Administration Guide

Amazon WorkSpaces Applications



Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon WorkSpaces Applications: Administration Guide

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

What Is Amazon WorkSpaces Applications? 1
Features 2
Key Concepts 4
How to Get Started 6
Accessing 7
Setting Up 8
Sign up for an AWS account 8
Get Started: Set Up With Sample Applications 9
Step 1: Set Up a Sample Stack, Choose an Image, and Configure a Fleet 9
Step 2: Provide Access to Users 17
Resources 18
Provide agents access to WorkSpaces 20
Why WorkSpaces for AI agents? 20
Use cases 20
How it works 22
Agents and LLMs 23
Best practices for computer-use agents 24
MCP Server 29
Overview 30
Prerequisites 30
Session behavior 36
Connecting 39
Agent identity 47
Connection modes 70
Session cleanup 71
Available tools 73
Computer use interfaces 80
Tool forwarding 81
Monitoring 85
Session evidence 86
Use AgentCore 101
Get Started 102
Write agent code with the Agent Toolkit for AWS 103
Prerequisites 104

Step 1: Create a stack or image builder with agent access enabled 104
Step 2: Generate a streaming URL 107
Step 3: Connect your agent to the MCP server 108
Step 4: Verify agent activity 110
Step 5: Change agent access settings or disable agent access 110
Step 6: Clean up resources 112
Networking and Access 114
Internet Access 114
VPC Requirements 116
VPC Setup Recommendations 116
Configure a VPC with Private Subnets and a NAT Gateway 119
Configure a VPC with a Public Subnet 128
Use the Default VPC and Public Subnet 132
Amazon S3 VPC Endpoints 134
Connections to Your VPC 136
Network Interfaces 136
Management Network Interface IP Address Range and Ports 136
Customer Network Interface Ports 137
User Connections to WorkSpaces Applications 138
Bandwidth Recommendations 139
IP Address and Port Requirements 140
Allowed Domains 141
Image Builders 157
Launch an Image Builder 158
Connect to an Image Builder 162
Console (Web Connection) 163
Streaming URL (Client or Web Connection) 164
Image Builder Actions 166
Instance Metadata for Image Builders 168
Install AMD Driver on Graphics Design Instances 169
Base Image and Managed Image Update Release Notes 170
Images 195
Default Settings and Application Launch Performance 196
Creating Default Application and Windows Settings 196
Optimizing the Launch Performance of Your Applications 197
Manage Agent Versions 198

Create an Image That Uses the Latest Version of the Agent ... 198
Create an Image That Uses a Specific Version of the Agent ... 199
Create an Image That Uses a Newer Version of the Agent ... 201
Unified graphics images ... 202
Agent Release Notes ... 203
Tutorial: Create a Custom Image by Using the Console ... 235
Step 1: Install Applications on the Image Builder ... 236
Step 2: Create an WorkSpaces Applications Application Catalog ... 237
Step 3: Create Default Application and Windows Settings ... 238
Step 4: Test Applications ... 240
Step 5: Optimize Applications ... 241
Step 6: Finish Creating Your Image ... 241
Step 7 (Optional): Tag and Copy an Image ... 243
Step 8: Clean Up ... 244
Administer Your Images ... 245
Delete a Private Image ... 246
Copy an Image That You Own to Another Region ... 246
Share an Image That You Own With Another Account ... 247
Stop Sharing an Image That You Own ... 249
Keep Your Image Up-to-Date ... 249
Windows Update and Antivirus Software ... 254
Programmatically Create a New Image ... 257
Manage License Included Applications ... 257
Import Image ... 261
Export Image ... 271
Create Your Image Programmatically ... 273
Default Application and Windows Settings ... 274
Launch Performance of Your Applications ... 275
Process Overview ... 276
Image Assistant CLI Operations ... 277
Create Your Linux-Based Images ... 287
Creating Default Application Settings for Your Users ... 288
Creating Default Environment Variables for Your Linux Users ... 291
Optimizing the Launch Performance of Your Linux Applications ... 292
Creating Session Scripts ... 293
Using the Image Assistant CLI Tool for Linux ... 294

Enabling and Disabling Webcam Support ... 294
Enabling and Disabling Heavy File Sync Mode for Home Folders ... 295
Tutorial: Create a Custom Linux-Based Image ... 296
Tutorial: Enable Japanese Support ... 304
Session Scripts to Manage Your Users' Streaming Experience ... 309
Run Scripts Before Streaming Sessions Begin ... 309
Run Scripts After Streaming Sessions End ... 314
Create and Specify Session Scripts ... 316
Session Scripts Configuration File ... 318
Using Windows PowerShell Files ... 321
Logging Session Script Output ... 321
Use Storage Connectors with Session Scripts ... 321
Enable Amazon S3 Bucket Storage for Session Script Logs ... 323
Use Session Scripts on Multi-Session Fleets ... 325
Applications Manager ... 328
App Blocks ... 328
Custom App Blocks ... 329
WorkSpaces Applications App Blocks ... 338
Unsupported Applications ... 350
App Block Builder ... 351
Create an App Block Builder ... 352
Connect to an App Block Builder ... 354
App Block Builder Actions ... 356
Applications ... 357
Store Application Icon, Setup Script, Session Script, and VHD in an S3 Bucket ... 359
Amazon S3 Bucket Permissions ... 359
Associate Applications to Elastic Fleets ... 362
Additional Resources ... 362
Fleets and Stacks ... 365
Session Context ... 365
Using Session Context to Pass Parameters to a Streaming Application ... 366
Fleet Types ... 368
Always-On and On-Demand Fleets ... 369
Elastic Fleets ... 370
Instance Families ... 370
Create a Fleet and Stack ... 376

Create a Fleet ... 376
Create a Stack ... 383
Provide Access to Users ... 389
Clean Up Resources ... 390
Customize a Fleet ... 391
Persist Environment Variables ... 392
Set Default File Associations ... 396
Disable Internet Explorer Enhanced Security Configuration ... 399
Change the Default Internet Explorer Home Page ... 401
User and Instance Metadata ... 405
Update a Fleet ... 407
Update a Fleet with a New Image ... 408
Manage Applications Associated to an Elastic Fleet ... 410
Fleet Auto Scaling ... 411
Scaling Concepts ... 412
Managing Fleet Scaling Using the Console ... 414
Managing Fleet Scaling Using the CLI ... 416
Additional Resources ... 424
Multi-Session Recommendations ... 425
Session Density ... 425
Image Management ... 428
Multi-session image optimization ... 428
Bring Your Own Windows desktop licenses ... 435
Overview ... 435
How it works ... 435
Prerequisites ... 436
Supported Windows Versions ... 436
Getting Started ... 437
Step 1: Enable your account for BYOL (First time only) ... 437
Step 2: Create a BYOL image in Amazon WorkSpaces ... 438
Step 3: Import BYOL Image into Amazon WorkSpaces Applications ... 438
Step 4: Customize your image (Optional) ... 439
Step 5: Create fleet and stream to users ... 439
Supported instance types ... 439
Region availability ... 440
Important considerations ... 441

Frequently asked questions ... 442
Application onboarding ... 443
Overview of the onboarding process ... 443
Part 1: Application compatibility validation ... 443
General compatibility ... 443
Multi-session compatibility ... 444
Native application mode compatibility ... 445
File handling and redirection ... 446
Printing ... 446
Local device interactions ... 447
Network performance ... 447
Multi-monitor support ... 447
Audio and video features ... 447
Validation environment ... 447
Part 2: Instance sizing and capacity planning ... 448
Choose an instance family ... 448
Size an instance for a single user ... 448
Size for multi-session fleets ... 449
Configure fleet scaling ... 450
Validate end-to-end with a pilot ... 450
Onboarding checklist ... 451
User Authentication ... 453
User Pools ... 453
User Pool End User Experience ... 454
Resetting a Forgotten Password ... 454
User Pool Administration ... 455
SAML 2.0 Integration ... 459
Example Authentication Workflow ... 460
Setting Up SAML ... 461
WorkSpaces Applications Integration with SAML 2.0 ... 477
Using Active Directory ... 481
Active Directory Domains ... 482
Before You Begin ... 484
Active Directory Domain Environment ... 484
Domain-Joined WorkSpaces Applications Streaming Instances ... 485
Group Policy Settings ... 486

Smart Card Authentication ... 487
Tutorial: Setting Up ... 487
Step 1: Create a Directory Config Object ... 488
Step 2: Create an Image by Using a Domain-Joined Image Builder ... 489
Step 3: Create a Domain-Joined Fleet ... 490
Step 4: Configure SAML 2.0 ... 491
Certificate-Based Authentication ... 492
Prerequisites ... 493
Enable Certificate-based Authentication ... 497
Manage Certificate-based Authentication ... 498
Enable Cross-account PCA Sharing ... 499
Administration ... 500
Granting Permissions to Create and Manage Active Directory Computer Objects ... 500
Finding the Organizational Unit Distinguished Name ... 502
Granting Local Administrator Rights on Image Builders ... 503
Updating the Service Account Used for Joining the Domain ... 505
Locking the Streaming Session When the User is Idle ... 505
Editing the Directory Configuration ... 507
Deleting a Directory Configuration ... 508
Configuring WorkSpaces Applications to Use Domain Trusts ... 508
Managing WorkSpaces Applications Computer Objects in Active Directory ... 509
Configuring Active Directory Sites and Services for WorkSpaces Applications ... 510
More Info ... 511
Add Custom Branding ... 512
Custom Branding Options ... 512
Adding Custom Branding ... 515
Custom Redirect URL and Feedback URL ... 517
Previewing Custom Branding Changes ... 517
Color Theme Palettes ... 517
Red ... 518
Light Blue ... 518
Blue ... 519
Pink ... 520
Embed Streaming Sessions ... 522
Prerequisites ... 522
Recommendations and Usage Considerations ... 523

Step 1: Specify a Host Domain 524
Step 2: Create a Streaming URL 525
Configuration Requirements for Using Custom Domains 525
Step 3: Download the Embedded Files 527
Step 4. Configure Your Website for Integration 528
Import the appstream-embed JavaScript File 528
Initialize and Configure the `AppStream.Embed` Interface Object 528
Examples for Hiding Items in the WorkSpaces Applications User Interface 530
Constants, Functions, and Events 531
Working with `HIDDEN_ELEMENTS` 532
Functions for the `AppStream.Embed` Object 533
Events for Embedded WorkSpaces Applications Streaming Sessions 534
Examples for Adding Event Listeners and Ending an Embedded WorkSpaces Applications Streaming Session 537
Administer Persistent Storage 539
Administer Home Folders 539
Files and Directories Associated with Compute-Intensive Applications 541
Enable Home Folders for Your WorkSpaces Applications Users 541
Administer Your Home Folders 542
Administer Google Drive 549
Enable Google Drive for Your WorkSpaces Applications Users 550
Disable Google Drive for Your WorkSpaces Applications Users 554
Controlling Google Drive Re-authentication 555
Administer OneDrive for Business 556
Enable OneDrive for Your WorkSpaces Applications Users 557
Disable OneDrive for Your WorkSpaces Applications Users 559
Controlling OneDrive Re-authentication 560
Administer Custom Shared Folders (SMB Network Drives) 561
Map Server Message Block (SMB) Network Drives 562
Enable Application Settings Persistence for Your Users 567
How Application Settings Persistence Works 567
Enabling Application Settings Persistence 570
Prerequisites for Enabling Application Settings Persistence 570
Best Practices for Application Settings Persistence 570
How to Enable Application Settings Persistence 571
Administer the VHDs for Your Users' Application Settings 571

Amazon S3 Bucket Storage ... 572
Reset a User's Application Settings ... 574
Enable Amazon S3 Object Versioning and Revert a User's Application Settings ... 574
Increase the Size of the Application Settings VHD ... 576
Enable Regional Settings for Your Users ... 579
Configure Default Regional Settings for Your Users ... 579
Specify a Default Time Zone ... 580
Specify a Default Display Language ... 584
Specify a Default System Locale ... 586
Specify a Default User Locale ... 588
Specify a Default Input Method ... 590
Configuring Chinese and Korean input methods on the image ... 592
Special Considerations for Application Settings Persistence ... 595
Special Considerations for Japanese Language Settings ... 595
Enable Your Users to Configure Their Regional Settings ... 596
Supported Locales ... 597
Enable Regional Settings for Your WorkSpaces Applications Users ... 598
Manage Application Entitlements ... 601
Attribute-Based Application Entitlements ... 601
Create Application Entitlements ... 601
SAML 2.0 Multi-Stack Application Catalog ... 604
Dynamic Application Framework ... 605
Example API Operations WorkFlow ... 606
Use the Dynamic Application Framework ... 607
Enable Dynamic App Providers ... 613
Test Dynamic App Providers ... 615
Additional Resources ... 616
Provide User Access ... 618
Supported Features ... 618
Provide Access Through a Web Browser ... 626
Requirements and Features ... 626
Configure a Connection Method for Your Users ... 630
Provide Access Through the Client ... 632
Requirements and Features ... 632
Install and Configure the Client ... 652
Tagging Your Resources ... 686

Tagging Basics ... 686
Tag Restrictions ... 687
Adding Tags during Resource Creation ... 688
Adding, Editing, and Deleting Tags ... 688
Using the API, SDK, or CLI ... 689
Monitoring and Reporting ... 692
Monitoring Resources ... 692
Viewing Fleet Usage Using the Console ... 693
Viewing instance and session metrics using the console ... 693
WorkSpaces Applications Metrics and Dimensions ... 695
Manage Multi-Session Fleet Instances ... 713
Create custom CloudWatch dashboards using CloudFormation templates ... 716
Monitor your WorkSpaces Applications health using the CloudWatch automatic dashboard ... 719
Usage Reports ... 720
Enable Usage Reports ... 721
Usage Reports Fields ... 724
Create Custom Reports ... 731
Logging WorkSpaces Applications API Calls ... 739
WorkSpaces Applications Information in CloudTrail ... 740
Example: WorkSpaces Applications Log File Entries ... 742
Security ... 745
Data Protection ... 746
Encryption at Rest ... 747
Encryption in Transit ... 747
Administrator Controls ... 748
Application Access ... 749
Identity and Access Management ... 751
Network Access ... 752
Access to WorkSpaces Applications Resources ... 752
Access to Application Auto Scaling ... 770
Access to the S3 Bucket for Home Folders and Application Settings Persistence ... 771
Access to Applications and Scripts on Streaming Instances ... 774
SELinux ... 779
Cookie-Based Authentication ... 779
Logging and Monitoring ... 780

Compliance Validation ... 782
Resilience ... 783
Infrastructure Security ... 783
Network Isolation ... 784
Isolation on Physical Hosts ... 785
Controlling Network Traffic ... 785
Interface VPC Endpoints ... 786
FIPS Endpoints ... 792
Security Groups ... 794
Update Management ... 795
Confused Deputy Prevention ... 796
Example: WorkSpaces Applications service role cross-service confused deputy prevention ... 797
Example: WorkSpaces Applications fleet machine role cross-service confused deputy prevention ... 798
Example: WorkSpaces Applications Elastic fleets session script Amazon S3 bucket policy cross-service confused deputy prevention ... 800
Example: WorkSpaces Applications Application Amazon S3 bucket policy cross-service confused deputy prevention ... 801
Security Best Practices ... 805
Securing Persistent Data ... 805
Endpoint Security and Antivirus ... 807
Network Exclusions ... 811
Securing an WorkSpaces Applications Session ... 812
Firewalls and Routing ... 813
Data Loss Prevention ... 813
Controlling egress traffic ... 814
Using AWS services ... 814
Troubleshooting ... 819
General Troubleshooting ... 819
SAML federation is not working. The user is not authorized to view WorkSpaces Applications applications. ... 820
After federating from an ADFS portal, my streaming session doesn't start. I am getting the error "Sorry connection went down". ... 820
I get an invalid redirect URI error. ... 820

My image builders and fleets never reach the running state. My DNS servers are in a Simple AD directory. 820
I've enabled application settings persistence for my users, but their persistent application settings aren't being saved or loaded. 821
I've enabled application settings persistence for my users, but for certain streaming applications, my users' passwords aren't persisting across sessions. 822
Google Chrome data is filling the VHD file that contains my users' persistent application settings. This is preventing their settings from persisting. How can I manage the Chrome profile? 822
I set up a custom domain for my embedded WorkSpaces Applications streaming sessions, but my WorkSpaces Applications streaming URLs aren't redirecting to my custom domain. 823
I launched an app on a smartcard-enabled WorkSpaces Applications fleet, and there are a limited number of certificates (or none) available to the app for authentication. 823
The Certification Propagation service isn't starting on my smartcard-enabled WorkSpaces Applications fleet. 825
I can't log in with my Active Directory username or password after SAML authentication. . 826
Troubleshooting Image Builders 826
I cannot connect to the internet from my image builder. 827
When I tried installing my application, I see an error that the operating system version is not supported. 827
I want to use a Windows PowerShell script to open my applications. 828
I want to make ClickOnce applications available to users. 828
When I connect to my image builder, I see a login screen asking me to enter Ctrl+Alt +Delete to log in. However, my local machine intercepts the key strokes. 829
When I switched between admin and test modes, I saw a request for a password. I don't know how to get a password. 829
I get an error when I add my installed application. 829
I accidentally quit a background service on the image builder and got disconnected. I am now unable to connect to that image builder. 830
The application fails to launch in test mode. 830
The application could not connect to a network resource in my VPC. 830
I customized my image builder desktop, but my changes are not available when connecting to a session after launching a fleet from the image I created. 830
My application is missing a command line parameter when launching. 830
I am unable to use my image with a fleet after installing an antivirus application. 831

My image creation failed. 831
The Image Assistant `create-image` operation failed with an error message that access to the PrewarmManifest.txt is denied 831
The image assistant reports that the SSM Agent version is below the recommended minimum 831
Troubleshooting Fleets 833
I tried to increase my fleet capacity, but the update isn't taking effect. 833
My applications won't work correctly unless I use the Internet Explorer defaults. How do I restore the Internet Explorer default settings? 835
I need to persist environment variables across my fleet instances. 836
I want to change the default Internet Explorer home page for my users. 836
When my users end a streaming session and then start a new one, they see a message that says no streaming resources are available. 836
Troubleshooting Active Directory 837
My image builders and fleet instances are stuck in the PENDING state. 837
My users aren't able to log in with the SAML application. 838
My fleet instances work for one user but don't cycle correctly. 838
My user Group Policy objects aren't being successfully applied. 838
My WorkSpaces Applications streaming instances aren't joining the Active Directory domain. 839
User login is taking a long time to complete on a domain-joined streaming session. 840
My users can't access a domain resource in a domain-joined streaming session, but they can access the resource from a domain-joined image builder. 840
My users receive the error “Certificate-Based Authentication not available” and are prompted to enter their domain password. Or users receive the error “Disconnected from session” when they are starting a session enabled with certificate-based authentication. . 841
I'm experiencing domain join failures after changing the Active Directory (AD) service account. 842
Troubleshooting WorkSpaces Applications User Issues 843
Enable advanced logging 843
My users' WorkSpaces Applications client installations fail, and they're getting a message stating that .NET Framework 4.6 is required. 845
My users' USB driver installations fail when they install the WorkSpaces Applications client, and now they can't use their USB devices with WorkSpaces Applications. 845
My WorkSpaces Applications client users are getting disconnected from their WorkSpaces Applications session after every 60 minutes. 846

My users can’t copy and paste between their local device and their streaming session. 847
Some keyboard shortcuts aren’t working for users during their streaming sessions. 849
My users' drawing tablets are not working with the streaming applications I deployed. 849
The Japanese language input method doesn't work for my users during their streaming sessions 850
My user sees an error about reaching the max number of streaming sessions when they try to launch an application from the application catalog. 853
My user sees a black screen or the desktop, and their application doesn’t launch on an Elastic fleet. No error appears. 853
Troubleshooting Persistent Storage Issues 853
My stack's home folders aren't working correctly. 854
My users can't access their home folder directory from one of our applications. 854
My users receive a “Device is not ready” error message when they access their home folder from one of our applications. 854
I removed or replaced a file in a user’s home folder in Amazon S3, but my users don’t see the changes in their home folder on the fleet instance during their streaming sessions. 855
Persistent storage isn't performing as expected. My users' files are taking longer than expected to save to persistent storage. 856
My users are getting errors that files are already in use when their files are not in use. 857
When a folder contains thousands of files, WorkSpaces Applications might take a long time to display the list of files. 857
Troubleshooting Notification Codes 858
Active Directory Internal Service 858
Active Directory Domain Join 858
Image Internal Service 862
Session Provisioning 862
Quotas 864
Guidance for WorkSpaces Applications Users 875
Access Methods and Clients 875
Web Browser Access 875
Client for Windows 887
Client for macOS 938
File Storage Options 952
Use Home Folders 953
Use Google Drive 954
Use OneDrive for Business 958

Use Custom Shared Network Folders 962
Regional Settings 962
Extension SDK Developer Guide 964
Prerequisites 964
Third-party vendor extensions 965
FAQs on supported protocols 966
What are the AWS JSON and Smithy RPC v2 CBOR protocols, and how do they differ? 966
How do I get started with the Smithy RPC v2 CBOR protocol? 966
What are the risks of enabling the Smithy RPC v2 CBOR protocol? 967
What if my solution does not support Smithy RPC v2 CBOR? 967
What AWS Regions support the Smithy RPC v2 CBOR protocol? 967
What latency improvements can I expect when upgrading to SDK versions that support Smithy RPC v2 CBOR? 967
Will the AWS JSON protocol be deprecated? 967
Where can I find more information about Smithy RPC v2 CBOR? 968
Document History 969
Earlier Updates 996

What Is Amazon WorkSpaces Applications?

Amazon WorkSpaces Applications is a fully managed application streaming service that provides users with instant access to their desktop applications from anywhere. WorkSpaces Applications manages the AWS resources required to host and run your applications, scales automatically, and provides access to your users on demand. WorkSpaces Applications provides users access to applications they need on the device of their choice, with a responsive, fluid user experience that is indistinguishable from natively installed applications.

With WorkSpaces Applications, you can easily add your existing desktop applications to AWS and enable your users to instantly stream them. Windows users can use either the WorkSpaces Applications client or an HTML5-capable web browser for application streaming. You can maintain a single version of each of your applications, which makes application management easier. Your users always access the latest versions of their applications. Your applications run on AWS compute resources, and data is never stored on users' devices, which means they always get a high performance, secure experience.

WorkSpaces Applications provides Model Context Protocol (MCP) tools that enable AI agents to interact with your desktop applications. Agents can type, click, take screenshots, and launch and switch between cataloged applications. You provide agents access to WorkSpaces in your stack and configure which capabilities agents can use. Agents authenticate using AWS IAM credentials. AWS CloudTrail logs agent activity, and you can monitor tool performance metrics in Amazon CloudWatch. For more information, see Get started providing agents with access to WorkSpaces Applications.

Unlike traditional on-premises solutions for desktop application streaming, WorkSpaces Applications offers pay-as-you-go pricing, with no upfront investment and no infrastructure to maintain. You can scale instantly and globally, ensuring that your users always have the best possible experience.

For more information, see WorkSpaces Applications.

Topics

- Features of Amazon WorkSpaces Applications
- Key Concepts of Amazon WorkSpaces Applications
- How to Get Started with Amazon WorkSpaces Applications

- Accessing Amazon WorkSpaces Applications

Features of Amazon WorkSpaces Applications

Using Amazon WorkSpaces Applications provides the following advantages:

Access desktop applications securely from any supported device

Your desktop applications can be accessed securely through an HTML5-capable web browser on Windows and Linux PCs, Macs, Chromebooks, iPads, and Android tablets. Or, for supported versions of Windows, the WorkSpaces Applications client can be used for application streaming.

Secure applications and data

Applications and data remain on AWS — only encrypted pixels are streamed to users. Applications run on an WorkSpaces Applications instance dedicated to each user so that compute resources are not shared. Applications can run inside your own virtual private cloud (VPC), and you can use Amazon VPC security features to control access. This enables you to isolate your applications and deliver them in a secure way.

Consistent, scalable performance

WorkSpaces Applications runs on AWS with access to compute capabilities not available on local devices, which means that your applications run with consistently high performance. You can instantly scale locally and globally, and ensure that your users always get a low-latency experience. Unlike on-premises solutions, you can quickly deploy your applications to the AWS region that is closest to your users, and start streaming with no incremental capital investment.

Integrate with your IT environment

Integrate with your existing AWS services and your on-premises environments. By running applications inside your VPCs, your users can access data and other resources that you have in AWS. This reduces the movement of data between AWS and your environment and provides a faster user experience.

Integrate with your existing Microsoft Active Directory environment. This enables you to use existing Active Directory governance, user experience, and security policies with your streaming applications.

Configure identity federation, which allows your users to access their applications using their corporate credentials. You can also allow authenticated access to your IT resources from applications running on WorkSpaces Applications.

Choose the fleet type that meets your needs

These are the types of fleets:

- Always-On — Streaming instances run all the time, even when no users are streaming applications and desktops. Streaming instances must be provisioned before a user is able to stream. The number of streaming instances provisioned is managed through auto scaling rules. For more information, see the section called “Fleet Auto Scaling”.

  When your users choose their application or desktop, they will start streaming instantly. You are charged the running instance fee for all streaming instances, even when no users are streaming.
- On-Demand — Streaming instances run only when users are streaming applications and desktops. Streaming instances not yet assigned to users are in a stopped state. Streaming instances must be provisioned before a user is able to stream. The number of streaming instances provisioned is managed through auto scaling rules. For more information, see the section called “Fleet Auto Scaling”.

  When your users choose their application or desktop, they will start streaming after a 1-2 minute wait. You are charged a lower stopped instance fee for streaming instances that are not yet assigned to users, and the running instance fee for streaming instances that are assigned to users.
- Elastic — The pool of streaming instances is managed by WorkSpaces Applications. When your users select their application or desktop to launch, they will start streaming after the app block has been downloaded and mounted to a streaming instance.

  You are charged the running instance fee for Elastic fleet streaming instances only for the duration of the streaming session, in seconds.

For more information, see Amazon WorkSpaces Applications Pricing.

Enable AI agents to interact with desktop applications

You enable agent access on your WorkSpaces Applications stack. When AI agents connect, they receive Model Context Protocol (MCP) tools to operate your desktop applications. Your agents can click buttons, enter text, and capture screenshots. This lets them complete multi-step workflows across your business applications. You configure which capabilities are available to agents, including computer input, computer vision, and screenshot storage. Agents authenticate by using standard AWS Identity and Access Management (IAM) credentials with SigV4 signing. No API keys or static secrets are required. Agent access works with Always-On and On-Demand

fleets. You can monitor agent activity through AWS CloudTrail. Operational metrics are available in CloudWatch.

Key Concepts of Amazon WorkSpaces Applications

To get the most out of WorkSpaces Applications, be familiar with the following concepts:

application

An *application* contains the information necessary to launch the application that you want to stream to your users. An application is associated with the resource that contains the files necessary to launch the application, such as an app block or image.

app block

An *app block* contains the application files that you want to stream to your users, and the details necessary to configure it.

app block builder

An *app block builder* is a virtual machine that you use to create an app block. You can launch and connect to an app block builder by using the WorkSpaces Applications console. After you connect to an appblock builder, you can install your application(s). App block builder packages your app contents, uploads it to an Amazon S3 bucket, and completes app block creation.

image builder

An *image builder* is a virtual machine that you use to create an image. You can launch and connect to an image builder by using the WorkSpaces Applications console. After you connect to an image builder, you can install, add, and test your applications, and then use the image builder to create an image. You can launch new image builders by using private images that you own.

image

An *image* contains applications that you can stream to your users, and default system and application settings to enable your users to get started with their applications quickly. AWS provides base images that you can use to create image builders to then create images that include your own applications. After you create an image, you can't change it. To add other applications, update existing applications, or change image settings, you must create a new image. You can copy your images to other AWS Regions or share them with other AWS accounts

in the same Region. your users, and default system and application settings to enable your users to get started with their applications quickly.

fleet

A *fleet* consists of fleet instances (also known as streaming instances) that run the applications and desktops that you specify.

stack

A *stack* consists of an associated fleet, user access policies, and storage configurations. You set up a stack to start streaming applications to users.

streaming instance

A *streaming instance* (also known as a fleet instance) is an EC2 instance that is made available to a single user for application streaming. After the user's session completes, the instance is terminated by EC2.

user pool

Use the *user pool* to manage users and their assigned stacks.

auto scaling rules

Auto scaling rules are schedule-based and usage-based policies that you can apply to an Always-On or On-Demand fleet to automatically manage the number of streaming instances available for users to stream from.

multi-session

A *multi-session* fleet allows you to provision more than one user session on a single fleet instance. The underlying infrastructure resources are shared across all of the user sessions.

> **Note**
>
> Multi-session is available only on Always-on and On-demand fleets powered by a Windows operating system. Multi-session is not available on Elastic fleets or the Linux operating system.
>
> Make sure you are using latest WorkSpaces Applications images for multi-session fleets. To keep your images are up-to-date, see the section called "Keep Your Image Up-to-Date". For details on supported images and WorkSpaces Applications agent versions for multi-session, see the section called "Base Image and Managed Image Update Release Notes".

agent session

An *agent session* is a WorkSpaces streaming session initiated by an AI agent. You create agent sessions by using the CreateStreamingURL API. The agent passes the streaming URL to the MCP server as a header on each request. You can monitor agent sessions through AWS CloudTrail, and you can view operational metrics in CloudWatch. You can optionally store screenshots captured during agent sessions in a customer-managed Amazon S3 bucket.

computer input

Computer input is a capability that allows agents to click buttons, enter text, and scroll on the desktop during a WorkSpaces streaming session. Computer input requires that you also enable computer vision.

computer vision

Computer vision is a capability that allows agents to see the desktop by taking screenshots during a WorkSpaces streaming session.

Model Context Protocol (MCP) server

The managed *MCP server* (`agentaccess-mcp`) provides AI agents with tools to click buttons, enter text, scroll, take screenshots of the desktop, and launch and switch between cataloged applications. Agents use these tools during a streaming session. Agents connect to the MCP server at a fixed endpoint using AWS IAM credentials with SigV4 signing. The MCP server is available when agent access is enabled on a stack.

How to Get Started with Amazon WorkSpaces Applications

If you are using WorkSpaces Applications for the first time, follow the Get Started with Amazon WorkSpaces Applications: Set Up With Sample Applications tutorial (available in the WorkSpaces Applications console). The Getting Started tutorial enables you to set up application streaming by using sample applications or your own applications. If you decide to start by using sample applications, you can always add your own applications later. For more information, see Amazon WorkSpaces Applications FAQs.

When you use the service for the first time, WorkSpaces Applications creates an AWS Identity and Access Management (IAM) role to create and manage WorkSpaces Applications resources on your behalf.

To run the Getting Started tutorial

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. Choose **Get Started**.
3. Select the option to learn more about WorkSpaces Applications resources.

Accessing Amazon WorkSpaces Applications

You can work with WorkSpaces Applications using any of the following interfaces:

AWS Management Console

The console is a browser-based interface to manage WorkSpaces Applications resources. For more information, see Get Started with Amazon WorkSpaces Applications: Set Up With Sample Applications.

AWS command line tools

AWS provides two sets of command line tools: the AWS Command Line Interface (AWS CLI) and the AWS Tools for Windows PowerShell. To use the AWS CLI to run WorkSpaces Applications commands, see Amazon WorkSpaces Applications Command Line Reference.

AWS SDKs

You can access WorkSpaces Applications from a variety of programming languages. The SDKs automatically take care of tasks such as the following:

- Setting up an WorkSpaces Applications stack or fleet
- Getting an application streaming URL to your stack
- Describing your resources

For more information, see Tools for Amazon Web Services.

Setting Up for Amazon WorkSpaces Applications

Complete the following tasks to get set up for Amazon WorkSpaces Applications.

Sign up for an AWS account

To get started with AWS, you need an AWS account. There are two ways to sign up for AWS. If you want complete control over account configuration and where your content is stored, or if you have regulated workloads such as HIPAA or FedRAMP, use Sign up for AWS (advanced). If you are a new user and want to use a login you already own to sign in to AWS and create an AWS environment with preconfigured defaults in select AWS Regions, use Sign up for AWS.

Get Started with Amazon WorkSpaces Applications: Set Up With Sample Applications

To stream your applications, Amazon WorkSpaces Applications requires an environment that includes a fleet that is associated with a stack, and at least one application image. This tutorial describes how to configure a sample WorkSpaces Applications environment for application streaming and give users access to that stream.

> **Note**
>
> For additional guidance in learning how to get started with WorkSpaces Applications, see the Amazon WorkSpaces Applications Getting Started Guide. This guide describes how to install and configure two applications, perform foundational administrative tasks using the WorkSpaces Applications console, and provision an Amazon Virtual Private Cloud by using a provided AWS CloudFormation template.

Tasks

- Step 1: Set Up a Sample Stack, Choose an Image, and Configure a Fleet
- Step 2: Provide Access to Users
- Resources

Step 1: Set Up a Sample Stack, Choose an Image, and Configure a Fleet

Before you can stream your applications, you need to set up a stack, choose an image that has applications installed, and configure a fleet. In this step, you use a template to help simplify these tasks.

To set up a sample stack, choose an image, and configure a fleet

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. Choose **Get Started** if you are new to the console, or **Quick Links** from the left navigation menu. Choose **Set up with sample apps**.

3. For **Step 1: Stack Details**, keep the default stack name or enter your own. Optionally, you can do the following:

   - **Display name** — Enter a name to display for the stack (maximum of 100 characters).
   - **Description**— Keep the default description or enter your own (maximum of 256 characters).
   - **Redirect URL** — Specify a URL to which users are redirected after their streaming sessions end.
   - **Feedback URL** — Specify a URL to which users are redirected after they click the **Send Feedback** link to submit feedback about their application streaming experience. If you do not specify a URL, this link is not displayed.
   - **Streaming Protocol Preference** — Specify the streaming protocol you'd like your stack to prefer, UDP or TCP. UDP is currently only supported in the Windows native client. For more information, see System Requirements and Feature Support (WorkSpaces Applications Client).
   - **Tags** — Choose **Add Tag**, and type the key and value for the tag. To add more tags, repeat this step as needed. For more information, see Tagging Your Amazon WorkSpaces Applications Resources.
   - **VPC Endpoints (Advanced)** — You can create a private link, which is an interface VPC endpoint (interface endpoint), in your virtual private cloud (VPC). To start creating the interface endpoint, select **Create VPC Endpoint**. Selecting this link opens the VPC console. To finish creating the endpoint, follow steps 3 through 6 in *To create an interface endpoint*, in Tutorial: Creating and Streaming from Interface VPC Endpoints.

     After you create the interface endpoint, you can use it to keep streaming traffic within your VPC.
   - **Embed WorkSpaces Applications (Optional)** — To embed an WorkSpaces Applications streaming session in a webpage, specify the domain to host the embedded streaming session. Embedded streaming sessions are only supported over HTTPS [TCP port 443].

     > **Note**
     >
     > You must meet prerequisites and perform additional steps to configure embedded WorkSpaces Applications streaming sessions. For more information, see Embed Amazon WorkSpaces Applications Streaming Sessions.

4. Choose **Next**.

5. For **Step 2: Choose Image**, a sample image is already selected. The image contains pre-installed open-source applications for evaluation purposes. Choose **Next**.
6. For **Step 3: Configure Fleet**, we recommend that you keep any default values that are provided. You can change most of these values after fleet creation.

   - **Choose instance type** — Choose the instance type that matches the performance requirements of your applications. All streaming instances in your fleet launch with the instance type that you select. For more information, see WorkSpaces Applications Instance Families.
   - **Fleet type** — Choose the fleet type that suits your use case. The fleet type determines its immediate availability and how you pay for it.
   - **Maximum session duration in minutes** — Choose the maximum amount of time that a streaming session can remain active. If users are still connected to a streaming instance five minutes before this limit is reached, they are prompted to save any open documents before being disconnected. After this time elapses, the instance is terminated and replaced by a new instance.
   - **Disconnect timeout in minutes** — Choose the amount of time that a streaming session should remain active after users disconnect. If users try to reconnect to the streaming instance after a disconnection or network interruption within this time interval, they are connected to the previous session. Otherwise, they are connected to a new session with a new instance. If you associate a stack with a fleet for which a redirect URL is specified, after users' streaming sessions end, the users are redirected to that URL.

     If a user ends the session by choosing **End Session** on the streaming session toolbar, the disconnect timeout does not apply. Instead, the user is prompted to save any open documents, and then immediately disconnected from the streaming instance.
   - **Idle disconnect timeout in minutes** — Choose the amount of time that users can be idle (inactive) before they are disconnected from their streaming session and the **Disconnect timeout in minutes** time interval begins. Users are notified before they are disconnected due to inactivity. If they try to reconnect to the streaming session before the time interval specified in **Disconnect timeout in minutes** has elapsed, they are connected to their previous session. Otherwise, they are connected to a new session with a new streaming instance. Setting this value to 0 disables it. When this value is disabled, users are not disconnected due to inactivity.

> **Note**
>
> Users are considered idle when they stop providing keyboard or mouse input during their streaming session. File uploads and downloads, audio in, audio out, and pixels changing do not qualify as user activity. If users continue to be idle after the time interval in **Idle disconnect timeout in minutes** elapses, they are disconnected.

- **Multiple user sessions** — Choose this option if you want to provision multiple user sessions on a single instance. By default, every unique user session is served by an instance (single-session).

> **Note**
>
> Multi-session is available only on Always-on and On-demand fleets powered by a Windows operating system. Multi-session is not available on Elastic fleets or the Linux operating system.
> Make sure you are using latest WorkSpaces Applications images for multi-session fleets. To keep your images are up-to-date, see the section called "Keep Your Image Up-to-Date". For details on supported images and WorkSpaces Applications agent versions for multi-session, see the section called "Base Image and Managed Image Update Release Notes".

- **Maximum sessions per instance** — Maximum number of user sessions on an instance. You must choose this value based on your end users' application performance needs. You can also adjust the maximum sessions per instance for a fleet after it is provisioned. In that case, the existing user sessions and instances will not be impacted, but the fleet will become consistent with the new value of maximum sessions per instance. The value must be between 2 and 50. Before setting this value for your fleet, see the section called "Multi-Session Recommendations".
- **Minimum capacity** — Choose a minimum number of instances for your fleet based on the minimum number of expected concurrent users. Every unique user session is served by an instance. For example, to have your stack support 100 concurrent users during low demand, specify a minimum capacity of 100. This ensures that 100 instances are running even if there are fewer than 100 users.
- **Maximum capacity** — Choose a maximum number of instances for your fleet based on the maximum number of expected concurrent users. Every unique user session is served by an

instance. For example, to have your stack support 500 concurrent users during high demand, specify a maximum capacity of 500. This ensures that up to 500 instances can be created on demand.

- **Minimum user sessions for fleet** — Choose a minimum number of user sessions for your fleet based on the minimum number of expected concurrent users. For example, to have your stack support 100 concurrent users during low demand, specify a minimum capacity of 100. This ensures that 100 user sessions are available even if there are fewer than 100 users.
- **Maximum user sessions for fleet** — Choose a maximum number of user sessions for your fleet based on the maximum number of expected concurrent users. For example, to have your stack support 500 concurrent users during high demand, specify a maximum capacity of 500. This ensures that up to 500 user sessions can be provisioned on demand.

> **Note**
>
> For a single-session fleet, one instance will be launched of every user session. However, for multi-session, the number of running instances depends on the maximum sessions per instance. You must provide the capacity in terms of user sessions. The service will decide how many instances are required based on your fleet type (multi-session or single-session) and maximum sessions per instance.

- **Scaling details** — Specify the scaling policies that WorkSpaces Applications uses to increase and decrease the capacity of your fleet. Note that the size of your fleet is limited by the minimum and maximum capacity that you specified. For more information, see Fleet Auto Scaling for Amazon WorkSpaces Applications.
- **IAM role (Advanced)** — When you apply an IAM role from your account to an WorkSpaces Applications fleet instance, you can make AWS API requests from the fleet instance without manually managing AWS credentials. To apply an IAM role, do either of the following:
  - To use an existing IAM role in your Amazon Web Services account, choose the role that you want to use from the **IAM role** list. The role must must be accessible from the fleet instance. For more information, see Configuring an Existing IAM Role to Use With WorkSpaces Applications Streaming Instances.
  - To create a new IAM role, choose **Create new IAM role** and follow the steps in How to Create an IAM Role to Use With WorkSpaces Applications Streaming Instances.

7. Choose **Next**.

8. For **Step 4: Configure Network**, a default VPC is provided. This VPC includes a default public subnet in each Availability Zone and an internet gateway that is attached to your VPC. The VPC also includes a default security group. To use the default VPC configuration, do the following:

   - Keep the **Default Internet Access** check box selected.

     When **Default Internet Access** is enabled, a maximum of 100 fleet instances is supported. If your deployment must support more than 100 concurrent users, use the NAT gateway configuration instead.

   - For **VPC**, keep the default VPC selected for your AWS Region.

     The default VPC name uses the following format: `vpc-`*`vpc-id`* `(No_default_value_Name)`.

   - For **Subnet 1** and **Subnet 2**, keep the default public subnets selected.

     The default subnet names use the following format: `subnet-`*`subnet-id`* `|` *`(IPv4 CIDR block)`* `| Default in` *`availability-zone`*.

   - For **Security groups**, keep the default security group selected.

     The default security group name uses the following format: `sg-`*`security-group-id`*`-default`.

9. For **Step 5: Enable Storage**, choose one or more of the following, then choose **Next**.

   > **Note**
   >
   > For multi-session fleets, Google Drive and OneDrive require an image that uses an WorkSpaces Applications agent released on or after June 29, 2026. Images using Update an Image by Using Managed WorkSpaces Applications Image Updates released on or after June 29, 2026 are also supported.

   - **Enable Home Folders** — By default, this setting is enabled. Keep the default setting. For information about requirements for enabling home folders, see Enable Home Folders for Your WorkSpaces Applications Users.
   - **Enable Google Drive** — Optionally, you can enable users to link their Google Drive for G Suite account to WorkSpaces Applications. You can enable Google Drive for accounts in G

Suite domains only, not for personal Gmail accounts. For information about requirements for enabling Google Drive, see Enable Google Drive for Your WorkSpaces Applications Users.

- **Enable OneDrive** — Optionally, you can enable users to link their OneDrive for Business account to WorkSpaces Applications. You can enable OneDrive for accounts in OneDrive domains only, not for personal accounts. For information about requirements for enabling OneDrive, see Enable OneDrive for Your WorkSpaces Applications Users.

10. For **Step 6: User Settings**, configure the following settings. When you're done, choose **Review**:

Clipboard, file transfer, print to local device, and authentication permissions options:

> **Note**
>
> **Print to local device** and **Smart card sign in for Active Directory** are currently not available for multi-session fleets.

- **Clipboard** — By default, users can copy and paste data between their local device and streaming applications. You can limit Clipboard options so that users can paste data to their remote streaming session only or copy data to their local device only. You can also disable Clipboard options entirely. Users can still copy and paste between applications in their streaming session. You can choose **Copy to local device character limit** or **Paste to remote session character limit** or both to limit the amount of data that users can copy or paste when using the clipboard, either in or out of their WorkSpaces Applications streaming session. The value can be between 1 and 20,971,520 (20 MB), and defaults to the maximum value when unspecified.
- **File transfer** — By default, users can upload and download files between their local device and streaming session. You can limit file transfer options so that users can upload files to their streaming session only or download files to their local device only. You can also disable file transfer entirely.

> **Important**
>
> If your users require WorkSpaces Applications file system redirection to access local drives and folders during their streaming sessions, you must enable both file upload and download. To use file system redirection, your users must have WorkSpaces

Applications client version 1.0.480 or later installed. For more information, see Enable File System Redirection for Your WorkSpaces Applications Users.

- **Print to local device** — By default, users can print to their local device from within a streaming application. When they choose **Print** in the application, they can download a .pdf file that they can print to a local printer. You can disable this option to prevent users from printing to a local device.
- **Password sign in for Active Directory** — Users can enter their Active Directory domain password to sign in to an WorkSpaces Applications streaming instance that is joined to an Active Directory domain.

  You can also enable **Smart card sign in for Active Directory**. At least one authentication method must be enabled.
- **Smart card sign in for Active Directory** — Users can use a smart card reader and smart card connected to their local computer to sign in to an WorkSpaces Applications streaming instance that is joined to an Active Directory domain.

  You can also enable **Password sign in for Active Directory**. At least one authentication method must be enabled.

> **Note**
>
> **Clipboard, file transfer, and print to local device settings** — These settings control only whether users can use WorkSpaces Applications data transfer features. If your image provides access to a browser, network printer, or other remote resource, your users might be able to transfer data to or from their streaming session in other ways.
> **Authentication settings** — These settings control only the authentication method that can be used for Windows sign in to an WorkSpaces Applications streaming instance (fleet or image builder). They do not control the authentication method that can be used for in-session authentication, after a user signs in to a streaming instance. For information about configuration requirements for using smart cards for Windows sign in and in-session authentication, see Smart Cards.

Application settings persistence options:

- **Enable Application Settings Persistence** — Users' application customizations and Windows settings are automatically saved after each streaming session and applied during the next session. These settings are saved to an Amazon Simple Storage Service (Amazon S3) bucket in your account, within the AWS Region in which application settings persistence is enabled.
- **Settings Group** — The settings group determines which saved application settings are used for a streaming session from this stack. If the same settings group is applied to another stack, both stacks use the same application settings. By default, the settings group value is the name of the stack.

> **Note**
>
> For information about requirements for enabling and administering application settings persistence, see Enable Application Settings Persistence for Your WorkSpaces Applications Users.

11. For **Step 7: Review**, confirm the details for the stack. To change the configuration for any section, choose **Edit** and make the needed changes. After you finish reviewing the configuration details, choose **Create**.
12. In the pricing acknowledgement dialog box, select the acknowledgement check box, and choose **Create**.
13. After the service sets up resources, the **Stacks** page appears. The status of your new stack appears as **Active** when it is ready to use.

Step 2: Provide Access to Users

After you create a stack with an associated fleet, you can provide access to users through the WorkSpaces Applications user pool, SAML 2.0 [single sign-on (SSO)], or the WorkSpaces Applications API. For more information, see User Pool Administration in Amazon WorkSpaces Applications and Amazon WorkSpaces Applications Integration with SAML 2.0.

> **Note**
>
> Users in the WorkSpaces Applications user pool can't be assigned to stacks with fleets that are joined to an Active Directory domain.

For this getting started exercise, you can use the WorkSpaces Applications user pool. This access method enables you to create and manage users by using a permanent login portal URL. To quickly test application streaming without setting up users, complete the following steps to create a temporary URL, also known as a streaming URL.

To provide access to users with a temporary URL

1. In the navigation pane, choose **Fleets**.
2. In the list of fleets, choose the fleet that is associated with the stack for which you want to create a streaming URL. Verify that the status of the fleet is **Running**.
3. In the navigation pane, choose **Stacks**. Select the stack, and then choose **Actions**, **Create Streaming URL**.
4. For **User id**, type the user ID. Choose an expiration time, which determines how long the generated URL is valid.
5. To view the user ID and URL, choose **Get URL**.
6. To copy the link to the clipboard, choose **Copy Link**.

After you provide your users with access to WorkSpaces Applications, they can start WorkSpaces Applications streaming sessions.

If you plan to use SAML 2.0 [single sign-on (SSO)] or the WorkSpaces Applications API to provide access to your users, you can make the WorkSpaces Applications client available to them. The WorkSpaces Applications client is a native application that is designed for users who require additional functionality during their WorkSpaces Applications streaming sessions. For more information, see Provide Access Through the WorkSpaces Applications Client.

Resources

For more information, see the following:

- Learn how to use the WorkSpaces Applications image builder to add your own applications and create images that you can stream to your users. For more information, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.
- Provide persistent storage for your session users by using WorkSpaces Applications home folders, Google Drive, and OneDrive. For more information, see Enable and Administer Persistent Storage for Your WorkSpaces Applications Users.

- Integrate your WorkSpaces Applications streaming resources with your Microsoft Active Directory environment. For more information, see Using Active Directory with WorkSpaces Applications.
- Control who has access to your WorkSpaces Applications streaming instances. For more information, see Identity and Access Management for Amazon WorkSpaces Applications, Amazon WorkSpaces Applications User Pools and Amazon WorkSpaces Applications Integration with SAML 2.0.
- Monitor your WorkSpaces Applications resources by using Amazon CloudWatch. For more information, see WorkSpaces Applications Metrics and Dimensions.
- Troubleshoot your WorkSpaces Applications streaming experience. For more information, see Troubleshooting.

Provide agents access to WorkSpaces

Amazon WorkSpaces Applications enables AI agents to connect to remote compute environments and operate desktop applications using a managed Model Context Protocol (MCP) server. You can also let human observers watch agent sessions in real time and stop them if needed. Use the topics in this section to set up, configure, and build agents that use Amazon WorkSpaces Applications.

Topics

- Why WorkSpaces for AI agents?
- Access the MCP server
- Get started providing agents with access to WorkSpaces Applications

Why WorkSpaces for AI agents?

WorkSpaces for AI agents gives AI agents a mouse, keyboard, and screen to operate desktop, legacy, and browser-based applications just like a person does, so agents can take on tasks that previously required a human operator. WorkSpaces for AI agents is the agent access feature of Amazon WorkSpaces Applications. This section helps you decide whether it fits your needs. It covers what kinds of work agents can do, how it works, which agents and models you can use, what to expect for performance and cost, and best practices for building and evaluating computer-use agents.

Topics

- Example agent use cases
- How it works
- Supported agents, frameworks, and LLMs
- Best practices for building and evaluating computer-use agents

Example agent use cases

With Amazon WorkSpaces Applications agent access, your agent works in a governed desktop session, operating in desktop applications, mainframe terminals, and browser sessions to complete tasks. This gives it access to the software your team already relies on. You bring your own agent, and agent access provides the session it runs in.

The agent runs under your IAM permissions, and you can review its activity through AWS CloudTrail logs and captured screenshots.

> **Note**
>
> You can monitor agent performance through CloudWatch metrics.

Back-office operations

Agents work through queues and batches to make judgment-heavy updates in the applications you rely on. Each case differs, so the agent reads the work, interprets it, and decides what to do.

- Data extraction from files and entry into ERP systems
- Insurance claims processing
- Transaction reconciliation and processing

Contact-initiated workflows

A customer or employee request kicks off an agent that completes the follow-up work across backend systems right away. You can invoke the agent through an API, so it starts the moment a request arrives, and your team can focus on higher-value work while the agent handles the busy work and reports back when it is done.

- CRM system contact data change
- Clinician update to an EHR system
- Payment terms change for billing

Software testing and IT operations

Agents run test cycles and recurring operational chores that teams otherwise do by hand or skip, reporting results back for review. This work runs against changing software, screens, and conditions. When a screen differs or an unexpected dialog appears, the agent adapts.

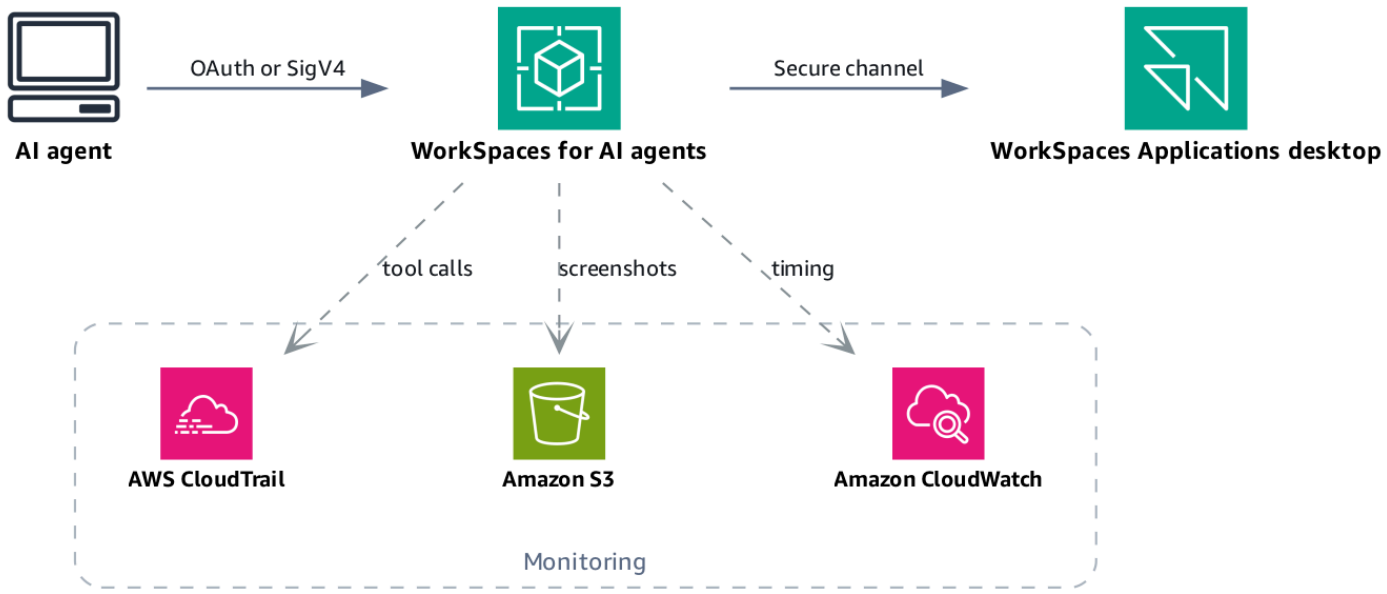
- Install, configure, and patch applications on an image builder before you create an image.
- Create default application settings on an image builder.

- Test applications and perform application-packaging checks as a standard (non-administrator) user on an image builder.
- Perform configuration management across desktops and images.
- Run nightly regression subsets against legacy applications.
- Run periodic access audits that capture screenshot evidence.
- Provide remote IT support by working directly on a user's desktop.
- Complete forms in multiple languages.

How it works

Amazon WorkSpaces Applications connects an AI agent to a desktop and gives it the tools to operate that desktop the way a person would. The AI agent is one that you bring, running in the framework or managed runtime of your choice. It connects to WorkSpaces for AI agents, the managed MCP server, using either OAuth or AWS IAM (SigV4) credentials. Through the MCP server, the agent then operates a WorkSpaces Applications desktop, which you prepare in advance, over a secure channel.

As the agent works, the MCP server records its activity to three monitoring services so that you can review every action: it logs each tool call to AWS CloudTrail, stores the screenshots it captures in Amazon S3, and reports the time each tool call takes to CloudWatch. The following diagram shows this flow.



The agent: you bring your own agent. This can be a coding agent such as Claude Code, Kiro, or Codex, or a custom-built agent that you develop with the framework or managed runtime of your choice. The agent connects to the managed MCP server with either OAuth or AWS IAM (SigV4) credentials and drives the work throughout the session.

The MCP interface: the agent reaches the desktop through a managed MCP server. WorkSpaces Applications gives the agent a set of tools that let it move and click the mouse, type with the keyboard, and view the screen through screenshots. You can forward additional tools through the same MCP server, so the agent can also read PDFs, access files on the desktop, or run shell commands alongside the built-in tools.

The desktop: the agent works on a WorkSpaces Applications desktop, a secure environment that runs from an image with the applications it needs. An agent can do two kinds of work. On a stack's streaming sessions, it works on a desktop that you prepared in advance. On an image builder, it can prepare the desktop itself: install and configure applications, then test them as a standard user before you create the image.

Monitoring: you have visibility into everything the agent does. Every tool call is logged in AWS CloudTrail, every screenshot can be stored in Amazon S3, and the time each tool call takes is available in CloudWatch.

Supported agents, frameworks, and LLMs

You choose which agent to use with agent access. You can run a coding agent on your local computer, or build your own agent that runs in the cloud. In both cases, the agent connects to the managed Model Context Protocol (MCP) server and uses it to operate the desktop.

AWS provides a sample agent to help you get started. For the sample code, see Sample code for WorkSpaces Applications agent access on GitHub.

Agents and frameworks

You can use a ready-to-run coding agent, build your own agent with a framework, or build and run an agent with a managed service. The following options work with agent access.

Option	Examples
Coding agents that run on your local computer	Claude Code, Kiro, Cursor, Codex

Option	Examples
Frameworks for building your own agent	LangChain, CrewAI, Strands Agents SDK
Managed services for building and running agents	Amazon Bedrock AgentCore, LangGraph Platform

Large language models

You bring your own large language model. Agent access does not require a specific one and works with any computer-use model. A computer-use model can operate a computer through vision and actions.

Computer-use models are available from multiple providers, including Anthropic, Google Gemini, and OpenAI. Because you select the model, you can match it to your task, cost, and quality needs.

With WorkSpaces Applications, you can present the desktop tools in the computer use tool schema that your model expects, which increases accuracy and lowers the cost for your agent to perform actions. You select the interface for each request by sending a request header to the MCP server. For more information, see Computer use interfaces.

Best practices for building and evaluating computer-use agents

A computer-use agent works in a loop: it looks at the screen, acts, and looks again. The following practices help you build agents that complete tasks reliably, in fewer steps, and at lower cost. They apply across workflows and model families. Results vary with your workflow, model, prompts, and environment, so measure each change on your own tasks.

Design the action loop

Most of an agent's effort goes into reading the screen and typing, not into moving between windows. Reduce reading errors and the number of steps before you tune navigation.

- **Give the agent text instead of pixels**. Reading exact values, such as names, email addresses, and account numbers, from a screenshot is a common source of errors. When the data has a text copy, provide it through a tool. For values in tables or columns, also have the agent view the image, because a text layer can mix up columns.

- **Put the instruction before the screenshot**. When the agent must find and click something, place the instruction ahead of the image in the request, so that the agent knows what to look for.
- **Start from a known layout**. Maximize each application with a keyboard shortcut before the agent works in it. A consistent layout makes repeated flows more predictable and reduces scrolling.
- **Prefer the keyboard**. Use keyboard shortcuts to scroll, maximize windows, and open menus, especially for small or hard-to-hit controls. Shortcuts don't depend on locating a control on the screen.
- **Confirm each action**. An application that is busy can ignore a keystroke. Wait briefly after each action, keep a few recent screenshots in view, and have the agent repeat an action that had no effect.
- **Use Tab only where it's safe**. The Tab key works well across plain text fields. Click into fields where the application reacts when a field loses focus, such as fields with autofill, dropdowns, or date pickers.
- **Switch, don't relaunch**. To move between running applications, use the taskbar or the application switcher. Relaunching an application adds steps.
- **Locate controls on the current screen**. Coordinates copied from an earlier screen go stale when a window moves. Use fixed coordinates only in a maximized, fixed layout.
- **Detect actions that change nothing**. When two actions in a row leave the screen unchanged, have the agent change approach. Limit repeated scrolling, so that a run doesn't spend its step budget on one action.
- **Batch actions with care**. A batch shows the agent only the screen after its last action, so end a batch on the one screen that the agent needs to read. Measure whether batching saves time on your full workflow before you rely on it. Return batch errors as text only, because some models reject an error result that includes an image.

Adapt to the model family

Model families operate a desktop differently. Build the agent loop once, and adapt the skill and the settings for each family. A skill tuned for one family usually needs corrections before it works with another.

- **Send a fresh screenshot after each change**. Some models keep acting on an old screenshot without signaling a problem. Attach a new screenshot after every action that changes the screen.

- **Handle coordinates explicitly**. Some models report pixel coordinates, and others report coordinates on a normalized grid that you must convert to pixels. State the screen size and the valid coordinate range in the prompt, and check that each click is within the screen.
- **Size screenshots to the model**. Match the screenshot resolution to the model's image limit. If you downscale a screenshot, scale the model's coordinates back up before you click.
- **Clear fields before typing**. On some models, typing appends to a field instead of replacing its value. Focus the field, select all, delete, type once, and read the value back.
- **Control parallel tool calls**. If a model sends several actions in one message regardless of the prompt, turn off parallel tool calls in the client. Run batched actions in order and stop at the first failure.
- **Keep reasoning effort low by default**. Simple tasks often succeed at low effort, and higher effort adds cost without improving accuracy. Raise the effort for branching or multi-step tasks.
- **Match the prompt length to the model**. Stronger models can work from a short goal. Add rules only where a model needs them, and shorten the prompt as models improve.
- **Keep reasoning between turns**. If you call a reasoning model through a framework or a compatibility API, confirm that the model's reasoning is carried from one turn to the next. A loop that drops it can stall.
- **Compare on cost per finished task**. Models differ in how many tokens they use per step, so price per token alone doesn't predict cost. Compare models with the tool surface that you plan to use.

Write skills and prompts

The skill, the instructions that you give the agent for a workflow, has a large effect on accuracy and cost. A poorly designed generic prompt can perform worse than no prompt at all.

- **Include a few working rules**. For example: remember what you read, verify before you move on, change approach after two failed attempts, and wait for the application to respond.
- **Give facts, not procedures**. Current models need less step-by-step guidance. Keep the skill short, and include the facts that the agent can't discover, such as file paths, field formats, and application quirks.
- **Define high-value fields**. State the exact format of fields such as identifiers and account numbers, and have the agent check each value against that format before typing.
- **Make verification required and targeted**. Agents can enter wrong values with full confidence, so don't leave verification to the agent's judgment. Check each value once, right after typing it,

and read back the few high-value fields before saving. A full re-check of every field adds steps without improving accuracy.

- **Correct errors in place**. When a value is wrong, have the agent fix that field instead of clearing and retyping the whole form.
- **Hand off cleanly to sub-agents**. Give a sub-agent complete instructions and ask for a report at the end, instead of checking its work as it goes.
- **Improve skills from your own runs**. Review the screenshots, actions, and reasoning of each run, change one rule at a time, and keep a change only if every run still succeeds. Divide the workflow into milestones, so that you can see where failed runs stop.

Manage context and cost

Cost depends mainly on the context that the model carries and on the number of steps in a run. Each call resends the history, so every extra step costs more than the last. The most effective way to lower cost is to choose the model and skill that finish the task in the fewest steps.

- **Reset context between tasks**. Start each independent task from a clean context and a known screen state.
- **Keep only recent screenshots**. Keep a few recent screenshots, such as the last three, and replace older images with a short text note. Keeping only one screenshot can increase cost, because the agent scrolls back to see what it lost.
- **Tell the agent what you remove**. Explain that older screenshots are removed on purpose, and have the agent record each value in its reply while the value is on the screen.

Prompt caching can lower cost substantially, but only when you build the payload carefully:

- Only add to the end of the conversation. Caching matches the payload from the start, so an edit to an earlier turn invalidates the cache.
- Remove old screenshots in batches, not on every turn, so that the cached prefix stays the same between removals.
- Don't rebuild the conversation on each turn. A payload that changes every turn can cost more with caching than without it.
- Confirm that caching engages. Models have a minimum prefix size for caching, and a smaller prefix isn't cached and doesn't return an error.

Measure results

- **Run enough trials**. Results vary between identical runs. Run many trials, report the spread, and use a statistical test, such as a paired test, before you conclude that one version is better.
- **Calculate cost from every token type**. Include fresh input, output, cache writes, and cache reads. Providers report cache tokens differently, so check how each one counts them. Report tokens alongside cost.
- **Report cost with the score**. A change can keep the same success rate although it increases cost. Report the success rate, cost, and steps together, and keep the cost of passing and failing runs separate.
- **Set the step limit with care**. Set the limit above what the workflow needs, and report how many runs reached it. A run that ran out of steps needs a different fix from a run that produced a wrong result. Also set time and cost limits, and support cancellation.
- **Change one thing at a time**. Hold everything else fixed, so that you can see the effect of each change.
- **Sort failures by cause**. Separate failures to read the screen from failures to act. A skill can improve how the agent acts, but not how well the model reads. Before you credit a rule with an improvement, check that the agent followed it.
- **Grade the outcome, not the agent's report**. Check the saved record and the final screen with an automatic evaluator that requires every field to match, and include a negative control. An agent can report a task as complete when it isn't.
- **Check your test setup**. Errors in the evaluator or the environment can look like agent failures. Log every prompt, screenshot, and action, so that you can trace failed runs. Don't publish results that you can't reproduce.

Oversight and safety

- **Keep a person in the loop**. Require confirmation before consequential or hard-to-reverse actions, such as payments, deletions, or submissions. Treat typing sensitive data into a form as sending it. Check for required confirmation before each action in a batch, and locate the target again after a person approves, because the screen might have changed.
- **Pass the agent's questions to a person**. Some models stop and ask for confirmation on their own, for example when the screen suggests a production system. Make sure that your loop can show the question to a person.

- **Limit what the agent can reach**. Run the agent in an isolated environment with minimal permissions. Keep sign-in credentials out of the agent's reach where possible, and limit internet access to an allowlist of domains.
- **Treat on-screen content as untrusted**. Text on a page or in a document can't grant permissions or change the task. Don't let instructions on the screen redirect the agent.
- **Keep the agent off the open web**. Put the information that the agent needs in the skill, instead of having it search the web, where it can encounter CAPTCHAs and untrusted content.

Related resources

Model providers publish computer-use guidance for their models. Their current documentation describes model-specific details, such as supported screen sizes, coordinate systems, and tools.

- Computer use tool in the Anthropic documentation
- Computer use in the OpenAI documentation
- Computer use in the Gemini API documentation

Access the MCP server

The WorkSpaces Applications MCP server is a fully managed service that provides AI agents with Model Context Protocol (MCP) tools to operate desktop applications during streaming sessions. Agents can click buttons, enter text, scroll, and take screenshots of the desktop.

Topics

- Overview
- Prerequisites
- Session behavior
- Connecting to the MCP server
- Agent identity
- Connection modes
- Session cleanup
- Available tools
- Computer use interfaces

- MCP tool forwarding
- Monitoring
- Session evidence
- Use Amazon Bedrock AgentCore

Overview

This section explains how to connect to and operate the WorkSpaces Applications Model Context Protocol (MCP) server, the endpoint your agent uses to work in a desktop session. For a conceptual overview of WorkSpaces for AI agents and agent access, see Why WorkSpaces for AI agents? and How it works.

An agent works with the MCP server in four stages:

1. **Authenticate**. The agent authenticates to the MCP server before it can connect.
2. **Initiate the connection**. The agent opens a session and retrieves the list of available tools.
3. **Make tool calls**. The agent calls tools to operate the desktop, such as moving the mouse, typing, or taking a screenshot.
4. **End the session**. The agent ends the session when the task is complete.

Throughout these stages, WorkSpaces Applications handles the connection between the MCP calls and the underlying desktop. The topics in this section describe each stage. To authenticate and open a connection, see Connecting to the MCP server. For the tools your agent can call, see Available tools. To end a session, see Session cleanup.

An agent can connect to a streaming session for a stack, or to an image builder. To connect to either, see Connecting to the MCP server.

You enable agent access when you create a stack (using the WorkSpaces Applications console, AWS CLI, or API) or an image builder (using the AWS CLI or API). Before you connect an agent, review Prerequisites.

Prerequisites

This topic covers what you need in place before an agent can connect: the AWS account and Amazon WorkSpaces Applications fleet that host the desktop session, and the permissions and

software your agent uses to reach it. Setting these up first lets your agent authenticate and operate the desktop without interruption.

Supported operating systems

Agent access supports the same operating systems as WorkSpaces Applications. How you get an image depends on the operating system.

Operating system	Image source	More information
Windows Server	AWS-provided base images, or bring your own image	WorkSpaces Applications Base Image and Managed Image Update Release Notes Import Image
Linux	AWS-provided base images	WorkSpaces Applications Base Image and Managed Image Update Release Notes
Windows 11	Bring Your Own License (BYOL): you supply the image	Bring Your Own Windows Desktop Licenses for Amazon WorkSpaces Applications

Windows 11 requires additional BYOL setup. For requirements, see Bring Your Own Windows Desktop Licenses for Amazon WorkSpaces Applications.

For image builders, agent access supports Windows Server and Windows 11 (BYOL). Linux image builders aren't supported.

Account, fleet, and image builder requirements

To set up agent access on a stack, you need the following:

- An AWS account with permissions to create and manage WorkSpaces Applications stacks.
- An active WorkSpaces Applications fleet (Always-On or On-Demand). If you haven't created a fleet, see Get Started with Amazon WorkSpaces Applications: Set Up With Sample Applications.

- A stack associated with a fleet. Agents cannot connect to a stack without an associated fleet.
- If you plan to enable screenshot storage, an Amazon S3 bucket. The bucket policy must grant the AppStream service principal access to list buckets. The agent that connects to the Model Context Protocol (MCP) server must have the `s3:PutObject` permission on the bucket.
- A fleet running the latest version of the WorkSpaces Applications Agent. This is the streaming software on the image, not the AI agent you build. For more information about updating it, see Keep Your Amazon WorkSpaces Applications Image Up-to-Date. For information about managing agent versions, see Managing agent versions.

To use agent access on an image builder, you need the following:

- An AWS account with permissions to create and manage WorkSpaces Applications images.
- An image builder in the `RUNNING` state. An agent can't connect to an image builder that isn't running.
- If you plan to enable screenshot storage, an Amazon S3 bucket. The bucket policy must grant the AppStream service principal access to list buckets. The agent that connects to the Model Context Protocol (MCP) server must have the `s3:PutObject` permission on the bucket.
- An image builder whose image uses a WorkSpaces Applications agent released on or after October 25, 2026. If you create the image builder from a WorkSpaces Applications base image, use a base image released on or after November 4, 2026. If the image uses an earlier version, the request to enable agent access fails with an `IncompatibleImageException` error.

  To check an image, run the following command. In the response, `AppstreamAgentVersion` shows the agent version and `PublicBaseImageReleasedDate` shows when the base image was released.

  ```
  aws appstream describe-images --names your-image-name
  ```

  If the image is older, use a newer base image, or update the agent on your image. For more information, see WorkSpaces Applications Agent Release Notes, WorkSpaces Applications Base Image and Managed Image Update Release Notes, and Update an Image by Using Managed WorkSpaces Applications Image Updates.

Agent and permissions requirements

To build and connect your agent, you need the following:

- AWS credentials (environment variables, AWS profile, EC2 instance role, or Lambda execution role) with the following IAM permissions. The agent uses these credentials directly to sign requests to the MCP server:

```
{
    "Sid": "MCP",
    "Effect": "Allow",
    "Action": ["agentaccess-mcp:*"],
    "Resource": "*"
},
{
    "Sid": "AppStream",
    "Effect": "Allow",
    "Action": ["appstream:CreateStreamingURL", "appstream:DescribeFleets",
 "appstream:CreateImageBuilderStreamingURL"],
    "Resource": "*"
}
```

The `agentaccess-mcp:*` wildcard includes the following actions:

- `agentaccess-mcp:InvokeMcp`: Initialize sessions and discover tools
- `agentaccess-mcp:GetScreenshot`: Capture screen state
- `agentaccess-mcp:LeftClick`: Perform left mouse click
- `agentaccess-mcp:DoubleClick`: Perform double click
- `agentaccess-mcp:TripleClick`: Perform triple click
- `agentaccess-mcp:RightClick`: Perform right mouse click
- `agentaccess-mcp:MiddleClick`: Perform middle mouse click
- `agentaccess-mcp:TypeText`: Type text string
- `agentaccess-mcp:KeyPress`: Press key or key combination
- `agentaccess-mcp:HoldKey`: Hold key for duration
- `agentaccess-mcp:Scroll`: Scroll at coordinates
- `agentaccess-mcp:MovePointer`: Move pointer to coordinates
- `agentaccess-mcp:LeftClickDrag`: Perform left click drag
- `agentaccess-mcp:LeftMouseDown`: Press and hold left mouse button
- `agentaccess-mcp:LeftMouseUp`: Release left mouse button

- `agentaccess-mcp:CheckConnectionStatus`: Check connection status of a streaming session
- `agentaccess-mcp:CallForwardedTool`: Invoke a forwarded tool on a remote instance
- `agentaccess-mcp:LaunchApplication`: Launch a cataloged application in the current session
- `agentaccess-mcp:ToggleAppSwitcher`: Open or close the app-switcher overlay
- `agentaccess-mcp:SwitchUser`: On an image builder, disconnect the current built-in user and log in as another built-in user (`ADMINISTRATOR`, `TEMPLATE_USER`, or `TEST_USER`)
- `agentaccess-mcp:GetCurrentUser`: Return the built-in user that's logged in to an image builder

  The `agentaccess-mcp:*` wildcard also allows `SwitchUser` and `GetCurrentUser` on every image builder, as any of the three built-in users. To limit an agent on image builders, name the actions and use the condition keys, as in the example policy later in this topic.

  The agent needs the same `agentaccess-mcp` permissions whether it connects to an image builder with a streaming URL or with the image builder name.

- An MCP-compatible agent framework that can send Streamable HTTP requests to the MCP endpoint, authorized in one of two ways:
  - **Signature Version 4 (SigV4) signing** – The agent signs each request to the MCP server with its AWS credentials.
  - **OAuth bearer token** – The agent uses OAuth 2.1 authorization through AWS Sign-In to retrieve a short-lived access token, and passes it as a bearer token on each request.

  For more information about both authorization methods, see Request authorization. The Strands Agents SDK provides native MCP client support, or you can use any framework with the `mcp-proxy-for-aws` transport.
- Python 3.10 or later on the machine where your agent runs. Your agent can run on any operating system, independent of the streamed desktop.

To enable or disable agent access, or change its settings on an existing image builder, you need the `appstream:UpdateImageBuilder` permission. This permission is separate from the agent's permissions.

On an image builder, the agent logs in as one of three built-in Windows users: `ADMINISTRATOR`, `TEMPLATE_USER`, or `TEST_USER`. For more information about these users, see Built-in users on image builders. To control which image builders an agent can use and which built-in users it can log in as, use the following IAM condition keys:

- `agentaccess-mcp:ImageBuilderArn`: The ARN of the image builder.
- `agentaccess-mcp:LocalUser`: The built-in user that the agent logs in as.

The `LocalUser` condition key is present only on requests that log a user in. Use the `StringEqualsIfExists` operator with it, as in the example policy later in this topic. With `StringEquals`, other requests, such as `tools/list`, are denied.

We recommend that you restrict the agent to the least-privileged built-in user that its tasks need. For example, an agent that only tests applications can log in as `TEST_USER`.

The following example policy allows an agent to use one image builder and to log in only as `TEST_USER`. The second statement lists the tools that the agent can use on that image builder. Add the other actions that your agent needs.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ConnectAndSwitchOnlyAsTestUser",
            "Effect": "Allow",
            "Action": ["agentaccess-mcp:InvokeMcp", "agentaccess-mcp:SwitchUser"],
            "Resource": "*",
            "Condition": {
                "ArnEquals": { "agentaccess-mcp:ImageBuilderArn":
 "arn:aws:appstream:region:account-id:image-builder/image-builder-name" },
                "StringEqualsIfExists": { "agentaccess-mcp:LocalUser": "TEST_USER" }
            }
        },
        {
            "Sid": "DriveTheDesktopOnThatBuilder",
            "Effect": "Allow",
            "Action": [
                "agentaccess-mcp:GetCurrentUser",
                "agentaccess-mcp:CheckConnectionStatus",
                "agentaccess-mcp:GetScreenshot",
                "agentaccess-mcp:LeftClick",
```

```
                "agentaccess-mcp:TypeText",
                "agentaccess-mcp:KeyPress"
            ],
            "Resource": "*",
            "Condition": {
                "ArnEquals": { "agentaccess-mcp:ImageBuilderArn":
 "arn:aws:appstream:region:account-id:image-builder/image-builder-name" }
            }
        }
    ]
}
```

Limitations

WorkSpaces Applications agent access configurations have the following limitations:

- VPC endpoints are not supported.
- Multi-session fleets are not supported.
- Elastic fleets are not supported.
- A single session supports only a single AI agent connection.

Agent access on image builders has the following additional limitations:

- Domain-joined image builders are not supported.
- App block builders, which you use with elastic fleets, aren't supported.
- An image builder supports only one AI agent connection at a time.
- Changes to agent access settings on the image builder cannot be made while the image builder is running. Stop the image builder first to make an update, and it takes effect on the next session.
- Either the AI agent or a human can provide input on an image builder, but not both at the same time. To switch between agent and human input, stop the image builder, change the setting to enable or disable agent access, and restart the image builder.

Session behavior

When you enable agent access on a stack or an image builder, you configure how the agent's desktop session behaves, including how the agent sees and controls the desktop, where

screenshots and files are stored, the screen layout, and whether people can watch and stop the agent. You set these options on the stack or on the image builder. Most options work the same way on both. Home folders and application settings persistence apply only to stacks.

If you make changes to agent access configurations on a stack, changes apply to new streaming sessions. On an image builder, you change the configuration while the image builder is stopped, and they apply at the next session. For steps, see Step 5: Change agent access settings or disable agent access.

Computer input and vision

Agent access gives you two options for computer use. You set them up on the stack or on the image builder. For other agent options, see MCP tool forwarding and User control mode.

- **Computer input**: during a streaming session, agents click buttons, enter text, and scroll on the desktop.
- **Computer vision**: during a streaming session, agents take screenshots to see the desktop.

Use these rules to choose what to turn on:

- For your agent to use computer input, computer vision must be turned on. The agent must see the desktop to use the keyboard and mouse.
- You can leave computer input and vision off if you only use tool forwarding. For details, see MCP tool forwarding.

Screenshot storage

You can optionally configure screenshot storage. When you enable it, the Model Context Protocol (MCP) server stores the screenshots captured during agent sessions in an Amazon S3 bucket. The MCP server uses the connecting agent's credentials to upload screenshots to the bucket, so the agent must have the `s3:PutObject` permission on the bucket. To encrypt screenshots in your bucket, see Configuring default encryption for an Amazon S3 bucket.

Desktop screen layout

You configure the screen resolution and image format for the agent's desktop session:

- **Screen resolution**: the display resolution for the agent's desktop session. The supported resolution is 1280x720.

- **Screen image type**: the image format for agent screen captures. You can choose PNG or JPEG.

Home folder storage

This option applies only to stacks.

You can enable home folders so that agent files are saved to an Amazon S3 bucket in your AWS account. The Amazon WorkSpaces Applications fleet associated with the stack must allow access to Amazon S3 through the internet or an Amazon S3 VPC endpoint. For more information, see Enable and administer home folders for your users.

Application settings persistence

This option applies only to stacks.

You can optionally enable application settings persistence. This setting is available on Windows and is not supported on Linux. When enabled, your agent's application customizations and Windows settings are saved after each streaming session and applied during the next session. These settings are saved to an Amazon S3 bucket in your AWS account. The Amazon WorkSpaces Applications fleet associated with the stack must allow access to Amazon S3 through the internet or an Amazon S3 VPC endpoint. For more information, see Enable application settings persistence for your users.

User control mode

You can enable user control mode to allow users to observe and interact with agent sessions in real time. When user control mode is enabled, observers connect to the session through their browser using a streaming URL. They see a real-time view of the desktop as the agent interacts with it.

In `VIEW_STOP` mode, a stop button appears at the top of the observer's screen. The observer can use this button to stop the agent if needed. On a stack, after an agent is stopped, it must start a new session to resume work. On an image builder, agent access returns only after you stop and start the image builder.

On an image builder with agent access enabled, a person who opens the image builder in a browser can only watch the agent. They can't log in or interact with the desktop. To work on the image builder yourself, stop it, disable agent access, and start it again. To disable agent access, use the `AttributesToDelete` parameter with the `AGENT_ACCESS_CONFIG` value.

Set the `UserControlMode` attribute in the `AgentAccessConfig` to one of the following values:

- `VIEW_ONLY`: users can view and observe agent actions as they happen.
- `VIEW_STOP`: users can view agent actions and stop the agent if needed.
- `DISABLED`: users cannot view or stop the agent session.

Connecting to the MCP server

This topic explains how to connect an agent to the Amazon WorkSpaces Applications Model Context Protocol (MCP) server. Connecting involves negotiating an MCP protocol version with the server and authorizing each request. This topic covers the endpoint the agent calls, how it authorizes its requests, and how it identifies a streaming session.

Agents connect to the MCP server at the following endpoint:

```
https://agentaccess-mcp.region.api.aws/mcp
```

Replace *region* with your Region code. For the supported Regions and their endpoints, see WorkSpaces Applications MCP endpoints.

The MCP server is hosted in the AWS cloud and uses Streamable HTTP as its transport protocol. You don't need to install or maintain any server components.

> **Note**
>
> Before an agent can connect, the fleet associated with your stack must be running and have available capacity. To connect to an image builder, the image builder must be in the `RUNNING` state.

Supported MCP protocol versions

The MCP server supports the following versions of the Model Context Protocol. Your MCP client negotiates the version with the server when it connects.

Version	Type
2025-03-26	Stateful

Version	Type
2025-06-18	Stateful
2025-11-25	Stateful
2026-07-28	Stateless

Version 2026-07-28 is the stateless MCP protocol. With the stateless protocol, the agent ends the session with the `close_session` tool. For more information about the tool, see Available tools. With the earlier MCP protocol versions, you end the session by closing the MCP client in your agent code, which sends an HTTP `DELETE` request. For more information, see Session cleanup.

Request authorization

You can authorize requests in one of two ways: SigV4 signing or OAuth 2.1 authorization through AWS Sign-In.

SigV4 signing

To use Signature Version 4 (SigV4), sign each request using IAM credentials with the service name `agentaccess-mcp`. The following Python example shows the general connection pattern using `mcp-proxy-for-aws`:

```
from mcp_proxy_for_aws import aws_iam_streamablehttp_client

async with aws_iam_streamablehttp_client(
    endpoint="https://agentaccess-mcp.region.api.aws/mcp",
    aws_service="agentaccess-mcp",
    aws_region="region",
    headers={
        # Fleet-type-specific headers (see the following subsections)
    },
    metadata={
        # Fleet-type-specific metadata (see the following subsections)
    },
) as (read, write, _):
    # Use read/write streams with your MCP client
    ...
```

For other languages, write your own SigV4 signing logic for outgoing MCP requests or use a library that supports SigV4 signing. For more information about `mcp-proxy-for-aws`, see mcp-proxy-for-aws on GitHub.

OAuth 2.1 authorization

As an alternative to signing requests with IAM credentials, the MCP server supports OAuth 2.1 authorization through AWS Sign-In, the same way as the AWS MCP Server. AWS Sign-In acts as the authorization server. It issues short-lived, scoped access tokens, and the MCP server accepts these tokens as bearer tokens. OAuth uses your existing IAM identities, permissions, and organizational controls. When you authorize an agent, it can access the MCP server on your behalf without gaining extra permissions. AWS still checks every request against your IAM policies.

OAuth-compatible MCP clients, such as Kiro, Claude Code, Cursor, Codex, and Gemini, can connect directly to the MCP server endpoint without local proxy software. AWS Sign-In supports two authorization models:

- **Interactive authorization** – This model is for agents that can open a browser. The client registers through dynamic client registration and redirects you to an AWS Sign-In page to authenticate and approve access. AWS Sign-In then issues and refreshes the access tokens.
- **Non-interactive authorization** – This model is for agents that already have AWS credentials and run without a browser (headless authorization). The agent exchanges its SigV4-signed credentials for a short-lived access token by using the OAuth client credentials grant.

OAuth authorization for the MCP server works the same way as it does for the AWS MCP Server, using the MCP server endpoint (`https://agentaccess-mcp.region.api.aws/mcp`). For setup, required IAM permissions, and governance and monitoring options, see OAuth 2.1 authentication for AWS MCP Server in the *Agent Toolkit for AWS User Guide* and Sign-In with OAuth 2.0 in the *AWS Sign-In User Guide*.

Session identification

How you identify the streaming session depends on your fleet type, or on whether the agent connects to an image builder.

Note

At any given time, only one agent can connect to a unique session. Named users, specified through the `UserId` parameter, can have only one active session per fleet at a time. To run multiple agents concurrently, each agent must connect to its own unique session.

Connecting with non-domain-joined fleets

For non-domain-joined fleets, you can connect in one of two ways. Provide the stack and user ID as headers, and the MCP server creates the streaming session for you. This suits statically configured MCP clients, because they can hold a fixed pair of headers. However, they cannot generate a streaming URL or replace one when it expires. Alternatively, generate a streaming URL yourself and pass it as a header. Use this approach when you need control over session creation, such as setting the validity period.

To connect with the stack and user ID, pass the following headers on every request:

- `x-amzn-agentaccess-workspaces-applications-stack`. The stack name or ARN.
- `x-amzn-agentaccess-workspaces-applications-user-id`. The user ID for the session.

Note

The `x-amzn-agentaccess-workspaces-applications-user-id` header is optional. If you omit it, the MCP server derives the user identity from the authenticated session. It uses the AWS `SourceIdentity` when one is present, and the caller's IAM principal otherwise. `SourceIdentity` preserves a user's original corporate identity when they assume an IAM role. This holds even when roles are chained through an external identity provider (IdP).

```
from mcp_proxy_for_aws import aws_iam_streamablehttp_client

# Connect to the MCP server with the stack and user ID
async with aws_iam_streamablehttp_client(
    endpoint="https://agentaccess-mcp.region.api.aws/mcp",
    aws_service="agentaccess-mcp",
    aws_region="region",
```

```
    headers={
        "x-amzn-agentaccess-workspaces-applications-stack": "stack-name-or-arn",
        "x-amzn-agentaccess-workspaces-applications-user-id": "user-id",
    },
) as (read, write, _):
    ...
```

To connect with a streaming URL instead, generate one using the `CreateStreamingURL` API and pass it as the `x-amzn-agentaccess-streaming-session-url` header on every request. No agent-specific parameters are required. The stack's agent access configuration determines the agent behavior.

```
import boto3
from mcp_proxy_for_aws import aws_iam_streamablehttp_client

# Generate streaming URL
appstream = boto3.client("appstream", region_name="region")
response = appstream.create_streaming_url(
    StackName="stack-name",
    FleetName="fleet-name",
    UserId="user-id",
)
streaming_url = response["StreamingURL"]

# Connect to MCP server
async with aws_iam_streamablehttp_client(
    endpoint="https://agentaccess-mcp.region.api.aws/mcp",
    aws_service="agentaccess-mcp",
    aws_region="region",
    headers={
        "x-amzn-agentaccess-streaming-session-url": streaming_url,
    },
) as (read, write, _):
    ...
```

For more information about the CreateStreamingURL API, see CreateStreamingURL in the *Amazon AppStream 2.0 API Reference*.

Connecting with domain-joined fleets

When agents access domain-joined streaming instances, the connection must be federated through a Security Assertion Markup Language (SAML) provider. This requirement applies to both traditional and agent sessions. For agent sessions, Certificate-Based Authentication is required.

> **Note**
>
> Domain join is supported on Windows only.

Because domain-joined streaming instances require access through SAML, your MCP client must provide a signed SAML assertion instead of a streaming URL. Encoded SAML assertions exceed HTTP header size limits. To stay within those limits, use the `metadata` field in `mcp-proxy-for-aws`:

```
from mcp_proxy_for_aws import aws_iam_streamablehttp_client

# saml_response: your signed, base64-encoded SAML assertion
# stack_arn: the ARN of the AppStream stack for the AD user

async with aws_iam_streamablehttp_client(
    endpoint="https://agentaccess-mcp.region.api.aws/mcp",
    aws_service="agentaccess-mcp",
    aws_region="region",
    metadata={
        "saml_response": saml_response,
        "stack_arn": stack_arn,
    },
) as (read, write, _):
    ...
```

> **Note**
>
> The `metadata` parameter was added in `mcp-proxy-for-aws` version 1.6.1. Earlier versions cannot inject the `_meta` field without additional development. To upgrade, run `pip install -U mcp-proxy-for-aws`.

For more information about setting up SAML federation with WorkSpaces Applications, see Setting up SAML in the *Amazon WorkSpaces Applications Administration Guide*. For more information and a complete working example, see the sample-code-for-workspaces-agent-access repository in AWS Samples on GitHub.

Connecting with image builders

For image builders, you can connect in the same two ways as for non-domain-joined fleets. Provide the image builder name as a header, and the MCP server creates the streaming session for you. Alternatively, generate a streaming URL yourself and pass it as a header.

With either method, you must also specify which built-in user the agent logs in as. Unlike the user ID header for fleets, this header is required and has no default.

To connect with the image builder name, pass the following headers on every request:

- `x-amzn-agentaccess-workspaces-applications-image-builder-name`. The image builder name.
- `x-amzn-agentaccess-local-user`. The built-in user that the agent logs in as: `ADMINISTRATOR`, `TEMPLATE_USER`, or `TEST_USER`. This header is required. For more information about these users, see Built-in users on image builders. To switch users after the agent connects, see User tools.

```
from mcp_proxy_for_aws import aws_iam_streamablehttp_client
# Connect to the MCP server with the image builder name and built-in user
async with aws_iam_streamablehttp_client(
    endpoint="https://agentaccess-mcp.region.api.aws/mcp",
    aws_service="agentaccess-mcp",
    aws_region="region",
    headers={
        "x-amzn-agentaccess-workspaces-applications-image-builder-name": "image-
builder-name",
        "x-amzn-agentaccess-local-user": "ADMINISTRATOR",
    },
    timeout=120,
) as (read, write, _):
    ...
```

To connect with a streaming URL instead, generate one using the `CreateImageBuilderStreamingURL` API. Pass it as the `x-amzn-agentaccess-streaming-`

`session-url` header, together with the `x-amzn-agentaccess-local-user` header, on every request.

```
import boto3
from mcp_proxy_for_aws import aws_iam_streamablehttp_client
# Generate streaming URL
appstream = boto3.client("appstream", region_name="region")
response = appstream.create_image_builder_streaming_url(
    Name="image-builder-name",
)
streaming_url = response["StreamingURL"]
# Connect to MCP server
async with aws_iam_streamablehttp_client(
    endpoint="https://agentaccess-mcp.region.api.aws/mcp",
    aws_service="agentaccess-mcp",
    aws_region="region",
    headers={
        "x-amzn-agentaccess-streaming-session-url": streaming_url,
        "x-amzn-agentaccess-local-user": "TEST_USER",
    },
    timeout=120,
) as (read, write, _):
    ...
```

Tip

The first connection to an image builder takes longer than a connection to a fleet. Set your MCP client timeout to at least 120 seconds.

Note

Only one agent can connect to an image builder at a time.

For more information about the `CreateImageBuilderStreamingURL` API, see CreateImageBuilderStreamingURL in the *Amazon AppStream 2.0 API Reference*.

Compatible frameworks

You can connect from any MCP-compatible framework that supports Streamable HTTP and SigV4 signing. For the agents, frameworks, and tools that work with agent access, see Supported agents, frameworks, and LLMs.

Agent identity

When an AI agent works in a streaming session on behalf of a person, you need to know two things: which agent acted, and which user it acted for. With *agent identity*, you use the identity provider (IdP) that you already use for your workforce, such as Microsoft Entra ID or Okta, to answer both questions. Each agent has its own identity in your IdP, unique to that agent and separate from the identity of any user. Your IdP signs the identities of the agent and of the user, so both are cryptographically verified rather than reported by the agent itself. AWS carries both identities through to the Model Context Protocol (MCP) server that controls the streaming session, and to the records in AWS CloudTrail.

With agent identity, you can do the following:

- **Attribute every action**. CloudTrail records which agent acted and which user it acted for, so that you can answer audit questions without having to correlate logs from different systems.
- **Reuse your IdP controls**. The user signs in through your IdP, so the multi-factor authentication, conditional access, and consent policies that you already enforce also apply when an agent acts for the user.
- **Control access in IAM policies**. Both identities arrive as IAM session tags, so you can write policies that require them, or that allow only specific agents to reach specific stacks.
- **Avoid long-lived credentials**. The agent receives temporary AWS credentials for each session. It doesn't store a password or an AWS access key.

How agent identity works

Agent identity builds on two standard mechanisms: the OAuth 2.0 on-behalf-of (OBO) token exchange that your IdP supports, and IAM web identity federation. It works as follows:

1. The user signs in to your IdP through the agent, and the agent receives a token for that user.
2. The agent exchanges the user's token at your IdP for a new access token that represents the agent acting on behalf of the user. This is the OBO exchange. Your IdP adds claims to the new token that identify both the agent and the user.

3. The agent calls the AWS Security Token Service (AWS STS) `AssumeRoleWithWebIdentity` operation with the access token. AWS STS validates the token against the IAM OpenID Connect (OIDC) identity provider that you registered, and returns temporary credentials for an IAM role. The claims in the token become session tags and the source identity of the role session.
4. The agent signs its requests to the MCP server with those credentials. The MCP server uses the source identity as the user of the streaming session.
5. CloudTrail records the role assumption with both identities. It also records the user on the MCP `initialize` request and, if you log data events, on every tool call that follows. Tool calls are data events, which CloudTrail doesn't log by default. For more information, see Monitoring and Logging data events in the *CloudTrail User Guide*.

The agent doesn't supply the identity values itself. They come only from the token that your IdP signed, which is why you can rely on them for authorization and audit.

Identity attributes

Configure your IdP to add the following claims to the access token that it issues in the OBO exchange. Session tag keys are case sensitive, so use the same spelling in your IdP claims and in your IAM policies.

Identity claims for agent identity

Session tag or claim	Value	Use
`AgentAccessAgentId`	The agent's immutable object ID in your IdP	Identifies the agent in policies and audit records
`AgentAccessAgentName` (optional)	The agent's display name	Display only. Display names can change, so don't use this tag in policies.
`AgentAccessUserId`	The user's short name	Identifies the user that the agent acts for
Source identity	The user's short name	Becomes the user of the streaming session, and appears on every later CloudTrail event

Session tags and the source identity do different jobs, so you need both:

- **Session tags** are available to IAM policy evaluation for the whole role session, and CloudTrail records them on the `AssumeRoleWithWebIdentity` event. Later events don't repeat them.
- **The source identity** can't be changed after the session is created, and CloudTrail records it on every event that the session produces, including tool calls when you log data events.

For more information, see Pass session tags in AWS STS and Monitor and control actions taken with assumed roles in the *IAM User Guide*.

Choose a setup for your fleet type

The setup depends on whether your fleet is joined to an Active Directory domain:

- **Non-domain-joined fleets**. The agent runs one OBO exchange and calls `AssumeRoleWithWebIdentity`. The MCP server creates the streaming session for the user named by the source identity. For more information, see Agent identity for non-domain-joined fleets.
- **Domain-joined fleets**. Windows must log the user on to the domain, so the agent also obtains a SAML assertion from your IdP. The assertion carries the user's Active Directory attributes, and certificate-based authentication uses them to log on without a password. For more information, see Agent identity for domain-joined fleets.

The examples in these topics use Microsoft Entra ID. You can use any IdP that supports OIDC, an OBO or token exchange grant, and custom claims in access tokens, such as Okta. For domain-joined fleets, your IdP must also be able to return a SAML assertion from the OBO or token exchange.

Security considerations

- **Protect your IdP configuration**. AWS trusts the claims that your IdP signs, and can't check that your claim mapping is correct. Treat the claim mapping and the IAM OIDC identity provider as security-critical configuration. Restrict who can change them, and review changes with the same care as a role trust policy.
- **Scope each role to specific stacks**. Grant the WorkSpaces Applications permissions in the role only on the stacks that the agent needs. This is where you authorize the agent, as distinct from the user it acts for. Use a separate role for each agent that needs different access.

- **Require both identities**. Use `Null` conditions in the role trust policy so that a token that is missing either identity can't be used to assume the role. Put the conditions in a statement that doesn't include `sts:SetSourceIdentity`. Otherwise, the role assumption fails. For an example, see Step 2: Configure IAM.
- **Use a signed client assertion**. Have the agent authenticate to your IdP with a private key JWT instead of a client secret. If you keep the private key in AWS KMS, the agent never holds the key itself.
- **Revoke access at any time**. Disabling the user in your IdP, or changing the role trust policy, prevents new sessions. It doesn't affect temporary credentials that AWS STS already issued, which stay valid until they expire at the end of the role session. To stop an agent that already has credentials, remove permissions from the role or revoke the role's active sessions. After the credentials expire, the agent needs a new OBO exchange to continue. For more information, see Revoke IAM role temporary security credentials in the *IAM User Guide*.

Agent identity for non-domain-joined fleets

For a non-domain-joined fleet, the agent runs one on-behalf-of (OBO) exchange at your identity provider (IdP), exchanges the resulting token for temporary AWS credentials, and connects to the Model Context Protocol (MCP) server with those credentials. The MCP server creates the streaming session for the user named in the source identity.

To set up agent identity, complete the following steps:

1. Step 1: Configure your identity provider
2. Step 2: Configure IAM
3. Step 3: Connect the agent
4. Step 4: Verify agent identity

Prerequisites

- A stack with agent access turned on, associated with a running non-domain-joined fleet. For more information, see Prerequisites.
- An IdP that supports OpenID Connect (OIDC), an OBO or token exchange grant, and custom claims in access tokens. Examples include Microsoft Entra ID, Okta, Ping Identity, Auth0, and Keycloak.

- Permissions to register applications and configure claims in your IdP, and to create IAM identity providers and roles in the AWS account that owns the stack.
- A current version of the AWS CLI.
- (Optional) A CloudTrail trail that logs data events for agent access, if you want tool calls recorded. Without it, CloudTrail records the role assumption and the MCP `initialize` request, but not the tool calls that follow. For more information, see Monitoring.

Step 1: Configure your identity provider

In your IdP, you create two applications and configure the claims that identify the agent and the user.

- **The AWS federation application** represents AWS as the resource that the agent requests access to. It exposes a scope that the agent requests in the OBO exchange, such as `access_as_agent`. You configure the identity claims on this application, so that one configuration covers every agent.
- **The agent application** represents the agent. The agent uses its client ID and credentials to run the OBO exchange. Create one agent application for each agent that needs a distinct identity.

Microsoft Entra ID

To configure Microsoft Entra ID for agent identity

1. Register the AWS federation application. Set an Application ID URI, such as `api://`*`federation-app-client-id`*, expose a scope named `access_as_agent`, and grant admin consent for it.
2. Register the agent application as a standard app registration. Add a delegated API permission for the `access_as_agent` scope of the federation application, and grant admin consent.
3. Add a certificate credential to the agent application. The agent signs a client assertion with the matching private key each time it calls the token endpoint. To keep the private key out of the agent, you can create an asymmetric AWS KMS key for signing, build an X.509 certificate for its public key, and sign assertions with the AWS KMS `Sign` operation.
4. Create a claims-mapping policy with the following definition, and assign it to the service principal of the AWS federation application. Because the policy is on the resource application, `application` claims resolve to whichever agent calls it, so one policy covers every agent.

```
{
    "ClaimsMappingPolicy": {
        "Version": 1,
        "IncludeBasicClaimSet": "true",
        "ClaimsSchema": [
            { "Source": "application", "ID": "objectid",
              "JwtClaimType": "https://aws.amazon.com/tags/principal_tags/
AgentAccessAgentId" },
            { "Source": "application", "ID": "displayname",
              "JwtClaimType": "https://aws.amazon.com/tags/principal_tags/
AgentAccessAgentName" },
            { "Source": "user", "ID": "mailnickname",
              "JwtClaimType": "https://aws.amazon.com/tags/principal_tags/
AgentAccessUserId" },
            { "Source": "user", "ID": "mailnickname",
              "JwtClaimType": "https://aws.amazon.com/source_identity" }
        ]
    }
}
```

For users that are synchronized from Active Directory, you can use `onpremisessamaccountname` instead of `mailnickname`. The `onpremises` attributes are empty for cloud-only users, and Entra omits a claim that has no value instead of returning an error.

5. On the AWS federation application, set `acceptMappedClaims` to `true` and `requestedAccessTokenVersion` to 2. Without `acceptMappedClaims`, Entra rejects the token request.

For more information, see Microsoft identity platform and OAuth 2.0 On-Behalf-Of flow and Claims customization in the Microsoft Entra documentation.

Okta

You can use Okta if it meets the requirements in Prerequisites. Configure Okta to add the claims listed in Identity attributes to the access token that it issues in the token exchange.

Other identity providers

If your IdP can add arbitrary claims to access tokens, you can emit all the session tags in one nested claim instead of one claim for each tag. Tag values must be arrays, even when there is only one value.

```
{
    "https://aws.amazon.com/tags": {
        "principal_tags": {
            "AgentAccessAgentId": ["agent-id"],
            "AgentAccessUserId": ["user-name"]
        }
    },
    "https://aws.amazon.com/source_identity": "user-name"
}
```

Transitive tag keys aren't required, because this flow doesn't chain roles. For more information about the claim formats, see Pass session tags in AWS STS in the *IAM User Guide*.

Step 2: Configure IAM

In the AWS account that owns the stack, register your IdP and create the role that the agent assumes.

To configure IAM for agent identity

1. Register your IdP as an IAM OIDC identity provider. The URL must match the `iss` claim of the token exactly. For Microsoft Entra ID, the URL is `https://login.microsoftonline.com/`*`tenant-id`*`/v2.0`.

   ```
   aws iam create-open-id-connect-provider \
       --url https://login.microsoftonline.com/tenant-id/v2.0 \
       --client-id-list agent-1-client-id agent-2-client-id
   ```

   For `--client-id-list`, specify the client ID of each agent application. If you created more than one agent application, include all of their client IDs. AWS STS rejects a token from an agent whose client ID isn't on the list. To add an agent later, use the `aws iam add-client-id-to-open-id-connect-provider` command.

2. Create a role with the following trust policy. The `Null` conditions reject any token that doesn't carry both identities. They're in a separate statement from `sts:SetSourceIdentity`, because session tags aren't part of the request for that action, so the conditions would deny it.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AssumeRoleWithBothIdentities",
            "Effect": "Allow",
            "Principal": {
                "Federated": "arn:aws:iam::111122223333:oidc-
provider/login.microsoftonline.com/tenant-id/v2.0"
            },
            "Action": [
                "sts:AssumeRoleWithWebIdentity",
                "sts:TagSession"
            ],
            "Condition": {
                "StringEquals": {
                    "login.microsoftonline.com/tenant-id/v2.0:aud": "agent-client-
id"
                },
                "Null": {
                    "aws:RequestTag/AgentAccessAgentId": "false",
                    "aws:RequestTag/AgentAccessUserId": "false"
                }
            }
        },
        {
            "Sid": "SetSourceIdentity",
            "Effect": "Allow",
            "Principal": {
                "Federated": "arn:aws:iam::111122223333:oidc-
provider/login.microsoftonline.com/tenant-id/v2.0"
            },
            "Action": "sts:SetSourceIdentity",
            "Condition": {
                "StringEquals": {
                    "login.microsoftonline.com/tenant-id/v2.0:aud": "agent-client-
id"
                }
            }
```

```
        }
    ]
}
```

For Microsoft Entra ID, the `:aud` condition key matches the `azp` claim of the token, which is the client ID of the agent application. This limits the role to that agent. To let several agents use the same role, list each client ID in the `:aud` condition.

3. Attach a permissions policy that allows the agent to call the MCP server and to start streaming sessions on the stack.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "agentaccess-mcp:*",
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "appstream:ListAssociatedFleets",
                "appstream:DescribeFleets",
                "appstream:CreateStreamingURL"
            ],
            "Resource": "*"
        }
    ]
}
```

4. (Optional) To require that the user tag matches the user of the streaming session, attach the following statement to the role. It denies MCP requests from any session whose `AgentAccessUserId` tag differs from its source identity. We recommend this statement when one role serves more than one user.

```
{
    "Sid": "UserTagMustMatchSourceIdentity",
    "Effect": "Deny",
    "Action": "agentaccess-mcp:*",
    "Resource": "*",
    "Condition": {
```

```
        "Null": { "aws:PrincipalTag/AgentAccessUserId": "false" },
        "StringNotEquals": {
            "aws:PrincipalTag/AgentAccessUserId": "${aws:SourceIdentity}"
        }
    }
}
```

Step 3: Connect the agent

The agent runs the OBO exchange, assumes the role, and connects to the MCP server with the resulting credentials. The following Python example uses Microsoft Entra ID. In it, `user_token` is the token that the agent received when the user signed in, and `client_assertion()` returns a client assertion signed with the agent's private key.

```
import boto3
import requests
from botocore.credentials import Credentials
from mcp_proxy_for_aws import aws_iam_streamablehttp_client

REGION = "region"

# 1. On-behalf-of exchange: the agent acts for the signed-in user
obo = requests.post(
    "https://login.microsoftonline.com/tenant-id/oauth2/v2.0/token",
    data={
        "grant_type": "urn:ietf:params:oauth:grant-type:jwt-bearer",
        "client_id": "agent-client-id",
        "client_assertion_type": "urn:ietf:params:oauth:client-assertion-type:jwt-
bearer",
        "client_assertion": client_assertion(),
        "assertion": user_token,
        "requested_token_use": "on_behalf_of",
        "scope": "api://federation-app-client-id/.default",
    },
).json()

# 2. Exchange the token for temporary AWS credentials.
#    The session tags and source identity come from the token.
creds = boto3.client("sts", region_name=REGION).assume_role_with_web_identity(
    RoleArn="arn:aws:iam::111122223333:role/agent-role",
    RoleSessionName="agent-session",
    WebIdentityToken=obo["access_token"],
```

```
)["Credentials"]

# 3. Connect to the MCP server. The user comes from the source identity,
#    so the user ID header isn't needed.
async with aws_iam_streamablehttp_client(
    endpoint=f"https://agentaccess-mcp.{REGION}.api.aws/mcp",
    aws_service="agentaccess-mcp",
    aws_region=REGION,
    credentials=Credentials(
        creds["AccessKeyId"], creds["SecretAccessKey"], creds["SessionToken"]
    ),
    headers={
        "x-amzn-agentaccess-workspaces-applications-stack": "stack-name-or-arn",
    },
) as (read, write, _):
    ...
```

The MCP server uses the source identity as the user of the streaming session. The user name must be 2 to 32 characters long. You don't need the `x-amzn-agentaccess-workspaces-applications-user-id` header. If you send it, it must match the source identity, or the request is rejected. For more information about the connection options, see Connecting with non-domain-joined fleets.

The credentials expire at the end of the role session. To continue after they expire, the agent runs a new OBO exchange and assumes the role again, so your IdP makes a new access decision each time.

The OBO exchange requires a valid token for the user. When the user's token expires, the agent must get a new one before it can run another OBO exchange, either by using a refresh token or by asking the user to sign in again. One sign-in doesn't let the agent run indefinitely. How long the agent can continue without the user signing in again depends on the token lifetime and refresh token policies in your IdP.

Step 4: Verify agent identity

In CloudTrail, confirm that both identities reached AWS:

1. Find the `AssumeRoleWithWebIdentity` event for the role. Under `requestParameters`, `principalTags` contains `AgentAccessAgentId` and `AgentAccessUserId`, and `sourceIdentity` contains the user. The event can take several minutes to appear in CloudTrail.
2. Find the `Initialize` event, which has the event source `agentaccess-mcp.amazonaws.com`. Under `userIdentity.sessionContext`, `sourceIdentity` contains the user, and

`webIdFederationData.federatedProvider` names your OIDC identity provider. The session tags don't appear on this event. That's expected, and it's why the source identity is required.

3. To confirm that the user is recorded on the agent's actions, find a `CallReadOnlyTool` or `CallPrivilegedTool` event (`eventCategory` is `Data`). Under `userIdentity.sessionContext`, `sourceIdentity` contains the user.

> **Note**
>
> Tool call events appear only if your trail logs data events for agent access. Without data events, you see only management events, such as `Initialize`, and not the agent's tool calls.

```
aws cloudtrail lookup-events --region region \
    --lookup-attributes AttributeKey=EventSource,AttributeValue=agentaccess-
mcp.amazonaws.com \
    --query 'Events[0].CloudTrailEvent' --output text
```

The `lookup-events` command returns management events only, such as `AssumeRoleWithWebIdentity` and `Initialize`. It never returns tool call events, because those are data events. To find tool call events, search the log files that your trail delivers to Amazon S3, or send the trail events to CloudWatch Logs and query them there.

Your IdP sign-in logs also record the OBO exchange against the agent's client ID and the user. Together with CloudTrail, they show which agent acted for which user, and when. For more information about agent access events, see Monitoring.

Agent identity for domain-joined fleets

On a domain-joined fleet, Windows must log the user on to the Active Directory domain, and an agent can't type a password. Instead, certificate-based authentication logs the user on with a short-lived certificate that WorkSpaces Applications issues for the session. To issue that certificate, WorkSpaces Applications needs the user's Active Directory identity, which your identity provider (IdP) supplies in a SAML assertion.

The agent therefore runs two on-behalf-of (OBO) exchanges from the same user token:

- **An OIDC exchange**. The agent exchanges the resulting token for temporary AWS credentials and signs its Model Context Protocol (MCP) requests with them. This is the same exchange as for non-domain-joined fleets, and it carries the agent and user identities into CloudTrail.
- **A SAML exchange**. The agent passes the resulting SAML assertion to the MCP server. WorkSpaces Applications uses it to establish the streaming session and to issue the logon certificate for the user's Active Directory account.

The steps in this topic use Microsoft Entra ID. To use another IdP, it must be able to return a SAML assertion from the OBO or token exchange, with the attributes listed in Step 2: Configure the SAML attributes.

To set up agent identity, complete the following steps:

1. Step 1: Create a SAML application in your identity provider
2. Step 2: Configure the SAML attributes
3. Step 3: Configure IAM for SAML
4. Step 4: Connect the agent
5. Step 5: Verify agent identity

Prerequisites

- A stack with agent access turned on, associated with a running domain-joined fleet.
- Certificate-based authentication turned on for the directory configuration, and working for human users. If a user can't log on with certificate-based authentication, an agent can't either. For more information, see Certificate-Based Authentication.
- The OIDC configuration for agent identity, which you complete in Step 1: Configure your identity provider and Step 2: Configure IAM for non-domain-joined fleets. The agent uses that configuration for its AWS credentials.
- (Optional) A CloudTrail trail that logs data events for agent access, if you want tool calls recorded. For more information, see Monitoring.

Step 1: Create a SAML application in your identity provider

Create a SAML application for AWS that the agent can request an assertion from in an OBO exchange. This application is separate from the AWS federation application that you created for the OIDC exchange.

To create the SAML application in Microsoft Entra ID

1. Register an application with the reply URL `https://signin.aws.amazon.com/saml`, create its service principal with SAML as the single sign-on mode, and add `urn:amazon:webservices` as an identifier URI.

   If another application in your tenant already uses `urn:amazon:webservices`, add a fragment, such as `urn:amazon:webservices#`*`app-name`*. AWS ignores the fragment when it matches the audience. Use the same value as the scope prefix in the OBO exchange.
2. Add an application-specific token signing certificate to the service principal. Entra can't issue a SAML assertion in an OBO exchange without one.
3. Expose a scope named `access_as_agent` on the SAML application, and grant the agent application consent to it. This consent is separate from the consent for the scope on the AWS federation application. Without it, the SAML exchange fails even when the OIDC exchange works.
4. Create a claims-mapping policy with the following definition. It emits the attributes in Step 2: Configure the SAML attributes. Replace the role ARN, the IAM SAML identity provider ARN, and the domain name with your own values.

```
{
    "ClaimsMappingPolicy": {
        "Version": 1,
        "IncludeBasicClaimSet": "false",
        "ClaimsSchema": [
            { "Source": "user", "ID": "userprincipalname",
              "SamlClaimType": "http://schemas.xmlsoap.org/ws/2005/05/identity/
claims/nameidentifier" },
            { "Value": "arn:aws:iam::111122223333:role/role-
name,arn:aws:iam::111122223333:saml-provider/AgentAccessSAML",
              "SamlClaimType": "https://aws.amazon.com/SAML/Attributes/Role" },
            { "Source": "user", "ID": "userprincipalname",
              "SamlClaimType": "https://aws.amazon.com/SAML/Attributes/
RoleSessionName" },
            { "Source": "user", "ID": "mailnickname",
```

```
              "SamlClaimType": "https://aws.amazon.com/SAML/Attributes/
SourceIdentity" },
            { "Source": "application", "ID": "objectid",
              "SamlClaimType": "https://aws.amazon.com/SAML/Attributes/
PrincipalTag:AgentAccessAgentId" },
            { "Source": "application", "ID": "displayname",
              "SamlClaimType": "https://aws.amazon.com/SAML/Attributes/
PrincipalTag:AgentAccessAgentName" },
            { "Source": "user", "ID": "mailnickname",
              "SamlClaimType": "https://aws.amazon.com/SAML/Attributes/
PrincipalTag:AgentAccessUserId" },
            { "Source": "user", "ID": "userprincipalname",
              "SamlClaimType": "https://aws.amazon.com/SAML/Attributes/
PrincipalTag:UserPrincipalName" },
            { "Source": "user", "ID": "onpremisesecurityidentifier",
              "SamlClaimType": "https://aws.amazon.com/SAML/Attributes/
PrincipalTag:ObjectSid" },
            { "Value": "corp.example.com",
              "SamlClaimType": "https://aws.amazon.com/SAML/Attributes/
PrincipalTag:Domain" }
        ]
    }
}
```

The `Domain` entry is optional. Include it only when the UPN suffix differs from the domain name.

5. Create the policy and assign it to the service principal of the SAML application with the Microsoft Graph API. You can't create or assign a claims-mapping policy in the Microsoft Entra admin center. For example, with the Azure CLI:

```
# Create the policy. In the request body, pass the definition as a single JSON
 string:
# { "displayName": "AgentAccessSAMLClaims", "definition": [ "policy-definition" ] }
az rest --method POST \
    --url https://graph.microsoft.com/v1.0/policies/claimsMappingPolicies \
    --headers Content-Type=application/json \
    --body @claims-policy-request.json

# Assign the policy to the service principal of the SAML application
az rest --method POST \
    --url 'https://graph.microsoft.com/v1.0/servicePrincipals/saml-sp-object-id/
claimsMappingPolicies/$ref' \
```

```
    --headers Content-Type=application/json \
    --body '{"@odata.id": "https://graph.microsoft.com/v1.0/policies/
claimsMappingPolicies/policy-id"}'
```

6. Set `acceptMappedClaims` to `true` on the SAML application. Without it, Entra rejects the token request.

```
az rest --method PATCH \
    --url https://graph.microsoft.com/v1.0/applications/saml-app-object-id \
    --headers Content-Type=application/json \
    --body '{"api": {"acceptMappedClaims": true}}'
```

For more information about SAML federation with WorkSpaces Applications, see Setting Up SAML.

Step 2: Configure the SAML attributes

The SAML assertion must carry the agent and user identities, and the Active Directory attributes that WorkSpaces Applications uses to build the logon certificate. Except for the name ID, each attribute name starts with `https://aws.amazon.com/SAML/Attributes/`.

SAML attributes for agent identity on domain-joined fleets

Attribute	Value
Name ID	The user's UPN, in the format *user*@*ad-dns-name* . Source it from the user's `userprincipalname` attribute. Entra ignores a constant value for the name ID. WorkSpaces Applications writes the name ID into the subject alternative name of the logon certificate.
`Role`	The role ARN and the IAM SAML identity provider ARN, separated by a comma
`RoleSessionName`	The user's UPN
`SourceIdentity`	The user's short name
`PrincipalTag:AgentAccessAgentId`	The agent's immutable object ID

Attribute	Value
`PrincipalTag:AgentAccessAgentName`	(Optional) The agent's display name
`PrincipalTag:AgentAccessUserId`	The user's short name
`PrincipalTag:UserPrincipalName`	The user's UPN
`PrincipalTag:ObjectSid`	The security identifier (SID) of the Active Directory user, such as `S-1-5-21-...`. WorkSpaces Applications adds it to the certificate as the strong certificate mapping that domain controllers require.
`PrincipalTag:Domain`	(Optional) The fully qualified domain name of the Active Directory domain. Use it when the UPN suffix differs from the domain name.

Keep the following in mind when you configure the attributes:

- In Microsoft Entra ID, set `IncludeBasicClaimSet` to `false` in the claims-mapping policy. Additional claims cause AWS Sign-In to show an interactive role selection page, which an agent can't complete.
- Use the same values for `SourceIdentity` and `PrincipalTag:AgentAccessUserId` as in the OIDC exchange, such as the user's `mailnickname` for the short name. The streaming session takes the user from the SAML assertion, and the MCP events in CloudTrail take it from the OIDC token. If the values differ, the two name different users.
- For users synchronized from Active Directory, source `PrincipalTag:ObjectSid` from `onpremisesecurityidentifier`. For other users, that attribute is empty, and the logon certificate can't be mapped to the account.

> **Important**
>
> For testing only, you can emit the SID of an Active Directory account as a constant value. A constant value in a claims-mapping policy applies to every user of the application,

so every user would get the same SID. Use a constant only for a single test user, or with a separate SAML application and policy for each user. Don't use a constant SID in production. To read the SID, run `Get-ADUser -Identity user-name -Properties objectSid` on a domain-joined computer.

- You don't need `TransitiveTagKeys`, because this flow doesn't chain roles. If you include it, emit one attribute value for each tag key. A single comma-separated value isn't a valid tag key, and the role assumption fails.

Step 3: Configure IAM for SAML

To configure IAM for the SAML exchange

1. Download the federation metadata of the SAML application from your IdP, and create an IAM SAML identity provider from it.

```
aws iam create-saml-provider --name AgentAccessSAML \
    --saml-metadata-document file://saml-metadata.xml
```

2. Add the following statements to the trust policy of the role that you created for the OIDC exchange. The role then trusts both identity providers. As in the OIDC trust policy, the `Null` conditions are in a separate statement from `sts:SetSourceIdentity`.

```
{
    "Sid": "AssumeRoleWithSAMLBothIdentities",
    "Effect": "Allow",
    "Principal": {
        "Federated": "arn:aws:iam::111122223333:saml-provider/AgentAccessSAML"
    },
    "Action": [
        "sts:AssumeRoleWithSAML",
        "sts:TagSession"
    ],
    "Condition": {
        "StringEquals": { "SAML:aud": "https://signin.aws.amazon.com/saml" },
        "Null": {
            "aws:RequestTag/AgentAccessAgentId": "false",
            "aws:RequestTag/AgentAccessUserId": "false"
        }
    }
```

```
},
{
    "Sid": "SAMLSetSourceIdentity",
    "Effect": "Allow",
    "Principal": {
        "Federated": "arn:aws:iam::111122223333:saml-provider/AgentAccessSAML"
    },
    "Action": "sts:SetSourceIdentity",
    "Condition": {
        "StringEquals": { "SAML:aud": "https://signin.aws.amazon.com/saml" }
    }
}
```

3. Allow the role to stream from the stack. Without this permission, the streaming session can't be established.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "appstream:Stream",
            "Resource": "arn:aws:appstream:region:111122223333:stack/stack-name"
        }
    ]
}
```

Step 4: Connect the agent

The agent runs both exchanges from the same user token. It signs MCP requests with the credentials from the OIDC exchange, as shown in Step 3: Connect the agent, and passes the SAML assertion and the stack ARN in the `_meta` field of every request. SAML assertions are too large for HTTP headers.

Microsoft Entra ID returns the assertion from the OBO exchange as a bare SAML assertion, encoded in base64url. AWS Sign-In expects a SAML response, so the agent wraps the assertion in an unsigned `samlp:Response` element. AWS Sign-In validates the signature of the assertion inside it.

```
import base64
import json
import uuid
```

```
from datetime import datetime, timezone

import requests
from botocore.auth import SigV4Auth
from botocore.awsrequest import AWSRequest
from botocore.credentials import Credentials

REGION = "region"
MCP_URL = f"https://agentaccess-mcp.{REGION}.api.aws/mcp"

# 1. SAML on-behalf-of exchange, from the same user token as the OIDC exchange
obo_saml = requests.post(
    "https://login.microsoftonline.com/tenant-id/oauth2/v2.0/token",
    data={
        "grant_type": "urn:ietf:params:oauth:grant-type:jwt-bearer",
        "client_id": "agent-client-id",
        "client_assertion_type": "urn:ietf:params:oauth:client-assertion-type:jwt-
bearer",
        "client_assertion": client_assertion(),
        "assertion": user_token,
        "requested_token_use": "on_behalf_of",
        "requested_token_type": "urn:ietf:params:oauth:token-type:saml2",
        "scope": "urn:amazon:webservices/.default",
    },
).json()

# 2. Decode the base64url assertion and wrap it in a SAML response
raw = obo_saml["access_token"]
assertion_xml = base64.urlsafe_b64decode(raw + "=" * (-len(raw) % 4)).decode()
now = datetime.now(timezone.utc).strftime("%Y-%m-%dT%H:%M:%SZ")
saml_response = base64.b64encode((
    '<samlp:Response xmlns:samlp="urn:oasis:names:tc:SAML:2.0:protocol" '
    f'ID="_{uuid.uuid4()}" Version="2.0" IssueInstant="{now}" '
    'Destination="https://signin.aws.amazon.com/saml">'
    '<samlp:Status><samlp:StatusCode '
    'Value="urn:oasis:names:tc:SAML:2.0:status:Success"/></samlp:Status>'
    f'{assertion_xml}</samlp:Response>'
).encode()).decode()

# 3. Send signed MCP requests with the OIDC credentials (creds),
#    with the SAML assertion and the stack ARN in _meta
credentials = Credentials(
    creds["AccessKeyId"], creds["SecretAccessKey"], creds["SessionToken"]
)
```

```
auth_meta = {
    "aws.agentaccess/workspacesApplicationsSamlAssertion": saml_response,
    "aws.agentaccess/workspacesApplicationsStackArn":
 "arn:aws:appstream:region:111122223333:stack/stack-name",
}
session_id = None

def mcp_request(method, params=None, request_id=1):
    """Send one JSON-RPC request. _meta is added before the body is signed."""
    params = dict(params or {})
    params["_meta"] = auth_meta
    body = json.dumps(
        {"jsonrpc": "2.0", "id": request_id, "method": method, "params": params}
    ).encode()
    headers = {"Content-Type": "application/json"}
    if session_id:
        headers["Mcp-Session-Id"] = session_id
    request = AWSRequest(method="POST", url=MCP_URL, data=body, headers=headers)
    SigV4Auth(credentials, "agentaccess-mcp", REGION).add_auth(request)
    # Send exactly the bytes that were signed
    response = requests.post(MCP_URL, data=body, headers=dict(request.headers))
    response.raise_for_status()
    return response

init = mcp_request("initialize", {
    "protocolVersion": "2025-03-26",
    "capabilities": {},
    "clientInfo": {"name": "my-agent", "version": "1.0"},
})
session_id = init.headers["Mcp-Session-Id"]

tools = mcp_request("tools/list", request_id=2).json()["result"]["tools"]
screenshot = mcp_request(
    "tools/call", {"name": "agentaccess___screenshot", "arguments": {}}, request_id=3
).json()
```

The desktop can take a few moments to connect after `initialize`. If the first `tools/list` or tool call fails, wait and retry. To end the session, send a signed `DELETE` request to the endpoint with the `Mcp-Session-Id` header.

The example follows these rules. Follow them in your own client too:

- Include `_meta` on every request, not only on `initialize`. If it's missing from a later request, the streaming session is lost.
- Compute the SigV4 signature over the exact bytes that you send, including `_meta`. If you change the body after you sign it, the request fails with a signature mismatch or a `Content-Length` error.
- The `initialize` response returns an `Mcp-Session-Id` header. Send it on every later request. Without it, the service returns HTTP 400 with the message `Missing Mcp-Session-Id header`.
- Send the SAML assertion and the stack ARN only in `_meta`. The service rejects a request that also sends them as headers.

The stack ARN must be in the same AWS account as the credentials that sign the request.

The SAML assertion is valid for a short time. When the agent starts a new session, it runs both exchanges again. Both exchanges require a valid token for the user, so the agent must refresh the user's token when it expires, as described in Step 3: Connect the agent for non-domain-joined fleets.

The example signs every request with the same static credentials, so requests fail after the role session expires. Before that happens, run both exchanges again and sign later requests with the new credentials and SAML assertion.

Step 5: Verify agent identity

The following signals confirm that agent identity works on a domain-joined fleet. Each one proves something different.

1. **Both identities reached AWS**. In CloudTrail, the `AssumeRoleWithSAML` event for the role contains `AgentAccessAgentId`, `AgentAccessUserId`, `UserPrincipalName`, and `ObjectSid` in `principalTags`, and the user in `sourceIdentity`. The `ObjectSid` tag shows that the security identifier for the logon certificate arrived. The event can take several minutes to appear in CloudTrail.
2. **The user is recorded on agent actions**. Events with the event source `agentaccess-mcp.amazonaws.com` show the user in `userIdentity.sessionContext.sourceIdentity`. These events belong to the OIDC role session, because the agent signs MCP requests with the OIDC credentials. The SAML assertion establishes the streaming session but doesn't sign AWS requests.

> **Note**
>
> Agent actions (tool calls, recorded as `CallReadOnlyTool` or `CallPrivilegedTool`) are data events. They appear in CloudTrail only if your trail logs data events for agent access. Without data events, you see only management events, such as `Initialize`, and not the agent's tool calls. For more information, see Monitoring.

3. **Windows logged on with the certificate**. On the domain controller, the Security log shows event 4768 with pre-authentication type 16 and a certificate thumbprint, followed by events 4769 and 4624. Event 4771 indicates a failed certificate logon.
4. **The agent can act**. A successful `tools/list` request shows that the session is established. A successful `tools/call`, such as taking a screenshot, shows that the agent can act on the desktop.

Troubleshooting

Troubleshooting agent identity on domain-joined fleets

Symptom	Possible cause	Solution
The SAML exchange fails, but the OIDC exchange works	The agent application doesn't have consent to the `access_as_agent` scope of the SAML application, or the SAML application has no application-specific signing certificate.	Grant the consent and add the certificate, as described in Step 1: Create a SAML application in your identity provider.
The MCP request fails with `403 Access denied by IAM policy`	The role doesn't allow `agentaccess-mcp` actions.	Add the permissions in Step 2: Configure IAM.
The streaming session fails with `INVALID_AUTH_POLICY`	The role doesn't allow `appstream:Stream` on the stack, or the stack ARN in the request isn't the one in the policy.	Allow `appstream:Stream` on the stack, as described in Step 3: Configure IAM for SAML, and send the same stack ARN in `_meta`.

Symptom	Possible cause	Solution
The session fails with `CBA certificate fetch failed`	WorkSpaces Applications can't issue the logon certificate from your private CA, because certificate-based authentication isn't working for the directory.	Confirm that certificate-based authentication works for a human user. For more information, see Certificate-Based Authentication.

Connection modes

You can control how your agent waits for the desktop session to become available by setting the `x-amzn-agentaccess-connect-mode` header on your Model Context Protocol (MCP) requests.

> **Note**
>
> Connection modes apply to both non-domain-joined and domain-joined fleets. Set the `x-amzn-agentaccess-connect-mode` header alongside whichever authentication mechanism your fleet type uses. For non-domain-joined fleets, that is either the stack and user ID headers or a streaming-URL header. For domain-joined fleets, it is the SAML assertion metadata.

The following modes are available:

- **BLOCKING** (default): the MCP server waits until the desktop connection is fully established before responding. When `tools/list` returns, all tools are immediately available.
- **POLLING**: the MCP server responds immediately without waiting for the desktop connection. Initially, only the `connection_status` tool is available. Your agent polls this tool until the connection is established, at which point the full tool set becomes available.

Use POLLING mode when you want your agent to perform other work while waiting for the desktop connection, or when you need more control over connection timeout behavior.

The following example shows how to use POLLING mode:

```
# Pass the header when creating the MCP connection
headers = {
```

```
    "x-amzn-agentaccess-streaming-session-url": streaming_url,  # non-domain-joined
 fleets
    "x-amzn-agentaccess-connect-mode": "POLLING",
}

# After initialize, tools/list returns immediately with connection_status
tools = await session.list_tools()
# tools = [connection_status]

# Poll connection_status until the desktop is ready
while True:
    result = await session.call_tool("connection_status", {})
    status = json.loads(result.content[0].text)
    if status["state"] == "CONNECTED":
        break
    time.sleep(2)

# Now tools/list returns the full set (screenshot, left_click, type_text, and so on)
tools = await session.list_tools()
```

Clean up: To avoid ongoing streaming charges, end the WorkSpaces Applications streaming session when your agent finishes its work. Disconnect the agent so the session closes, and do not leave idle sessions running.

Session cleanup

When your agent finishes its work, it ends its session with the Model Context Protocol (MCP) server. Ending the session, and optionally expiring the streaming session, releases the underlying streaming instance so the fleet's autoscaling policy can reclaim idle capacity and avoid unnecessary cost. How the agent ends the session, and what happens to the underlying WorkSpaces Applications streaming session, depends on the MCP protocol version. For the supported versions, see Supported MCP protocol versions. On an image builder, ending the MCP session doesn't log out the built-in users. Only stopping the image builder ends its streaming session. For more information, see Image builder sessions.

Stateless MCP protocol

With the stateless MCP protocol (version 2026-07-28), the agent ends the session by calling the `close_session` tool. This closes the MCP session. By default, `close_session` does not expire the WorkSpaces Applications streaming session; the streaming instance is preserved for reuse until the session's disconnect timeout is reached. To also expire the streaming session, the agent sets

`expire_streaming_session` to `true` when it calls `close_session`. Expiring the streaming session terminates the streaming instance and triggers the fleet's configured autoscaling policy. For more information about the tool, see Available tools. For more information about the disconnect timeout, see Create a Fleet in Amazon WorkSpaces Applications. For more information about autoscaling, see Fleet Auto Scaling for Amazon WorkSpaces Applications.

Earlier MCP protocol versions

With the earlier MCP protocol versions, you end the session by closing the MCP client in your agent code. When the MCP client lifecycle ends, the client sends an HTTP `DELETE` request to the MCP server. Some agent frameworks, such as Strands, close the client for you when the session ends. You control whether ending the session also expires the streaming session by setting the `x-amzn-agentaccess-expire-streaming-session-on-delete` header on your MCP requests. The following values are available:

- `true` – When the MCP client sends the HTTP `DELETE` request, the MCP server expires the WorkSpaces Applications streaming session as part of cleanup.
  - On a fleet, this terminates the streaming instance and triggers the fleet's configured autoscaling policy. For more information, see Fleet Auto Scaling for Amazon WorkSpaces Applications.
  - On an image builder, this header doesn't end the streaming session. The streaming session and the built-in users that the agent logged in keep running until you stop the image builder. Stopping the image builder ends its streaming session and logs out all users.
- `false` (default) – The streaming session continues running.
  - On a fleet, the session runs until the disconnect timeout is reached. For more information about the disconnect timeout, see Create a Fleet in Amazon WorkSpaces Applications.
  - On an image builder, the session runs until you stop the image builder. For more information, see Image Builder Actions.

> **Note**
>
> By default, an `mcp-proxy-for-aws` MCP client handles the `DELETE` request for you when you properly end the client lifecycle.

Image builder sessions

On an image builder, also keep the following in mind:

- Ending the MCP session doesn't end the image builder session. The MCP session closes after about 13 minutes of inactivity, or when your agent closes the session. After an idle close, the next tool call reconnects as the last logged-in user.
- When the image builder session ends, all users are logged out.
- An image builder is billed while it's running, even after its streaming session ends. Stop the image builder when your agent is done. For more information, see Image Builder Actions.

Available tools

The Model Context Protocol (MCP) server provides the following tools, grouped into mouse, keyboard, screen, application, sequencing, session, and user tools. Agents use them to operate the desktop during a streaming session. All tool names use the `agentaccess___` prefix.

The tools available to an agent depend on whether the session is on a stack or on an image builder. Application tools are available only on stacks, and user tools are available only on image builders.

Mouse tools

`left_click`

Perform a left click at the given coordinates.

Parameters: `x` (required), `y` (required), `modifiers` (optional, for example `ctrl` or `ctrl+shift`).

`double_click`

Perform a double click at the given coordinates.

Parameters: `x` (required), `y` (required), `modifiers` (optional).

`triple_click`

Perform a triple click at the given coordinates.

Parameters: `x` (required), `y` (required), `modifiers` (optional).

`right_click`

Perform a right click at the given coordinates.

Parameters: `x` (required), `y` (required), `modifiers` (optional).

`middle_click`

Perform a middle click at the given coordinates.

Parameters: `x` (required), `y` (required), `modifiers` (optional).

`left_click_drag`

Perform a left click drag from start coordinates to end coordinates.

Parameters: `start_x` (required), `start_y` (required), `end_x` (required), `end_y` (required).

`left_mouse_down`

Press and hold the left mouse button at the given coordinates.

Parameters: `x` (required), `y` (required), `modifiers` (optional).

`left_mouse_up`

Release the left mouse button at the given coordinates.

Parameters: `x` (required), `y` (required), `modifiers` (optional).

`move_pointer`

Move the pointer to the given coordinates.

Parameters: `x` (required), `y` (required).

`scroll`

Scroll the mouse wheel at the given coordinates.

Parameters: `x` (required), `y` (required), `scroll_direction` (required, one of `Up`, `Down`, `Left`, or `Right`), `scroll_amount` (required, in ticks, where 120 ticks equals one wheel notch), `modifiers` (optional).

Keyboard tools

`type_text`

Type text by simulating keyboard events for each character.

Parameters: `text` (required, up to 10,000 characters).

`key`

Press a key or key combination.

Parameters: `keys` (required, a single key or combination joined by +, for example `a`, `ctrl+c`, or `ctrl+shift+s`).

`hold_key`

Hold a key or key combination for a specified duration.

Parameters: `keys` (required), `duration` (required, 1 to 30 seconds).

Screen tools

`screenshot`

Capture a screenshot of the desktop. The returned image dimensions define the coordinate space for all mouse tools.

Parameters: `include_cursor` (optional, defaults to `false`).

Application tools

Application tools appear in the agent's tool list only when applications are registered in the fleet image's catalog. If the image has no registered applications, these tools are not available.

`launch_application`

Launch an application from the ones available in the fleet image's catalog.

Parameters: `id` (required). The application identifier.

The tool's own description lists the identifiers of the available applications. The MCP server populates this description from the fleet image's catalog (for example, `firefox`, `notepad`

++, or eclipse). Use one of the listed identifiers. An id that is not in the catalog returns isError: true with SERVER_ERROR: Invalid input. IDs are set when the image is built. You cannot determine an ID from the application name.

Launching an application does not replace what is already running. The previously launched application stays open. That is why the app switcher exists.

toggle_app_switcher

Open or close the app-switcher overlay.

Parameters: none.

This tool only toggles the overlay. To bring a particular application to the foreground, follow up with a left_click on its thumbnail, using coordinates read from a screenshot. No tool reports which application is currently in the foreground, so an agent that needs to know must capture a screenshot and look.

> **Note**
>
> The app switcher is available only when the fleet's **Stream view** is set to **Application**, which displays only the windows of applications that users open. It is not available when **Stream view** is set to **Desktop**, because WorkSpaces Applications then provides the standard desktop and taskbar for the operating system.

Sequencing tools

Use these tools to run several actions in one call, and to pace them, instead of making a separate tool call for each step.

execute_sequence

Run several tool calls in one call, in order. Each item in the actions list is an object with a type that names the tool to run. Write the type the same way the tool appears in tools/list, for example agentaccess___left_click or agentaccess___wait. Include that tool's own arguments as sibling keys, not nested. For example:

```
[{"type": "agentaccess___type_text", "text": "hello"},
 {"type": "agentaccess___key", "keys": "Tab"}]
```

The sequence stops at the first failure. You cannot use `screenshot` or `execute_sequence` inside a sequence. If the sequence succeeds, the MCP server pauses briefly and returns a final screenshot. If the last action in the sequence is a `wait`, the server skips this pause. You can end a sequence with a `wait` to set your own delay after a slow action, such as opening a window. If an action fails or the sequence reaches the time limit, the MCP server returns text only. It includes guidance to take a screenshot before the next action.

A sequence has the following limits. If a sequence exceeds any limit, the MCP server rejects it before any action runs:

- A sequence can include at most 20 actions.
- Each `wait` lasts 0.001 to 5 seconds, and all waits together last at most 10 seconds.
- All `type_text` text together totals at most 10,000 characters.
- The entire sequence stops after 60 seconds.

Parameters: `actions` (required, a list of action objects).

The following example request fills in a form and submits it. It enters text in three fields, presses Enter to submit, and ends with a `wait` so the page has time to update:

```
{
  "jsonrpc": "2.0",
  "method": "tools/call",
  "id": 5,
  "params": {
    "name": "agentaccess___execute_sequence",
    "arguments": {
      "actions": [
        {"type": "agentaccess___left_click", "x": 450, "y": 320},
        {"type": "agentaccess___type_text", "text": "Jane Doe"},
        {"type": "agentaccess___key", "keys": "Tab"},
        {"type": "agentaccess___type_text", "text": "Seattle"},
        {"type": "agentaccess___key", "keys": "Tab"},
        {"type": "agentaccess___type_text", "text": "Please add me to the
 newsletter."},
        {"type": "agentaccess___key", "keys": "Enter"},
        {"type": "agentaccess___wait", "duration": 2}
      ]
    }
  }
}
```

`wait`

Pause for `duration` seconds without interacting with the desktop. Use `wait` on its own, or as an action inside `execute_sequence` to pace actions that need time to settle.

Parameters: `duration` (required, 0.001 to 5 seconds).

Session tools

The `close_session` tool is part of the stateless MCP protocol (version 2026-07-28). With the earlier MCP protocol versions, you end the session by closing the MCP client in your agent code, which sends an HTTP `DELETE` request instead. For more information, see Session cleanup.

Note

To watch an agent while it works, code the agent to open a streaming URL in a browser. The `get_session_info` tool doesn't return a streaming URL, but its values let your agent's code create one:

- On a stack, call the `CreateStreamingURL` API with the `stack`, `fleet`, and `userId` values from `get_session_info`.
- On an image builder, call the `CreateImageBuilderStreamingURL` API with the image builder name. The name is the last part of the `imageBuilder` ARN from `get_session_info`.

Then read the `streaming_url` value from the `_meta` field of your MCP client, and open it in a browser. The full URL contains session credentials, so keeping it in `_meta` keeps it out of the model's context.

`get_session_info`

Return information about the current session.

Parameters: none.

Returns the following identifier fields as the tool result: `mcpSessionId`, `stack`, `userId`, `workSpacesSessionId`, and `fleet`.

On an image builder session, returns `mcpSessionId`, `workSpacesSessionId`, and `imageBuilder` (the ARN of the image builder).

`close_session`

The agent calls this tool to end the current session. This closes the MCP session. By default, it does not expire the WorkSpaces Applications streaming session, and the streaming instance is preserved for reuse.

Parameters: `expire_streaming_session` (optional, Boolean, default `false`). The following values are available:

- `true` – The MCP server also expires the WorkSpaces Applications streaming session.
  - On a fleet, this terminates the streaming instance and triggers the fleet's configured autoscaling policy. For more information, see Fleet Auto Scaling for Amazon WorkSpaces Applications.
  - On an image builder, this doesn't end the streaming session. The streaming session and the built-in users that the agent logged in keep running until you stop the image builder. Stopping the image builder ends its streaming session and logs out all users.
- `false` (default) – The streaming session continues running.
  - On a fleet, the session runs until the disconnect timeout is reached. For more information about the disconnect timeout, see Create a Fleet in Amazon WorkSpaces Applications.
  - On an image builder, the session runs until you stop the image builder. For more information, see Image Builder Actions.

User tools

User tools control which built-in user is logged in to the image builder. The logged-in built-in user determines which actions the agent can perform, because each built-in user has different permissions on the image builder. These tools are available only on image builders. For more information about connecting as a built-in user, see Connecting with image builders.

`switch_user`

Disconnect the current built-in user's Windows session and log in another built-in user. The previous user's session keeps running.

Parameters: `user` (required, one of `ADMINISTRATOR`, `TEMPLATE_USER`, or `TEST_USER`). These are the same values as the `x-amzn-agentaccess-local-user` header.

`get_current_user`

Return the built-in user that is logged in to the image builder.

Parameters: none.

Returns the logged-in built-in user.

Computer use interfaces

By default, an agent uses the native WorkSpaces Applications tools, such as `left_click` and `type_text`. With a model-specific computer use interface, you present those tools in the model's own schema instead. The Model Context Protocol (MCP) server converts each call into the underlying desktop actions. You select the interface for each request by sending an HTTP header to the MCP server.

Different providers train their computer use models to call tools by different names and shapes. For example, to click at a point, the native tools use a `left_click` tool. One provider's model calls a single `computer` tool with an action of `left_click`. Another provider's model calls a `computer` tool with an action of `click`.

When you select the computer use interface that matches your model, the model uses the tool names and schema it was trained on. This increases accuracy and lowers the cost for your agent to perform actions.

Select an interface

You select a computer use interface by sending the `x-amzn-agentaccess-computer-use-interface-id` HTTP header on the `tools/list` request and on every `tools/call` request. The header name is case-insensitive, and the value is case-sensitive.

```
x-amzn-agentaccess-computer-use-interface-id: ANTHROPIC_COMPUTER_20251124
```

You send the value only as this HTTP header. The interface is selected per request and isn't stored. A change to the header takes effect on the next `tools/list` or `tools/call` request.

If you omit the header or send it empty, the agent uses the native WorkSpaces Applications tools described in Available tools. If you send a value that isn't one of the supported values, the MCP server returns the JSON-RPC error `-32602`.

Selecting an interface replaces only the standard desktop tools. Session tools, forwarded tools, and application catalog tools are unchanged.

Supported interfaces

The following table lists the computer use interfaces that WorkSpaces Applications supports.

Computer use interface	Header value
Default (standard desktop tools)	Omit the header
Anthropic computer use (2025-01-24)	`ANTHROPIC_COMPUTER_20250124`
Anthropic computer use (2025-11-24)	`ANTHROPIC_COMPUTER_20251124`
Anthropic computer toolset (2026-08-01)	`ANTHROPIC_COMPUTER_TOOLSET_20260801`
OpenAI computer use v1	`OPENAI_COMPUTER_V1`
Gemini computer use 3.x desktop v1	`GOOGLE_COMPUTER_USE_3X_DESKTOP_V1`

MCP tool forwarding

With Model Context Protocol (MCP) tool forwarding, agents can work directly with the applications and operating system running on the desktop. They make direct MCP calls instead of using computer-use tools. When you turn on tool forwarding, the MCP server forwards the tools set up on the WorkSpaces Applications session to your agent.

Setting up tool forwarding

To set up MCP tool forwarding:

Step 1: Enable tool forwarding

Turn on the `FORWARD_MCP_TOOLS` agent action through the API or console settings.

Step 2: Add the configuration file to the fleet's image

WorkSpaces Applications creates each streaming instance from the fleet's image, so the configuration file must be part of that image. Add it when you build or update the image. A file placed on a running instance does not carry over to new sessions. The service looks for the file at the following location, which depends on the platform:

Platform	Configuration file path
Windows instances	`C:\ProgramData\NICE\dcv\mcp_server_redirection_config.json`
Linux instances	`/etc/dcv/mcp_server_redirection_config.json`

Tool forwarding works the same way on image builders. Add the same configuration file, at the same location, to the image that you launch the image builder from.

Step 3: Define the MCP servers in the configuration file

The configuration file is JSON with a single top-level `mcpServers` object. Each key is a unique name that you choose for a server. Each value specifies how to launch that server.

The following example is for a Windows instance:

```
{
    "mcpServers": {
        "filesystem": {
            "command": "C:/path/to/python.exe",
            "args": ["C:/mcpServerPath/filesystem.py", "C:/UserName/Documents"]
        },
        "weather": {
            "command": "C:/Program Files/my-mcp/weather.exe"
        }
    }
}
```

Field	Required	Type	Description
`command`	Yes	String	Absolute path to the executable to launch.

Field	Required	Type	Description
`args`	No	Array of strings	Arguments passed to the executable.

Step 4: Verify tool availability

If the configuration file is present, the service connects to the MCP servers it lists and forwards their tools. Forwarded tools appear when the agent lists its available tools.

Step 5: Test the configuration

To confirm that tool forwarding works, run the `mcpforwardingtester` script. You can find it in the sample-code-for-workspaces-agent-access repository on GitHub. The script tests the MCP setup on your WorkSpace. It reports which tools are ready and flags any issues.

> **Note**
>
> Both IAM access and the service setting must be enabled for tool forwarding to work. IAM permissions do not override the service setting.

MCP tool forwarding considerations

Note the following considerations when you configure MCP tool forwarding:

- **Transport is standard I/O (stdio) only.** Each entry must launch a process that speaks MCP over its standard input and output. Remote HTTP or SSE MCP endpoints are not supported. To use a remote endpoint, wrap it in a local stdio server.
- **Only `command` and `args` are supported.** There is no field for environment variables or working directory. Each server inherits the environment of the streaming session and runs as the session user. Use absolute paths for `command` and for any path arguments.
- **On Windows instances, use forward slashes in paths.** For example, use `C:/Program Files/my-mcp/server.exe`. JSON treats the backslash as an escape character, so a Windows-style path written with single backslashes is invalid. Windows accepts forward slashes for absolute paths, which avoids the need to escape every separator as `\\`.

How forwarded tools appear to the agent

To prevent collisions between servers, the MCP server renames each forwarded tool in the agent's tool list using the following pattern:

```
forwarded___server-name___original-tool-name
```

The *server-name* is the key from your configuration file. For example, a `get_forecast` tool from the `weather` server is listed as `forwarded___weather___get_forecast`. When the agent calls the forwarded name, the MCP server routes the request to the original tool on the owning server. Agent code that matches on tool names must expect this prefix.

IAM permissions for tool forwarding

The IAM action for calling forwarded tools is `CallForwardedTool`. To limit access to specific stacks, use the `StackArn` condition key:

```
{
  "Effect": "Allow",
  "Action": ["agentaccess-mcp:InvokeMcp", "agentaccess-mcp:CallForwardedTool"],
  "Resource": "*",
  "Condition": {
    "ArnLike": {
      "agentaccess-mcp:StackArn": "arn:aws:appstream:region:account-id:stack/stack-
name"
    }
  }
}
```

Image builder sessions use the `ImageBuilderArn` condition key instead of `StackArn`. To allow access only to specific image builders, use the `ImageBuilderArn` condition key:

```
{
  "Effect": "Allow",
  "Action": ["agentaccess-mcp:InvokeMcp", "agentaccess-mcp:CallForwardedTool"],
  "Resource": "*",
  "Condition": {
    "ArnEquals": {
      "agentaccess-mcp:ImageBuilderArn": "arn:aws:appstream:region:account-id:image-
builder/image-builder-name"
```

```
    }
  }
}
```

Monitoring

You can monitor agent activity through three services: AWS CloudTrail for session events, CloudWatch for operational metrics, and Amazon S3 for stored screenshots.

- **AWS CloudTrail**: logs the following agent session activity:
  - When an agent connects to a session
  - The tools an agent calls
  - When a session ends

  Tool calls are recorded as data events, one for each `agentaccess-mcp` action the agent invokes (for example, `GetScreenshot`, `LeftClick`, or `TypeText`). For the full list of actions, see Agent and permissions requirements. Because these are data events, you must set up a trail to log them. For more information, see Logging data events in the *CloudTrail User Guide*.
- **CloudWatch**: the following operational metrics are available in CloudWatch for agent sessions:
  - `Invocations`
  - `Latency`
  - `ClientErrors`
  - `ServerErrors`
  - `McpSessionStart`
  - `McpSessionDuration`
- **Amazon S3**: if you configure screenshot storage, screenshots captured during agent sessions are available in the Amazon S3 bucket that you specify. Screenshots are stored with the following key format:

  ```
  agentaccess/screenshots/year=YYYY/month=MM/day=DD/session-id/timestamp.png
  ```

  The *`session-id`* in the path is the WorkSpaces Applications streaming session ID.

Session evidence

Session evidence is the set of capabilities that you use to track what happened during an agent session. Session evidence gives you an audit-ready record of what the AI agent did, delivered to your own Amazon S3 bucket.

You can choose either of the following evidence options:

- **Screenshots and CloudWatch logs**. Store the screenshots that the agent takes during a session in your Amazon S3 bucket, and review agent activity in CloudWatch. For more information, see Screenshot storage and Monitoring.
- **[Public Preview] Replay Session**. Combine the screenshots taken before each agent action into a replayable session video, paired with an AI-generated transcript and an HTML report. The rest of this topic describes Replay Session. Replay Session is available on stacks and not on image builder.

How Replay Session works

Replay Session is turned off by default. An administrator turns it on for each stack and names an Amazon S3 bucket to receive the recordings. When Replay Session is turned on, it works as follows:

1. During the session, just before the agent performs each action, Replay Session captures a frame of the whole screen and records metadata about the action, such as the tool that the agent called, its arguments, the tool response, and a timestamp.
2. Replay Session stores the frames and their metadata in storage that AWS owns, encrypted at rest. For more information, see Encryption at rest for Replay Session.
3. After the session ends, Replay Session uses Amazon Bedrock to analyze the frames and metadata, and produces the following artifacts:
   - A session video built from the frames.
   - A transcript of the session.
   - An HTML report.
4. Replay Session delivers the artifacts to your Amazon S3 bucket. Processing and delivery continue for a few minutes after the session ends.

Frames, metadata, tool responses, and finished artifacts that are awaiting delivery remain in AWS-owned storage for up to 15 days, and then expire. This retention period applies only to the copy

that Replay Session holds during processing. The artifacts in your bucket are retained according to the lifecycle configuration of your bucket, which you control.

What Replay Session records and how it is processed

Before you turn on Replay Session, review what it records and how it processes that content. By turning on Replay Session, you accept the following:

- **Whole-screen capture**. Each frame records the entire screen as it appears just before the agent acts. Anything visible on the screen at that moment is in the recording.
- **Typed text**. Text that the agent types is saved in the `arguments` field of the per-action metadata, and can appear in the transcript.
- **No redaction**. Replay Session doesn't redact sensitive data from frames, metadata, or artifacts.
- **Processing by Amazon Bedrock**. Replay Session sends frames and metadata to Amazon Bedrock to produce the AI-generated transcript. This content is processed in AWS Regions other than the Region where the session runs. For details, see Cross-Region inference for Replay Session.

Cross-Region inference for Replay Session

Review this section before you turn on Replay Session.

Replay Session uses Amazon Bedrock cross-Region inference to produce the AI-generated transcript. In-Region inference for the model that Replay Session uses isn't available in any of the AWS Regions that support Replay Session. As a result, inference requests, including the frames and metadata from your sessions, can be routed to and processed in AWS Regions other than the Region where your session runs:

- **Geographic cross-Region inference**. Where Amazon Bedrock offers an inference profile for the geography of your Region, requests are routed only to Regions within that geography, such as North America, Europe, or Australia.
- **Global cross-Region inference**. Where no inference profile exists for the geography of your Region, requests can be routed to commercial AWS Regions worldwide.

The following table shows how requests are routed for each Region where Replay Session is available.

Cross-Region inference routing for Replay Session

Session Region	Inference Regions
US East (N. Virginia) (us-east-1) US East (Ohio) (us-east-2) US West (Oregon) (us-west-2) Canada (Central) (ca-central-1)	US East (N. Virginia) (us-east-1) US East (Ohio) (us-east-2) US West (Oregon) (us-west-2) Canada (Central) (ca-central-1)
Europe (Frankfurt) (eu-central-1) Europe (Ireland) (eu-west-1) Europe (London) (eu-west-2) Europe (Milan) (eu-south-1) Europe (Paris) (eu-west-3) Europe (Spain) (eu-south-2)	Europe (Frankfurt) (eu-central-1) Europe (Ireland) (eu-west-1) Europe (London) (eu-west-2) Europe (Milan) (eu-south-1) Europe (Paris) (eu-west-3) Europe (Spain) (eu-south-2) Europe (Stockholm) (eu-north-1)
Asia Pacific (Sydney) (ap-southeast-2)	Asia Pacific (Sydney) (ap-southeast-2) Asia Pacific (Melbourne) (ap-southeast-4)
Asia Pacific (Malaysia) (ap-southeast-5) Asia Pacific (Mumbai) (ap-south-1) Asia Pacific (Seoul) (ap-northeast-2)	Commercial AWS Regions

Session Region	Inference Regions
Asia Pacific (Singapore) (ap-southeast-1) Asia Pacific (Tokyo) (ap-north east-1) South America (São Paulo) (sa-east-1)	

For the current list of destination Regions in each geography, and for the Regions that a global inference profile can route to, see Supported Regions and models for inference profiles and Increase throughput with cross-Region inference in the *Amazon Bedrock User Guide*.

Amazon Bedrock doesn't retain your content after it serves an inference request, and doesn't use your content to train models. For more information, see Data protection in the *Amazon Bedrock User Guide*.

> **Important**
>
> You are responsible for determining whether cross-Region inference is consistent with your regulatory and data residency obligations. If you require that inference data remain within a specific Region or geography, don't turn on Replay Session in a Region that routes requests outside that Region or geography.

Destination bucket requirements

The Amazon S3 bucket that receives Replay Session artifacts must meet the following requirements:

- The bucket must be in the same AWS account as the principal that creates or updates the stack.
- The IAM identity that your agent uses to connect to the Model Context Protocol (MCP) service must have the `s3:PutObject` permission on the bucket. Replay Session writes the artifacts as that identity.

The artifacts in your bucket are encrypted with the default encryption settings of the bucket. To use your own AWS KMS key for the artifacts in your bucket, see Encrypting delivered artifacts in your bucket.

Replay Session doesn't add a signature or checksum to the artifacts that it delivers. To protect the artifacts from being changed or deleted, we recommend that you turn on versioning and Amazon S3 Object Lock for the bucket. For more information, see Retaining multiple versions of objects with S3 Versioning and Locking objects with Object Lock in the *Amazon S3 User Guide*.

Turn on Replay Session

To turn on Replay Session, set the following fields in the `AgentAccessConfig` when you call `CreateStack` or `UpdateStack`:

- `ReplaySessionEnabled`. Set to `true` to record agent sessions on the stack.
- `S3BucketArn`. The ARN of the Amazon S3 bucket that receives the artifacts.
- `KmsKeyArn` (optional). The key ARN of a customer managed key that encrypts the recording while Replay Session processes it. For more information, see Encryption at rest for Replay Session.

The following example creates a stack with Replay Session turned on:

```
aws appstream create-stack \
    --name your-stack-name \
    --agent-access-config '{
        "Settings": [
            {"AgentAction": "COMPUTER_VISION", "Permission": "ENABLED"},
            {"AgentAction": "COMPUTER_INPUT", "Permission": "ENABLED"}
        ],
        "ScreenResolution": "W_1280xH_720",
        "ScreenImageFormat": "PNG",
        "S3BucketArn": "arn:aws:s3:::amzn-s3-demo-bucket",
        "ReplaySessionEnabled": true,
        "KmsKeyArn": "arn:aws:kms:us-
east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab"
    }'
```

The identity that your agent uses must keep its Amazon S3 permissions, and its AWS KMS permissions if you use a customer managed key, for the whole session and until processing finishes

a few minutes after the session ends. If the identity loses these permissions before processing finishes, the recording can't be delivered.

Encryption at rest for Replay Session

Replay Session encrypts your recordings at rest by default. While a session is recorded and processed, Replay Session holds the screen frames, their metadata, and the finished artifacts that are waiting for delivery in storage that AWS owns. By default, this content is encrypted at rest with Amazon S3 managed keys (SSE-S3). You don't need to take any action.

Optionally, you can provide a customer managed AWS Key Management Service (AWS KMS) key. When you do, Replay Session uses the AWS Encryption SDK to encrypt each frame, metadata object, tool response record, and artifact checkpoint under your key before it stores the object. You create, own, and manage the customer managed key, and you have full control over it. For more information, see AWS KMS concepts in the *AWS Key Management Service Developer Guide*.

AWS KMS charges apply for customer managed keys. For more information, see AWS KMS pricing.

Topics

- How Replay Session uses a customer managed key
- Encrypting delivered artifacts in your bucket
- Configuring a customer managed key for Replay Session
- Monitoring Replay Session interaction with AWS KMS

How Replay Session uses a customer managed key

Replay Session uses your customer managed key to encrypt the following content while a session is recorded and processed:

- **Screen frames**. The frames captured during the session.
- **Frame metadata**. The per-action metadata for each frame.
- **Tool responses**. The result that each tool call returns.
- **Artifact checkpoints**. The transcript, session video, and HTML report that Replay Session produces after the session, while they wait for delivery.

The customer managed key protects only this content in AWS-owned storage. It isn't used for the artifacts that Replay Session delivers to your bucket. For those artifacts, see Encrypting delivered artifacts in your bucket.

Principals that use the key

Every AWS KMS request that Replay Session makes runs as your own identity, with your own permissions. Two of your principals are involved:

- **The principal that configures the stack**. The IAM identity that calls `CreateStack` or `UpdateStack`. It validates the key when you set or change it.
- **The principal that uses the stack**. The IAM identity that starts the agent session, which is the identity that your agent uses to connect to the Model Context Protocol (MCP) service. It makes every other AWS KMS request. This principal isn't necessarily the person whose screen activity is recorded.

Replay Session doesn't use grants, service roles, service-linked roles, or roles that you provide, and you don't add a service principal to your key policy.

AWS KMS operations

Replay Session makes the following AWS KMS requests with your customer managed key. Each request arrives through one of two services, which you can use as the value of the `kms:ViaService` condition key:

- **Stack configurer** requests come from the principal that configures the stack, through `appstream.region.amazonaws.com`.
- **Stack user** requests come from the principal that uses the stack, through `agentaccess-mcp.region.amazonaws.com`.

AWS KMS requests that Replay Session makes

Operation	When	Caller	Purpose
`DescribeKey`	When you call `CreateStack` or `UpdateStack` with `KmsKeyArn`	Stack configurer	Validate the key

Operation	When	Caller	Purpose
`Decrypt`	Once per session, at the first recorded action	Stack user	Confirm at the start of the session that processing will be able to decrypt the recording. No data is decrypted.
`GenerateDataKey`	During the session, three times for each action: once for the frame, once for the metadata object, and once for the tool response record	Stack user	Encrypt each object before it is stored
`Decrypt`	After the session, three times for each action: once for the frame, once for the metadata object, and once for the tool response record	Stack user	Read the recording for processing
`GenerateDataKey`	After the session, at least three times per session	Stack user	Encrypt the transcript, video, and report checkpoints
`Decrypt`	After the session, once for each object, only if processing fails	Stack user	Deliver the raw frames, metadata, tool response records, and any completed artifact checkpoints to your bucket instead of the finished artifacts, so that no captured evidence is lost

Replay Session never calls `Encrypt`, `ReEncrypt*`, or `CreateGrant`. Replay Session doesn't cache data keys, so if you disable the key or remove access to it, the change takes effect at the next request.

AWS KMS requests continue for a few minutes after a session ends, while the recording is processed. As a result, you can see `Decrypt` and `GenerateDataKey` events in CloudTrail after the agent has disconnected.

Request quotas and using a dedicated key

Each recorded action makes three `GenerateDataKey` requests during the session and three `Decrypt` requests after it: one each for the frame, the metadata object, and the tool response record. These requests count toward the AWS KMS request quotas for your AWS account, which all keys in the account and Region share. For more information, see Request quotas in the *AWS Key Management Service Developer Guide*.

We recommend that you use a dedicated key for Replay Session. That way, key policy changes that you make for other workloads can't stop Replay Session from recording.

When the key becomes unavailable

If you disable the key, schedule it for deletion, or change its policy so that Replay Session can no longer use it, the following happens:

- **During a session**. The next frame capture fails, and the agent's action fails with it. Every later action in that session also fails, because Replay Session doesn't let an agent act without recording the action. Changes to the key or its policy therefore control whether agents can act on stacks that use the key.
- **After a session**. If the key becomes unavailable before processing finishes, the recording can't be processed or delivered, and it can't be recovered.

The principal that uses the stack must keep the `kms:Decrypt` and `kms:GenerateDataKey` permissions on the key, and the `s3:PutObject` permission on the destination bucket, for the whole session and until processing finishes. Keep the key enabled, and keep its permissions in place, until every session that used it has finished processing.

Encrypting delivered artifacts in your bucket

Replay Session writes the delivered artifacts to your Amazon S3 bucket as the principal that uses the stack. The artifacts are encrypted with the default encryption of the bucket, which can be SSE-S3. The customer managed key that you set for Replay Session isn't used for these artifacts.

To encrypt the delivered artifacts with your own key, do the following:

1. Set SSE-KMS with an Amazon S3 Bucket Key as the default encryption for the bucket. For more information, see Setting default server-side encryption behavior for Amazon S3 buckets and Reducing the cost of SSE-KMS with Amazon S3 Bucket Keys in the *Amazon S3 User Guide*.
2. Give the principal that uses the stack the following permissions on the bucket's KMS key, through Amazon S3:
   - `kms:GenerateDataKey`, to write each artifact.
   - `kms:Decrypt`, to write artifacts larger than 8 MB. Larger artifacts are uploaded in multiple parts, and multipart uploads to an SSE-KMS bucket require `kms:Decrypt`.

Configuring a customer managed key for Replay Session

Replay Session supports a symmetric encryption KMS key that is enabled and is in the same Region as the stack. Keys in custom key stores and external key stores are supported. You can choose a Multi-Region key. The primary key or replica key that you choose must be in the same Region as the stack. Replay Session doesn't support the multi-Region features of the key.

To encrypt Replay Session recordings with a customer managed key

1. Create a symmetric encryption KMS key in the same Region as the stack, or choose an existing one. For more information, see Create a symmetric encryption KMS key in the *AWS Key Management Service Developer Guide*.
2. Add key policy statements that allow the principal that uses the stack to encrypt and read recordings, and allow the principal that configures the stack to validate the key. For details, see Permissions for the customer managed key.
3. Attach IAM policies to both principals.
4. Set `KmsKeyArn` in the `AgentAccessConfig` when you call `CreateStack` or `UpdateStack`. For details, see Setting the key on a stack.
5. Verify in CloudTrail that the expected `DescribeKey`, `GenerateDataKey`, and `Decrypt` events appear. For details, see Monitoring Replay Session interaction with AWS KMS.

Permissions for the customer managed key

Key policies control access to your customer managed key. Every customer managed key must have exactly one key policy, which contains statements that determine who can use the key and how they can use it. For more information, see KMS key access and permissions in the *AWS Key Management Service Developer Guide*.

Which policies you need depends on where the key is:

- If the key is in the same account as the stack, and its key policy includes the default statement that allows IAM policies to grant access to the key, the IAM policies in this section are enough.
- If the key is in another account, or its key policy doesn't allow IAM policies to grant access, also add the key policy statements in this section.

Key policy statements

The following example statements give the principal that uses the stack (`ReplaySessionUserRole`) and the principal that configures the stack (`ReplaySessionAdminRole`) the permissions they need. Add these statements to your existing key policy. Don't replace the existing statements.

```
{
    "Sid": "AllowAgentAccessDirectKMS",
    "Effect": "Allow",
    "Principal": { "AWS": "arn:aws:iam::111122223333:role/ReplaySessionUserRole" },
    "Action": [
        "kms:GenerateDataKey",
        "kms:Decrypt"
    ],
    "Resource": "*",
    "Condition": {
        "StringEquals": {
            "kms:ViaService": "agentaccess-mcp.us-east-1.amazonaws.com"
        },
        "StringLike": {
            "kms:EncryptionContext:aws:agentaccess-mcp:session_id": "*"
        }
    }
},
{
    "Sid": "AllowAppstreamToDescribeKey",
    "Effect": "Allow",
    "Principal": { "AWS": "arn:aws:iam::111122223333:role/ReplaySessionAdminRole" },
    "Action": "kms:DescribeKey",
    "Resource": "*",
    "Condition": {
        "StringEquals": {
            "kms:ViaService": "appstream.us-east-1.amazonaws.com"
```

```
        }
    }
}
```

- `AllowAgentAccessDirectKMS` lets the principal that uses the stack encrypt and read the recording. The `kms:ViaService` condition limits these operations to requests that come through Replay Session. The `kms:EncryptionContext` condition requires the session encryption context to be present, with any value, because the value is a session ID that is assigned when the session starts.
- `AllowAppstreamToDescribeKey` lets the principal that configures the stack validate the key when it creates or updates the stack. This statement has no encryption context condition, because `DescribeKey` doesn't take an encryption context.

Because each statement names a specific role as its principal, you don't need a `kms:CallerAccount` condition. If more than one role starts sessions or configures the stack, list each role in the matching statement.

IAM policy for the principal that uses the stack

Attach the following statement to the IAM identity that your agent uses to start sessions.

```
{
    "Sid": "AllowReplaySessionFrameEncryption",
    "Effect": "Allow",
    "Action": [
        "kms:GenerateDataKey",
        "kms:Decrypt"
    ],
    "Resource": "arn:aws:kms:us-
east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab",
    "Condition": {
        "StringEquals": {
            "kms:ViaService": "agentaccess-mcp.us-east-1.amazonaws.com"
        }
    }
}
```

IAM policy for the principal that configures the stack

Attach the following statements to the IAM identity that calls `CreateStack` or `UpdateStack`.

```
{
    "Sid": "AllowReplaySessionKeyValidation",
    "Effect": "Allow",
    "Action": "kms:DescribeKey",
    "Resource": "arn:*:kms:*:*:key/*"
},
{
    "Sid": "AllowReplaySessionConsoleKeyListing",
    "Effect": "Allow",
    "Action": "kms:ListAliases",
    "Resource": "*"
}
```

These statements don't include a `kms:ViaService` condition, because the console lists your keys with `kms:ListAliases` and validates the key that you select with `kms:DescribeKey` directly, without a `kms:ViaService` value. If you configure stacks only through the API, you can remove the `AllowReplaySessionConsoleKeyListing` statement, limit `Resource` to your key, and add a `kms:ViaService` condition with the value `appstream.region.amazonaws.com`.

For more information, see Specifying KMS keys in IAM policy statements and Troubleshooting key access in the *AWS Key Management Service Developer Guide*.

Using a key in another account

To use a key in another account, configure the stack through the API. Both the key policy and the IAM policy of the principal that configures the stack must allow `kms:DescribeKey`.

Setting the key on a stack

Set the key with the `KmsKeyArn` field in the `AgentAccessConfig` of the `CreateStack` and `UpdateStack` API operations. You can set `KmsKeyArn` whether or not `ReplaySessionEnabled` is `true`.

`KmsKeyArn` accepts a key ARN only. If you provide an alias ARN, an alias name, or a key ID, the request is rejected before Replay Session contacts AWS KMS.

When `KmsKeyArn` is present, the service calls `DescribeKey` as the principal that makes the request, to confirm that the key exists, is enabled, is a symmetric encryption KMS key, and is in the same Region as the stack.

The following example adds a customer managed key to an existing stack:

```
aws appstream update-stack \
    --name your-stack-name \
    --agent-access-config '{
        "S3BucketArn": "arn:aws:s3:::amzn-s3-demo-bucket",
        "ReplaySessionEnabled": true,
        "KmsKeyArn": "arn:aws:kms:us-
east-1:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab"
    }'
```

Changing the key

To change the key, call `UpdateStack` with a new `KmsKeyArn`. The service validates the new key the same way that it validates a key when you create a stack. Configuration is fixed for each session, so sessions that are already running keep the key that they started with, and new sessions use the new key. Replay Session doesn't re-encrypt content.

Content recorded under the previous key stays encrypted under that key. Keep the previous key enabled, and keep its key policy and IAM permissions in place, until every session that started under it has finished processing.

Scoping down access to the key

Encryption context

An encryption context is an optional set of key-value pairs that contain additional contextual information about the data. AWS KMS uses the encryption context as additional authenticated data to support authenticated encryption. The encryption context also appears in CloudTrail logs, so you can use it to identify how your key is being used.

Every `GenerateDataKey` and `Decrypt` request that Replay Session makes includes the following encryption context pair. The key is `aws:agentaccess-mcp:session_id`, and the value is the ID of the session being recorded.

```
"encryptionContext": {
    "aws:agentaccess-mcp:session_id": "a1b2c3d4-5678-90ab-cdef-EXAMPLE11111"
}
```

The `aws:` prefix is part of the literal name of the encryption context key. It isn't a global IAM condition key.

Requests made through the AWS Encryption SDK, which are all requests except the `Decrypt` request at the start of the session, also include an `aws-crypto-public-key` encryption context pair that the AWS Encryption SDK adds.

When you write your own encryption context conditions, require the `aws:agentaccess-mcp:session_id` key and allow the `aws-crypto-public-key` key. Session IDs are assigned when each session starts, and every session on the stack uses the same key, so you can't name a specific session ID in a policy. Instead, require that the session ID key is present, as the `AllowAgentAccessDirectKMS` statement does with `"StringLike": { "kms:EncryptionContext:aws:agentaccess-mcp:session_id": "*" }`. Don't use a condition that requires the encryption context to contain only the session ID pair, or requests that the AWS Encryption SDK makes are denied.

kms:ViaService

The kms:ViaService condition key limits use of the key to requests that come from a specific AWS service. Replay Session uses the following values:

- `agentaccess-mcp.region.amazonaws.com` for `GenerateDataKey` and `Decrypt` requests made to record, process, and deliver the recording.
- `appstream.region.amazonaws.com` for `DescribeKey` requests made to validate the key.

Confused deputy protection

The `aws:SourceArn` and `aws:SourceAccount` condition keys don't apply to Replay Session. Every AWS KMS request runs as your own identity, and Replay Session doesn't use a service role, so there is no cross-service request for these keys to constrain. The key policy scopes access by naming the two roles.

Monitoring Replay Session interaction with AWS KMS

When you use a customer managed key with Replay Session, you can use CloudTrail to track the requests that Replay Session sends to AWS KMS. For more information, see Logging AWS KMS API calls with AWS CloudTrail in the *AWS Key Management Service Developer Guide*.

Look for the following events:

- `DescribeKey`, made by the principal that configures the stack when the stack is created or updated, with the `kms:ViaService` value `appstream.region.amazonaws.com`.

- `GenerateDataKey` and `Decrypt`, made by the principal that uses the stack, with the `kms:ViaService` value `agentaccess-mcp.region.amazonaws.com` and the `aws:agentaccess-mcp:session_id` encryption context. These events continue for a few minutes after the agent disconnects.

Events appear in the trail of the account that makes the request. For a key in another account, the events also appear in the account that owns the key.

Use Amazon Bedrock AgentCore

You can build and run agents with Amazon Bedrock AgentCore that operate desktop and mainframe applications through WorkSpaces Applications agent access. AgentCore provides a managed environment to create, test, and run your agent, and agent access gives each agent a secure Windows or Linux desktop with built-in monitoring, auditing, and governance. There are no extra components to build: an AgentCore agent connects to the managed MCP server the same way any other agent does.

You can start in the AgentCore Harness Playground and run your first agent without writing code. The playground runs the same managed loop as the production harness. After you test an agent, you can promote it, along with its model, system prompt, tools, skills, memory, and limits, to a versioned harness that runs behind an endpoint and supports policy controls. AgentCore Gateway and AgentCore Identity connect agents to your identity provider, while each agent's desktop stays inside your corporate network. After you deploy, you can use AgentCore evaluations and cost management features to improve accuracy and reduce cost.

Connect an agent to a desktop from AgentCore

To connect an agent to a WorkSpaces Applications desktop, add a remote MCP server tool to your agent in the Harness Playground, then configure it for agent access:

1. Set the stack and user ID headers so that the agent connects to your non-domain-joined stack. For more information about identifying a session, see Session identification.
2. Set the outbound authorization to either IAM or OAuth. When you use IAM, the request is signed with SigV4 so that WorkSpaces Applications can verify it. For more information about the authorization options, see Request authorization.

The agent connects to the MCP server at the endpoint described in Connecting to the MCP server. The other steps in the Harness Playground, such as adding tools and setting your prompts, are the same as for any other agent.

Get started providing agents with access to WorkSpaces Applications

To enable AI agents to operate desktop applications through Amazon WorkSpaces Applications, you create a stack with access enabled for agents, generate a streaming URL, and connect your agent to the managed MCP server. You can also enable agent access on an image builder, so that an agent can install and test applications before you create an image.

If you are setting up agent access for the first time, you can use the Build Your First Agent quick-start experience. You can also follow the steps in this topic to configure agent access manually. The quick-start experience is available in the GitHub repository — sample-code-for-workspaces-agent-access.

This tutorial takes approximately 15 minutes to complete.

> **Important**
>
> The resources you create in this tutorial might result in charges to your AWS account. Delete the stack and stop the fleet when you are done to avoid ongoing charges.

In this tutorial, you complete the following tasks:

- Step 1: Create a stack or image builder with agent access enabled
- Step 2: Generate a streaming URL
- Step 3: Connect your agent to the MCP server
- Step 4: Verify agent activity
- Step 5: Change agent access settings or disable agent access
- Step 6: Clean up resources

Write agent code with the Agent Toolkit for AWS

If you build your agent with an AI coding assistant, you can use the Agent Toolkit for AWS. It helps you set up agent access and build your first agent.

The toolkit uses the AWS MCP Server. This managed server gives agents safe access to AWS through the Model Context Protocol (MCP). Agents can run AWS CLI commands, search AWS docs, and follow ready-made skills. They connect through one secure endpoint. You also get CloudWatch metrics and IAM access controls. To set up the server for your agent, see Setting up the AWS MCP Server in the *Agent Toolkit for AWS User Guide*.

The toolkit includes the following skills for WorkSpaces Applications:

`setting-up-workspaces-applications`

Sets up a WorkSpaces Applications desktop for an AI agent or for people. Your agent asks about your workflow one question at a time. It then recommends image, fleet, and stack settings, and shows you a plan that you can approve to create the resources. It can also build a custom image with an image builder.

`amazon-workspaces-agent-access`

Connects your agent to the MCP server and helps it operate the desktop, following best practices for agent access. It also helps you fix connection errors. The skill includes three sample workflows that you can run in your own AWS account to see agent access working.

After you set up the AWS MCP Server, your agent can find and load these skills when it needs them. To install the skills for the AI coding agents on your system instead, run the following commands. The `add-skill` command requires AWS CLI version 2.35.0 or later.

```
aws agent-toolkit add-skill \
    --skill-name setting-up-workspaces-applications

aws agent-toolkit add-skill \
    --skill-name amazon-workspaces-agent-access
```

To get started, ask your agent "Help me complete the getting started WorkSpaces agent access tutorial".

Prerequisites

Before you begin, review the agent access requirements. For the account, fleet, permissions, and agent prerequisites, see Prerequisites.

Step 1: Create a stack or image builder with agent access enabled

Create a WorkSpaces Applications stack or image builder with agent access enabled to allow AI agents to interact with desktop applications.

Create a stack using the AWS Management Console

To create a stack with agent access

1. Open the WorkSpaces Applications console.
2. In the left navigation pane, choose **Stacks**, then choose **Create Stack**.
3. On the **Stack details** page (step 1 of 4), under **AI agent access**, select **Enable AI agent access**. Choose **Next**.
4. On the **Enable storage** page (step 2 of 4), optionally enable **Home folders** to allow your agent to save files to an Amazon S3 bucket in your AWS account. The fleet associated with this stack must allow access to Amazon S3 through the internet or a Amazon VPC endpoint for Amazon S3. Choose **Next**.
5. On the **Edit agent settings** page (step 3 of 4), configure the following:

   - **Enable computer input** — Allow agents to choose buttons, enter text, and scroll on the desktop. If you enable computer input, you must also enable computer vision.
   - **Enable computer vision** — Allow agents to see the desktop.
   - **Screenshot storage** — Configure where agent screenshots are stored during streaming sessions. If enabled, provide an Amazon S3 bucket that you have permissions to write to.
   - **Screen resolution** — Select the display resolution for the agent streaming environment (1280x720).
   - **Screen image type** — Select the image format for agent screen captures (PNG or JPEG).
   - **Application settings persistence** — Optionally enable this to save your agent's application customizations and Windows settings between sessions. Settings are saved to an Amazon S3 bucket in your AWS account. This setting is available on Windows and is not supported on Linux.

- **Enable MCP tool forwarding** — Allow agents to interact with applications and the desktop operating system through direct MCP calls rather than using computer use tools. When enabled, MCP tools configured on the WorkSpaces application session are forwarded to the agent. Tool forwarding also requires a configuration file on the image. For the platform-specific path and setup steps, see Setting up tool forwarding.
- **Enable user control mode** — Allow users to observe agent sessions in real time through their browser. Observers see a live view of the desktop as the agent works. In VIEW_STOP mode, observers can stop the agent using a button at the top of the screen. After the agent stops, it must start a new session to resume.

> **Note**
>
> You must enable at least one of computer input or computer vision.

Choose **Next**.

6. On the **Review and Create** page (step 4 of 4), review your settings and choose **Create Stack**.

Create a stack using the AWS CLI

Run the following command to create a stack with agent access enabled:

```
aws appstream create-stack \
    --name your-stack-name \
    --agent-access-config '{
        "Settings": [
            {"AgentAction": "COMPUTER_VISION", "Permission": "ENABLED"},
            {"AgentAction": "COMPUTER_INPUT", "Permission": "ENABLED"},
            {"AgentAction": "FORWARD_MCP_TOOLS", "Permission": "ENABLED"}
        ],
        "ScreenResolution": "W_1280xH_720",
        "ScreenImageFormat": "PNG"
    }'
```

To also enable screenshot storage, add the `S3BucketArn` and `ScreenshotsUploadEnabled` parameters:

```
aws appstream create-stack \
```

```
    --name your-stack-name \
    --agent-access-config '{
        "Settings": [
            {"AgentAction": "COMPUTER_VISION", "Permission": "ENABLED"},
            {"AgentAction": "COMPUTER_INPUT", "Permission": "ENABLED"},
            {"AgentAction": "FORWARD_MCP_TOOLS", "Permission": "ENABLED"}
        ],
        "ScreenResolution": "W_1280xH_720",
        "ScreenImageFormat": "PNG",
        "S3BucketArn": "arn:aws:s3:::your-bucket-name",
        "ScreenshotsUploadEnabled": true
    }'
```

After you create the stack, associate it with a fleet. Agents cannot connect to a stack that does not have an associated fleet.

```
aws appstream associate-fleet \
    --stack-name your-stack-name \
    --fleet-name your-fleet-name
```

Enable agent access on an image builder using the AWS CLI

With agent access on an image builder, an agent can install, configure, and test applications on the image builder. You then create an image from the image builder and use that image for the fleet that you associate with your stack. If you don't have an image for your fleet yet, you can complete this step first.

You enable agent access when you create the image builder. Domain-joined image builders aren't supported. For more information, see Limitations.

Run the following command to create an image builder with agent access enabled:

```
aws appstream create-image-builder \
    --name your-image-builder-name \
    --instance-type your-instance-type \
    --image-name your-base-image-name \
    --agent-access-config '{
        "Settings": [
            {"AgentAction": "COMPUTER_VISION", "Permission": "ENABLED"},
            {"AgentAction": "COMPUTER_INPUT", "Permission": "ENABLED"},
            {"AgentAction": "FORWARD_MCP_TOOLS", "Permission": "ENABLED"}
```

```
        ],
        "ScreenResolution": "W_1280xH_720",
        "ScreenImageFormat": "PNG",
        "UserControlMode": "VIEW_ONLY",
        "ScreenshotsUploadEnabled": false
    }'
```

Instead of `--image-name`, you can specify the base image with `--image-arn`.

The `Settings`, `ScreenResolution`, and `ScreenImageFormat` members are required, and at least one action in `Settings` must be `ENABLED`. `UserControlMode` is optional. For a description of each setting, see Session behavior.

To confirm that agent access is enabled, run the following command. The response includes an `AgentAccessConfig` object with your settings.

```
aws appstream describe-image-builders \
    --names your-image-builder-name
```

Step 2: Generate a streaming URL

Generate a streaming URL for your stack or image builder. You don't need agent-specific parameters. The stack's or image builder's agent access configuration determines the agent-specific behavior.

Stack using the AWS Management Console

To generate a streaming URL using the console

1. Open the WorkSpaces Applications console.
2. In the left navigation pane, choose **Stacks**, then choose the stack you created with agent access enabled.
3. Choose the **Actions** button, and in the dropdown, select **Create Streaming URL**. Your stack must be selected for this option to be available.
4. In the **UserID** section, enter a user. You can enter `TestUser` if you are testing.
5. In the **URL expiration** section, select the time you want the URL to be valid. A shorter time is recommended. The default is 30 minutes, which is recommended for testing.
6. Choose **GetURL** and copy the URL generated.

Stack using the AWS CLI

Create a streaming URL using the standard WorkSpaces Applications `CreateStreamingURL` API. Run the following command to generate a streaming URL:

```
aws appstream create-streaming-url \
    --stack-name your-stack-name \
    --fleet-name your-fleet-name \
    --user-id your-agent-id \
    --validity 3600
```

The response includes a `StreamingURL` that you pass to your agent in the next step. The URL is valid for the duration specified by the `--validity` parameter.

> **Note**
>
> At any given time, only one agent can connect to a unique session. Named users, specified through the `UserId` parameter, can have only one active session per fleet at a time. To run multiple agents concurrently, each agent must connect to its own unique session.

Image builder using the AWS CLI

If you use an image builder, you can generate a streaming URL with the following command. Or you can skip this step and connect with the image builder name. For both options, see Connecting with image builders.

```
aws appstream create-image-builder-streaming-url \
    --name your-image-builder-name
```

Step 3: Connect your agent to the MCP server

Your agent connects to the managed MCP server at the following fixed endpoint:

`https://agentaccess-mcp.region.api.aws/mcp`

Replace *region* with your Region code. For the supported Regions and their endpoints, see WorkSpaces Applications MCP endpoints.

The connection uses SigV4 signing with the service name `agentaccess-mcp`. Your agent signs each request using the AWS credentials configured in the prerequisites with the `agentaccess-mcp:*` permissions. You pass the streaming URL from Step 2 as a header on every MCP request.

The following example shows how to establish the connection using mcp-proxy-for-aws:

```
aws_iam_streamablehttp_client(
    endpoint="https://agentaccess-mcp.region.api.aws/mcp",
    aws_service="agentaccess-mcp",
    aws_region="region",
    headers={
        "x-amzn-agentaccess-streaming-session-url": streaming_url,
    },
)
```

After the agent connects, it can use the MCP tools to enter text, choose buttons, and take screenshots of the desktop.

If you connect to an image builder, you must also pass the `x-amzn-agentaccess-local-user` header with the built-in user that the agent logs in as. For the headers and examples, see Connecting with image builders.

Built-in users on image builders

An image builder has three built-in Windows users. The agent works as one of them:

- `ADMINISTRATOR`: The local administrator. Use this user to install software and create the image.
- `TEMPLATE_USER`: A standard user whose profile becomes the default profile for your users.
- `TEST_USER`: A standard user for testing applications.

An image builder doesn't log in a user automatically. The agent chooses the user when it connects, and it can switch to another user with the `switch_user` tool. For more information, see User tools.

For example, an agent can prepare and test an application as follows:

1. Connect as `ADMINISTRATOR`.
2. Install the application.
3. Call `switch_user` to switch to `TEST_USER`.

4. Open the application and check that it works.

Step 4: Verify agent activity

You can verify agent activity using the following AWS services:

- **AWS CloudTrail** — AWS CloudTrail logs agent session events. Open the AWS CloudTrail console to view agent activity. Tool calls are logged as data events, which requires a trail configured to log data events.
- **CloudWatch** — CloudWatch provides operational metrics for agent sessions. Open the CloudWatch console to view metrics.
- **Amazon S3** — If you enabled screenshot storage, Amazon S3 stores the screenshots in the bucket you specified during stack configuration.

Step 5: Change agent access settings or disable agent access

This step changes the saved agent access settings on a stack or an image builder. To switch the Windows user that is logged in to an image builder, the agent uses the `switch_user` tool instead. For more information, see User tools.

To change agent access settings on a stack

1. Run the following command with only the fields that you want to change. This example changes the user control mode.

   ```
   aws appstream update-stack \
       --name your-stack-name \
       --agent-access-config '{
           "UserControlMode": "VIEW_STOP"
       }'
   ```

   Fields that you don't include keep their current values. After the update, at least one action must be enabled.
2. Start a new streaming session. The changes apply to new sessions, not to sessions that are already running.

To disable agent access on a stack, run the following command:

```
aws appstream update-stack \
    --name your-stack-name \
    --attributes-to-delete AGENT_ACCESS_CONFIG
```

You can change agent access settings on an image builder only while it is stopped.

To change agent access settings on an image builder

1. Stop the image builder.

   ```
   aws appstream stop-image-builder \
       --name your-image-builder-name
   ```

2. Run the following command until the `State` value in the response is `STOPPED`.

   ```
   aws appstream describe-image-builders \
       --names your-image-builder-name
   ```

3. Run the following command with only the fields that you want to change. This example enables MCP tool forwarding.

   ```
   aws appstream update-image-builder \
       --name your-image-builder-name \
       --agent-access-config '{
           "Settings": [
               {"AgentAction": "FORWARD_MCP_TOOLS", "Permission": "ENABLED"}
           ]
       }'
   ```

   The same rules apply as for stacks. If you set `ScreenshotsUploadEnabled` to `false`, the `S3BucketArn` value is also removed.

4. Start the image builder. The changes apply at the next session.

   ```
   aws appstream start-image-builder \
       --name your-image-builder-name
   ```

To disable agent access on an image builder, stop the image builder as described in the previous procedure, and then run the following command. After you start the image builder again, you can work on it yourself in a browser.

```
aws appstream update-image-builder \
    --name your-image-builder-name \
    --attributes-to-delete AGENT_ACCESS_CONFIG
```

Step 6: Clean up resources

To avoid ongoing charges, delete the stack you created in this tutorial. You must stop the fleet and disassociate it from the stack before you can delete the stack. Optionally, you can also delete the fleet.

If you created an image builder, stop it when you are done to avoid ongoing charges. You can also delete it. For the commands, see Using the AWS CLI.

Using the AWS Management Console

To clean up resources

1. Open the WorkSpaces Applications console.
2. In the left navigation pane, choose **Fleets**.
3. Select the fleet associated with the stack. Choose **Actions**, **Stop**. Wait for the fleet to stop.
4. In the left navigation pane, choose **Stacks**.
5. Select the stack you created, and choose **Actions**, **Disassociate Fleet**.
6. With the stack still selected, choose **Actions**, **Delete**.
7. (Optional) To delete the fleet, in the left navigation pane, choose **Fleets**. Select the fleet and choose **Actions**, **Delete**.

Using the AWS CLI

Run the following commands to clean up resources:

```
aws appstream stop-fleet \
    --name your-fleet-name

aws appstream disassociate-fleet \
    --stack-name your-stack-name \
    --fleet-name your-fleet-name

aws appstream delete-stack \
```

```
    --name your-stack-name
```

(Optional) To also delete the fleet after it has stopped:

```
aws appstream delete-fleet \
    --name your-fleet-name
```

If you created an image builder, stop it when you are done. A stopped image builder isn't billed for the instance.

```
aws appstream stop-image-builder \
    --name your-image-builder-name
```

(Optional) To delete the image builder:

```
aws appstream delete-image-builder \
    --name your-image-builder-name
```

Networking and Access for Amazon WorkSpaces Applications

The following topics provide information about enabling users to connect to WorkSpaces Applications streaming instances (fleet instances) and enabling your WorkSpaces Applications fleets, image builders, and app block builders to access network resources and the internet.

Contents

- Internet Access
- Configure a VPC for WorkSpaces Applications
- Using Amazon S3 VPC Endpoints for WorkSpaces Applications Features
- Amazon WorkSpaces Applications Connections to Your VPC
- User Connections to Amazon WorkSpaces Applications

Internet Access

If your fleets, app block builders, and image builders require internet access, you can enable internet access in several ways. When you choose a method for enabling internet access, consider the number of users your deployment must support and your deployment goals. For example:

- If your deployment must support more than 100 concurrent users, configure a VPC with private subnets and a NAT gateway.
- If your deployment supports fewer than 100 concurrent users, you can configure a new or existing VPC with a public subnet.
- If your deployment supports fewer than 100 concurrent users and you are new to WorkSpaces Applications and want to get started using the service, you can use the default VPC, public subnet, and security group.

The following sections provide more information about each of these deployment options.

- Configure a VPC with Private Subnets and a NAT Gateway (recommended) — With this configuration, you launch your fleets, app block builders, and image builders in a private subnet and configure a NAT gateway in a public subnet in your VPC. Your streaming instances are assigned a private IP address that is not directly accessible from the internet.

In addition, unlike configurations that use the **Default Internet Access** option for enabling internet access, the NAT configuration is not limited to 100 fleet instances. If your deployment must support more than 100 concurrent users, use this configuration.

You can create and configure a new VPC to use with a NAT gateway, or add a NAT gateway to an existing VPC.

- Configure a New or Existing VPC with a Public Subnet — With this configuration, you launch your fleets, app block builders, and image builders in a public subnet and enable **Default Internet Access**. When you enable this option, WorkSpaces Applications uses the internet gateway in your Amazon VPC public subnet to provide the internet connection. Your streaming instances are assigned a public IP address that is directly accessible from the internet. You can create a new VPC or configure an existing one for this purpose.

  > **Note**
  >
  > When **Default Internet Access** is enabled, a maximum of 100 fleet instances is supported. If your deployment must support more than 100 concurrent users, use the NAT gateway configuration instead.

- Use the Default VPC, Public Subnet, and Security Group — If you are new to WorkSpaces Applications and want to get started using the service, you can launch your fleets, app block builders, and image builders in a default public subnet and enable **Default Internet Access**. When you enable this option, WorkSpaces Applications uses the internet gateway in your Amazon VPC public subnet to provide the internet connection. Your streaming instances are assigned a public IP address that is directly accessible from the internet.

  Default VPCs are available in Amazon Web Services accounts created after 2013-12-04.

  The default VPC includes a default public subnet in each Availability Zone and an internet gateway that is attached to your VPC. The VPC also includes a default security group.

  > **Note**
  >
  > When **Default Internet Access** is enabled, a maximum of 100 fleet instances is supported. If your deployment must support more than 100 concurrent users, use the NAT gateway configuration instead.

Configure a VPC for WorkSpaces Applications

When you set up WorkSpaces Applications, you must specify the virtual private cloud (VPC) and the subnets in which to launch your fleet instances and image builders. Fleets require at least two subnets in different Availability Zones; image builders require at least one subnet. A VPC is a virtual network in your own logically isolated area within the Amazon Web Services Cloud. A subnet is a range of IP addresses in your VPC.

When you configure your VPC for WorkSpaces Applications, you can specify either public or private subnets, or a mix of both types of subnets. A public subnet has direct access to the internet through an internet gateway. A private subnet, which doesn't have a route to an internet gateway, requires a Network Address Translation (NAT) gateway or NAT instance to provide access to the internet.

Contents

- VPC Setup Recommendations
- Configure a VPC with Private Subnets and a NAT Gateway
- Configure a New or Existing VPC with a Public Subnet
- Use the Default VPC, Public Subnet, and Security Group

VPC Setup Recommendations

When you create a fleet, or launch an image builder or app block builder, you specify the VPC and one or more subnets to use. You can provide additional access control to your VPC by specifying security groups.

The following recommendations can help you configure your VPC more effectively and securely. In addition, they can help you configure an environment that supports effective fleet scaling. With effective fleet scaling, you can meet current and anticipated WorkSpaces Applications user demand, while avoiding unecessary resource usage and associated costs.

Overall VPC Configuration

- Make sure that your VPC configuration can support your fleet scaling needs.

  As you develop your plan for fleet scaling, keep in mind that one user requires one fleet instance. Therefore, the size of your fleet determines the number of users who can stream concurrently. For this reason, for each instance type that you plan to use, make sure that the number of fleet

instances that your VPC can support is greater than the number of anticipated concurrent users for the same instance type.

- Make sure that your WorkSpaces Applications account quotas (also referred to as limits) are sufficient to support your anticipated demand. To request a quota increase, you can use the Service Quotas console at https://console.aws.amazon.com/servicequotas/. For information about default WorkSpaces Applications quotas, see Amazon WorkSpaces Applications Service Quotas.
- If you plan to provide your streaming instances (fleet instances, app block builder, or image builders) with access to the internet, we recommend that you configure a VPC with two private subnets for your streaming instances and a NAT gateway in a public subnet.

  The NAT gateway lets the streaming instances in your private subnets connect to the internet or other AWS services. However, it prevents the internet from initiating a connection with those instances. In addition, unlike configurations that use the **Default Internet Access** option for enabling internet access, the NAT configuration supports more than 100 fleet instances. For more information, see Configure a VPC with Private Subnets and a NAT Gateway.

Elastic Network Interfaces

- WorkSpaces Applications creates as many elastic network interfaces (network interfaces) as the maximum desired capacity of your fleet. By default, the limit for network interfaces per Region is 5000.

  When planning capacity for very large deployments, for example, thousands of streaming instances, consider the number of EC2 instances that are also used in the same Region.

Subnets

- If you are configuring more than one private subnet for your VPC, configure each in a different Availability Zone. Doing so increases fault tolerance and can help prevent insufficient capacity errors. If you use two subnets in the same AZ, you might run out of IP addresses, because WorkSpaces Applications will not use the second subnet.
- Make sure that the network resources required for your applications are accessible through both of your private subnets.
- Configure each of your private subnets with a subnet mask that allows for enough client IP addresses to account for the maximum number of expected concurrent users. In addition, allow

for additional IP addresses to account for anticipated growth. For more information, see VPC and Subnet Sizing for IPv4.

- If you are using a VPC with NAT, configure at least one public subnet with a NAT Gateway for internet access, preferably two. Configure the public subnets in the same Availability Zones where your private subnets reside.

  To enhance fault tolerance and reduce the chance of insufficient capacity errors for large WorkSpaces Applications fleet deployments, consider extending your VPC configuration into a third Availability Zone. Include a private subnet, public subnet, and NAT gateway in this additional Availability Zone.
- If you enable auto assign IPV6 option for your subnet then Elastic Network Interface for your instances will be auto assigned with a global IPV6 address. For more information, see modify-subnet.
- Enabling default internet access is only applicable for IPV4 addresses in IPv4 only or dual-stack subnets. To allow internet access for IPV6 addresses add an internet gateway or egress only gateway. For more information, see egress-only-internet-gateway.

> **Note**
>
> By default IPV6 addresses are globally addressable. If your subnet has an internet gateway and appropriate subnet groups and acl allowing IPV6 traffic rules your streaming instances can be connected to the internet with IPV6 address.

Security Groups

- Use security groups to provide additional access control to your VPC.

  Security groups that belong to your VPC let you control the network traffic between WorkSpaces Applications streaming instances and network resources required by applications. These resources may include other AWS services such as Amazon RDS or Amazon FSx, license servers, database servers, file servers, and application servers.
- Make sure that the security groups provide access to the network resources that your applications require.

For more information about configuring security groups for WorkSpaces Applications, see Security Groups in Amazon WorkSpaces Applications. For general information about security groups, see Security Groups for Your VPC in the *Amazon VPC User Guide*.

Configure a VPC with Private Subnets and a NAT Gateway

If you plan to provide your streaming instances (fleet instances, app block builders, and image builders) with access to the internet, we recommend that you configure a VPC with two private subnets for your streaming instances and a NAT gateway in a public subnet. You can create and configure a new VPC to use with a NAT gateway, or add a NAT gateway to an existing VPC. For additional VPC configuration recommendations, see VPC Setup Recommendations.

The NAT gateway lets the streaming instances in your private subnets connect to the internet or other AWS services, but prevents the internet from initiating a connection with those instances. In addition, unlike configurations that use the **Default Internet Access** option for enabling internet access for WorkSpaces Applications streaming instances, this configuration is not limited to 100 fleet instances.

For information about using NAT Gateways and this configuration, see NAT Gateways and VPC with Public and Private Subnets (NAT) in the *Amazon VPC User Guide*.

Contents

- Create and Configure a New VPC
- Add a NAT Gateway to an Existing VPC
- Enable Internet Access for Your Fleet, Image Builder, or App Block Builder in Amazon WorkSpaces Applications

Create and Configure a New VPC

This topic describes how to use the VPC wizard to create a VPC with a public subnet and one private subnet. As part of this process, the wizard creates an internet gateway and a NAT gateway. It also creates a custom route table associated with the public subnet and updates the main route table associated with the private subnet. The NAT gateway is automatically created in the public subnet of your VPC.

After you use the wizard to create the initial VPC configuration, you'll add a second private subnet. For more information about this configuration, see VPC with Public and Private Subnets (NAT) in the *Amazon VPC User Guide*.

> **Note**
>
> If you already have a VPC, complete the steps in Add a NAT Gateway to an Existing VPC instead.

Contents

- Step 1: Allocate an Elastic IP Address
- Step 2: Create a New VPC
- Step 3: Add a Second Private Subnet
- Step 4: Verify and Name Your Subnet Route Tables

Step 1: Allocate an Elastic IP Address

Before you create your VPC, you must allocate an Elastic IP address in your WorkSpaces Applications Region. You must first allocate an Elastic IP address for use in your VPC, and then associate it with your NAT gateway. For more information, see Elastic IP Addresses in the *Amazon VPC User Guide*.

> **Note**
>
> Charges may apply to Elastic IP addresses that you use. For more information, see Elastic IP Addresses on the Amazon EC2 pricing page.

Complete the following steps if you don't already have an Elastic IP address. If you want to use an existing Elastic IP address, verify that it's not currently associated with another instance or network interface.

To allocate an Elastic IP address

1. Open the Amazon EC2 console at https://console.aws.amazon.com/ec2/.
2. In the navigation pane, under **Network & Security**, choose **Elastic IPs**.

3. Choose **Allocate New Address**, and then choose **Allocate**.
4. Note the Elastic IP address.
5. In the upper right of the **Elastic IPs** pane, click the X icon to close the pane.

Step 2: Create a New VPC

Complete the following steps to create a new VPC with a public subnet and one private subnet.

To create a new VPC

1. Open the Amazon VPC console at https://console.aws.amazon.com/vpc/.
2. In the navigation pane, choose **VPC Dashboard**.
3. Choose **Launch VPC Wizard**.
4. In **Step 1: Select a VPC Configuration**, choose **VPC with Public and Private Subnets**, and then choose **Select**.
5. In **Step 2: VPC with Public and Private Subnets**, configure the VPC as follows:
   - For **IPv4 CIDR block**, specify an IPv4 CIDR block for the VPC.
   - For **IPv6 CIDR block**, keep the default value, **No IPv6 CIDR Block**.
   - For **VPC name**, type a unique name for the VPC.
6. Configure the public subnet as follows:
   - For **Public subnet's IPv4 CIDR**, specify the CIDR block for the subnet.
   - For **Availability Zone**, keep the default value, **No Preference**.
   - For **Public subnet name**, type a name for the subnet; for example, `AppStream2 Public Subnet`.
7. Configure the first private subnet as follows:
   - For **Private subnet's IPv4 CIDR**, specify the CIDR block for the subnet. Make a note of the value that you specify.
   - For **Availability Zone**, select a specific zone and make a note of the zone that you select.
   - For **Private subnet name**, type a name for the subnet; for example, `AppStream2 Private Subnet1`.
   - For the remaining fields, where applicable, keep the default values.

8. For **Elastic IP Allocation ID**, click in the text box and select the value that corresponds to the Elastic IP address that you created. This address is assigned to the NAT gateway. If you don't have an Elastic IP address, create one by using the Amazon VPC console at https://console.aws.amazon.com/vpc/.
9. For **Service endpoints**, if an Amazon S3 endpoint is required for your environment, specify one. An S3 endpoint is required to provide users with access to home folders or to enable application settings persistence for your users in a private network.

   To specify an Amazon S3 endpoint, do the following:

   a. Choose **Add Endpoint**.
   b. For **Service**, select the entry in the list that ends with "s3" (the `com.amazonaws.region.s3` entry that corresponds to the Region in which the VPC is being created).
   c. For **Subnet**, choose **Private subnet**.
   d. For **Policy**, keep the default value, **Full Access**.
10. For **Enable DNS hostnames**, keep the default value, **Yes**.
11. For **Hardware tenancy**, keep the default value, **Default**.
12. Choose **Create VPC**.
13. Note that it takes several minutes to set up your VPC. After the VPC is created, choose **OK**.

Step 3: Add a Second Private Subnet

In the previous step (Step 2: Create a New VPC), you created a VPC with one public subnet and one private subnet. Perform the following steps to add a second private subnet. We recommend that you add a second private subnet in a different Availability Zone than your first private subnet.

1. In the navigation pane, choose **Subnets**.
2. Select the first private subnet that you created in the previous step. On the **Description** tab, below the list of subnets, make a note of the Availability Zone for this subnet.
3. On the upper left of the subnets pane, choose **Create Subnet**.
4. For **Name tag**, type a name for the private subnet; for example, `AppStream2 Private Subnet2`.
5. For **VPC**, select the VPC that you created in the previous step.

6. For **Availability Zone**, select an Availability Zone other than the one you are using for your first private subnet. Selecting a different Availability Zone increases fault tolerance and helps prevent insufficient capacity errors.
7. For **IPv4 CIDR block**, specify a unique CIDR block range for the new subnet. For example, if your first private subnet has an IPv4 CIDR block range of `10.0.1.0/24`, you could specify a CIDR block range of `10.0.2.0/24` for the new private subnet.
8. Choose **Create**.
9. After your subnet is created, choose **Close**.

Step 4: Verify and Name Your Subnet Route Tables

After you've created and configured your VPC, complete the following steps to specify a name for your route tables, and to verify that:

- The route table associated with the subnet in which your NAT gateway resides includes a route that points internet traffic to an internet gateway. This ensures that your NAT gateway can access the internet.
- The route tables associated with your private subnets are configured to point internet traffic to the NAT gateway. This enables the streaming instances in your private subnets to communicate with the internet.

1. In the navigation pane, choose **Subnets**, and select the public subnet that you created; for example, `WorkSpaces Applications Public Subnet`.
    a. On the **Route Table** tab, choose the ID of the route table; for example, `rtb-12345678`.
    b. Select the route table. Under **Name**, choose the edit icon (the pencil), and type a name (for example, `appstream2-public-routetable`), and then select the check mark to save the name.
    c. With the public route table still selected, on the **Routes** tab, verify that there is one route for local traffic and another route that sends all other traffic to the internet gateway for the VPC. The following table describes these two routes:

Destination	Target	Description
Public subnet IPv4 CIDR Block (for example, 10.0.0/20)	Local	All traffic from the resources destined for IPv4 addresses within the public subnet IPv4 CIDR block is routed locally within the VPC.
Traffic destined to all other IPv4 addresses (for example, 0.0.0.0/0	Outbound (`igw-`*`ID`*)	Traffic destined for all other IPv4 addresses is routed to the internet gateway (identified by `igw-`*`ID`*) that was created by the VPC Wizard.

2. In the navigation pane, choose **Subnets**, and select the first private subnet that you created (for example, `AppStream2 Private Subnet1`).

   a. On the **Route Table** tab, choose the ID of the route table.

   b. Select the route table. Under **Name**, choose the edit icon (the pencil), and enter a name (for example, `appstream2-private-routetable`), and then choose the check mark to save the name.

   c. On the **Routes** tab, verify that the route table includes the following routes:

Destination	Target	Description
Public subnet IPv4 CIDR Block (for example, 10.0.0/20)	Local	All traffic from the resources destined for IPv4 addresses within the public subnet IPv4 CIDR block is routed locally within the VPC.
Traffic destined to all other IPv4 addresses (for example, 0.0.0.0/0	Outbound (`nat-`*`ID`*)	Traffic destined for all other IPv4 addresses is routed to the NAT gateway (identified by `nat-`*`ID`*).
Traffic destined for S3 buckets (applicable if you specified an S3 endpoint)	Storage (`vpce-`*`ID`*)	Traffic destined for S3 buckets is routed to the S3 endpoint (identifi ed by `vpce-`*`ID`*).

Destination	Target	Description
[pl-*ID* (com.amazonaws. *region*.s3)]		

3. In the navigation pane, choose **Subnets**, and select the second private subnet that you created (for example, `AppStream2 Private Subnet2`).
4. On the **Route Table** tab, verify that the route table is the private route table (for example, `appstream2-private-routetable`). If the route table is different, choose **Edit** and select this route table.

Next Steps

To enable your fleet instances, app block builders, and image builders to access the internet, complete the steps in Enable Internet Access for Your Fleet, Image Builder, or App Block Builder in Amazon WorkSpaces Applications.

Add a NAT Gateway to an Existing VPC

If you have already configured a VPC, complete the following steps to add a NAT gateway to your VPC. If you need to create a new VPC, see Create and Configure a New VPC.

To add a NAT gateway to an existing VPC

1. To create your NAT gateway, complete the steps in Creating a NAT Gateway in the *Amazon VPC User Guide*.
2. Verify that your VPC has at least one private subnet. We recommend that you specify two private subnets from different Availability Zones for high availability and fault tolerance. For information about how to create a second private subnet, see Step 3: Add a Second Private Subnet.
3. Update the route table associated with one or more of your private subnets to point internet-bound traffic to the NAT gateway. This enables the streaming instances in your private subnets to communicate with the internet. To do so, complete the steps in Configure route tables.

Next Steps

To enable your fleet instances, app block builders, and image builders to access the internet, complete the steps in Enable Internet Access for Your Fleet, Image Builder, or App Block Builder in Amazon WorkSpaces Applications.

Enable Internet Access for Your Fleet, Image Builder, or App Block Builder in Amazon WorkSpaces Applications

After your NAT gateway is available on a VPC, you can enable internet access for your fleet, image builder, and app block builder.

> **Note**
>
> When working with IPv6 only subnets, default internet access cannot be enabled. You'll need to set up an Egress-Only Internet Gateway and configure route table to allow outbound internet traffic. For more information check the steps. You also need to enable auto-assign IPv6 addresses for your subnets. The Egress-Only gateway handles outbound internet traffic only so if you need inbound access, you'll still need a regular internet gateway. You can find more details about this in egress only internet gateway documentation. Ensure appropriate security and networking controls exist to prevent accidental inbound/outbound access for your subnets.

Topics

- Enable Internet Access for Your Fleet in Amazon WorkSpaces Applications
- Enable Internet Access for Your Image Builder in Amazon WorkSpaces Applications
- Enable Internet Access for Your App Block Builder in Amazon WorkSpaces Applications

Enable Internet Access for Your Fleet in Amazon WorkSpaces Applications

You can enable internet access either when you create the fleet or later.

To enable internet access at fleet creation

1. Complete the steps in Create a Fleet in Amazon WorkSpaces Applications up to **Step 4: Configure Network**.
2. Choose a VPC with a NAT gateway.

3. If the subnet fields are empty, select a private subnet for **Subnet 1** and, optionally, another private subnet for **Subnet 2**. If you don't already have a private subnet in your VPC, you may need to create a second private subnet.
4. Continue with the steps in Create a Fleet in Amazon WorkSpaces Applications.

To enable internet access after fleet creation by using a NAT gateway

1. In the navigation pane, choose **Fleets**.
2. Select a fleet and verify that the state is **Stopped**.
3. Choose **Fleet Details**, **Edit**, and choose a VPC with a NAT gateway.
4. Choose a private subnet for **Subnet 1** and, optionally, another private subnet for **Subnet 2**. If you don't already have a private subnet in your VPC, you may need to create a second private subnet.
5. Choose **Update**.

You can test your internet connectivity by starting your fleet, and then connecting to your streaming instance and browsing to the internet.

Enable Internet Access for Your Image Builder in Amazon WorkSpaces Applications

If you plan to enable internet access for your image builder, you must do so when you create the image builder.

To enable internet access for an image builder

1. Complete the steps in Launch an Image Builder to Install and Configure Streaming Applications, up to **Step 3: Configure Network**.
2. Choose the VPC with a NAT gateway.
3. If **Subnet** is empty, select a subnet.
4. Continue with the steps in Launch an Image Builder to Install and Configure Streaming Applications.

Enable Internet Access for Your App Block Builder in Amazon WorkSpaces Applications

If you plan to enable internet access for your app block builder, you must do so when you create the app block builder.

To enable internet access for an app block builder

1. Complete the steps in the section called "Create an App Block Builder" up to **Step 2: Configure Network**.
2. Choose the VPC with a NAT gateway.
3. If **Subnet** is empty, select a subnet.
4. Continue with the steps in the section called "Create an App Block Builder".

Configure a New or Existing VPC with a Public Subnet

If you created your Amazon Web Services account after 2013-12-04, you have a default VPC in each AWS Region that includes default public subnets. However, you may want to create your own nondefault VPC or configure an existing VPC to use with WorkSpaces Applications. This topic describes how to configure a nondefault VPC and public subnet to use with WorkSpaces Applications.

After you configure your VPC and public subnet, you can provide your streaming instances (fleet instances and image builders) with access to the internet by enabling the **Default Internet Access** option. When you enable this option, WorkSpaces Applications enables internet connectivity by associating an Elastic IP address to the network interface that is attached from the streaming instance to your public subnet. An Elastic IP address is a public IPv4 address that is reachable from the internet. For this reason, we recommend that you instead use a NAT gateway to provide internet access to your WorkSpaces Applications instances. In addition, when **Default Internet Access** is enabled, a maximum of 100 fleet instances is supported. If your deployment must support more than 100 concurrent users, use the NAT gateway configuration instead.

For more information, see the steps in Configure a VPC with Private Subnets and a NAT Gateway. For additional VPC configuration recommendations, see VPC Setup Recommendations.

Contents

- Step 1: Configure a VPC with a Public Subnet
- Step 2: Enable Default Internet Access Your Fleet, Image Builder, or App Block Builder

Step 1: Configure a VPC with a Public Subnet

You can configure your own non-default VPC with a public subnet by using either of the following methods:

- Create a New VPC with a Single Public Subnet
- Configure an Existing VPC

> **Note**
>
> When working with IPv6 only subnets, default internet access cannot be enabled. You'll need to set up an Egress-Only Internet Gateway and configure route table to allow outbound internet traffic. For more information check the steps. You also need to enable auto-assign IPv6 addresses for your subnets. The Egress-Only gateway handles outbound internet traffic only so if you need inbound access, you'll still need a regular internet gateway. You can find more details about this in egress only internet gateway documentation.

Create a New VPC with a Single Public Subnet

When you use the VPC wizard to create a new VPC, the wizard creates an internet gateway and a custom route table that is associated with the public subnet. The route table routes all traffic destined for an address outside the VPC to the internet gateway. For more information about this configuration, see VPC with a Single Public Subnet in the *Amazon VPC User Guide*.

1. Complete the steps in Step 1: Create the VPC in the *Amazon VPC User Guide* to create your VPC.
2. To enable your fleet instances and image builders to access the internet, complete the steps in Step 2: Enable Default Internet Access Your Fleet, Image Builder, or App Block Builder.

Configure an Existing VPC

If you want to use an existing VPC that does not have a public subnet, you can add a new public subnet. In addition to a public subnet, you must also have an internet gateway attached to your VPC and a route table that routes all traffic destined for an address outside the VPC to the internet gateway. To configure these components, complete the following steps.

1. To add a public subnet, complete the steps in Creating a Subnet in Your VPC. Use the existing VPC that you plan to use with WorkSpaces Applications.

If your VPC is configured to support IPv6 addressing, the **IPv6 CIDR block** list displays. Select **Don't assign Ipv6**.

2. To create and attach an internet gateway to your VPC, complete the steps in Creating and Attaching an Internet Gateway.
3. To configure your subnet to route internet traffic through the internet gateway, complete the steps in Creating a Custom Route Table. In step 5, for **Destination**, use IPv4 format (`0.0.0.0/0`).
4. To enable your fleet instances and image builders to access the internet, complete the steps in Step 2: Enable Default Internet Access Your Fleet, Image Builder, or App Block Builder.

Step 2: Enable Default Internet Access Your Fleet, Image Builder, or App Block Builder

After you configure a VPC that has a public subnet, you can enable the **Default Internet Access** option for your fleet and image builder.

Enable Default Internet Access for a Fleet

You can enable the **Default Internet Access** option when you create the fleet, or later.

> **Note**
>
> For fleet instances that have the **Default Internet Access** option enabled, the limit is 100.

To enable internet access at fleet creation

1. Complete the steps in Create a Fleet in Amazon WorkSpaces Applications up to **Step 4: Configure Network**.
2. Select the **Default Internet Access** check box.
3. If the subnet fields are empty, select a subnet for **Subnet 1** and, optionally, **Subnet 2**.
4. Continue with the steps in Create a Fleet in Amazon WorkSpaces Applications.

To enable internet access after fleet creation

1. In the navigation pane, choose **Fleets**.

2. Select a fleet and verify that its state is **Stopped**.
3. Choose **Fleet Details**, **Edit**, then select the **Default Internet Access** check box.
4. Choose a subnet for **Subnet 1** and, optionally, **Subnet 2**. Choose **Update**.

You can test internet connectivity by starting your fleet, creating a stack, associating the fleet with a stack, and browsing the internet within a streaming session for stack. For more information, see Create an Amazon WorkSpaces Applications Fleet and Stack.

Enable Default Internet Access for an Image Builder

After you configure a VPC that has a public subnet, you can enable the **Default Internet Access** option for your image builder. You can do so when you create the image builder.

To enable internet access for an image builder

1. Complete the steps in Launch an Image Builder to Install and Configure Streaming Applications up to **Step 3: Configure Network**.
2. Select the **Default Internet Access** check box.
3. If **Subnet 1** is empty, select a subnet.
4. Continue with the steps in Launch an Image Builder to Install and Configure Streaming Applications.

Enable Default Internet Access for an App Block Builder

After you configure a VPC that has a public subnet, you can enable the **Default Internet Access** option for your app block builder. You can do so when you create the app block builder.

To enable internet access for an app block builder

1. Follow the steps in the section called “Create an App Block Builder”, up to **Step 2: Configure Network**.
2. Select the **Default Internet Access** check box.
3. If **Subnet** is empty, select a subnet.
4. Continue with the steps in the section called “Create an App Block Builder”.

Use the Default VPC, Public Subnet, and Security Group

Your Amazon Web Services account, if it was created after 2013-12-04, has a default VPC in each AWS Region. The default VPC includes a default public subnet in each Availability Zone and an internet gateway that is attached to your VPC. The VPC also includes a default security group. If you are new to WorkSpaces Applications and want to get started using the service, you can keep the default VPC and security group selected when you create a fleet, create an app block builder, or launch an image builder. Then, you can select at least one default subnet.

> **Note**
>
> If your Amazon Web Services account was created before 2013-12-04, you must create a new VPC or configure an existing one to use with WorkSpaces Applications. We recommend that you manually configure a VPC with two private subnets for your fleets, app block builders, and image builders and a NAT gateway in a public subnet. For more information, see Configure a VPC with Private Subnets and a NAT Gateway. Alternatively, you can configure a non-default VPC with a public subnet. For more information, see Configure a New or Existing VPC with a Public Subnet.

To use the default VPC, subnet, and security group for a fleet

1. Complete the steps in Create a Fleet in Amazon WorkSpaces Applications up to **Step 4: Configure Network**.
2. In **Step 4: Configure Network**, do the following:

   - To enable your fleet instances to access the internet, select the **Default Internet Access** check box.

     > **Note**
     >
     > For fleet instances that have the **Default Internet Access** option enabled, the limit is 100.

   - For **VPC**, choose the default VPC for your AWS Region.

     The default VPC name uses the following format: `vpc-`*`vpc-id`* `(No_default_value_Name)`.

- For **Subnet 1**, choose a default public subnet and make a note of the Availability Zone.

  The default subnet names use the following format: `subnet-`*`subnet-id`* `|` `(`*`IPv4 CIDR block`*`) | Default in` *`availability-zone`*.
- Optionally, for **Subnet 2**, choose a default subnet in a different Availability Zone.
- For **Security groups**, select the default security group.

  The default security group name uses the following format: `sg-`*`security-group-id`*`-default`

3. Continue with the steps in Create a Fleet in Amazon WorkSpaces Applications.

Complete the following steps to use the default VPC, subnet, and security group for an image builder.

To use the default VPC, subnet, and security group for an image builder

1. Follow the steps in Launch an Image Builder to Install and Configure Streaming Applications up to **Step 3: Configure Network**.
2. In **Step 4: Configure Network**, do the following:

   - To enable your image builder to access the internet, select the **Default Internet Access** check box.
   - For **VPC**, choose the default VPC for your AWS Region.

     The default VPC name uses the following format: `vpc-`*`vpc-id`* `(No_default_value_Name)`.
   - For **Subnet 1**, choose a default public subnet.

     The default subnet names use the following format: `subnet-`*`subnet-id`* `|` `(`*`IPv4 CIDR block`*`) | Default in` *`availability-zone`*.
   - For **Security groups**, select the default security group.

     The default security group name uses the following format: `sg-`*`security-group-id`*`-default`
3. Continue with the steps in Launch an Image Builder to Install and Configure Streaming Applications.

Complete the following steps to use the default VPC, subnet, and security group for an app block builder.

To use the default VPC, subnet, and security group for an app block builder

1. Follow the steps in the section called "Create an App Block Builder", up to **Step 2: Configure Network**.
2. In **Step 2: Configure Network**, do the following:
   - To enable your image builder to access the internet, select the **Default Internet Access** check box.
   - For **VPC**, choose the default VPC for your AWS Region.

     The default VPC name uses the following format: `vpc-`*`vpc-id`* `(No_default_value_Name)`.
   - For **Subnet 1**, choose a default public subnet.

     The default subnet names use the following format: `subnet-`*`subnet-id`* `|` `(`*`IPv4 CIDR block`*`) | Default in` *`availability-zone`*.
   - For **Security groups**, select the default security group.

     The default security group name uses the following format: `sg-`*`security-group-id`*`-default`
3. Continue with the steps in the section called "Create an App Block Builder".

Using Amazon S3 VPC Endpoints for WorkSpaces Applications Features

When you enable Application Settings Persistence or Home folders on a stack, WorkSpaces Applications uses the VPC you specify for your fleet to provide access to Amazon Simple Storage Service (Amazon S3) buckets. For Elastic fleets, WorkSpaces Applications will use the VPC to access the Amazon S3 bucket containing applications assigned to the fleet's app block. To enable WorkSpaces Applications access to your private S3 endpoint, attach the following custom policy to your VPC endpoint for Amazon S3. For more information about private Amazon S3 endpoints, see VPC Endpoints and Endpoints for Amazon S3 in the *Amazon VPC User Guide*.

JSON

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Sid": "Allow-AppStream-to-access-S3-buckets",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:sts::111122223333:assumed-role/
AmazonAppStreamServiceAccess/AppStream2.0"
            },
            "Action": [
                "s3:ListBucket",
                "s3:GetObject",
                "s3:PutObject",
                "s3:DeleteObject",
                "s3:GetObjectVersion",
                "s3:DeleteObjectVersion"
            ],
            "Resource": [
                "arn:aws:s3:::appstream2-36fb080bb8-*",
                "arn:aws:s3:::appstream-app-settings-*",
                "arn:aws:s3:::appstream-logs-*"
            ]
        },
        {
            "Sid": "Allow-AppStream-ElasticFleetstoRetrieveObjects",
            "Effect": "Allow",
            "Principal": "*",
            "Action": "s3:GetObject",
            "Resource": "arn:aws:s3:::bucket-with-application-or-app-block-
objects/*",
            "Condition": {
                "StringEquals": {
                    "aws:PrincipalServiceName": "appstream.amazonaws.com"
                }
            }
        }
    ]
}
```

Amazon WorkSpaces Applications Connections to Your VPC

To enable WorkSpaces Applications connectivity to network resources and the internet, configure your streaming instances as follows.

Topics

- Network Interfaces in Amazon WorkSpaces Applications
- Management Network Interface IP Address Range and Ports in Amazon WorkSpaces Applications
- Customer Network Interface Ports in Amazon WorkSpaces Applications

Network Interfaces in Amazon WorkSpaces Applications

Each WorkSpaces Applications streaming instance has the following network interfaces:

- The customer network interface provides connectivity to the resources within your VPC, as well as the internet, and is used to join the streaming instance to your directory.
- The management network interface is connected to a secure WorkSpaces Applications management network. It is used for interactive streaming of the streaming instance to a user's device, and to allow WorkSpaces Applications to manage the streaming instance.

WorkSpaces Applications selects the IP address for the management network interface from the following private IP address range: 198.19.0.0/16. Do not use this range for your VPC CIDR or peer your VPC with another VPC with this range, as this might create a conflict and cause streaming instances to be unreachable. Also, do not modify or delete any of the network interfaces attached to a streaming instance, as this might also cause the streaming instance to become unreachable.

Management Network Interface IP Address Range and Ports in Amazon WorkSpaces Applications

The management network interface IP address range is 198.19.0.0/16. The following ports must be open on the management network interface of all streaming instances:

- Inbound TCP on port 8300. This is used for establishment of the streaming connection.
- Inbound TCP on ports 8000 and 8443. These are used for management of the streaming instance by WorkSpaces Applications.

- Inbound UDP on port 8300. This is used for establishment of the streaming connection over UDP.
- Outbound TCP on port 1688. This is used for Microsoft license included application activation on fleet streaming instances.

Limit the inbound range on the management network interface to 198.19.0.0/16.

Under normal circumstances, WorkSpaces Applications correctly configures these ports for your streaming instances. If any security or firewall software is installed on a streaming instance that blocks any of these ports, the streaming instance may not function correctly or may be unreachable.

Do not disable IPv6. If you disable IPv6, WorkSpaces Applications will not function correctly. For information about configuring IPv6 for Windows, see Guidance for configuring IPv6 in Windows for advanced users.

> **Note**
>
> WorkSpaces Applications relies on the DNS servers within your VPC to return a non-existent domain (NXDOMAIN) response for local domain names that don't exist. This enables the WorkSpaces Applications-managed network interface to communicate with the management servers.
> When you create a directory with Simple AD, AWS Directory Service creates two domain controllers that also function as DNS servers on your behalf. Because the domain controllers don't provide the NXDOMAIN response, they can't be used with WorkSpaces Applications.

Customer Network Interface Ports in Amazon WorkSpaces Applications

Follow the guidance below for customer network interface ports.

- For internet connectivity, the following ports must be open to all destinations. If you are using a modified or custom security group, you need to add the required rules manually. For more information, see Security Group Rules in the *Amazon VPC User Guide*.
  - TCP 80 (HTTP)
  - TCP 443 (HTTPS)
  - UDP 8433

- If you join your streaming instances to a directory, the following ports must be open between your WorkSpaces Applications VPC and your directory controllers.
  - TCP/UDP 53 - DNS
  - TCP/UDP 88 - Kerberos authentication
  - UDP 123 - NTP
  - TCP 135 - RPC
  - UDP 137-138 - Netlogon
  - TCP 139 - Netlogon
  - TCP/UDP 389 - LDAP
  - TCP/UDP 445 - SMB
  - TCP 1024-65535 - Dynamic ports for RPC

  For a complete list of ports, see Active Directory and Active Directory Domain Services Port Requirements in the Microsoft documentation.
- All streaming instances require that port 80 (HTTP) be open to IP address `169.254.169.254` to allow access to the EC2 metadata service. The IP address range `169.254.0.0/16` is reserved for WorkSpaces Applications service usage for management traffic. Failure to exclude this range might result in streaming issues.

User Connections to Amazon WorkSpaces Applications

Users can connect to WorkSpaces Applications streaming instances through the default public internet endpoint, or by using an interface VPC endpoint (interface endpoint) that you create in your virtual private cloud (VPC). For more information, see Tutorial: Creating and Streaming from Interface VPC Endpoints.

By default, WorkSpaces Applications is configured to route streaming connections over the public internet. Internet connectivity is required to authenticate users and deliver the web assets that WorkSpaces Applications requires to function. To allow this traffic, you must allow the domains listed in Allowed Domains.

> **Note**
>
> For user authentication, WorkSpaces Applications supports user pools, Security Assertion Markup Language 2.0 (SAML 2.0), and the CreateStreamingURL API action. For more information, see User Authentication.

The following topics provide information about how to enable user connections to WorkSpaces Applications.

Contents

- Bandwidth Recommendations
- IP Address and Port Requirements for WorkSpaces Applications User Devices
- Allowed Domains

Bandwidth Recommendations

To optimize the performance of WorkSpaces Applications, make sure that your network bandwidth and latency can sustain your users' needs.

WorkSpaces Applications uses NICE Desktop Cloud Visualization (DCV) to enable your users to securely access and stream your applications over varying network conditions. To help reduce bandwidth consumption, NICE DCV uses H.264-based video compression and encoding. During streaming sessions, the visual output of applications is compressed and streamed to your users as an AES-256 encrypted pixel stream over HTTPS. After the stream is received, it is decrypted and output to your users' local screen. When your users interact with their streaming applications, the NICE DCV protocol captures their input and sends it back to their streaming applications over HTTPS.

Network conditions are constantly measured during this process and information is sent back to WorkSpaces Applications. WorkSpaces Applications dynamically responds to changing network conditions by changing the video and audio encoding in real time to produce a high-quality stream for a wide variety of applications and network conditions.

The recommended bandwidth and latency for WorkSpaces Applications streaming sessions depends on the workload. For example, a user who works with graphic-intensive applications to perform computer-aided design tasks will require more bandwidth and lower latency than a user who works with business productivity applications to write documents.

The following table provides guidance on the recommended network bandwidth and latency for WorkSpaces Applications streaming sessions based on common workloads.

For each workload, the bandwidth recommendation is based on what an individual user might require at a specific point in time. The recommendation does not reflect the bandwidth required for sustained throughput. When only a few pixels change on the screen during a streaming session, the sustained throughput is much lower. Although users who have less bandwidth available can still stream their applications, the frame rate or image quality may not be optimal.

Workload	Description	Bandwidth recommended per user	Recommended maximum roundtrip latency
Line of business applications	Document writing applications, database analysis utilities	2 Mbps	< 150 ms
Graphics applications	Computer-aided design and modeling applications, photo and video editing	5 Mbps	< 100 ms
High fidelity	High-fidelity datasets or maps across multiple monitors	10 Mbps	< 50 ms

IP Address and Port Requirements for WorkSpaces Applications User Devices

WorkSpaces Applications users' devices require outbound access on port 443 (TCP) and port 8433 (UDP) when using the internet endpoints, and if you are using DNS servers for domain name resolution, port 53 (UDP).

- Port 443 is used for HTTPS communication between WorkSpaces Applications users' devices and streaming instances when using the internet endpoints. Typically, when end users browse the web during streaming sessions, the web browser randomly selects a source port in the high range for streaming traffic. You must ensure that return traffic to this port is allowed.

  > **Note**
  >
  > WorkSpaces Applications uses WebSockets on port 443.

- Port 8433 is used for UDP HTTPS communication between WorkSpaces Applications users' devices and streaming instances when using the internet endpoints. This is currently only supported in the Windows native client. UDP is not supported if you are using VPC endpoints.

  > **Note**
  >
  > Streaming through interface VPC endpoints requires additional ports. For more information, see Tutorial: Creating and Streaming from Interface VPC Endpoints.

- Port 53 is used for communication between WorkSpaces Applications users' devices and your DNS servers. The port must be open to the IP addresses for your DNS servers so that public domain names can be resolved. This port is optional if you are not using DNS servers for domain name resolution.

WorkSpaces Applications client for Windows (version 1.2.1581 or greater), Mac (version 1.2.0 or greater) and Web Browser access automatically prefer to connect on IPv6 network than over IPv4 network. In case of adverse network conditions such as network latency the clients will fall back to IPV4 network.

Allowed Domains

Streaming Instances

For WorkSpaces Applications users to access streaming instances, you must allow the following domain on the network from which users initiate access to the streaming instances.

- Streaming Gateway: *.amazonappstream.com

> **Note**
>
> Instead of using a wildcard to allowlist all streaming gateways, you can create a VPC endpoint and allowlist only that specific endpoint. For more information, see the section called "Interface VPC Endpoints".

For IPV4 support, you must allow the following domain on the network from which users initiate access to the streaming instances. It is region based and follows the following format: `*.streaming.{region}.appstream2.amazonappstream.com` and `*.dcv-streaming.{region}.appstream2.amazonappstream.com`. If its a FIPS compliant region, it will also need another endpoint with the format `*.streaming.{region}.appstream2-fips.amazonappstream.com` and `*.dcv-streaming.{region}.appstream2-fips.amazonappstream.com`. Check table below.

> **Note**
>
> `*.streaming.*` URLs are not supported for the Asia Pacific (Malaysia), Asia Pacific (Osaka), Canada West (Calgary), Europe (Milan), Europe (Spain), Europe (Zurich), and Israel (Tel Aviv) Regions. For these Regions, allow only the `*.dcv-streaming.*` URLs listed in the table below.

Region	Domain
US East (N. Virginia)	*.streaming.us-east-1.appstream2.amazonappstream.com *.dcv-streaming.us-east-1.appstream2.amazonappstream.com *.streaming.us-east-1.appstream2-fips.amazonappstream.com *.dcv-streaming.us-east-1.appstream2-fips.amazonappstream.com
US East (Ohio)	*.streaming.us-east-2.appstream2.amazonappstream.com *.dcv-streaming.us-east-2.appstream2.amazonappstream.com

Region	Domain
US West (Oregon)	*.streaming.us-west-2.appstream2.amazonappstream.com *.dcv-streaming.us-west-2.appstream2.amazonappstream.com *.streaming.us-west-2.appstream2-fips.amazonappstream.com *.dcv-streaming.us-west-2.appstream2-fips.amazonappstream.com
Asia Pacific (Malaysia)	*.dcv-streaming.ap-southeast-5.appstream2.amazonappstream.com
Asia Pacific (Mumbai)	*.streaming.ap-south-1.appstream2.amazonappstream.com *.dcv-streaming.ap-south-1.appstream2.amazonappstream.com
Asia Pacific (Seoul)	*.streaming.ap-northeast-2.appstream2.amazonappstream.com *.dcv-streaming.ap-northeast-2.appstream2.amazonappstream.com
Asia Pacific (Singapore)	*.streaming.ap-southeast-1.appstream2.amazonappstream.com *.dcv-streaming.ap-southeast-1.appstream2.amazonappstream.com
Asia Pacific (Sydney)	*.streaming.ap-southeast-2.appstream2.amazonappstream.com *.dcv-streaming.ap-southeast-2.appstream2.amazonappstream.com

Region	Domain
Asia Pacific (Tokyo)	*.streaming.ap-northeast-1.appstream2.amazonappstream.com *.dcv-streaming.ap-northeast-1.appstream2.amazonappstream.com
Asia Pacific (Osaka)	*.dcv-streaming.ap-northeast-3.appstream2.amazonappstream.com
Canada (Central)	*.streaming.ca-central-1.appstream2.amazonappstream.com *.dcv-streaming.ca-central-1.appstream2.amazonappstream.com
Canada West (Calgary)	*.dcv-streaming.ca-west-1.appstream2.amazonappstream.com *.dcv-streaming.ca-west-1.appstream2-fips.amazonappstream.com
Europe (Frankfurt)	*.streaming.eu-central-1.appstream2.amazonappstream.com *.dcv-streaming.eu-central-1.appstream2.amazonappstream.com
Europe (London)	*.streaming.eu-west-2.appstream2.amazonappstream.com *.dcv-streaming.eu-west-2.appstream2.amazonappstream.com
Europe (Ireland)	*.streaming.eu-west-1.appstream2.amazonappstream.com *.dcv-streaming.eu-west-1.appstream2.amazonappstream.com
Europe (Paris)	*.streaming.eu-west-3.appstream2.amazonappstream.com *.dcv-streaming.eu-west-3.appstream2.amazonappstream.com
Europe (Milan)	*.dcv-streaming.eu-south-1.appstream2.amazonappstream.com

Region	Domain
Europe (Spain)	*.dcv-streaming.eu-south-2.appstream2.amazonappstream.com
Europe (Zurich)	*.dcv-streaming.eu-central-2.appstream2.amazonappstream.com
AWS GovCloud (US-East)	*.streaming.us-gov-east-1.appstream2.amazonappstream.com *.dcv-streaming.us-gov-east-1.appstream2.amazonappstream.com *.streaming.us-gov-east-1.appstream2-fips.amazonappstream.com *.dcv-streaming.us-gov-east-1.appstream2-fips.amazonappstream.com
AWS GovCloud (US-West)	*.streaming.us-gov-west-1.appstream2.amazonappstream.com *.dcv-streaming.us-gov-west-1.appstream2.amazonappstream.com *.streaming.us-gov-west-1.appstream2-fips.amazonappstream.com *.dcv-streaming.us-gov-west-1.appstream2-fips.amazonappstream.com
South America (São Paulo)	*.streaming.sa-east-1.appstream2.amazonappstream.com *.dcv-streaming.sa-east-1.appstream2.amazonappstream.com
Israel (Tel Aviv)	*.dcv-streaming.il-central-1.appstream2.amazonappstream.com

For IPV6 support, you must allow the following domain on the network from which users initiate access to the streaming instances. It is region based and follows the following format: `*.streaming.appstream2.{region}.on.aws` and `*.dcv-streaming.appstream2.`

`{region}.on.aws`. If its a FIPS compliant region, it will also need another endpoint with the format `*.streaming.appstream2-fips.{region}.on.aws` and `*.dcv-streaming.appstream2-fips.{region}.on.aws`. Check table below.

In order to use IPV6 address your base images must be updated to the images published on September 05, 2025 or later. For more information check managed image updates.

Region	Domain
US East (N. Virginia)	*.streaming.appstream2.us-east-1.on.aws *.dcv-streaming.appstream2.us-east-1.on.aws *.streaming.appstream2-fips.us-east-1.on.aws *.dcv-streaming.appstream2-fips.us-east-1.on.aws
US East (Ohio)	*.streaming.appstream2.us-east-2.on.aws *.dcv-streaming.appstream2.us-east-2.on.aws
US West (Oregon)	*.streaming.appstream2.us-west-2.on.aws *.dcv-streaming.appstream2.us-west-2.on.aws *.streaming.appstream2-fips.us-west-2.on.aws *.dcv-streaming.appstream2-fips.us-west-2.on.aws
Asia Pacific (Malaysia)	*.dcv-streaming.appstream2.ap-southeast-5.on.aws
Asia Pacific (Mumbai)	*.streaming.appstream2.ap-south-1.on.aws *.dcv-streaming.appstream2.ap-south-1.on.aws
Asia Pacific (Seoul)	*.streaming.appstream2.ap-northeast-2.on.aws *.dcv-streaming.appstream2.ap-northeast-2.on.aws
Asia Pacific (Singapore)	*.streaming.appstream2.ap-southeast-1.on.aws *.dcv-streaming.appstream2.ap-southeast-1.on.aws

Region	Domain
Asia Pacific (Sydney)	*.streaming.appstream2.ap-southeast-2.on.aws *.dcv-streaming.appstream2.ap-southeast-2.on.aws
Asia Pacific (Tokyo)	*.streaming.appstream2.ap-northeast-1.on.aws *.dcv-streaming.appstream2.ap-northeast-1.on.aws
Asia Pacific (Osaka)	*.dcv-streaming.appstream2.ap-northeast-3.on.aws
Canada (Central)	*.streaming.appstream2.ca-central-1.on.aws *.dcv-streaming.appstream2.ca-central-1.on.aws
Canada West (Calgary)	*.dcv-streaming.appstream2.ca-west-1.on.aws *.dcv-streaming.appstream2-fips.ca-west-1.on.aws
Europe (Frankfurt)	*.streaming.appstream2.eu-central-1.on.aws *.dcv-streaming.appstream2.eu-central-1.on.aws
Europe (London)	*.streaming.appstream2.eu-west-2.on.aws *.dcv-streaming.appstream2.eu-west-2.on.aws
Europe (Ireland)	*.streaming.appstream2.eu-west-1.on.aws *.dcv-streaming.appstream2.eu-west-1.on.aws
Europe (Paris)	*.streaming.appstream2.eu-west-3.on.aws *.dcv-streaming.appstream2.eu-west-3.on.aws
Europe (Milan)	*.dcv-streaming.appstream2.eu-south-1.on.aws
Europe (Spain)	*.dcv-streaming.appstream2.eu-south-2.on.aws
Europe (Zurich)	*.dcv-streaming.appstream2.eu-central-2.on.aws

Region	Domain
AWS GovCloud (US-East)	*.streaming.appstream2.us-gov-east-1.on.aws *.dcv-streaming.appstream2.us-gov-east-1.on.aws *.streaming.appstream2-fips.us-gov-east-1.on.aws *.dcv-streaming.appstream2-fips.us-gov-east-1.on.aws
AWS GovCloud (US-West)	*.streaming.appstream2.us-gov-west-1.on.aws *.dcv-streaming.appstream2.us-gov-west-1.on.aws *.streaming.appstream2-fips.us-gov-west-1.on.aws *.dcv-streaming.appstream2-fips.us-gov-west-1.on.aws
South America (São Paulo)	*.streaming.appstream2.sa-east-1.on.aws *.dcv-streaming.appstream2.sa-east-1.on.aws
Israel (Tel Aviv)	*.dcv-streaming.appstream2.il-central-1.on.aws

User Authentication

One or more of the following domains must be allowed to enable user authentication. You must allow the domains and subdomains that correspond to the Regions where WorkSpaces Applications is deployed.

Region	Domain
US East (N. Virginia)	*.appstream2.us-east-1.aws.amazon.com
US East (Ohio)	*.appstream2.us-east-2.aws.amazon.com
US West (Oregon)	*.appstream2.us-west-2.aws.amazon.com
Asia Pacific (Malaysia)	*.appstream2.ap-southeast-5.aws.amazon.com

Region	Domain
Asia Pacific (Mumbai)	*.appstream2.ap-south-1.aws.amazon.com
Asia Pacific (Seoul)	*.appstream2.ap-northeast-2.aws.amazon.com
Asia Pacific (Singapore)	*.appstream2.ap-southeast-1.aws.amazon.com
Asia Pacific (Sydney)	*.appstream2.ap-southeast-2.aws.amazon.com
Asia Pacific (Tokyo)	*.appstream2.ap-northeast-1.aws.amazon.com
Asia Pacific (Osaka)	*.appstream2.ap-northeast-3.aws.amazon.com
Canada (Central)	*.appstream2.ca-central-1.aws.amazon.com
Canada West (Calgary)	*.appstream2.ca-west-1.aws.amazon.com
Europe (Frankfurt)	*.appstream2.eu-central-1.aws.amazon.com
Europe (London)	*.appstream2.eu-west-2.aws.amazon.com
Europe (Ireland)	*.appstream2.eu-west-1.aws.amazon.com
Europe (Milan)	*.appstream2.eu-south-1.aws.amazon.com
Europe (Paris)	*.appstream2.eu-west-3.aws.amazon.com
Europe (Spain)	*.appstream2.eu-south-2.aws.amazon.com
Europe (Zurich)	*.appstream2.eu-central-2.aws.amazon.com
AWS GovCloud (US-East)	*.appstream2.us-gov-east-1.amazonaws-us-gov.com
AWS GovCloud (US-West)	*.appstream2.us-gov-west-1.amazonaws-us-gov.com
South America (São Paulo)	*.appstream2.sa-east-1.aws.amazon.com
Israel (Tel Aviv)	*.appstream2.il-central-1.aws.amazon.com

> **Note**
>
> If your users use a network proxy to access streaming instances, disable any proxy caching for the user auth domains in the table and the session gateway, *.amazonappstream.com.

For SAML 2.0 [single sign-on (SSO)] authentication, you must allow these three additional endpoints:

1. AWS Sign-in endpoint
2. AWS Sign-in region endpoints where WorkSpaces Applications is available
3. WorkSpaces Applications relay state region endpoints

The following table lists the AWS Sign-in endpoint.

AWS Sign-in endpoint
`signin.aws.amazon.com`

The following table lists the AWS Sign-in region endpoints for the Regions where WorkSpaces Applications is available.

AWS Sign-in region endpoints

Region	Domain
US East (N. Virginia)	`us-east-1.signin.aws.amazon.com`
US East (Ohio)	`us-east-2.signin.aws.amazon.com`
US West (Oregon)	`us-west-2.signin.aws.amazon.com`
Asia Pacific (Malaysia)	`ap-southeast-5.signin.aws.amazon.com`
Asia Pacific (Mumbai)	`ap-south-1.signin.aws.amazon.com`
Asia Pacific (Seoul)	`ap-northeast-2.signin.aws.amazon.com`

Region	Domain
Asia Pacific (Singapore)	ap-southeast-1.signin.aws.amazon.com
Asia Pacific (Sydney)	ap-southeast-2.signin.aws.amazon.com
Asia Pacific (Tokyo)	ap-northeast-1.signin.aws.amazon.com
Canada (Central)	ca-central-1.signin.aws.amazon.com
Europe (Frankfurt)	eu-central-1.signin.aws.amazon.com
Europe (Ireland)	eu-west-1.signin.aws.amazon.com
Europe (London)	eu-west-2.signin.aws.amazon.com
Europe (Milan)	eu-south-1.signin.aws.amazon.com
Europe (Paris)	eu-west-3.signin.aws.amazon.com
Europe (Spain)	eu-south-2.signin.aws.amazon.com
AWS GovCloud (US-East)	us-gov-east-1.signin.amazonaws-us-gov.com
AWS GovCloud (US-West)	signin.amazonaws-us-gov.com
South America (São Paulo)	sa-east-1.signin.aws.amazon.com
Israel (Tel Aviv)	il-central-1.signin.aws.amazon.com

The following table lists the WorkSpaces Applications relay state region endpoints.

WorkSpaces Applications relay state region endpoints

Region	Domain
US East (N. Virginia)	appstream2.euc-sso.us-east-1.aws.amazon.com

Region	Domain
	(FIPS) appstream2.euc-sso-fips.us-east-1.aws.amazon.com
US East (Ohio)	appstream2.euc-sso.us-east-2.aws.amazon.com
US West (Oregon)	appstream2.euc-sso.us-west-2.aws.amazon.com (FIPS) appstream2.euc-sso-fips.us-west-2.aws.amazon.com
Asia Pacific (Malaysia)	appstream2.euc-sso.ap-southeast-5.aws.amazon.com
Asia Pacific (Mumbai)	appstream2.euc-sso.ap-south-1.aws.amazon.com
Asia Pacific (Seoul)	appstream2.euc-sso.ap-northeast-2.aws.amazon.com
Asia Pacific (Singapore)	appstream2.euc-sso.ap-southeast-1.aws.amazon.com
Asia Pacific (Sydney)	appstream2.euc-sso.ap-southeast-2.aws.amazon.com
Asia Pacific (Tokyo)	appstream2.euc-sso.ap-northeast-1.aws.amazon.com
Canada (Central)	appstream2.euc-sso.ca-central-1.aws.amazon.com
Europe (Frankfurt)	appstream2.euc-sso.eu-central-1.aws.amazon.com
Europe (Ireland)	appstream2.euc-sso.eu-west-1.aws.amazon.com

Region	Domain
Europe (London)	`appstream2.euc-sso.eu-west-2.aws.amazon.com`
Europe (Milan)	`appstream2.euc-sso.eu-south-1.aws.amazon.com`
Europe (Paris)	`appstream2.euc-sso.eu-west-3.aws.amazon.com`
Europe (Spain)	`appstream2.euc-sso.eu-south-2.aws.amazon.com`
AWS GovCloud (US-East)	`appstream2.euc-sso.us-gov-east-1.amazonaws-us-gov.com` (FIPS) `appstream2.euc-sso-fips.us-gov-east-1.amazonaws-us-gov.com` **Note** For more information about using WorkSpaces Applications in AWS GovCloud (US) Regions, see Amazon WorkSpaces Applications in the *AWS GovCloud (US) User Guide*.

Region	Domain
AWS GovCloud (US-West)	`appstream2.euc-sso.us-gov-west-1.amazonaws-us-gov.com` (FIPS) `appstream2.euc-sso-fips.us-gov-west-1.amazonaws-us-gov.com` **Note** For more information about using WorkSpaces Applications in AWS GovCloud (US) Regions, see Amazon WorkSpaces Applications in the *AWS GovCloud (US) User Guide*.
South America (São Paulo)	`appstream2.euc-sso.sa-east-1.aws.amazon.com`
Israel (Tel Aviv)	`appstream2.euc-sso.il-central-1.aws.amazon.com`

WorkSpaces Applications MCP endpoints

The following table lists the WorkSpaces Applications Model Context Protocol (MCP) endpoints.

WorkSpaces Applications MCP endpoints

Region	Domain
Asia Pacific (Tokyo)	`agentaccess-mcp.ap-northeast-1.api.aws`
Asia Pacific (Seoul)	`agentaccess-mcp.ap-northeast-2.api.aws`
Asia Pacific (Mumbai)	`agentaccess-mcp.ap-south-1.api.aws`
Asia Pacific (Singapore)	`agentaccess-mcp.ap-southeast-1.api.aws`
Asia Pacific (Sydney)	`agentaccess-mcp.ap-southeast-2.api.aws`

Region	Domain
Asia Pacific (Malaysia)	`agentaccess-mcp.ap-southeast-5.api.aws`
Canada (Central)	`agentaccess-mcp.ca-central-1.api.aws`
Europe (Frankfurt)	`agentaccess-mcp.eu-central-1.api.aws`
Europe (Milan)	`agentaccess-mcp.eu-south-1.api.aws`
Europe (Spain)	`agentaccess-mcp.eu-south-2.api.aws`
Europe (Ireland)	`agentaccess-mcp.eu-west-1.api.aws`
Europe (London)	`agentaccess-mcp.eu-west-2.api.aws`
Europe (Paris)	`agentaccess-mcp.eu-west-3.api.aws`
South America (São Paulo)	`agentaccess-mcp.sa-east-1.api.aws`
US East (N. Virginia)	`agentaccess-mcp.us-east-1.api.aws`
US East (Ohio)	`agentaccess-mcp.us-east-2.api.aws`
US West (Oregon)	`agentaccess-mcp.us-west-2.api.aws`

AWS IP address ranges

AWS publishes its current IP address ranges, including the ranges that the Session Gateway and CloudFront domains may resolve to, in JSON format. For information about how to download the .json file and view the current ranges, see AWS IP Address Ranges in the Amazon Web Services General Reference. Or, if you are using AWS Tools for Windows PowerShell, you can access the same information by using the `Get-AWSPublicIpAddressRange` cmdlet. For more information, see Querying the Public IP Address Ranges for AWS.

For WorkSpaces Applications users that are accessing Elastic fleets, you must allow access to the domain for the Amazon Simple Storage Service (S3) bucket that contains the application icon.

Note

If your S3 bucket has a “.” character in the name, the domain used is https://s3.<AWS Region>.amazonaws.com. If your S3 bucket does not have a “.” character in the name, the domain used is https://<*bucket name*>.s3.<*AWS Region*>.amazonaws.com.

Image Builders

Amazon WorkSpaces Applications uses EC2 instances to stream applications. You launch instances from base images, called *image builders*, which WorkSpaces Applications provides. To create your own custom image, you connect to an image builder instance, install and configure your applications for streaming, and then create your image by creating a snapshot of the image builder instance.

When you launch an image builder, you choose:

- An instance type — WorkSpaces Applications provides different instance types with various compute, memory, and graphics configurations. The instance type must align with the instance family you need. For more information, see WorkSpaces Applications Instance Families.
- An operating system — WorkSpaces Applications provides the following operating systems:
  - Windows 11
  - Windows Server 2025 Base
  - Windows Server 2022 Base
  - Windows Server 2019 Base
  - Windows Server 2016 Base
  - Red Hat Enterprise Linux 8
  - Rocky Linux 8 (Rocky Linux from CIQ)
- The subnet and security groups to use — Make sure that the subnet and security groups provide access to the network resources that your applications require. Typical network resources required by applications may include licensing servers, database servers, file servers, and application servers.

Contents

- Launch an Image Builder to Install and Configure Streaming Applications
- Connect to an Image Builder in Amazon WorkSpaces Applications
- Image Builder Actions
- Instance Metadata for WorkSpaces Applications Image Builders
- Install AMD Driver on Graphics Design Instances
- WorkSpaces Applications Base Image and Managed Image Update Release Notes

Launch an Image Builder to Install and Configure Streaming Applications

To install and configure applications to stream to your users, you start by launching an image builder instance as described in the following procedure.

> **Important**
>
> After you launch an image builder and it is running, your account may incur nominal charges. For more information, see WorkSpaces Applications Pricing.

> **Note**
>
> When you create an image builder with the AWS CLI or API, you can enable agent access so that AI agents can connect to the image builder. For more information, see Enable agent access on an image builder using the AWS CLI.

To launch an image builder

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. You can launch the image builder in the following ways:
   - If a welcome screen appears displaying two options (**Try it now** and **Get started**), choose **Get started**, **Custom set up**.

     For information about these two options, see Amazon WorkSpaces Applications FAQs.
   - If a welcome screen does not appear, choose **Quick links** in the left navigation pane, then **Custom set up**.
   - Alternatively, choose **Images** in the left navigation pane, then the **Image Builder** tab, **Launch Image Builder**.
3. For **Step 1: Choose Image**, choose a base image. If you are launching the image builder for the first time, you can use one of the latest base images released by AWS (selected by default). For a list of the latest versions of base images released by AWS, see WorkSpaces Applications Base Image and Managed Image Update Release Notes. If you have already created images,

or you want to update applications in an existing image, you can select one of your existing images. Be sure to select an image that aligns with the instance family that you need. For more information, see WorkSpaces Applications Instance Families.

> **Note**
>
> To create a unified graphics image, select the unified graphics base image (the AWS base image with the name prefix `AppStream-Graphics-NV-`). You can launch the image builder on any supported graphics family. You can also launch the image builder from one of your existing generation-specific graphics images; when you create the image, WorkSpaces Applications converts it to a unified graphics image as long as a graphics driver is detected. In either case, the resulting unified graphics image supports graphics instance families based on the installed graphics driver version. For supported driver versions, see Install NVIDIA GRID drivers. For more information, see Unified graphics images.

Choose **Next**.

4. For **Step 2: Configure Image Builder**, configure the image builder by doing the following:

   - **Name**: Type a unique name identifier for the image builder.
   - **Display name (optional)**: Type a name to display for the image builder (maximum of 100 characters).
   - **Tags (optional)**: Choose **Add Tag**, and type the key and value for the tag. To add more tags, repeat this step. For more information, see Tagging Your Amazon WorkSpaces Applications Resources.
   - **Instance Type**: Select the instance type for the image builder. Choose a type that matches the performance requirements of the applications that you plan to install. For more information, see WorkSpaces Applications Instance Families.
     - You can use stream.* instance types for images with `type = "native"` and `type = "BYOL"`. To use any of the following instance types, you must use the Import Image option to import an EC2 AMI and create an image with `type = "custom"`.
       - GeneralPurpose.*
       - MemoryOptimized.*
       - ComputeOptimized.*

  - Accelerated.*

- Configure **storage** volumes for Image Builder. By default, the storage volume matches image volume size, with a service default of 200 GB included in the hourly instance rate. You can customize your Image Builder storage capacity from 200 GB up to 500 GB based on your requirements.

  > **Note**
  >
  > Note: Storage volume size cannot be set below the image volume size. Storage capacity can be increased up to 500 GB, with additional charges applying to any storage beyond the included 200 GB. These charges apply to Image Builders in RUNNING state.

- **Available license included applications**: Select one or more license included applications that you want to install on your image.
  - Already installed license included applications display a checked checkbox next to them, and the status says **Installed**.
  - If you want to uninstall an already installed license included application, uncheck the checkbox next to it.
  - After connecting to the Image Builder, you can see apps you selected installed there.
- **VPC Endpoints (Advanced)**: You can create an interface VPC endpoint (interface endpoint), in your virtual private cloud (VPC). To start creating the interface endpoint, select **Create VPC Endpoint**. Selecting this link opens the VPC console. To finish creating the endpoint, follow steps 3 through 6 in *To create an interface endpoint*, in Tutorial: Creating and Streaming from Interface VPC Endpoints.

  After you create the interface endpoint, you can use it to keep streaming traffic within your VPC.
- **WorkSpaces Applications Agent**: This section displays only if you are not using the latest base image from AWS or a custom image that uses the latest version of the agent.

  The WorkSpaces Applications agent software runs on your streaming instances, enabling your users to connect to and stream their applications. Starting December 7, 2017, your streaming instances can be automatically updated with the latest WorkSpaces Applications agent software. This capability helps to ensure that your image builder includes the latest features, performance improvements, and security updates that are available from AWS.

You can enable automatic updates of the WorkSpaces Applications agent by creating a new image from any base image published by AWS on or after December 7, 2017. If the image that you are launching your image builder from doesn't use the latest version of the WorkSpaces Applications agent, we recommend that you select the option to launch your image builder with the latest agent.

- **IAM role (Advanced)**: When you apply an IAM role from your account to an WorkSpaces Applications image builder, you can make AWS API requests from the image builder instance without manually managing AWS credentials. To apply an IAM role to the image builder, do either of the following:
  - To use an existing IAM role in your Amazon Web Services account, choose the role that you want to use from the **IAM role** list. The role must must be accessible from the image builder. For more information, see Configuring an Existing IAM Role to Use With WorkSpaces Applications Streaming Instances.
  - To create a new IAM role, choose **Create new IAM role** and follow the steps in How to Create an IAM Role to Use With WorkSpaces Applications Streaming Instances.

5. Choose **Next**.
6. For **Step 3: Configure Network**, do the following:

- To add internet access for the image builder in a VPC with a public subnet, choose **Default Internet Access**. If you are providing internet access by using a NAT gateway, leave **Default Internet Access** unselected. For more information, see Internet Access.
- For **VPC** and **Subnet 1**, choose a VPC and at least one subnet. For increased fault tolerance, we recommend that you choose two subnets in different Availability Zones. For more information, see Configure a VPC with Private Subnets and a NAT Gateway.

  If you don't have your own VPC and subnet, you can use the default VPC or create your own. To create your own, choose the **Create a new VPC** and **Create new subnet** links to create them. Choosing these links opens the Amazon VPC console. After you create your VPC and subnets, return to the WorkSpaces Applications console and choose the refresh icon to the left of the **Create a new VPC** and **Create new subnet** links to display them in the list. For more information, see Configure a VPC for WorkSpaces Applications.
- For **Security group(s)**, choose up to five security groups to associate with this image builder. If you don't have your own security group and you don't want to use the default security group, choose the **Create new security group** link to create one. After you create your subnets in the Amazon VPC console, return to the WorkSpaces Applications console and

choose the refresh icon to the left of the **Create new security group** link to display them in the list. For more information, see Security Groups in Amazon WorkSpaces Applications.

7. For **Active Directory Domain (Optional)**, expand this section to choose the Active Directory configuration and organizational unit in which to place your streaming instance computer objects. Ensure that the selected network access settings enable DNS resolvability and communication with your directory. For more information, see Using Active Directory with WorkSpaces Applications.
8. Choose **Review** and confirm the details for the image builder. To change the configuration for any section, choose **Edit** and make the needed changes.
9. After you finish reviewing the configuration details, choose **Launch**.

> **Note**
>
> If an error message notifies you that you don't have sufficient limits (quotas) to create the image builder, submit a limit increase request through the Service Quotas console at https://console.aws.amazon.com/servicequotas/. For more information, see Requesting a quota increase in the *Service Quotas User Guide*.

10. During the image builder creation process, the status of the image builder displays as **Pending** while WorkSpaces Applications prepares the necessary resources. Choose the **Refresh** icon periodically to update the image builder status. After the status changes to **Running**, the image builder is ready to use and you can create a custom image.

Next Steps

Next, install and configure your applications for streaming, and then create an image by creating a snapshot of the image builder instance. For more information, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

Connect to an Image Builder in Amazon WorkSpaces Applications

You can connect to an image builder by doing either of the following:

- Using the WorkSpaces Applications console (for web connections only)
- Creating a streaming URL (for web or WorkSpaces Applications client connections)

> **Note**
>
> If the image builder that you want to connect to is joined to an Active Directory domain and your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection.

Contents

- WorkSpaces Applications Console (Web Connection)
- Streaming URL (WorkSpaces Applications Client or Web Connection)

WorkSpaces Applications Console (Web Connection)

To use the WorkSpaces Applications console to connect to an image builder through a web browser, complete the following steps.

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the left navigation pane, choose **Images**, **Image Builder**.
3. In the list of image builders, choose the image builder to which you want to connect. Verify that the status of the image builder is **Running**, and choose **Connect**.

   For this step to work, you might need to configure your browser to allow pop-ups from https://stream.<aws-region>.amazonappstream.com/.
4. Log in to the image builder by doing either of the following:

   - If your image builder is powered by Windows, Red Hat Enterprise Linux, or Rocky Linux, on the **Local User** tab, choose one of the following:
     - **Administrator** — Choose **Administrator** to install your applications on the image builder and create an image, or to perform any other tasks that require local administrator permissions.
     - **Template User (Windows only)** — Choose **Template User** to create default application and Windows settings.
     - **Test User** — Choose **Test User** to open your applications and verify their settings.
   - If your image builder is powered by Windows, Red Hat Enterprise Linux, or Rocky Linux, it's joined to an Active Directory domain, and you require access to resources that are managed

by Active Directory to install your applications, choose the **Directory User** tab. Then, enter the credentials for a domain account that has local administrator permissions on the image builder.

> **Note**
>
> Smart card sign in isn't supported for connections through a web browser. Instead, you must create a streaming URL and use the WorkSpaces Applications client. For information about smart card sign in, see Smart Cards.

- If your image builder is powered by Linux, you are automatically logged in as

  the **ImageBuilderAdmin** user in the Linux GNOME desktop and have root admin privileges.

Streaming URL (WorkSpaces Applications Client or Web Connection)

You can create a streaming URL to connect to an image builder through a web browser or the WorkSpaces Applications client. Unlike a streaming URL that you create to enable user access to a fleet instance, which is valid for a maximum of seven days, by default, a streaming URL that you create to access an image builder expires after one hour. To set a different expiration time, you must generate the streaming URL by using the CreateStreamingURL API action.

If the image builder that you want to connect to is joined to an Active Directory domain and your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection.

> **Note**
>
> The streaming URL provides direct access to the image builder instance. Manage the streaming URL properly and do not share it with unintended users.

> **Note**
>
> Native application mode is not supported for WorkSpaces Applications client connections to image builders. If you use the WorkSpaces Applications client to connect to an image builder and the **Start in native application mode** check box is selected, an WorkSpaces

Applications error notification displays, stating that your session was switched to classic mode.

You can create a streaming URL in any of the following ways:

- WorkSpaces Applications console
- The CreateStreamingURL API action
- The create-streaming-url AWS CLI command

To create a streaming URL and connect to the image builder by using the WorkSpaces Applications console, complete the steps in the following procedure.

To create a streaming URL and connect to the image builder by using the WorkSpaces Applications console

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the navigation pane, choose **Images**, **Image Builder**.
3. In the list of image builders, choose the image builder to which you want to connect. Verify that the status of the image builder is **Running**.
4. Choose **Actions**, **Create streaming URL**.
5. Do one of the following:
   - To connect to the image builder through the WorkSpaces Applications client, choose **Launch in Client**. When you choose this option, the WorkSpaces Applications client sign-in page is prepopulated with the streaming URL.
   - To connect to the image builder through a web browser, choose **Launch in Browser**. When you choose this option, a web browser opens with the address bar prepopulated with the streaming URL.
6. After you create the streaming URL and connect to the image builder, log in to the image builder by doing either of the following:
   - If your image builder is Windows-based and not joined to an Active Directory domain, on the **Local User** tab, choose one of the following:

- **Administrator** — Choose **Administrator** to install your applications on the image builder and create an image, or to perform any other tasks that require local administrator permissions.
- **Template User** — Choose **Template User** to create default application and Windows settings.
- **Test User** — Choose **Test User** to open your applications and verify their settings.

- If your image builder is Windows-based and joined to an Active Directory domain and you require access to resources that are managed by Active Directory to install your applications, choose the **Directory User** tab, and enter the credentials for a domain account that has local administrator permissions on the image builder.

  > **Note**
  >
  > If you're using the WorkSpaces Applications client, you can enter either your Active Directory domain password and choose **Password sign in**, or select **Choose a smart card** and provide your smart card PIN when prompted.
  > If you're using a web browser, you must enter your Active Directory domain password. Smart card sign in is supported only for WorkSpaces Applications client connections to streaming instances.

- If your image builder is Linux-based, you are automatically logged in as

  the **ImageBuilderAdmin** user in the Linux GNOME desktop and have root admin privileges.

Image Builder Actions

You can perform the following actions on an image builder, depending on the current state (status) of the image builder instance.

Delete

Permanently delete an image builder.

The instance must be in a **Stopped** state.

Connect

Connect to a running image builder. This action starts a desktop streaming session with the image builder to install and add applications to the image, and create an image.

The instance must be in a **Running** state.

Start

Start a stopped image builder. A running instance is billed to your account.

The instance must be in a **Stopped** state.

Stop

Stop a running image builder. A stopped instance is not billed to your account.

The instance must be in a **Running** state.

Update

Change the description and display name of an image builder by using the `UpdateImageBuilder` API operation or the `aws appstream update-image-builder` AWS CLI command.

You can use this operation to enable or disable agent access on an image builder, and to change agent access configuration settings. For more information, see Step 5: Change agent access settings or disable agent access.

The instance must be in a **Stopped** state.

Manage license included applications

Install or uninstall license included applications.

- Select one or more license included applications that you want to install on your image.
  - Already installed license included applications display a checked checkbox next to them, and the status says **Installed**.
  - If you want to uninstall an already installed license included application, uncheck the checkbox next to it.
  - If one or more applications are in **Failed to install**, **Installation timed out**, or **Installation skipped** status, the service retries the installation after you select **Apply Changes**. If you want to uninstall these applications, uncheck the checkbox next to them.
  - If one or more applications are in **Failed to uninstall**, **Uninstallation timed out**, or **Uninstallation skipped** status, and you want to uninstall these applications, uncheck the checkbox next to them
- Select **Apply Changes**.

- Verify the **Manage license included applications summary** and select **Confirm and deploy updates**.
- After connecting to the Image Builder, you can see the apps you selected already installed there.

None of these actions can be performed on an instance in any of the following intermediate states:

- **Pending**
- **Snapshotting**
- **Stopping**
- **Starting**
- **Deleting**
- **Updating**
- **Pending Qualification**
- **Pending Syncing Apps**
- **Syncing Apps**

Instance Metadata for WorkSpaces Applications Image Builders

WorkSpaces Applications image builder instances have instance metadata available through Windows environment variables. You can use the following environment variables in your applications and scripts to modify your environment based on the image builder instance details.

Environment Variable	Context	Description
AppStream _Image_Arn	Machine	The ARN of the image that was used to create the streaming instance.
AppStream _Instance_Type	Machine	The instance type of the streaming instance. For example, **`stream.standard.medium`** .
AppStream _Resource_Type	Machine	The type of WorkSpaces Applications resource. The value is either **`fleet`** or **`imagebuilder`** .

Environment Variable	Context	Description
AppStream _Resource _Name	Machine	The name of the image builder.

On Linux image builders, environment variables are exported through the script at **/etc/profile.d/appstream_system_vars.sh**. To access the environment variables, you can explicitly source this file in your application.

Install AMD Driver on Graphics Design Instances

If you need to update the AMD driver on your Windows Image Builder that is using a Graphics Design instance, you can either use the latest WorkSpaces Applications Graphics Design base images, or download the AMD driver and install it on your Image Builder. If you need to update the AMD driver for an existing image of the Graphics Design instance family, you can use managed WorkSpaces Applications image updates. For more information, see the section called "Update an Image by Using Managed WorkSpaces Applications Image Updates".

The AMD driver download is available to AWS customers only. By downloading, you agree to use the downloaded software only to build images for use with WorkSpaces Applications Graphics Design instances using AMD FirePro S7150x2 Server GPU hardware. Upon installation of the software, you are bound by the terms of the AMD Software End User License Agreements.

The latest AMD driver version for Graphics Design instances is version 24.20.13028.5012.

Before you begin, make sure that you meet the following prerequisites:

- Configure default credentials for the AWS Tools for Windows PowerShell on your Windows instance. For more information, see Getting Started with the AWS Tools for Windows PowerShell.
- IAM users must have the permissions granted by the **AmazonS3ReadOnlyAccess** policy.

To install the AMD driver on your Image Builder

1. Connect to your Windows Image Builder instance and open a PowerShell window as an Administrator.

2. Download the drivers from Amazon S3 to your desktop using the following PowerShell commands:

```
$Bucket = "appstream2-driver-patches"
$LocalPath = "$home\Desktop\AMD"
$Objects = Get-S3Object -BucketName $Bucket -Region us-east-1
foreach ($Object in $Objects) {
   $LocalFileName = $Object.Key
   if ($LocalFileName -ne '' -and $Object.Size -ne 0) {
     $LocalFilePath = Join-Path $LocalPath $LocalFileName
     Copy-S3Object -BucketName $Bucket -Key $Object.Key -LocalFile $LocalFilePath -
Region us-east-1
   }
 }
```

3. Unzip the downloaded driver file and run the installer using the following PowerShell commands:

```
Expand-Archive $LocalFilePath -DestinationPath $home\Desktop -Verbose
$Driverdir = Get-ChildItem $home\Desktop\ -Directory -Filter "*210426a-366782C*"
Write-Host $Driverdir
pnputil /add-driver $home\Desktop\$Driverdir\Packages\Drivers\Display\WT6A_INF
\*inf /install
```

4. Follow the instructions to install the driver and reboot your instance as required.
5. To verify that the GPU is working properly, check **Device Manager**. You should see **AMD MxGPU** listed as a display adapter, with the latest driver version.

WorkSpaces Applications Base Image and Managed Image Update Release Notes

Amazon WorkSpaces Applications provides base images to help you create images that include your own applications. Base images are Amazon Machine Images (AMIs) that contain software configurations specific to the operating system. For WorkSpaces Applications, each base image

includes the WorkSpaces Applications agent and the latest version of one of the following operating systems:

> **Important**
>
> Operating system versions that are no longer supported by the vendor are not guaranteed to work and are not supported by AWS Support.

- Windows Server 2025 Base — Available on the following image types: Base, Graphics G4dn, Graphics G5, Graphics G6, Graphics G7, and Graphics NV
- Windows Server 2022 Base — Available on the following image types: Base, Graphics G4dn, Graphics G5, Graphics G6, Graphics G7, and Graphics NV
- Windows Server 2019 Base — Available on the following image types: Base, Graphics G4dn and Graphics G5
- Windows Server 2016 Base — Available on the following image types: Base, Graphics G4dn and Graphics G5
- Red Hat Enterprise Linux 8 – Available on the following image types: Base, Graphics G4dn, Graphics G5, Graphics G6, Graphics G7, and Graphics NV
- Rocky Linux 8 – Available on the following image types: Base, Graphics G4dn, Graphics G5, Graphics G6, Graphics G7, and Graphics NV

After you create your own image that includes your own applications, you are responsible for installing and maintaining the updates for the operating system, your applications, and their dependencies. WorkSpaces Applications provides an automated way to update your image using managed WorkSpaces Applications image updates. With managed image updates, you select the image that you want to update. WorkSpaces Applications creates an image builder in the same AWS account and Region to install the updates and create the new image. After the new image is created, you can test it on a pre-production fleet before updating your production fleets or sharing the image with other AWS accounts. For more information, see "Keep Your WorkSpaces Applications Image Up-to-Date" in Administer Your Amazon WorkSpaces Applications Images.

For information about the latest WorkSpaces Applications agent, see WorkSpaces Applications Agent Release Notes.

The following table lists the latest released images.

Image type	Image name
Base	• AppStream-WinServer2025-08-31-2026 • AppStream-WinServer2022-08-31-2026 • AppStream-WinServer2019-08-31-2026 • AppStream-WinServer2016-11-10-2025 • AppStream-RHEL8-08-31-2026 • AppStream-RockyLinux8-08-31-2026
Graphics G4dn	• AppStream-Graphics-G4dn-WinServer2025-08-31-2026 • AppStream-Graphics-G4dn-WinServer2022-08-31-2026 • AppStream-Graphics-G4dn-WinServer2019-08-31-2026 • AppStream-Graphics-G4dn-WinServer2016-11-10-2025 • AppStream-Graphics-G4dn-RHEL8-08-31-2026 • AppStream-Graphics-G4dn-RockyLinux8-08-31-2026
Graphics G5	• AppStream-Graphics-G5-WinServer2025-08-31-2026 • AppStream-Graphics-G5-WinServer2022-08-31-2026 • AppStream-Graphics-G5-WinServer2019-08-31-2026 • AppStream-Graphics-G5-WinServer2016-11-10-2025 • AppStream-Graphics-G5-RHEL8-08-31-2026 • AppStream-Graphics-G5-RockyLinux8-08-31-2026
Graphics G6	• AppStream-Graphics-G6-WinServer2025-08-31-2026 • AppStream-Graphics-G6-WinServer2022-08-31-2026 • AppStream-Graphics-G6-RHEL8-08-31-2026 • AppStream-Graphics-G6-RockyLinux8-08-31-2026
Graphics G7	• AppStream-Graphics-G7-WinServer2025-08-31-2026 • AppStream-Graphics-G7-WinServer2022-08-31-2026

Image type	Image name
	• AppStream-Graphics-G7-RHEL8-08-31-2026 • AppStream-Graphics-G7-RockyLinux8-08-31-2026
Graphics NV (unified)	• AppStream-Graphics-NV-WinServer2025-09-22-2026 • AppStream-Graphics-NV-WinServer2022-09-22-2026 • AppStream-Graphics-NV-RHEL8-09-22-2026 • AppStream-Graphics-NV-RockyLinux8-09-22-2026
Sample apps	Amazon-AppStream2-Sample-Image-06-17-2024 For information about how to access this base image, see Get Started with Amazon WorkSpaces Applications: Set Up With Sample Applications.

The following table lists the software components for the latest released base images and the components that are available if you update your image using managed image updates. If the version is marked “latest”, the current stable software component available from the vendor will be installed. If the version is marked “not included”, managed image updates is not managing the component and the version will not be changed when you update your image.

The following table lists the software components for the latest released Windows, Rocky Linux, and Red Hat Enterprise Linux base images and Managed image updates.

Windows

Software component	Latest base images (August 31, 2026)	Managed image updates (August 31, 2026)
Amazon AWS (AvsCamera) Driver	1.0.27.0	1.0.23.0
Amazon CloudWatch Agent	1.300072.0b1766	1.300069.0b1529
SSM Agent	3.3.4851.0	3.3.4851.0
NICE DCV Virtual Display	2026.0-22938	2026.0-22938

Software component	Latest base images (August 31, 2026)	Managed image updates (August 31, 2026)
AMD Driver for Graphics Design instances	24.20.13028.7002	24.20.13028.7002
AppStream 2.0 Agent	LATEST	--
AWS Command Line Interface (AWS CLI)	2.36.35 (Windows Server 2022/2025) 1.46.1 (Windows Server 2019) 1.40.24 (Windows Server 2016)	Not included
Firefox	144 (Windows Server 2016/2019)	Not included
Microsoft Message Queuing (MSMQ)	Installed with Windows Server	Installed with Windows Server
NVIDIA Graphics Driver for G4dn, G5 and G6 instances	596.86 (Windows Server 2022/2025) 539.19 (Windows Server 2019) 512.78 (Windows Server 2016)	596.36 (Windows Server 2022/2025) 539.19 (Windows Server 2019) 512.78 (Windows Server 2016)
Process monitor	4.01	Latest
Quality Windows Audio/Video Experience (qWAVE)	Installed with Windows Server	Installed with Windows Server

Software component	Latest base images (August 31, 2026)	Managed image updates (August 31, 2026)
Visual C++ redistributable packages	Microsoft Visual C++ 2013 Redistributable (x64) - 12.0.40664.0 Microsoft Visual C++ 2015-2022 Redistributable (x64) - 14.44.35211	Microsoft Visual C++ 2013 Redistributable (x64) - 12.0.30501 Microsoft Visual C++ 2015-2022 Redistributable (x64) - 14.44.35211
Windows Server updates	Latest as of August 12, 2026 (KB5120242)	Latest
WinSCard Filter Driver	1.50.1.0	1.0.19.0
Paravirtual (PV) driver	8.6.0	8.6.1
ENA driver	2.11.0.0	2.11.0
AWS NVMe driver	1.8.2	1.8.2

Rocky Linux

Software component	Latest base images (August 31, 2026)	Managed image updates (August 31, 2026)
AWS Command Line Interface (AWS CLI)	2.36.27	2.36.27
Amazon CloudWatch Agent	1.300071.0b1720-1	1.300071.0b1720-1
SSM Agent	3.3.4851.0-1	3.3.4851.0-1
NICE DCV Server AppStream	2026.0.24079-1	2026.0.24079-1
Cloud-init	23.4-78_10.11.0.2	23.4-78_10.11.0.2
Kernel	4.18.0-553.156.1	4.18.0-553.156.1

Software component	Latest base images (August 31, 2026)	Managed image updates (August 31, 2026)
NVIDIA Graphics Driver for G4dn, G5, G6, and G7 instances	595.91.07	595.91.07
Cuda Version	13.2	13.2

Red Hat Enterprise Linux

Software component	Latest base images (August 31, 2026)	Managed image updates (August 31, 2026)
AWS Command Line Interface (AWS CLI)	2.36.27	2.36.27
Amazon CloudWatch Agent	1.300071.0b1720-1	1.300071.0b1720-1
SSM Agent	3.3.4851.0-1	3.3.4851.0-1
NICE DCV Server AppStream	2026.0.24079-1	2026.0.24079-1
Cloud-init	23.4-78_10.11	23.4-78_10.11
Kernel	4.18.0-553.156.1	4.18.0-553.156.1
NVIDIA Graphics Driver for G4dn, G5, G6, and G7 instances	595.91.07	595.91.07
Cuda Version	13.2	13.2

Important

The following public images are deprecated and therefore no longer available from AWS:

- 2016/2019/2022 Windows images released before May 30, 2025

- Images for the Graphics Desktop, Graphics Design, and Graphics Pro instance families
- Amazon Linux 2 (AL2) base images, which are no longer available from AWS after 04/15/2026 due to End of Support for AL2 on WorkSpaces Applications

If you want to use an image for a multi-session fleet, the image must meet the following conditions:

- The image must be created from a base image released on or after June 12, 2023. Or, the image must be updated by using managed WorkSpaces Applications image updates released on or after September 6, 2023. For more information, see the section called "Update an Image by Using Managed WorkSpaces Applications Image Updates".
- The WorkSpaces Applications agent release version must be 09-06-2023 or later. For more information, see the section called "Manage Agent Versions".
- If you have updated your image using Managed WorkSpaces Applications Image updates, then the WorkSpaces Applications agent release version is not applicable. Your image must be updated using a Managed Image Update released on or after September 6, 2023. For more information, see the section called "Update an Image by Using Managed WorkSpaces Applications Image Updates".
- Multi-session fleets are supported only for Microsoft Server 2019, 2022, and 2025.

The following table describes all released base images.

Release	Platform	Image	Changes
08/31/2026	Red Hat Enterprise Linux	• Base • Graphics G4dn • Graphics G5 • Graphics G6 • Graphics G7	• Support added for the Graphics G7 instance family • NVIDIA Graphics Driver updated to 595.91.07 • CW Agent updated to 1.300071.0b1720-1 • SSM Agent updated to 3.3.4851.0-1 • Amazon DCV Server updated to 2026.0.24079-1

Release	Platform	Image	Changes
08/31/2026	Rocky Linux	• Base • Graphics G4dn • Graphics G5 • Graphics G6 • Graphics G7	• Support added for the Graphics G7 instance family • NVIDIA Graphics Driver updated to 595.91.07 • CW Agent updated to 1.300071.0b1720-1 • SSM Agent updated to 3.3.4851.0-1 • Amazon DCV Server updated to 2026.0.24079-1
04/16/2026	Windows	• Base • Graphics G4dn • Graphics G5 • Graphics G6	• Support added for the Agent Access feature • General bug fixes and improvements • CW Agent updated to 1.4.37922 • SSM Agent updated to 3.3.4121.0 • Amazon DCV Server updated to 2026.0-21918
04/16/2026	Red Hat Enterprise Linux	• Base • Graphics G4dn • Graphics G5 • Graphics G6	• General bug fixes and improvements • CW Agent updated to 1.300066.1b1374-1 • SSM Agent updated to 3.3.4177.0-1
04/16/2026	Rocky Linux	• Base • Graphics G4dn • Graphics G5 • Graphics G6	• General bug fixes and improvements • CW Agent updated to 1.300066.1b1374-1 • SSM Agent updated to 3.3.4177.0-1

Release	Platform	Image	Changes
02/18/2026	Red Hat Enterprise Linux	• Base • Graphics G4dn • Graphics G5 • Graphics G6	• 4K resolution support for non-accelerated instance types • CW Agent updated to 1.300064.0b1337 • SSM Agent updated to 3.3.3598.0 • NICE DCV Server AppStream updated to 2024.0.17598
02/18/2026	Rocky Linux	• Base • Graphics G4dn • Graphics G5 • Graphics G6	• 4K resolution support for non-accelerated instance types • CW Agent updated to 1.300064.0b1337 • SSM Agent updated to 3.3.3598.0 • NICE DCV Server AppStream updated to 2024.0.17598
12/18/2025	Windows	• Base • Graphics G4dn • Graphics G5 • Graphics G6	• Includes support for Windows Server 2025
11/10/2025	Windows	• Base • Graphics G4dn • Graphics G5 • Graphics G6	• General bug fixes and improvements • Updated to latest NVIDIA drivers • Includes new CloudWatch Agent 1.4.37917 • Include new SSM Agent 3.3.3050.0
11/10/2025	Red Hat Enterprise Linux	• Base • Graphics G4dn • Graphics G5 • Graphics G6	• General bug fixes and improvements • Updated to latest NVIDIA drivers • CW Agent updated to 1.300061.0b1289-1 • SSM Agent updated to 3.3.3270.0-1 • NICE DCV Server AppStream updated to 2024.0.17598-18

Release	Platform	Image	Changes
11/10/2025	Rocky Linux	• Base • Graphics G4dn • Graphics G5 • Graphics G6	• General bug fixes and improvements • Updated to latest NVIDIA drivers • CW Agent updated to 1.300061.0b1289-1 • SSM Agent updated to 3.3.3270.0-1 • NICE DCV Server AppStream updated to 2024.0.17598-18
09-05-2025	Red Hat Enterprise Linux	• Base • Graphics G4dn • Graphics G5 • Graphics G6	• Updated to latest NVIDIA drivers • Includes new CloudWatch Agent 1.300057.1b1167-1 • Include new SSM Agent 3.3.2958.0-1 • General bug fixes and improvements
09-05-2025	Rocky Linux	• Base • Graphics G4dn • Graphics G5 • Graphics G6	• Updated to latest NVIDIA drivers • Includes new CloudWatch Agent 1.300057.1b1167-1 • Include new SSM Agent 3.3.2958.0-1 • General bug fixes and improvements
05-30-2025	Windows	• Base • Graphics G4dn • Graphics G5	• Updated to latest NVIDIA drivers • Amazon DCV updated to version 2024.0-19143 • Includes new CloudWatch Agent 1.4.37911 • Include new SSM Agent 3.3.2299.0

Release	Platform	Image	Changes
05-30-2025	Red Hat Enterprise Linux	• Base • Graphics G4dn • Graphics G5	• Updated to latest NVIDIA drivers • Amazon DCV server updated to version 2024.0.17598-1 • Includes new CloudWatch Agent 1.300055.1-1 • Include new SSM Agent 3.3.2471.0-1
05-30-2025	Rocky Linux	• Base • Graphics G4dn • Graphics G5	• Updated to latest NVIDIA drivers • Amazon DCV server updated to version 2024.0.17598-1 • Includes new CloudWatch Agent 1.300055.1-1 • Include new SSM Agent 3.3.2471.0-1
02-11-2025	Windows	• Base • Graphics G4dn • Graphics G5	• Includes new ENA and NVMe drivers • Includes updated PV driver • Updating drivers and software with new versions available
12-19-2024	Rocky Linux	• Base • Graphics G4dn • Graphics G5	• Support for Rocky Linux 8
10-22-2024	Windows	• Base • Graphics Design • Graphics G4dn • Graphics Pro • Graphics G5	• Includes latest NVIDIA drivers for Windows • Includes Microsoft Visual C++ 2015-2022 Redistributable (x64) - 14.40.33816 • Includes Microsoft security updates up to October 9, 2024

Release	Platform	Image	Changes
07-30-2024	Red Hat Enterprise Linux	• Base • Graphics G4dn • Graphics G5	• Includes support for Red Hat Enterprise Linux 8
06-17-2024	Windows	• Base • Graphics Design • Graphics G4dn • Graphics Pro • Graphics G5	• Includes Microsoft security updates up to June 13, 2024 • Includes CloudWatch Agent 1.4.37896 • Includes SSM Agent 3.3.484.0 • Includes AWS Command Line Interface (AWS CLI) (WinServer 2016/2019) 1.33.9 • Includes AWS Command Line Interface (AWS CLI) (WinServer 2022) 2.16.9.0
05-08-2024	Windows	• Base • Graphics Design • Graphics G4dn • Graphics Pro • Graphics G5	• Includes Microsoft security updates up to May 2024 • Includes latest NVIDIA Graphics Driver (552.08) for Graphics Pro and G4dn instances for Windows Server 2016 and Windows Server 2019 • Includes CloudWatch Agent 1.4.37891 • Includes SSM Agent 3.3.131.0-1 • Includes AWS Command Line Interface (AWS CLI) 1.32.89 • Includes AWSVirtualSmartCardReader 1.0.0.59
05-08-2024	Linux	• Base • Graphics G4dn • Graphics Pro • Graphics G5	• Updated Linux to version 2.0.20240412.0. For more information, see Amazon Linux 2.0.20240412.0 release notes.

Release	Platform	Image	Changes
03-24-2024	Windows	• Base • Graphics Design • Graphics G4dn • Graphics Pro • Graphics G5	• Includes Microsoft security updates up to March 2024 • Includes latest NVIDIA Graphics Driver (551.61) for Graphics Pro and G4dn instances for Windows Server 2016 and Windows Server 2019 • Includes CloudWatch Agent 1.3.50742 • Includes SSM Agent 3.2.2303.0 • Includes AWS Command Line Interface (AWS CLI) 2.15.33.0 • Includes AWSVirtualSmartCardReader 1.0.0.59
03-24-2024	Linux	• Base • Graphics G4dn • Graphics Pro • Graphics G5	• Updated Linux to version 2.0.20240318.0. For more information, see 2.0.20240318.0.
01-26-2024	Windows	• Base • Graphics Design • Graphics G4dn • Graphics Pro • Graphics G5	• Includes Microsoft security updates up to January 2024
12-11-2023	Windows	• Base • Graphics G4dn • Graphics G5	• Add support for Windows Server 2022

Release	Platform	Image	Changes
11-13-2023	Windows	• Base • Graphics Design • Graphics G4dn • Graphics Pro • Graphics G5	• Includes Microsoft security updates up to November 2023
11-13-2023	Amazon Linux 2	• Base • Graphics G4dn • Graphics Pro • Graphics G5	• Updated Linux to version 2.0.20231101.0. For more information, see Amazon Linux 2.0.20231101.0 release notes.
06-12-2023	Windows	• Base • Graphics Design • Graphics G4dn • Graphics Pro	• Includes Microsoft security updates up to June 2023
06-11-2023	Amazon Linux 2	• Base • Graphics G4dn • Graphics Pro	• Updated Linux to version 2.0.20230530.0. For more information, see Amazon Linux 2 2.0.20230530.0 release notes.
03-29-2023	Windows	• Base • Graphics Design • Graphics G4dn • Graphics Pro	• Includes Microsoft security updates up to February 2023

Release	Platform	Image	Changes
03-15-2023	Amazon Linux 2	• Base • Graphics G4dn • Graphics Pro	• Updated Linux to version 2.0.20220805.0. For more information, see Amazon Linux 2 2.0.20230221.0 release notes. • Improves Webcam experience • Resolves an issue that prevents WorkSpaces Applications fleet instances from provisioning when the system cryptography is set to use FIPS-compliant algorithms
10-05-2022	Windows	• Base • Graphics Design • Graphics G4dn • Graphics Pro	• Includes Microsoft security updates up to September 13, 2022
09-21-2022	Amazon Linux 2	• Base • Graphics G4dn • Graphics Pro	• Updated Linux to version 2.0.20220805.0. For more information, see Amazon Linux 2.0.20220805.0 release notes. • Includes Image Assistant GUI • Includes webcam support
09-14-2022	Amazon Linux 2	• Graphics G4dn • Graphics Pro	• Includes NVIDIA Graphics Driver (510.85.02)
09-01-2022	Windows	• Graphics G4dn • Graphics Pro	• Includes NVIDIA Graphics Driver (473.47) for Windows Server 2012 R2 • Includes NVIDIA Graphics Driver (512.78) for Windows Server 2016 and Windows Server 2019

Release	Platform	Image	Changes
07-12-2022	Windows	• Base • Graphics Design • Graphics G4dn • Graphics Pro	• Includes Microsoft security updates up to June 14, 2022 • Includes latest AMD Driver (24.20.13 028.7002) for Graphics Design instances for Windows Server 2016 and Windows Server 2019 • Includes latest NVIDIA Graphics Driver (472.98) for Graphics Pro and G4dn instances for Windows Server 2012R2 • Includes latest NVIDIA Graphics Driver (511.65) for Graphics Pro and G4dn instances for Windows Server 2016 and Windows Server 2019 • Includes CloudWatch Agent 1.3.50742 • Includes SSM Agent 3.1.1575.0 • Includes AWS Command Line Interface (AWS CLI) 1.23.11
06-20-2022	Amazon Linux 2	• Base • Graphics G4dn • Graphics Pro	• Updated Linux to version 2.0.20220 426.0. For more information, see Amazon Linux 2.0.20220426.0 release notes.
03-03-2022	Windows	• Base • Graphics Design • Graphics G4dn • Graphics Pro	• Includes Microsoft security updates up to January 11, 2022
02-18-2022	Amazon Linux 2	• Base • Graphics G4dn • Graphics Pro	• Updated Linux to version 2.0.20211 223.0. For more information, see Amazon Linux 2.0.20211223.0 release notes. • Latest Linux base images

Release	Platform	Image	Changes
11-19-2021	Amazon Linux 2	• Base • Graphics G4dn • Graphics Pro	• Latest Linux base images, including blank screen fixes on small instance types
11-15-2021	Amazon Linux 2	• Base • Graphics G4dn • Graphics Pro	• Linux base images
10-08-2021	Windows	• Base • Graphics Design • Graphics G4dn • Graphics Pro • Sample apps	• Includes Microsoft security updates up to September 15, 2021 • AWS Tools for PowerShell updated to version 3.15.1398
07-19-2021	Windows	• Base • Graphics Design • Graphics G4dn • Graphics Pro	• Includes Microsoft Windows updates up to July 13, 2021
06-01-2021	Windows	• Base • Graphics Design • Graphics G4dn • Graphics Pro	• Includes Microsoft Windows updates up to April 14, 2021 • Includes AMD driver version 24.20.130 28.5012 for Graphics Design instances

Release	Platform	Image	Changes
12-28-2020	Windows	• Base • Graphics Design • Graphics G4dn • Graphics Pro	• Includes a driver that adds support for using smart cards. Smart cards can be used for Windows sign in, Active Directory-joined streaming instances , and in-session authentication for streaming applications • Includes Microsoft Windows updates up to December 9, 2020 • Includes AWS CLI version 1.18.138 • Includes NVIDIA Graphics Driver version 451.48 for Graphics Pro and Graphics G4dn instances
07-16-2020	Windows	• Base • Graphics Design • Graphics G4dn • Graphics Pro	• Includes Microsoft Windows updates up to June 9, 2020 • Includes AWS CLI version 1.18.86 • Includes NVIDIA Graphics Driver version 441.66 for Graphics Pro instances
04-22-2020	Windows	• Base (Windows Server 2019) • Graphics Design (Windows Server 2019) • Graphics G4dn (Windows Server 2019) • Graphics Pro (Windows Server 2019)	• Includes Microsoft Windows updates up to March 10, 2020 • Includes AWS CLI version 1.18.21 • Includes NVIDIA Graphics Driver version 441.66 for Graphics Pro instances

Release	Platform	Image	Changes
03-18-2020	Windows	• Base • Graphics Design • Graphics Pro	• Includes Microsoft Windows updates up to February 11, 2020 • Includes AWS CLI version 1.17.5 • Includes NVIDIA Graphics Driver version 412.16 for Graphics Pro instances
03-16-2020	Windows	• Graphics G4dn	• Adds support for Graphics G4dn instances based on the EC2 G4dn family (Windows Server 2012 R2) • Includes Microsoft Windows updates up to February 11, 2020 • Includes AWS CLI version 1.17.5
03-05-2020	Windows	• Graphics G4dn	• Adds support for Graphics G4dn instances based on the EC2 G4dn family (Windows Server 2016 and Windows Server 2019) • Includes Microsoft Windows updates up to February 11, 2020 • Includes AWS CLI version 1.17.5
01-13-2020	Windows	• Graphics Design	• Adds support for Windows Server 2019, with Microsoft Windows updates up to November 12, 2019
12-12-2019	Windows	• Base • Graphics Design • Graphics Pro	• Includes Microsoft Windows updates up to November 12, 2019 • Includes AWS CLI version 1.16.284 • Includes a new version of the SSM Agent (v2.3.760.0), which resolves an issue that prevented streaming instances from being provisioned

Release	Platform	Image	Changes
09-18-2019	Windows	• Base • Graphics Design • Graphics Pro	• Includes Microsoft Windows updates up to August 13, 2019 for all Base and Graphics Pro instances and for Graphics Design Windows Server 2012 R2. Graphics Design Windows Server 2016 instances already include this version. • Includes AWS CLI version 1.16.222 for all Base and Graphics Pro instances and Graphics Design Windows Server 2012 R2. Graphics Design Windows Server 2016 instances already include this version. • Includes a fix to prevent Windows Defender from being enabled by default on Windows Server 2016 and Windows Server 2019 image builder instances. For more information, see Windows Update and Antivirus Software on Amazon WorkSpaces Applications.
09-05-2019	Windows	• Graphics Design	• Adds support for Windows Server 2016 • Includes Microsoft Windows updates up to August 13, 2019 • Includes AWS CLI version 1.16.222 • Includes AMD Driver version 24.20.130 28.3002 for Graphics Design instances (compatible with Windows Server 2016)
06-24-2019	Windows	• Base • Graphics Pro	• Adds support for Windows Server 2016 and Windows Server 2019

Release	Platform	Image	Changes
05-28-2019	Windows	• Base • Graphics Design • Graphics Pro	• Includes Microsoft Windows updates up to May 14, 2019
04-29-2019	Windows	• Base • Graphics Design • Graphics Pro	• Includes Microsoft Windows updates up to April 20, 2019 • Includes AWS CLI version 1.16.126 • Includes NVIDIA Graphics Driver 412.16 for Graphics Pro instances
01-22-2019	Windows	• Base • Graphics Design • Graphics Pro	• Includes Microsoft Windows updates up to December 10, 2018 • Includes AWS CLI version 1.16.84 • Includes NVIDIA Graphics Driver version 391.58 for Graphics Pro instances
06-12-2018	Windows	• Base • Graphics Design • Graphics Desktop • Graphics Pro	• Includes Microsoft Windows updates up to May 9, 2018 • Includes Windows PowerShell 5.1
05-02-2018	Windows	• Base • Graphics Design • Graphics Desktop • Graphics Pro	• Includes Microsoft Windows updates up to April 10, 2018 • Adds the following language packs: Japanese, Korean, Portuguese (Brazil), Thai, Chinese (Simplified), Chinese (Traditional)

Release	Platform	Image	Changes
03-19-2018	Windows	• Base • Graphics Design • Graphics Desktop • Graphics Pro	• Includes Microsoft Windows updates up to February 23, 2018 • Includes the following language packs: German, French, Italian, Spanish, Dutch • Resolves intermittent issues with using Microsoft Visio and Microsoft Project applications during streaming sessions
01-24-2018	Windows	• Base • Graphics Design • Graphics Desktop • Graphics Pro	• Includes Microsoft Windows updates up to January 5, 2018 • Includes Microsoft Windows updates for the Spectre and Meltdown vulnerabilities • Enables a default profile to be created on image builders and used for the AWS Command Line Interface (CLI) during streaming sessions
01-01-2018	Windows	• Base • Graphics Design • Graphics Desktop • Graphics Pro	• Resolves an issue with connectivity to WorkSpaces Applications instances
12-07-2017	Windows	• Base • Graphics Design • Graphics Desktop • Graphics Pro	• Includes Microsoft Windows updates up to November 19, 2017 • Adds support for managed WorkSpaces Applications agent updates
11-13-2017	Windows	• Base	• Resolves an issue with Microsoft Office 365 applications not working during streaming sessions • Includes Microsoft Windows updates up to October 11, 2017

Release	Platform	Image	Changes
09-05-2017	Windows	• Base • Graphics Design • Graphics Desktop • Graphics Pro	• New Graphics Design instance family • Support for On-Demand fleets • Updated approach for session context • Includes Microsoft Windows updates up to August 9, 2017 • Resolves an intermittent issue with applications not coming to the foreground • Resolves an intermittent issue with applications not appearing in tile view
07-25-2017	Windows	• Graphics Desktop • Graphics Pro	• New Graphics Desktop and Graphics Pro instance families • Adds support for 2 K resolution
07-24-2017	Windows	• Base	• Includes Microsoft Windows updates up to July 13, 2017 • Adds support for Microsoft Active Directory domains
06-20-2017	Windows	• Base • Sample apps	• Optimizes application launch performan ce • Resolves an issue with applications not displaying in tile view • Resolves an issue with applications displaying in tile view only • Resolves an issue with applications displaying multiple times in tile view • Resolves an issue with recently launched application windows not appearing in the foreground • Resolves an issue with page margins when printing

Release	Platform	Image	Changes
05-18-2017	Windows	• Base • Sample apps	• Adds support for WorkSpaces Applications home folders • Includes Microsoft Windows updates up to May 16, 2017 • Resolves an intermittent network issue that affects internet connections from streaming instances • Resolves an issue with application tiles not functioning correctly

Images

You can create Amazon WorkSpaces Applications images that contain applications you can stream to your users and default system and application settings to enable your users to get started with those applications quickly. However, after you create an image, you can't change it. To add other applications, update existing applications, or change image settings, you must start and reconnect to the image builder that you used to create the image. If you deleted that image builder, launch a new image builder that is based on your image. Then make your changes and create a new image. For more information, see Launch an Image Builder to Install and Configure Streaming Applications and Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

Images that are available to you are listed in the **Image Registry** in the WorkSpaces Applications console. They are categorized as public, private, or shared. You can use any of these image types to launch an image builder and set up an WorkSpaces Applications fleet. Shared images are owned by other Amazon Web Services accounts and shared with you. Permissions set on images that are shared with you may limit what you can do with those images. For more information, see Administer Your Amazon WorkSpaces Applications Images.

Contents

- Default Application and Windows Settings and Application Launch Performance in Amazon WorkSpaces Applications
- Manage WorkSpaces Applications Agent Versions
- Unified graphics images
- WorkSpaces Applications Agent Release Notes
- Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console
- Administer Your Amazon WorkSpaces Applications Images
- Create Your Amazon WorkSpaces Applications Image Programmatically by Using the Image Assistant CLI Operations
- Create Your Linux-Based Images
- Use Session Scripts to Manage Your Amazon WorkSpaces Applications Users' Streaming Experience

Default Application and Windows Settings and Application Launch Performance in Amazon WorkSpaces Applications

You can create default application and Windows settings to enable your users to get started with their applications quickly, so that they won't need to create or configure the settings themselves.

WorkSpaces Applications optimizes the launch performance of your applications for your users' streaming sessions. To ensure that all of the required files are included in this process, you may need to manually add certain files and folders to the optimization manifest.

Contents

- Creating Default Application and Windows Settings for Your WorkSpaces Applications Users
- Optimizing the Launch Performance of Your Applications in Amazon WorkSpaces Applications

Creating Default Application and Windows Settings for Your WorkSpaces Applications Users

Application customizations and Windows settings that are saved to the Windows user profile folder or the user registry hive can be set as defaults. When you save the default settings by using the **Template User** in Image Assistant, WorkSpaces Applications replaces the Windows default user profile with the profile that you configure. The Windows default user profile is then used to create the initial settings for users in the fleet instance. If the application or Windows settings that you configure don't work in the fleet, confirm that they are saved in the Windows user profile. For more information, see Step 3: Create Default Application and Windows Settings in Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

Default settings that you can create and configure include:

- Application preferences, including a browser home page, toolbar customizations, and security settings.
- Application data settings, including browser bookmarks and connection profiles.
- Windows experience settings, including displaying file name extensions and hidden folders.

Additionally, you can modify or disable Internet Explorer security settings, such as Enhanced Security Configuration (ESC). For more information, see Disable Internet Explorer Enhanced Security Configuration in Amazon WorkSpaces Applications.

Optimizing the Launch Performance of Your Applications in Amazon WorkSpaces Applications

When you create an image, WorkSpaces Applications requires that you optimize the launch performance of your applications for your users' streaming sessions. When your applications are opened during this process, make sure that they use the initial components required by your users. Doing so ensures that these components are captured by the optimization process. In some cases, not all of the files required for the optimizations are detected. Examples of such files would be plug-ins or components that aren't opened in the image builder. To ensure that all of the files needed for your application are captured, you can include them in the optimization manifest. Adding files to the optimization manifest may increase the time it takes for fleet instances to be created and made available for users. Doing so, however, reduces the time it takes for the application to be launched the first time on the fleet instance.

To optimize all the files in a folder, open PowerShell and use the following PowerShell command:

```
dir -path "C:\Path\To\Folder\To\Optimize" -Recurse -ErrorAction SilentlyContinue |
 %{$_.FullName} | Out-File "C:\ProgramData\Amazon\Photon\Prewarm\PrewarmManifest.txt" -
encoding UTF8 -append
```

By default, Image Assistant replaces the application optimization manifest each time the Image Assistant **Optimize** step runs. You must run the PowerShell command to optimize all files in a folder:

- Each time after the **Optimize** step runs.
- Before you choose **Disconnect and create image** on the Image Assistant **Review** page.

Alternatively, you can specify the optimization manifest on a per-application basis by using the Image Assistant command line interface (CLI) operations. When you specify the optimization manifest by using the Image Assistant CLI operations, WorkSpaces Applications merges the specified application optimization manifest with the files identified by the Image Assistant **Optimize** step. For more information, see Create Your Amazon WorkSpaces Applications Image Programmatically by Using the Image Assistant CLI Operations.

Manage WorkSpaces Applications Agent Versions

The WorkSpaces Applications agent is software that runs on your streaming instances and enables users to stream applications. When you create a new image, the **Always use latest agent version** option is selected by default. When this option is selected, new image builders or fleet instances that are launched from your image always use the latest WorkSpaces Applications agent version. You might want to control agent updates to ensure compatibility with your software or to qualify the updated environment before you deploy it for your end users.

The following procedures describe how to manage WorkSpaces Applications agent versions.

Contents

- Create an Image That Always Uses the Latest Version of the WorkSpaces Applications Agent
- Create an Image That Uses a Specific Version of the WorkSpaces Applications Agent
- Create an Image That Uses a Newer Version of the WorkSpaces Applications Agent

Create an Image That Always Uses the Latest Version of the WorkSpaces Applications Agent

When your images are configured to always use the latest WorkSpaces Applications agent version, your streaming instances are automatically updated with the latest features, performance improvements, and security updates that are available from AWS when a new agent version is released.

> **Note**
>
> In some cases, a new WorkSpaces Applications agent version might conflict with your software. We recommend that you qualify the new WorkSpaces Applications agent version before deploying it to your production fleets.

To create an image that always uses the latest version of the WorkSpaces Applications agent

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. Do either of the following:

- If you have an image builder that you want to use to create the image, start the image builder and then connect to it. If the image builder is not running the latest version of the WorkSpaces Applications agent, you are prompted to choose whether to start the image builder with the latest agent. Make sure that this option is selected, choose **Start**, and then connect to the image builder.
- If you do not have an image builder that you want to use to create the image, launch a new image builder. In **Step 1: Choose Image**, choose an AWS base image or a custom image. In **Step 2: Configure Image Builder**, if the image that you choose is not running the latest version of the WorkSpaces Applications agent, the **WorkSpaces Applications** section displays. In the **Agent version** list, select the latest agent version. Complete the remaining steps to create the image builder, and then connect to it. For more information, see Launch an Image Builder to Install and Configure Streaming Applications.

3. On the image builder desktop, open Image Assistant and follow the steps to create your new image. For the **Configure Image** step, make sure that **Always use the latest agent version** is selected. For more information, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

   If you decide later to not always use the latest version of the WorkSpaces Applications agent, you must create a new image and clear that option.

4. Create a new fleet or modify an existing one. When you configure the fleet, select the new image that you created. For more information, see Create an Amazon WorkSpaces Applications Fleet and Stack.
5. Create a new stack or modify an existing one and associate it with your fleet.

Create an Image That Uses a Specific Version of the WorkSpaces Applications Agent

You may want to control WorkSpaces Applications agent updates rather than always using the latest version so that you can test for compatibility first. To ensure that the version of the WorkSpaces Applications agent you use is compatible with your streaming applications, you can create an image that uses a specific version of the agent software. Then perform your qualification tests in a separate fleet before deploying to your production fleet.

When you create the image, make sure that the **Always use latest agent version** option is not selected. Doing so pins your image to the version of the WorkSpaces Applications agent that you

selected when you launched the image builder, rather than always using the latest version. After you finish your qualification tests, you can update your production fleet with the image.

To create an image that uses a specific version of the WorkSpaces Applications agent

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. Do either of the following:
   - If you have an image builder that you want to use to create the image, start the image builder and then connect to it.
   - If you do not have an image builder that you want to use to create the image, launch a new image builder. In **Step 1: Choose Image**, choose an AWS base image or a custom image. In **Step 2: Configure Image Builder**, if the image that you choose is not running the latest version of the WorkSpaces Applications agent, the **WorkSpaces Applications** section displays. In the **Agent version** list, do not select the latest agent version. Complete the remaining steps to create the image builder, and then connect to it. For more information, see Launch an Image Builder to Install and Configure Streaming Applications.
3. On the image builder desktop, open Image Assistant and follow the steps to create your new image. For the **Configure Image** step in Image Assistant, make sure that **Always use the latest agent version** is not selected. For more information, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

   If you decide later to always use the latest version of the WorkSpaces Applications agent, you must create a new image and select that option.
4. Create a new fleet or modify an existing one. When you configure the fleet, select the new image that you created. For more information, see Create an Amazon WorkSpaces Applications Fleet and Stack.
5. Create a new stack or modify an existing one and associate it with your fleet.
6. Connect to your fleet and test your applications for compatibility.

Create an Image That Uses a Newer Version of the WorkSpaces Applications Agent

If you pin your image to a specific WorkSpaces Applications agent version, you must update to a newer version by creating a new image. This approach lets you test each agent update for compatibility first, and then update your fleet incrementally.

When you create the image, make sure that the **Always use latest agent version** option is not selected. After you create your image, perform your qualification tests in a separate fleet before deploying to your production fleet. After you finish your qualification tests, you can update your production fleet with the image.

To create an image that uses a newer version of the WorkSpaces Applications agent

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. Do either of the following:
   - If you have an image builder that you want to use to create the image, start the image builder and then connect to it. If the image builder is not running the latest version of the WorkSpaces Applications agent, you are prompted to choose whether to start the image builder with the latest agent. Make sure that this option is selected, choose **Start**, and then connect to the image builder.
   - If you do not have an image builder that you want to use to create the image, launch a new image builder. In **Step 1: Choose Image**, choose an AWS base image or a custom image. In **Step 2: Configure Image Builder**, if the image that you choose is not running the latest version of the WorkSpaces Applications agent, the **WorkSpaces Applications** section displays. In the **Agent version** list, select the latest agent version. Complete the remaining steps to create the image builder, and then connect to it. For more information, see Launch an Image Builder to Install and Configure Streaming Applications.
3. On the image builder desktop, open Image Assistant and follow the steps to create your new image. For the **Configure Image** step in Image Assistant, make sure that **Always use the latest agent version** is not selected. For more information, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

   If you decide later to always use the latest version of the WorkSpaces Applications agent, you must create a new image and select that option.

4. Create a new fleet or modify an existing one. When you configure the fleet, select the new image that you created. For more information, see Create an Amazon WorkSpaces Applications Fleet and Stack.
5. Create a new stack or modify an existing one and associate it with your fleet.
6. Connect to your fleet and test your applications for compatibility.

Unified graphics images

A unified graphics image is a single graphics image that can support multiple graphics instance families, including Graphics G4dn, G5, G6, Gr6, G6f, Gr6f, and G7. Previously, each graphics family used a separate base image. The AWS base image for a unified graphics image uses the name prefix `AppStream-Graphics-NV-` and is available for Windows Server 2025 and 2022, Red Hat Enterprise Linux 8, and Rocky Linux 8.

With a unified graphics image, you can do the following:

- Launch an image builder on any supported graphics family from the same image.
- Create a fleet on any supported graphics family from the same image.
- Change a fleet's instance type across supported graphics families. For more information, see Update a Fleet with a New Image in Amazon WorkSpaces Applications.

There are three ways to use a unified graphics image:

1. Use the `AppStream-Graphics-NV-`*`OperatingSystem`*`-`*`MM-DD-YYYY`* base images or create your own images from them.
2. Update an existing graphics image through managed image updates. Managed image updates convert your existing graphics image to a unified graphics image. For more information, see the section called “Update an Image by Using Managed WorkSpaces Applications Image Updates”.
3. Create an image from an existing graphics image by using the image builder. As long as a graphics driver is detected on the source image, WorkSpaces Applications converts the image to a unified graphics image. The unified image supports graphics instance families based on the installed driver version. For more information about creating an image, see Launch an Image Builder to Install and Configure Streaming Applications. For details about driver version support, see Install NVIDIA GRID drivers.

> **Note**
>
> Unified graphics images are not supported for custom or BYOL image types.

WorkSpaces Applications Agent Release Notes

The Amazon WorkSpaces Applications agent software runs on your streaming instances, enabling end users to connect to and start their streaming applications. Starting December 7, 2017, your streaming instances can be automatically updated with the latest features, performance improvements, and security updates that are available from AWS. Before December 7, 2017, agent updates were included with new base image releases.

To use the latest WorkSpaces Applications agent software, you need to rebuild your images by using new base images published by AWS on or after December 7, 2017. When you do this, the option to enable automatic updates of the agent is selected by default in the Image Assistant. We recommend that you leave this option selected so that any new image builder or fleet instance that is launched from your image always uses the latest version of the agent. For more information, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

The following table describes the latest updates that are available in released versions of the WorkSpaces Applications agent for Windows instances.

Amazon WorkSpaces Applications agent version	Changes
10-25-2026	• Improves streaming session reliability • General bug fixes and improvements
10-20-2026	• Adds support for printer preference redirection
10-07-2026	• General bug fixes and improvements
10-01-2026	• General bug fixes and improvements
08-31-2026	• Adds support for Graphics NV images • General bug fixes and improvements

Amazon WorkSpaces Applications agent version	Changes
08-04-2026	• Adds support for the Microsoft Smart Card Key Storage Provider (KSP) • General bug fixes and improvements
07-03-2026	• General bug fixes and improvements
06-29-2026	• Adds support for Google Drive and OneDrive storage connectors on multi-session fleets
06-22-2026	• General bug fixes and improvements
06-14-2026	• Adds support for Agent Access (general availability) • Improves session logon latency • Adds a collection of new system resource metrics related to GPU, CPU, memory, and disk • General bug fixes and improvements
05-19-2026	• Adds support for Windows 11 Desktop Operating System (BYOL) • General bug fixes and improvements
04-23-2026	• Adds support for Agent Access (public preview) • General bug fixes and improvements
03-30-2026	• Resolves an idle disconnect issue for streaming sessions • 4K resolution support for WorkSpaces Pools • General bug fixes and improvements
02-09-2026	• Security patch for multisession S3 storage functiona lity

Amazon WorkSpaces Applications agent version	Changes
02-04-2026	• Enhancements for automatic time zone redirection • Improves App Switcher functionality with enhanced window management • 4K resolution support for non-accelerated instance types • General bug fixes and improvements
12-06-2025	• Adds support for Windows Server 2025 • Support for performance metrics for single session • Adds additional performance metrics for multi session • General bug fixes and improvements
10-02-2025	• Support for Microsoft license included applications
09-30-2025	• Support for local file redirection on Multi-Session fleets • General bug fixes and improvements
07-15-2025	• Resolves rendering issues that affect certain applicati ons and improves window management for a seamless user experience • Expanded support for custom paper sizes in printing scenarios to accommodate diverse business needs • Various performance optimizations and reliability enhancements across the platform

Amazon WorkSpaces Applications agent version	Changes
05-29-2025	• Resolves an issue with quick links access to home folder / temporary folder on multi-session fleets • Resolves an issue with certificate validation handling for certificate-based authentication on multi-session WorkSpaces Applications • Improves the display resolution quality • General bug fixes and improvements
03-05-2025	• Resolves an instance launch issue with Windows FIPS • Resolves an issue with Application Settings Persisten ce with private Amazon S3 interface endpoint • General bug fixes and improvements
02-07-2025	• Adds support for multi-session certification-based authentication • Adds performance enhancements for multi-session fleets • Fixes an issue where certain legacy apps that use User Account Control (UAC) Virtualization encounter permissions issues when using Application Settings Persistence
10-31-2024	• Bug fixes and improvements for multi-session fleets
10-21-2024	• Resolves an issue with WorkSpaces Applications agent version 09-18-2024 that caused instance launch failures when used with Microsoft Visual C++ Redistrib utable (versions newer than 14.38.33135.0)

Amazon WorkSpaces Applications agent version	Changes
09-18-2024	• Regional settings support for multi-session fleets. End users can customize time zone, locale, and input method for their streaming sessions. • Printer redirection support for multi-session fleets. End users can redirect print jobs from their streaming application to a printer that is connected to their local computer, including any network printers that the users have mapped. • Automatic time zone redirection for application and desktop streaming sessions for Windows (single session or multi-session). WorkSpaces Applicati ons administrators can enable or disable time zone redirection for their end users' streaming sessions. Once enabled, end users see the same time zone settings on their local devices and WorkSpaces Applications session. • New default maximum size for user profile VHD is 5 GB for Always-On and On-Demand fleets • MSIX Application support for Always-On and On-Demand fleet with Desktop View • General bug fixes and improvements
05-21-2024	• Support for audio in for multi-session fleets • Stability improvement for app view fleets • Support for active directory trust relationships in AD-joined multi-session fleets • General bug fixes and improvements

Amazon WorkSpaces Applications agent version	Changes
04-15-2024	• Improves streaming resiliency when Application Settings Persistence is enabled • Adds support for Seamless mode/Native application mode for multi-session fleets • Improves mouse cursor end user experience in multi-session streaming
01-17-2024	• Adds support for audio out on multi-session fleets • Adds support for session scripts on multi-session fleets • Improves provisioning resilience on multi-session fleets
12-07-2023	• Adds support for Windows Server 2022 • Improves streaming performance on Windows Server 2019 • Adds AWS CLI v2 support • Adds keyboard support to switch between applicati ons • Solves an issue with certificate-based authentication when Windows session is locked • Note: Windows Server 2012 R2 reached end of support on October 10th, 2023. For better streaming experience support, upgrade to Windows Server 2016, Windows Server 2019, or Windows Server 2022.
09-06-2023	• Adds support for multi-session fleets • Improves instance and session provisioning • Resolves an issue with copy/paste functionality • Requires the following software components: • Microsoft .NET Framework Runtime — 4.7.2

Amazon WorkSpaces Applications agent version	Changes
05-30-2023	• Improves instance provisioning resilience
05-08-2023	• Resolves an issue with a shutdown warning in fleet instances for Windows 2016 and Windows 2012 R2
04-13-2023	• Resolves an issue with the streaming session being stuck in a connecting state
03-21-2023	• Resolves an issue with application freezing • Resolves an issue with physical smartcard authentic ation failure • Resolves an issue with home folders not working with FIPS enabled on Windows • Improves instance provisioning resilience • Improves performance with physical smartcard logon time for Windows Server 2019
10-13-2022	• Improves performance with agents • Resolves issue with DCV physical smartcard
06-20-2022	• Adds backwards compatibility for the USB string filter file location on old images • Improves instance provisioning resilience • Improves session connection reliability
03-14-2022	• Resolves an issue with regional settings not updating

Amazon WorkSpaces Applications agent version	Changes

02-21-2022	• Resolves issue with Microsoft OneDrive copying larger files
	• Improves agent robustness on small instance types
	• Works with the following software components. For more information, see the section called "Base Image and Managed Image Update Release Notes".
	  • Amazon SSM Agent — 3.0.1295.0
	  • Amazon WDDM Hook Driver — 1.0.0.56 (Windows Server 2012 R2)
	  • NICE DCV Virtual Display — 1.0.34.0 (Windows Server 2016/2019)
	  • EC2Config service (Windows Server 2012 R2 only) — 4.9.4500

Amazon WorkSpaces Applications agent version	Changes

12-20-2021	• Resolves an issue with the mouse disappearing when using the native client
	• Resolves an issue of storage unmount time for session termination
	• Resolves issues with system crashes on Graphics instances running Windows Server 2016
	• Added support for Windows Server instances when system cryptography group policy is enabled. For more information, see System cryptography.
	• Added the ability to toggle file system caching
	• Works with the following software components. For more information, see the section called “Base Image and Managed Image Update Release Notes”.
	  • Amazon SSM Agent — 3.0.1295.0
	  • Amazon WDDM Hook Driver — 1.0.0.56 (Windows Server 2012 R2)
	  • NICE DCV Virtual Display — 1.0.34.0 (Windows Server 2016/2019)
	  • EC2Config service (Windows Server 2012 R2 only) — 4.9.4500

Amazon WorkSpaces Applications agent version	Changes

10-19-2021	• Resolves an issue preventing users from streaming when the Microsoft Windows printer service is disabled
	• Resolves an issue where language pack installation doesn't complete successfully
	• Resolves an issue with the S3 Home Folder where folders and files are being changed to all uppercase
	• Works with the following software components. For more information, see the section called “Base Image and Managed Image Update Release Notes”.
	  • Amazon SSM Agent — 3.0.1295.0
	  • Amazon WDDM Hook Driver — 1.0.0.56 (Windows Server 2012 R2)
	  • NICE DCV Virtual Display — 1.0.34.0 (Windows Server 2016/2019)
	  • EC2Config service (Windows Server 2012 R2 only) — 4.9.4500

Amazon WorkSpaces Applications agent version	Changes

08-02-2021	• Updated USB driver to include important fixes
	• Resolves an issue where the customer's local machine's caps lock state becomes out of sync with the caps lock state of the remote machine
	• Works with the following software components. For more information, see the section called "Base Image and Managed Image Update Release Notes".
	  • Amazon SSM Agent — 3.0.1295.0
	  • Amazon WDDM Hook Driver — 1.0.0.56 (Windows Server 2012 R2)
	  • NICE DCV Virtual Display — 1.0.34.0 (Windows Server 2016/2019)
	  • EC2Config service (Windows Server 2012 R2 only) — 4.9.4419.0

07-01-2021	• Incremental agent release for Managed Image Updates. For more information, see Update an Image by Using Managed WorkSpaces Applications Image Updates. • Includes changes from the 06-25-2021 agent.

Amazon WorkSpaces Applications agent version	Changes
06-25-2021	• Resolved various networking issues • Resolved an issue where local group policies were overridden • Resolved an issue where files were failing to be created if they were contained in parent directori es that did not exist after attempting to fetch from OneDrive and Google cloud storage • Resolved an issue where session scripts failed to run at the end of a session • Added support for webcam redirection in the web client
05-17-2021	• Enables real-time audio-video (AV) feature by default • Fixes an output of the Image Assistant CLI command to be valid JSON • Fixes an issue causing instance provisioning failures due to internal timeouts • Amazon SSM Agent, Amazon WDDM Hook Driver, and EC2Config service versions remain the same as the previous agent version release

Amazon WorkSpaces Applications agent version	Changes

03-04-2021	• Resolves issues with smart card authentication that cause connection failures. The connection failures occur when users close and reopen a streaming session multiple times
	• Resolves an issue that causes right-click menu items in Microsoft Office applications to be unavailable
	• Resolves an issue that causes multiple storage connector processes for OneDrive and Google Drive to appear in Task Manager
	• Resolves an issue that prevents files larger than 2 GB from being downloaded from Google Drive
	• Resolves an intermittent issue that causes provision ing delays for WorkSpaces Applications fleet instances that are joined to a Microsoft Active Directory domain
	• Works with these software components:
	  • Amazon SSM Agent — 3.0.431.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service (Windows Server 2012 R2 only) — 4.9.4279.0

Amazon WorkSpaces Applications agent version	Changes

12-17-2020	• Resolves an issue that causes the application settings persistence VHD file to not download to the WorkSpaces Applications fleet streaming instance
	• Resolves an issue that causes local printer redirecti on to stop working during WorkSpaces Applicati ons streaming sessions. This issue might occur when Microsoft KB4571694 is installed on the WorkSpaces Applications image builder or fleet streaming instance
	• Resolves an issue that causes the Image Assistant `update-default-profile` command line interface (CLI) operation to return an error when attempting to reference a local Microsoft Windows user as the source for the default user profile
	• Resolves an issue that prevents WorkSpaces Applicati ons fleet instances from provisioning when the system cryptography is set to use FIPS-compliant algorithms
	• Resolves an issue that prevents icons from appearing on users' local computer taskbar during streaming sessions in native application mode
	• Adds support for files shared by Microsoft SharePoin t to the OneDrive for Business persistent storage connector
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.1319.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service (Windows Server 2012 R2 only) — 4.9.4222.0

Amazon WorkSpaces Applications agent version	Changes

01-04-2021	• Adds support for using a smart card for Windows sign in to Active Directory-joined streaming instances and in-session authentication for streaming applications
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.1319.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service (Windows Server 2012 R2 only) — 4.9.4222.0

Amazon WorkSpaces Applications agent version	Changes

12-17-2020	• Resolves an issue that causes the application settings persistence VHD file to not download to the WorkSpaces Applications fleet streaming instance
	• Resolves an issue that causes local printer redirecti on to stop working during WorkSpaces Applicati ons streaming sessions. This issue might occur when Microsoft KB4571694 is installed on the WorkSpaces Applications image builder or fleet streaming instance
	• Resolves an issue that causes the Image Assistant `update-default-profile` command line interface (CLI) operation to return an error when attempting to reference a local Microsoft Windows user as the source for the default user profile
	• Resolves an issue that prevents WorkSpaces Applicati ons fleet instances from provisioning when the system cryptography is set to use FIPS-compliant algorithms
	• Resolves an issue that prevents icons from appearing on users' local computer taskbar during streaming sessions in native application mode
	• Adds support for files shared by Microsoft SharePoin t to the OneDrive for Business persistent storage connector
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.1319.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service (Windows Server 2012 R2 only) — 4.9.4222.0

Amazon WorkSpaces Applications agent version	Changes

10-08-2020	• Resolves an issue that causes users to receive an internal error notification when they connect to WorkSpaces Applications streaming sessions
	• Resolves an issue that causes intermittent copy and paste failures during WorkSpaces Applications streaming sessions
	• Resolves an issue that causes application icons to not appear on the taskbar during WorkSpaces Applicati ons streaming sessions in native application mode
	• Resolves an issue that causes the application catalog to appear empty when users reconnect to WorkSpaces Applications after an idle disconnect
	• Improves the download speed between WorkSpaces Applications home folders and WorkSpaces Applicati ons fleet instances
	• Works with these software components:
	  • Amazon SSM Agent — 3.0.161.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service (Windows Server 2012R2 only) — 4.9.4222.0

Amazon WorkSpaces Applications agent version	Changes

09-01-2020	• Resolves an issue that causes Graphics Design instances to not display the correct resolution
	• Resolves an issue that causes a white screen when using the WorkSpaces Applications client in native application mode to stream Microsoft Remote Desktop
	• Resolves an issue that causes a streaming applicati on to freeze when minimized. This issue occurs when using the WorkSpaces Applications client in native application mode
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.1319.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service (Windows Server 2012R2 only) — 4.9.4222.0

07-30-2020	• Adds support for printer redirection to the WorkSpace s Applications client for Windows
	• Resolves an issue that causes file downloads for files greater than 5 GB to stop, and then fail
	• Improves clipboard performance when using Microsoft Office 2016 plug-ins
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.1319.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service (Windows Server 2012R2 only) — 4.9.4222.0

Amazon WorkSpaces Applications agent version	Changes

05-27-2020	• Resolves an issue that prevents some applications from being resized, moved, or maximized when users stream in native application mode using the WorkSpaces Applications client for Windows
	• Resolves an intermittent issue with downloading utility software. The issue may prevent image builders and fleet instances from being provisioned
	• Resolves an intermittent issue with certain language settings that may prevent image builders and fleet instances from being provisioned
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.701.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service (Windows Server 2012R2 only) — 4.9.3519.0

04-20-2020	• Resolves an issue that causes streaming sessions to fail when session scripts are run
	• Improves performance when IAM roles are used
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.701.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service (Windows Server 2012R2 only) — 4.9.3519.0

Amazon WorkSpaces Applications agent version	Changes

02-19-2020	• Adds support for native application mode. For more information, see Native Application Mode
	• Adds support for the **Desktop** stream view
	• Improves interprocess communication between WorkSpaces Applications components
	• Resolves an issue that caused streaming instances to fail to be provisioned
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.701.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service (Windows Server 2012R2 only) — 4.9.3519.0

01-13-2020	• For persistent storage with Google Drive for G Suite, *Team Drives* have been renamed to *Shared Drives*
	• Resolves an issue that causes slow provisioning for streaming instances in Active Directory environments that have many users
	• Resolves an issue with accessing applications from the application switcher when the fleet user is an administrator
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.701.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service (Windows Server 2012R2 only) — 4.9.3519.0

Amazon WorkSpaces Applications agent version	Changes

11-13-2019	• WorkSpaces Applications assemblies are now signed, including executables and installer packages
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.701.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.3519.0

10-08-2019	• Modifies the WorkSpaces Applications storage connector to no longer bypass the system proxy server
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.701.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.3519.0

09-23-2019	• Resolves an issue that occurs when launching applicati ons that start child processes
	• Resolves an issue with directory traversal
	• Resolves an issue that causes the WorkSpaces Applications agent to stop functioning, which prevents interaction with applications
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.701.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.3519.0

Amazon WorkSpaces Applications agent version	Changes

09-03-2019	• Adds support for applying IAM roles to WorkSpaces Applications streaming instances. For more informati on, see Using an IAM Role to Grant Permissions to Applications and Scripts Running on WorkSpaces Applications Streaming Instances
	• Adds support for specifying tags when creating WorkSpaces Applications images programmatically with a command line interface
	• Modifies the WorkSpaces Applications storage connector to bypass the system proxy server when mounting storage
	• Resolves an issue that prevented .lnk files from being specified in Image Assistant
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.612.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.3429

08-08-2019	• Adds support for WorkSpaces Applications file system redirection. For more information, see Enable File System Redirection for Your WorkSpaces Applications Users
	• Adds support for three new locales: English-United Kingdom (en-GB), English-Canada (en-CA), and English-Australia (en-AU)
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.612.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.3429

Amazon WorkSpaces Applications agent version	Changes

07-26-2019	• Adds support for creating and managing WorkSpaces Applications images programmatically with a command line interface. For more information, see Create Your Amazon WorkSpaces Applications Image Programmatically by Using the Image Assistant CLI Operations.
	• Image creation is no longer blocked when automatic Windows updates are enabled on an image builder. However, a message notifies administrators that automatic Windows updates will be disabled on the fleet in this case (that is, automatic Windows updates won't be enabled on fleet instances).
	• Disables Windows updates when a fleet instance starts
	• Users in the Administrators group are no longer disabled when an image builder instance starts
	• Users in the Administrators group are now disabled rather than deleted when an image builder instance starts
	• Resolves an issue that prevents the streaming resolution from resizing when network connections change
	• Resolves a race condition that prevents the streaming resolution from resizing when application settings persistence is enabled
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.612.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.3429

Amazon WorkSpaces Applications agent version	Changes

06-19-2019	• Adds support for Windows Server 2016 and Windows Server 2019 base images
	• WorkSpaces Applications session scripts are now terminated after the configured timeout is exceeded
	• Resolves an issue where streaming instances may not be provisioned if the locale is changed
	• Includes a change to block image creation when automatic Windows updates are enabled on an image builder
	• Resolves an issue where streaming instances may take a long time to stop if the storage connector mount fails
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.612.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.3429

Amazon WorkSpaces Applications agent version	Changes

05-07-2019	• Adds support for subscribing to WorkSpaces Applications usage reports. For more information, see WorkSpaces Applications Usage Reports.
	• Adds support for configuring the amount of time that users can be idle (inactive) before they are disconnected from their streaming session. For more information, see "Create a Fleet" in Create an Amazon WorkSpaces Applications Fleet and Stack.
	• Resolves an issue with using Amazon S3 buckets for home folder and application settings persistence with an Amazon S3 virtual private gateway
	• Includes a change to block image creation when automatic Windows updates are enabled on an image builder
	• Resolves an issue with persistent storage drives (home folders, OneDrive, and Google Drive) intermittently disappearing from the **My Files** dialog box
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.542.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.3289

04-02-2019	• Resolves an issue with session scripts and storage connector mounting
	• Resolves a minor issue with instance provisioning
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.344.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.3067

Amazon WorkSpaces Applications agent version	Changes

03-07-2019	• Adds support for gestures on touch-enabled iPads, Android tablets, and Windows devices
	• Resolves an issue with switching users in an image builder instance
	• Resolves an intermittent issue with instance reservati ons
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.344.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.3067

Amazon WorkSpaces Applications agent version	Changes

01-22-2019	• Adds support for using on-instance session scripts to run your own custom scripts when specific events occur in users' streaming sessions
	• Adds support for adding tags to the following WorkSpaces Applications resource types during resource creation: image builders, images, fleets, and stacks
	• Includes a fix removing storage connector log files from the application settings persistence Virtual Hard Disk (VHD) file
	• Prevents image creation when the display language is changed from English and the AWS Command Line Interface (AWS CLI) version is earlier than 1.16.36. For more information, see "Special Considerations for Japanese Language Settings" in Configure Default Regional Settings for Your WorkSpaces Applications Users.
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.344.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.3067

01-08-2019	• Improves the instance provisioning time for base images dated 01-08-2019
	• Works with these software components:
	  • Amazon SSM Agent — 2.3.344.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.3067

Amazon WorkSpaces Applications agent version	Changes

12-19-2018	• Resolves an issue with dynamic applications not being added to the application catalog
	• Works with these software components:
	  • Amazon SSM Agent — 2.2.619.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.2644

12-17-2018	• The WorkSpaces Applications client now supports a multiple-monitor experience for streaming instances that use a Graphics Design instance type
	• Resolves an issue with the temporary drive being visible on fleet instances that use a Graphics Desktop or Memory Optimized instance type
	• Works with these software components:
	  • Amazon SSM Agent — 2.2.619.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.2644

Amazon WorkSpaces Applications agent version	Changes

12-04-2018	• Adds support for using a Japanese keyboard with web clients that run on Windows
	• Adds support for using the WorkSpaces Applicati ons dynamic application framework APIs to build a dynamic app provider
	• Resolves an issue with streaming the same session concurrently on multiple tabs or browsers
	• Includes a fix to make home folders, Google Drive, and OneDrive read-only until mounting is completed
	• Improves the mount time for home folders that are stored on fleet instances connected to an Amazon S3 VPC endpoint
	• Works with these software components:
	  • Amazon SSM Agent — 2.2.619.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.2644

11-14-2018	• Adds support for launching streaming sessions using the WorkSpaces Applications Windows client
	• Resolves an issue with opening applications that use environment variables for the fleet user name
	• Works with these software components:
	  • Amazon SSM Agent — 2.2.619.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.2644

Amazon WorkSpaces Applications agent version	Changes

10-30-2018	• Resolves an issue with mounting home folders that are larger than 1 GB when application settings persistence is enabled
	• Resolves an issue with image creation when IPv6 is disabled
	• Session information is now provided as environment variables within streaming instances
	• Works with these software components:
	  • Amazon SSM Agent — 2.2.619.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.2644

10-24-2018	• Includes a fix to display more than 1,000 files in the Amazon S3 home folders directory
	• Works with these software components:
	  • Amazon SSM Agent — 2.2.619.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.2644

10-01-2018	• Improves the performance of application settings persistence
	• Includes a fix to unhide all drives on a fleet instance, except Drive C and Drive D, during user streaming sessions that are launched from the instance
	• Resolves an issue with accessing minimized applicati on subwindows from the application switcher
	• Works with these software components:
	  • Amazon SSM Agent — 2.2.619.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.2644

Amazon WorkSpaces Applications agent version	Changes

08-29-2018	• Adds support for application settings persistence
	• Resolves an issue with copying and pasting large amounts of data between applications within an WorkSpaces Applications streaming session
	• Resolves an issue with accessing unresponsive applications from the application switcher
	• Works with these software components:
	  • Amazon SSM Agent — 2.2.619.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.2644

07-26-2018	• Adds support for OneDrive persistent storage
	• Resolves an issue with saving Visio files to home folders and Google Drive
	• Works with these software components:
	  • Amazon SSM Agent — 2.2.619.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.2644

06-19-2018	• Resolves an issue with optimizing images for applicati on launch
	• Works with these software components:
	  • Amazon SSM Agent — 2.2.619.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.2644

Amazon WorkSpaces Applications agent version	Changes

06-06-2018	• Adds support for regional settings and default application and Windows settings
	• Works with these software components:
	  • Amazon SSM Agent — 2.2.619.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.2644

05-31-2018	• Adds support for Google Drive persistent storage
	• Works with these software components:
	  • Amazon SSM Agent — 2.2.392.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.2586

05-21-2018	• Adds support for administrative controls for data transfer
	• Adds support for the Safari browser on macOS X
	• Works with these software components:
	  • Amazon SSM Agent — 2.2.392.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.2586

03-19-2018	• Resolves an issue with minimizing the application window in certain environments
	• Works with these software components:
	  • Amazon SSM Agent — 2.2.160.0
	  • Amazon WDDM Hook Driver — 1.0.0.56
	  • EC2Config service — 4.9.2400.0

Amazon WorkSpaces Applications agent version	Changes

01-24-2018	• Resolves an issue with the Alt Graph key not working on certain keyboard layouts
	• Works with these software components:
	  • Amazon SSM Agent — 2.2.93.0
	  • Amazon WDDM Hook Driver — 1.0.0.50
	  • EC2Config service — 4.9.2262.0

12-07-2017	• Resolves issues with using ALT key combinations
	• Resolves an issue with file uploads from local computers to streaming sessions
	• Works with these software components:
	  • Amazon SSM Agent — 2.2.93.0
	  • Amazon WDDM Hook Driver — 1.0.0.21
	  • EC2Config service — 4.9.2218.0

Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console

This tutorial describes how to create WorkSpaces Applications images that are based on Microsoft Windows Server operating systems. If you want to create custom images that are based on the Rocky Linux or Red Hat Enterprise Linux operating systems, see the section called “Tutorial: Create a Custom Linux-Based Image”.

In this tutorial, you will learn how to create a custom Amazon WorkSpaces Applications image that contains applications you can stream to your users, and default application and Windows settings to enable your users to get started with their applications quickly. To complete this tutorial, you must already have an image builder. If you don't have an image builder, see Launch an Image Builder to Install and Configure Streaming Applications.

> **Important**
>
> This tutorial includes information that applies to the latest base image release. For more information, see WorkSpaces Applications Base Image and Managed Image Update Release Notes.

Contents

- Step 1: Install Applications on the Image Builder
- Step 2: Create an WorkSpaces Applications Application Catalog
- Step 3: Create Default Application and Windows Settings
- Step 4: Test Applications
- Step 5: Optimize Applications
- Step 6: Finish Creating Your Image
- Step 7 (Optional): Tag and Copy an Image
- Step 8: Clean Up

Step 1: Install Applications on the Image Builder

In this step, you connect an image builder and install your applications on the image builder.

> **Important**
>
> To complete this step, you must log into the image builder with the local **Administrator** account or a domain account that has local administrator permissions. Do not rename or delete the local built-in **Administrator** account. If you rename or delete this account, the image builder will not start and image creation will fail.

To install applications on the image builder

1. Connect to the image builder by doing either of the following:
   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

> **Note**
>
> If the image builder that you want to connect to is joined to an Active Directory domain and your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection. For information about smart card sign in, see Smart Cards.

2. Install applications from an application website or other download source. Install the applications you want before proceeding to the next step.

> **Note**
>
> Download and install applications only from sites that you trust.

If an application requires the Windows operating system to restart, let it do so. Before the operating system restarts, you are disconnected from your image builder. After the restart is complete, connect to the image builder again, then finish installing the application.

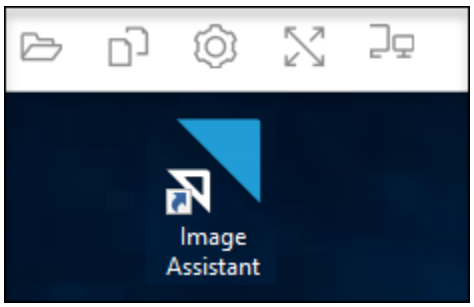
Step 2: Create an WorkSpaces Applications Application Catalog

In this step, create an WorkSpaces Applications application catalog by specifying applications (*.exe*), batch scripts (*.bat*), and application shortcuts (*.lnk*) for your image. For each application that you plan to stream, you can specify the name, display name, executable file to launch, and icon to display. If you choose an application shortcut, these values are prepopulated for you.

> **Important**
>
> To complete this step, you must be logged into the image builder with the local **Administrator** account or a domain account that has local administrator permissions.

To create an WorkSpaces Applications application catalog

1. From the image builder desktop, open Image Assistant. Image Assistant guides you through the image creation process.



2. In **1. Add Apps**, choose **+ Add App**, and navigate to the location of the application, script, or shortcut to add. Choose **Open**.
3. In the **App Launch Settings** dialog box, keep or change the default settings for **Name**, **Display Name**, and **Icon Path**. Optionally, you can specify launch parameters (additional arguments passed to the application when it is launched) and a working directory for the application. When you're done, choose **Save**.

   The **Display Name** and **Icon Path** settings determine how your application name and icon appear in the application catalog. The catalog displays to users when they sign in to an WorkSpaces Applications streaming session.
4. Repeat steps 2 and 3 for each application in Image Assistant and confirm that the applications appear on the **Add Apps** tab. When you're done, choose **Next** to continue using Image Assistant to create your image.

Step 3: Create Default Application and Windows Settings

In this step, you create default application and Windows settings for your WorkSpaces Applications users. Doing this enables your users to get started with applications quickly during their WorkSpaces Applications streaming sessions, without the need to create or configure these settings themselves.

> **Important**
>
> To complete this step, you must be logged into the image builder with the local **Template User** account or a domain user account that does not have local administrator permissions.

To create default application and Windows settings for your users

1. In Image Assistant, in **2. Configure Apps**, choose **Switch user**. This disconnects you from the current session and displays the login menu.
2. Do either of the following:
   - If your image builder is not joined to an Active Directory domain, on the **Local User** tab, choose **Template User**. This account enables you to create your default application and Windows settings.
   - If your image builder is joined to an Active Directory domain, choose **Directory User**, and log in as a domain user that does not have local administrator permissions.
3. From the image builder desktop, open Image Assistant, which displays the applications that you added when you created the application catalog.
4. Choose the application for which you want to create default application settings.
5. After the application opens, create these settings as needed.
6. When you're done, close the application, and return to Image Assistant.
7. If you specified more than one application in Image Assistant, repeat steps 4 through 6 for each application as needed.
8. If you want default Windows settings, create them now. When you're done, return to Image Assistant.
9. Choose **Switch user** and log in with the same account that you used to create the application catalog (an account that has local administrator permissions).
10. In Image Assistant, in **2. Configure Apps**, do either of the following:
    - If your image builder is not joined to an Active Directory domain, choose **Save settings**.
    - If your image builder is joined to an Active Directory domain, in the **Choose which user settings to copy** list, choose the same account that you used to log into the image builder when you created the default application and Windows settings, then choose **Save settings**.

      The **Choose which settings to copy** list displays any account that currently has settings saved on the image builder.
11. When you're done, choose **Next** to continue creating your image.

Step 4: Test Applications

In this step, verify that the applications you've added open correctly and perform as expected. To do so, start a new Windows session as a user who has the same permissions as your users.

> **Important**
>
> To complete this step, you must log in to the image builder with the **Test User** account or a domain account that does not have local administrator permissions.

To test your applications

1. In Image Assistant, in **3. Test**, do either of the following:

   - If your image builder is not joined to an Active Directory domain, choose **Switch user**.
   - If your image builder is joined to an Active Directory domain, you require a domain account to test your applications, and the user already has settings on the image builder, you must reset the application settings for that user. To do so, select the user from the **User to reset** list, and choose **Reset**. When you're done, choose **Switch user**.

     > **Note**
     >
     > If your image builder is new and no users have settings on the image builder, the list does not display any users.

2. Choose the user to use for testing by doing either of the following:

   - If your image builder is not joined to an Active Directory domain, choose **Test User**. This account enables you to test your applications by using the same policies and permissions as your users.
   - If your image builder is joined to an Active Directory domain, choose **Directory User**, specify the credentials for a domain account that does not have local administrator permissions, then choose **Log in**.

3. From the image builder desktop, open Image Assistant, which displays the applications that you specified when you created the application catalog.
4. Choose the application that you want to test, to confirm that it opens correctly and that any default application settings you created are applied.

5. After the application opens, test it as needed. When you're done, close the application and return to Image Assistant.
6. If you specified more than one application in Image Assistant, repeat steps 4 and 5 to test each application as needed.
7. When you're done, choose **Switch user**, then do either of the following:
   - If your image builder is not joined to an Active Directory domain, on the **Local User** tab, choose **Administrator**.
   - If your image builder is joined to an Active Directory domain and you logged in as a domain user with local administrator permissions to specify applications in Image Assistant, log in as that user.
8. Choose **Next** to continue creating your image.

Step 5: Optimize Applications

In this step, Image Assistant opens your applications one after another, identifies their launch dependencies, and performs optimizations to ensure that applications launch quickly. These are required steps that are performed on all applications in the list.

To optimize your applications

1. In Image Assistant, in **4. Optimize**, choose **Launch**.
2. WorkSpaces Applications automatically launches the first application in your list. When the application completely starts, provide any required input to perform the first run experience for the application. For example, a web browser may prompt you to import settings before it is completely up and running.
3. After you complete the first run experience and verify that the application performs as expected, choose **Continue**. If you added more than one application to your image, each application opens automatically. Repeat this step for each application as needed, leaving all applications running.
4. When you're done, the next tab in Image Assistant, **5. Configure Image**, automatically displays.

Step 6: Finish Creating Your Image

In this step, choose an image name and finish creating your image.

To create the image

1. Type a unique image name, and an optional image display name and description. The image name cannot begin with "Amazon," "AWS," or "AppStream."

   You can also add one or more tags to the image. To do so, choose **Add Tag**, and type the key and value for the tag. To add more tags, repeat this step. For more information, see Tagging Your Amazon WorkSpaces Applications Resources. When you're done, choose **Next**.

   > **Note**
   >
   > If you choose a base image that is published by AWS on or after December 7, 2017, the option **Always use the latest agent version** appears, selected by default. We recommend that you leave this option selected so that streaming instances that are launched from the image always use the latest version of the agent. If you disable this option, you can't enable it again after you finish creating the image. For information about the latest release of the WorkSpaces Applications agent, see WorkSpaces Applications Agent Release Notes.

2. In **6. Review**, verify the image details. To make changes, choose **Previous** to navigate to the appropriate Image Assistant tab, make your changes, and then proceed through the steps in Image Assistant as needed.
3. After you finish reviewing the image details, choose **Disconnect and Create Image**.
4. The remote session disconnects within a few moments. When the **Lost Connectivity** message appears, close the browser tab. While the image is created, the image builder status appears as **Snapshotting**. You cannot connect to the image builder until this process finishes.
5. Return to the console and navigate to **Images**, **Image Registry**. Verify that your new image appears in the list.

   While your image is being created, the image status in the image registry of the console appears as **Pending** and you cannot connect to it.
6. Choose the **Refresh** icon periodically to update the status. After your image is created, the image status changes to **Available** and the image builder is automatically stopped.

   To continue creating images, start the image builder and connect to it from the console, or create a new image builder.

Note

After you create your image, you are responsible for maintaining updates for the Windows operating system. To do so, you can use managed WorkSpaces Applications image updates. You are also responsible for maintaining updates for your applications and their dependencies. For more information, see Keep Your Amazon WorkSpaces Applications Image Up-to-Date.

To add other applications, update existing applications, or change image settings, you must start and reconnect to the image builder that you used to create the image. Or, if you deleted that image builder, launch a new image builder that is based on your image. Then, make your changes and create a new image.

Step 7 (Optional): Tag and Copy an Image

You can add one or more tags to an image during image creation or after you create an image. You can also copy the image within the same Region or to a new Region within the same Amazon Web Services account. Copying a source image results in an identical but distinct destination image. AWS does not copy any user-defined tags, however. Also, you can only copy custom images that you create, not the base images that are provided by AWS.

Note

You can copy up to two images at the same time to a destination. If the destination to which you are copying an image is at the image limit, you receive an error. To copy the image in this case, you must first remove images from the destination. After the destination is below the image quota (also referred to as limit), initiate the image copy from the source Region. For more information, see Amazon WorkSpaces Applications Service Quotas.

To add tags to an existing image

1. In the navigation pane, choose **Images**, **Image Registry**.
2. In the image list, select the image to which you want to add tags.
3. Choose **Tags**, choose **Add/Edit Tags**, choose **Add Tag**, specify the key and value for the tag, and then choose **Save**.

For more information, see Tagging Your Amazon WorkSpaces Applications Resources.

To copy an image

Copying an image across geographically diverse regions enables you to stream applications from multiple regions based on the same image. By streaming your applications in closer proximity to your users, you can improve your users' experience streaming applications with WorkSpaces Applications.

1. In the navigation pane, choose **Images**, **Image Registry**.
2. In the image list, select the image that you want to copy.
3. Choose **Actions**, **Copy**.
4. In the **Copy Image** dialog box, specify the following information, and then choose **Copy Image:**
   - For **Destination region**, choose the region to which to copy the new image.
   - For **Name**, specify a name that the image will have when it is copied to the destination.
   - For **Description** (optional), specify a description that the image will have when it is copied to the destination.
5. To check on the progress of the copy operation, return to the console and navigate to **Images**, **Image Registry**. Use the navigation bar to switch to the destination region (if applicable), and confirm that your new image appears in the list of images.

   The new image first appears with a status of **Copying** in the image registry of your console. After the image is successfully created, the status of the image changes to **Available**, which means that you can use the image to launch a stack and stream your applications.

Step 8: Clean Up

Finally, stop your running image builders to free up resources and avoid unintended charges to your account. We recommend stopping any unused, running image builders. For more information, see WorkSpaces Applications Pricing.

To stop a running image builder

1. In the navigation pane, choose **Images**, **Image Builders**, and select the running image builder instance.

2. Choose **Actions**, **Stop**.

Administer Your Amazon WorkSpaces Applications Images

Available images are listed in the **Image Registry** in the WorkSpaces Applications console, and categorized by visibility as follows:

- **Public** — Base images that are owned and made available by AWS. Base images include the latest Windows operating system and the WorkSpaces Applications agent software. You can use these base images to create new images that include your own applications. For information about the base images released by AWS, see WorkSpaces Applications Base Image and Managed Image Update Release Notes.
- **Private** — Images that you create and own, and that you have not shared with other AWS accounts.
- **Shared with others** — Images that you create and own, and that you have shared with one or more AWS accounts in the same AWS Region. When you share an image with another AWS account, you can specify whether the image can be used for an image builder (to create a new image), for a fleet, or both.
- **Shared with me** — Images that are created and owned by another AWS account in the same AWS Region, and that are shared with your AWS account. Depending on the permissions that the owner provided when sharing the image with your account, you can use this image for image builders, for fleets, or both.

Contents

- Delete a Private Image in Amazon WorkSpaces Applications
- Copy an Image That You Own to Another AWS Region in Amazon WorkSpaces Applications
- Share an Image That You Own With Another AWS Account in Amazon WorkSpaces Applications
- Stop Sharing an Image That You Own in Amazon WorkSpaces Applications
- Keep Your Amazon WorkSpaces Applications Image Up-to-Date
- Windows Update and Antivirus Software on Amazon WorkSpaces Applications
- Programmatically Create a New Image in Amazon WorkSpaces Applications
- Manage License Included Applications on Your Image in Amazon WorkSpaces Applications
- Import Image

- Export Image

Delete a Private Image in Amazon WorkSpaces Applications

You can delete your private images when you no longer need them. You can't delete an image that is used by fleets or shared with other AWS accounts. To delete an image that is used by fleets or shared, you must first remove the image from any fleets and remove all image sharing permissions. After you delete an image, you can't recover it.

To delete a private image

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the navigation pane, choose **Images**, **Image Registry**.
3. In the image list, select the private image you want to delete.
4. Choose **Actions**, **Delete**, then choose **Delete** again.

   The image is removed from the image registry and deleted.

Copy an Image That You Own to Another AWS Region in Amazon WorkSpaces Applications

> **Important**
>
> For Asia Pacific (Malaysia), Asia Pacific (Osaka), Canada West (Calgary), Europe (Milan), Europe (Spain), Europe (Zurich), and Israel (Tel Aviv) AWS Regions: Cross-region image copying is only supported for images with WorkSpaces Applications agent versions released on/after October 02, 2025, or images using managed updates released on/after September 05, 2025. Older versions are not eligible for copying between regions. Update your images to meet these minimum version requirements to enable cross-region copy functionality.

You can copy images that you own to another AWS Region. Using the same image across different AWS Regions can help simplify global deployments of your applications on WorkSpaces Applications. By deploying your applications in the AWS Regions that are geographically closest to your users, you can help provide your users with a more responsive experience.

To copy an image that you own to another AWS Region

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the navigation pane, choose **Images**, **Image Registry**.
3. In the image list, select the image that you want to copy to another AWS Region.
4. Choose **Actions**, **Copy**.
5. In the **Copy image** dialog box, in **Destination region**, select the AWS Region that you want to copy the image to.
6. Type a unique name and optionally, a description for the image in **Destination region**.
7. Choose **Copy Image**.

Share an Image That You Own With Another AWS Account in Amazon WorkSpaces Applications

WorkSpaces Applications images are a regional resource, so you can share an image that you own with other AWS accounts within the same AWS Region. Doing so can be helpful in several different scenarios. For example, if you separate your development and production resources by using different AWS accounts, you can create an image by using your development account. Then you can share the image with your production account. If your organization is an independent software vendor (ISV), you can share optimized images with your customers. Optimized images that have the required applications already installed and configured let your customers get started with your applications quickly, so that they won't need to install and configure those applications themselves.

When you share an image with another AWS account, you specify whether the destination account can use the image in a fleet or create new images by creating an image builder. You continue to own images that you share. This way, you can add, change, or remove permissions as needed for your shared images.

If you share an image with an account and grant the account fleet permissions, the shared image can be used to create or update fleets in that account. If you remove these permissions later, the account can no longer use the image. For fleets in the account that use the shared image, the desired capacity is set to 0, which prevents new fleet instances from being created. Existing sessions continue until the streaming session ends. For new fleet instances to be created, the fleet in that account must be updated with a valid image.

If you share an image with an account and grant the account image builder permissions, the shared image can be used to create image builders and images in that account. If you remove these permissions later, image builders and images that were created from your image are not affected.

> **Important**
>
> After you share an image with an account, you can't control image builders or images in the account that are created from your image. For this reason, grant image builder permissions to an account only if you want to enable the account to make a copy of your image, and retain access to the copy after you stop sharing your image.

To share an image that you own with another AWS account

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the navigation pane, choose **Images**, **Image Registry**.
3. In the image list, select the image that you want to share.
4. Choose **Actions**, **Share**.
5. In the **Share image** dialog box, choose **Add account**.
6. Type the 12-digit AWS account ID of the account that you want to share the image with, and then select whether the account can do one or both of the following:

   - Use the image to launch an image builder, if you want to create a new image.
   - Use the image with a fleet.

   To remove an account from the list of accounts that the image is shared with, in the row for the account you want to remove, choose the X icon to the right of the **Use for fleet** option.
7. To share the image with more AWS accounts, repeat step 6 for each account that you want to share the image with.
8. Choose **Share Image**.

To add or update image sharing permissions for an image that you own

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the navigation pane, choose **Images**, **Image Registry**.

3. In the image list, select the image that you want to change the permissions for.
4. Below the image list, choose the **Permissions** tab for the image you selected, then choose **Edit**.
5. In the **Edit image permissions** dialog box, select or clear one or both of the following image sharing options as needed for one or more AWS accounts. If you clear both options for an account, the image is no longer shared with that account.

   - Use the image to launch an image builder, if you want to create a new image.
   - Use the image with a fleet.

   To remove an account from the list of accounts that the image is shared with, in the row for the account you want to remove, choose the X icon to the right of the **Use for fleet** option.
6. To edit image sharing permissions for more AWS accounts, repeat step 5 for each account you want to update permissions for.
7. Choose **Update image sharing permissions**.

Stop Sharing an Image That You Own in Amazon WorkSpaces Applications

Follow these steps to stop sharing an image that you own with any other AWS account.

To stop sharing an image that you own with any other AWS account

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the navigation pane, choose **Images**, **Image Registry**.
3. In the image list, select the image that you want to change the permissions for.
4. Below the image list, choose the **Permissions** tab for the image you selected, then choose **Edit**.
5. In the **Edit image permissions** dialog box, in the row for all AWS accounts that the image is shared with, choose the X icon to the right of the **Use for fleet** option.
6. Choose **Update image sharing permissions**.

Keep Your Amazon WorkSpaces Applications Image Up-to-Date

You can keep your WorkSpaces Applications image up-to-date by doing either of the following:

- Update an Image by Using Managed WorkSpaces Applications Image Updates – This update method provides the latest operating system updates and driver updates, and the latest WorkSpaces Applications agent software.
- Update the WorkSpaces Applications Agent Software by Using Managed WorkSpaces Applications Agent Versions – This update method provides the latest WorkSpaces Applications agent software.

Update an Image by Using Managed WorkSpaces Applications Image Updates

WorkSpaces Applications provides an automated way to update your image with the latest operating system updates, license included application updates, driver updates, and WorkSpaces Applications agent software. With managed WorkSpaces Applications image updates, you select the image that you want to update. WorkSpaces Applications creates an image builder in the same AWS account and Region to install the updates and create the new image. After the new image is created, you can test it on a pre-production fleet before updating your production fleets or sharing the image with other AWS accounts.

Note

Managed WorkSpaces Applications Image Updates is available for Microsoft Windows Server, Red Hat Enterprise Linux, and Rocky Linux operating systems.

Note

When you update a generation-specific graphics image (for example, a Graphics G6 image), managed WorkSpaces Applications image updates produce a unified graphics image. WorkSpaces Applications selects an appropriate instance type to launch a temporary image builder, re-tags the image with expanded instance-family compatibility, and validates driver compatibility, which might require a driver update. After the update completes, the image shows compatibility with all supported graphics families. For more information, see Unified graphics images.

When available, the graphics driver version appears in the image details.

Note

After your new image is created, you're responsible for maintaining updates for the operating system. To do so, you can continue using managed WorkSpaces Applications image updates.

You are responsible for maintaining updates for the Amazon EC2 Windows Paravirtual (PV) driver, ENA driver, and AWS NVMe driver. For more information about how to update the drivers, see Manage device drivers for your EC2 instance.

You're also responsible for maintaining your applications and their dependencies. To add other applications, update existing applications, or change image settings, you must start and reconnect to the image builder that you used to create the image. Or, if you deleted that image builder, launch a new image builder that is based on your image. Then, make your changes and create a new image.

Prerequisites

The following are prerequisites and considerations for working with managed image updates.

- Make sure that your WorkSpaces Applications account quotas (also referred to as limits) are sufficient to support the creation of a new image builder and a new image. To request a quota increase, you can use the Service Quotas console at https://console.aws.amazon.com/servicequotas/. For information about default WorkSpaces Applications quotas, see Amazon WorkSpaces Applications Service Quotas.
- You must own the image that you update. You can't update an image that is shared with you.
- When WorkSpaces Applications creates an image builder to install the latest operating system updates, driver updates, and WorkSpaces Applications agent software, and creates the new image, you're charged for the image builder instance while it's updating.
- Supported images must be created from a base image released on 2017-07-24T00:00:00Z or later.
- English and Japanese are supported display languages. For more information, see Specify a Default Display Language.
- Use the latest version of SSM Agent. For version information, see the section called “Base Image and Managed Image Update Release Notes”. For installation information, see Manually install SSM Agent on EC2 instances for Windows Server.

How to Update an Image by Using Managed WorkSpaces Applications Image Updates

To update an WorkSpaces Applications image with the latest patches, driver updates, and WorkSpaces Applications agent software, perform the following steps.

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the navigation pane, choose **Images**, **Image Registry**.
3. In the image list, choose the image that you want to update. Verify that the status of the image is **Available**.
4. Choose **Actions**, **Update**.
5. In the **Update image** dialog box, do the following:
   - For **New image name**, enter an image name that is unique within the AWS account and Region. The image name can't begin with "Amazon," "AWS," or "AppStream."
   - For **New image display name**, you can optionally enter a name to display for the image.
   - For **New image description**, you can optionally provide a description for the image.
   - For **Tags**, you can choose **Add Tag**, and type the key and value for the tag. To add more tags, repeat this step. For more information, see Tagging Your Amazon WorkSpaces Applications Resources.
6. Choose **Update image**.

   If your current image is already up to date, a message notifies you.
7. In the navigation pane, choose **Images**, and then choose **Image Builder**.
8. In the list of image builders, verify that a new image builder appears in the **Updating** state. The name of the image builder includes a random 10-digit suffix.

   The image builder is the smallest size in the instance family that you chose for the new image in step 5. No subnet is specified because the image builder is not attached to your virtual private cloud (VPC).
9. Choose **Image Registry** and verify that your new image appears in the list.

   While your image is being created, the image status in the image registry of the console appears as **Creating**.
10. After your image is created, WorkSpaces Applications performs a qualification process to verify that the image works as expected.

During this time, the image builder, which is also used for this process, appears in the **Image Builder** list with a status of **Pending Qualification**.

11. After the qualification process successfully completes, a **Success** message appears at the top of the console and the image status in the image registry appears as **Available**.

    In addition, the image builder that WorkSpaces Applications created is deleted automatically.

    > **Note**
    >
    > Depending on the volume of operating system updates, it might take several hours for an image update to complete. If an issue prevents the image from being updated, a red icon with an exclamation point appears next to the image name, and the image status in the image registry appears as **Failed**. If this occurs, select the image, choose the **Notifications** tab, and review any error notifications. For more information, see the information in the Image Internal Service section of the documentation for troubleshooting notification codes.
    > If the qualification process is not successful, the image builder that WorkSpaces Applications created is still deleted automatically.

12. After WorkSpaces Applications creates the new image, test the image on a pre-production fleet. After you verify that your applications work as expected, update your production fleet with the new image.

Update the WorkSpaces Applications Agent Software by Using Managed WorkSpaces Applications Agent Versions

WorkSpaces Applications provides an automated way to update your image builder with newer WorkSpaces Applications agent software. Doing so enables you to create a new image whenever a new version of the agent is released. You can then test the image before updating your production fleets. For more information about how to manage the WorkSpaces Applications agent software, see Manage WorkSpaces Applications Agent Versions.

> **Note**
>
> You're responsible for installing and maintaining the updates for the Windows operating system, your applications, and their dependencies.

To keep your WorkSpaces Applications image updated with the latest Windows operating system updates, do one of the following:

- Install your applications on the latest base image each time a new image is released.
- Install the updates for the Windows operating system, your applications, and their dependencies on an existing image builder.
- Install the updates for the Windows operating system, your applications, and their dependencies on a new image builder from an existing image.

After you create a new image with the latest Windows operating system, applications and their dependencies, and the WorkSpaces Applications agent software, test the image on a development fleet. After you verify that your applications work as expected, update your production fleet with the new image.

Windows Update and Antivirus Software on Amazon WorkSpaces Applications

WorkSpaces Applications streaming instances are non-persistent. When a user streaming session ends, WorkSpaces Applications terminates the instance used by the session and, depending on your scaling policies, provisions a new instance to replace it in your fleet. All fleet instances are provisioned from the same image. Because images cannot be changed once created, all fleet instances used in user streaming sessions have only the Windows and application updates that were installed on the underlying image when the image was created. In addition, because a fleet instance used for a streaming session terminates at the end of the session, any updates made to Windows or to applications on the instance during the streaming session will not persist to future sessions by the same user or other users.

> **Note**
>
> If you enabled application settings persistence for your stack, WorkSpaces Applications persists Windows and application configuration changes made by a user to future sessions for the same user if those configuration changes are stored in the user's Windows profile. However, the application settings persistence feature persists only Windows and application configuration settings. It does not persist software updates to Windows or applications on the streaming instance.

For these reasons, WorkSpaces Applications takes the following approach to Windows Update and antivirus software on WorkSpaces Applications instances.

Windows Update

Windows Update is not enabled by default on WorkSpaces Applications base images. If you enable Windows Update on an image builder and then try to create an image, Image Assistant displays a warning and disables Windows Update during the image creation process. To ensure that your fleet instances have the latest Windows updates installed, we recommend that you install Windows updates on your image builder, create a new image, and update your fleet with the new image on a regular basis.

Antivirus Software

If you choose to install antivirus software on your image, we recommend that you do not enable automatic updates for the antivirus software. Otherwise, the antivirus software may attempt to update itself with the latest definition files or other updates during user sessions. This may affect performance. In addition, any updates made to the antivirus software will not persist beyond the current user session. To ensure that your fleet instances always have the latest antivirus updates, we recommend that you do either of the following:

- Update your image builder and create a new image on a regular basis (for example, by using the Image Assistant CLI operations).
- Use an antivirus application that delegates scanning or other operations to an always-up-to-date external server.

> **Note**
>
> Even if you do not enable automatic updates for your antivirus software, the antivirus software may perform hard drive scans or other operations that may impact the performance of your fleet instances during user sessions.

On WorkSpaces Applications Windows Server 2025/2022/2019/2016 base images published on or after September 10, 2019, Windows Defender is not enabled by default. On WorkSpaces Applications Windows Server 2016 and Windows Server 2019 base images published on June 24, 2019, Windows Defender is enabled by default.

To enable Windows Defender manually

If Windows Defender is not enabled on your base image, you can enable it manually. To do so, complete the following steps.

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the left navigation pane, choose **Images**, **Image Builder**.
3. Choose the image builder on which to enable Windows Defender, verify that it is in the **Running** state, and choose **Connect**.
4. Log in to the image builder with the local **Administrator** account or with a domain account that has local administrator permissions.
5. Open Registry Editor.
6. Navigate to the following location in the registry: **HKLM\SOFTWARE\Policies\Microsoft\Windows Defender\DisableAntiSpyware**.
7. To edit this registry key, double-click it, or right-click the registry key, and choose **Modify**.
8. In the **Edit DWORD (32-bit) Value** dialog box, in **Value data**, change **1** to **0**.
9. Choose **OK**.
10. Close Registry Editor.
11. Open the Microsoft Management Console (MMC) **Services** snap-in (`services.msc`).
12. In the list of services, do one of the following.

    If you are using Microsoft Windows Server 2022/2025, do either of the following:

    - Right-click **Microsoft Defender Antivirus Service**, and choose **Start**.
    - Double-click **Microsoft Defender Antivirus Service**, choose **Start** in the properties dialog box, and then choose **OK**.

    If you are using Microsoft Windows Server 2019 or 2016, do either of the following:

    - Right-click **Windows Defender Antivirus Service**, and choose **Start**.
    - Double-click **Windows Defender Antivirus Service**, choose **Start** in the properties dialog box, and then choose **OK**.
13. Close the **Services** snap-in.

Programmatically Create a New Image in Amazon WorkSpaces Applications

You can create WorkSpaces Applications images programmatically by connecting to an image builder and using the Image Assistant command line interface (CLI) operations. For more information, see Create Your Amazon WorkSpaces Applications Image Programmatically by Using the Image Assistant CLI Operations.

Manage License Included Applications on Your Image in Amazon WorkSpaces Applications

You can stream the following Microsoft license included applications using WorkSpaces Applications. You can install these applications on your Windows Image, use this custom image to create fleet(s), and then stream these applications. All of the following applications are available in 32-bit and 64-bit architecture:

- Microsoft Office LTSC Professional Plus 2021/2024
- Microsoft Visio LTSC Professional 2021/2024
- Microsoft Project Professional 2021/2024
- Microsoft Office LTSC Standard 2021/2024
- Microsoft Visio LTSC Standard 2021/2024
- Microsoft Project Standard 2021/2024

Important

- Microsoft Office, Visio, and Project must follow the same versions. For example, you can't mix 2021 applications with 2024 applications.
- Microsoft Office, Visio, and Project must follow the same architecture. For example, you can't mix 32-bit applications with 64-bit applications.
- Microsoft Office, Visio, and Project 2021 Standard/Professional versions are supported on Microsoft Windows Server 2019/2022/2025 and Windows 11. Microsoft Office, Visio, and Project 2024 Standard/Professional versions are supported on Microsoft Windows Server 2022/2025 and Windows 11.
- To enable this feature, you must use an WorkSpaces Applications Image Builder that uses an WorkSpaces Applications agent released on or after October 2, 2025. For more

information, see the section called “Manage Agent Versions” . Or, your image must use managed WorkSpaces Applications image updates released on or after October 3, 2025. For more information, see the section called “Keep Your Image Up-to-Date”.

- Outbound TCP on port 1688 must be open on the management network interface of all streaming instances.
- All users streaming through a fleet powered by an image with one or more licensed apps incur billing for these apps monthly, regardless of usage. The application entitlement feature doesn't restrict access for specific users.
- License included applications on Image Builder aren't activated because they are installed for administrative purposes. Activation occurs when users stream through a fleet instance.

Contents

- View the list of license included applications installed on your image
- View the list of license included applications on your image builder
- Install or uninstall license included applications
- Update license-included applications

View the list of license included applications installed on your image

View the list of license included applications installed on your image

To view the list of license included applications installed on your image, follow these steps.

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. Choose **Images** in the left navigation pane and the **Image Registry** tab.
3. Select an image, and choose **View Details**.
4. Review the list of all the installed applications under **License included applications**.

View the list of license included applications on your image builder

View the list of license included applications on your image builder

To view the list of license included applications on your image builder, follow these steps.

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. Choose **Images** in the left navigation pane and the **Image builder** tab.
3. Select an image builder, and choose **View details**.
4. Review the list of applications and their statuses under **License included applications**.

Install or uninstall license included applications

Install or uninstall license included applications

To install or uninstall one or more license included application(s) on your image, follow these steps.

1. Complete one of the following options:
   - Launch an image Builder and configure license included applications. For more information, see the section called “Launch an Image Builder”.
   - Manage license included applications on your image builder. For more information, see the section called “Attribute-Based Application Entitlements”.
2. When you have an image created with one or more license included applications, you can use this image to create fleets. Users connecting to this fleet can access these applications.

> **Important**
>
> All the users streaming through a fleet powered by an image with one or more licensed apps will incur billing for these apps monthly, regardless of usage. The application entitlement feature does not restrict access for specific users.

If you encounter failures during license included app installation or uninstallation, you will see a failure status on your Image Builder's details page. To troubleshoot these issues, we recommend connecting to your Image Builder and enabling verbose logging. For more information see How to enable Microsoft 365 Apps for enterprise logging. If the problem persists after reviewing the logs and troubleshooting, contact AWS Support for help.

Update license-included applications

Updates for license-included applications are disabled by default. You can enable and perform updates for these applications on an image builder that includes one or more of them or perform

the updates using Managed Image Updates. Updates on fleet instances remain disabled to prevent installation during session setup.

There are three options for updating Office applications:

- Enable and perform updates from the application menu on an image builder
- Enable and perform updates manually with PowerShell on an image builder
- Use Managed Image Updates

After updates complete, snapshot the image and use the new image for your fleets to continue using the updated applications.

Option 1: Enable and perform updates from the application menu

To enable updates for license-included applications from within an Office application, do the following:

1. On the image builder, open any license-included application (for example, Word or Excel).
2. Choose **File**, **Account**, **Update Options**, **Enable Updates**.
3. Once updates are enabled, choose **Update Options**, **Update Now** to start the updates.

Option 2: Enable and perform updates manually with PowerShell

To enable updates for license-included applications by using PowerShell, do the following:

1. On the image builder, open PowerShell as an administrator.
2. Enable updates through Click-to-Run:

   ```
   Set-ItemProperty -Path "HKLM:\SOFTWARE\Microsoft\Office\ClickToRun
   \Configuration" -Name UpdatesEnabled -Value True
   ```
3. Run the update:

   ```
   Start-Process -FilePath "C:\Program Files\Common Files\Microsoft Shared
   \ClickToRun\OfficeC2RClient.exe" -ArgumentList "/update user" -Wait -
   PassThru
   ```

The update can take several minutes to complete. A pop-up window displays the update status.

Option 3: Enable updates by using Managed Image Updates

You can also receive updates for Microsoft license-included applications through Managed Image Updates. For more information, see the section called “Update an Image by Using Managed WorkSpaces Applications Image Updates”.

Import Image

You can create WorkSpaces Applications images by importing your customized EC2 AMIs or WorkSpaces images. Here's how it works:

1. Customize your EC2 AMI using any preferred method including EC2 Image Builder.
2. Import your customized AMI or WorkSpaces images into WorkSpaces Applications to create a WorkSpaces Applications image
3. Optionally, use Image Builder for additional image customization

WorkSpaces Applications images created through EC2 AMI or WorkSpaces Windows Server image import are of `type = "custom"`. WorkSpaces Windows BYOL images import will create WorkSpaces Applications images with `type = "BYOL"`. WorkSpaces Applications provided images are of `type = "native"`.

You can use stream.* instance types for images with `type = "native"` and `type = "BYOL"`. To use any of the following instance types, you must import your AMI and create an image with `type = "custom"`.

- GeneralPurpose.*
- MemoryOptimized.*
- ComputeOptimized.*
- Accelerated.*

> **Important**
>
> WorkSpaces images with License included applications installed are not supported. You must uninstall those applications from the image using Manage Applications and configure those applications using Manage License Included Applications on Your Image in Amazon WorkSpaces Applications.

Prerequisites for EC2 AMI image import

All these prerequisites are important for a successful workflow execution. Supported EC2 AMI configurations and other mandatory requirements are listed below. Note these do not apply to importing WorkSpaces images, as that is managed by the WorkSpaces service.

Required AMI Properties

EBS

- Less or equal to 500GB size
  - You can import an AMI with < 200 GB, however, the imported image will use minimum 200GB.
- GP2
  - You can import an AMI with gp2 or gp3 EBS volume type, however, the imported image will use gp2.
- One volume per image
- `/dev/sda1` Root Device Name

- Image Type: Machine
- Architecture: x86_64
- Virtualization Type: HVM
- Boot Mode: UEFI
- TPM Support: v2.0. This is required, Refer to https://docs.aws.amazon.com/ec2/latest/windows-ami-reference/ami-windows-tpm.html#ami-windows-tpm-find to find a TPM enabled AMI.
- ENA Support: true
- Platform: Windows
- Platform Details: Windows

Operating System Properties

Windows Server 2022/2025 **Full Base**

- Windows Server **Core** is not supported
- Windows with SQL Server is not supported

Agents

- EC2 Launch V2
  - Installed by default in Windows Server 2022/2025

Library Support

- .NET Framework 4.8 or greater
  - Installed by default in Windows Server 2022/2025
- PowerShell 5.1 or greater
  - Installed by default in Windows Server 2022/2025

- Windows Features: Remote Desktop Services Licensing and Remote Desktop Services Session Host must not be installed
- Ports: Ports 8000, 8300, and 8443 must be unblocked and unoccupied
- Boot Mode: UEFI

If you want to use image with graphics instances such as Accelerated.g4dn, Accelerated.g5, Accelerated.g6, Accelerated.g6e, or Accelerated.g7 you must install proper GRID driver on your AMI. For more details please refer to Install NVIDIA GRID drivers. If the drivers are not setup correctly the streaming will work, however, graphics card may not be available.

Note

Unified graphics images are not supported for imported images. An imported graphics image is compatible with the specific Accelerated instance family that you configure it for. For more information about unified graphics images, see Unified graphics images.

Important

"Owner Account Id" of the AMI must be your AWS account id. You cannot import a public EC2 AMI.
Perform any Windows updates and disable automatic Windows updates before importing the image.
Import of encrypted EC2 AMIs is currently not supported

IAM Role Requirements

> **Important**
>
> Create an IAM role with the following permissions to use for image import:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowModifyImageAttributeWithTagCondition",
            "Effect": "Allow",
            "Action": "ec2:ModifyImageAttribute",
            "Resource": "*"
        },
        {
            "Sid": "AllowDescribeImages",
            "Effect": "Allow",
            "Action": "ec2:DescribeImages",
            "Resource": "*"
        }
    ]
}
```

Add the following trust relationship for this IAM role

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "Service": "appstream.amazonaws.com"
            },
            "Action": "sts:AssumeRole"
        }
    ]
}
```

Requirements for WorkSpaces image import

When importing an Amazon WorkSpaces image into WorkSpaces Applications, the following operating system and protocol requirements apply:

- Windows Server 2022 — Only WSP (WorkSpaces Streaming Protocol) images are supported. PCoIP images are not supported.
- Windows Server 2025 — All images are supported.
- Windows 11 — Version 24H2 or newer only.

To import an image

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.

   > **Note**
   >
   > Your console role will also need to include the `workspaces:DescribeWorkspaceImages` permission to import an Amazon WorkSpaces image.

2. In the left navigation pane, choose **Images** and then choose **Image registry**.
3. Choose **Import Image**.
4. **Image source** - choose from the below options based on what type of image you want to import
   - **AMI ID** - Enter an AMI ID for AMI that you would like to import to WorkSpaces Applications. You can also search for your AMI using this field.
   - **Amazon WorkSpaces Image** - Enter a WorkSpaces Image ID (starts with "wsi-"). You can also search for your WorkSpaces Image ID using this field.
5. **Image name** - Enter a unique name for the image that will be created because of import operation.
6. **Display name** *(Optional)* - Enter a display name for the image.
7. **Description** *(Optional)* – Enter a description for the image.
8. **IAM Role** - Select the IAM role that you have created for image import. For more details refer to the section called “IAM Role Requirements”.
   - Required for EC2 AMI import
   - Not applicable for WorkSpaces image import

9. **Manage WorkSpaces Applications agent** – Select this option if you want to always use the latest WorkSpaces Applications agent version, your streaming instances are automatically updated with the latest features, performance improvements, and security updates that are available from AWS when a new agent version is released.

10. **Runtime validation** *(Optional)*: Select this option and service will provision an instance with the image being imported and run streaming tests.

    - > **Note**
      >
      > These streaming tests will be executed in the background, you cannot connect to this instance via WorkSpaces Applications client.

    - We recommend using this option to get higher confidence that your image is suitable for WorkSpaces Applications.
    - You will be billed for the hourly price of that instance.
    - You can avoid running runtime validation if you are re-importing your AMI after making minor changes that may not affect the streaming test, and if runtime validation passed the last time, you imported this AMI.
    - **Choose instance type** *(Optional)*: Select the right instance family, type, and size for running the streaming test. It is recommended that you use the same instance which you are planning to use for fleet creation.

11. **Applications catalog and launch performance manifest** *(Optional)*: Provide details to create applications catalog for your end users and improve the launch performance of your applications.

    - **Application catalog**: To create an application catalog specify details about the applications installed your image. For each application that you plan to stream, you can specify the name, display name, executable file to launch, and icon to display.
    - **Launch performance**: Adding files to the application optimization manifest reduces the time that it takes for the application to launch for the first time on a new fleet instance. The optimization manifest is a line-delimited text file that is per application.

    To learn more refer to the section called “Applications Details”.

12. **Tags** *(Optional)* - Choose **Add Tag** and type the key and value for the tag. To add more tags, repeat this step. For more information, see Tagging Your Amazon WorkSpaces Applications Resources.

13 **Import Image** – Review all the information you have entered and choose **Import Image**. Service will run compatibility checks to make sure AMI is compatible with WorkSpaces Applications.

- If the static checks fail, you will receive an error straight away.
- If the static checks pass, your import request will be submitted and depending upon the options you have selected it could take 30-60 min to create a new WorkSpaces Applications image with `type = "custom"` or `type = "BYOL"`

Applications Details

Applications details contains information about pre-warm manifests and app catalog configurations.

Application PreWarm Manifests

When creating WorkSpaces Applications images you may specify applications to be made available to your users. To speed up the application's launch time you can prepare a PreWarm manifest. This is essentially a catalog of the files that your application needs to launch when users launch your application. During instance provisioning these files will be prepared ahead of session connections to speed up application launch times in user sessions.

Prewarm manifests must be pre-created on your AMI before being imported into the WorkSpaces Applications environment. You can choose to either create one common Prewarm manifest file or one per each application. This changes how you will import your AMI later.

Common Prewarm Manifest

For each application you wish to prewarm, launch the application and perform any initial interactions your users may perform. Then, use the following command targeting the directory where your applications data is stored.

```
dir -path "C:\Path\To\Folder\To\Optimize" -Recurse -ErrorAction SilentlyContinue |
 %{$_.FullName} | Out-File "C:\ProgramData\Amazon\Photon\Prewarm\PrewarmManifest.txt" -
encoding UTF8 -append
```

This will append the files to optimize for each application into the common `C:\\ProgramData\\Amazon\\Photon\\Prewarm\\PrewarmManifest.txt` file. There is no additional action needed to perform application prewarming. WorkSpaces Applications will look for the prewarm file at the above location and use it if it is present.

This process is optional and as the size of the prewarm manifest increases, fleet provisioning time will also increase. So take care to balance optimization with fleet provisioning.

Application Specific Manifests

During image import, you may wish to specify separate application manifest files per application for easier tracking of the prewarm assets per application. To do this perform the same steps as above, but instead of creating a common `C:\\ProgramData\\Amazon\\Photon\\Prewarm\\PrewarmManifest.txt` file, create a file per application on your AMI.

For each application you wish to prewarm, launch the application and perform any initial interactions your users may perform. Then, use the following command targeting the directory where your applications data is stored.

```
dir -path "C:\Path\To\Folder\To\Optimize" -Recurse -ErrorAction SilentlyContinue |
 %{$_.FullName} | Out-File "C:\Path\To\My\<ApplicationName>PreWarm.txt" -encoding UTF8
 -append
```

We will use these application prewarm files during the image import process. Again this is completely optional. You may choose to use this method, the Common Prewarm Manifest method, or no Prewarm manifest at all.

Application Catalog Configs

`AppCatalogConfig` which allows you to specify the applications you wish to register to your WorkSpaces Applications image during AMI import. The `AppCatalogConfig` is a JSON list of Application configuration objects of the following structure.

```
[
    {
        "Name": "Rufus", //Required and must be unique among the list of applications
        "DisplayName": "Rufus",
        "AbsoluteAppPath": "Rufus", //Required
        "AbsoluteIconPath": "Rufus",
        "AbsoluteManifestPath": "Rufus",
        "WorkingDirectory": "Rufus",
        "LaunchParameters": "Rufus"
    }

    ...
```

```
    // Up to 50 applications total
]
```

The only required fields per application are the `Name` and the `AbsoluteAppPath`. The details of each field as follows:

Name [**Required**]

- A given name for your application to identify it
- Between 1 and 100 characters
- Allowed characters regex `^[a-zA-Z0-9][a-zA-Z0-9_.-]{0,99}$`
- Must be unique in a given AppCatalogConfig

DisplayName

- The display name for a given application to display to users
- Between 0 and 100 characters
- Allowed characters regex `^[a-zA-Z0-9][a-zA-Z0-9_. -]{0,99}$`

AbsoluteAppPath [**Required**]

- The path to the executable to launch your application
  - This is the executable that will be launched when users select your application
- Between 1 and 32767 characters
  - This character length upper limit is to support extended file paths in Windows. Ensure your AMI and application is properly configured to support Windows extended file paths if using file paths larger than 260 characters.
- Use escaped file path strings such as
  - `"C:\\Windows\\System32\\notepad.exe"`

AbsoluteManifestPath

- Only applicable if you are using the section called “Application Specific Manifests”
- Path to prewarm manifest file for this application
- Between 0 and 32767 characters
  - This character length upper limit is to support extended file paths in Windows. Ensure your AMI and application is properly configured to support Windows extended file paths if using file paths larger than 260 characters.

- Use escaped file path strings such as
  - `"C:\\Path\\To\\PrewarmManifest.txt"`

AbsoluteIconPath

- Path to icon file on the AMI to use for the application.
  - This icon will be shown to users when streaming to this image.
  - If none is provided the icon will be derived from the executable itself.
  - Take care to select icon files with appropriately handled background transparency for a good client experience for your users
    - Use PNG images
- Between 1 and 32767 characters
  - This character length upper limit is to support extended file paths in Windows. Ensure your AMI and application is properly configured to support Windows extended file paths if using file paths larger than 260 characters.
- Use escaped file path strings such as
  - `"C:\\Path\\To\\ApplicationIcon.png"`

WorkingDirectory

- The working directory to launch your application in
- Between 0 and 32767 characters
  - This character length upper limit is to support extended file paths in Windows. Ensure your AMI and application is properly configured to support Windows extended file paths if using file paths larger than 260 characters.
- Use escaped file path strings such as
  - `"C:\\Path\\To\\Working\\Directory"`

LaunchParameters

- A string to use as the launch parameters for the executable specified in `AbsoluteAppPath`
- Between 0 and 1024 characters
- Use escaped strings with the full list of required launch parameters such as the following example which shows how you may use PowerShell scripts as your applications by using the PowerShell executable as your app with a script provided in the launch parameters
  - AbsoluteAppPath
    - `"C:\\Windows\\System32\\WindowsPowerShell\\v1.0\\powershell.exe"`

- LaunchParameters
  - `"-File \"C:\\Path\\To\\App\\Script.ps1\""`

Sample AppCatalogConfig

This is a bare bones example of an AppCatalogConfig for Notepad, Google Chrome, and Mozilla Firefox

```
[
    {
        "Name": "Notepad",
        "DisplayName": "Notepad",
        "AbsoluteAppPath": "C:\\Windows\\System32\\notepad.exe"
    },
    {
        "Name": "Chrome",
        "DisplayName": "Chrome",
        "AbsoluteAppPath": "C:\\Program Files\\Google\\Chrome\\Application\
\chrome.exe",
        "LaunchParameters": "https://www.amazon.com/"
    },
    {
        "Name": "Firefox",
        "DisplayName": "Firefox",
        "AbsoluteAppPath": "C:\\Program Files\\Mozilla Firefox\\firefox.exe",
        "LaunchParameters": "https://aws.amazon.com/"
    }
 ]
```

Export Image

You can export your images to create EC2 AMIs. Later you can Import Image those AMIs back to create WorkSpaces Applications images. This helps you to use your own AMI customization tools for customizing of your images.

> **Note**
>
> During export following components will be removed from your images
>
> - WorkSpaces Applications agent

- Microsoft license included applications, which were added using Image Builder
- Only Microsoft Windows Server 2022 and 2025 images can be exported.

IAM Role Requirements

Important

Create an IAM role with the following permissions to use for export import:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowCopyImage",
            "Effect": "Allow",
            "Action": "ec2:CopyImage",
            "Resource": "*"
        },
        {
            "Sid": "AllowDescribeImages",
            "Effect": "Allow",
            "Action": "ec2:DescribeImages",
            "Resource": "*"
        },
        {
            "Sid": "AllowCreateTags",
            "Effect": "Allow",
            "Action": "ec2:CreateTags",
            "Resource": "*"
        }
    ]
}
```

Add the following trust relationship for this IAM role

```
{
    "Version": "2012-10-17",
    "Statement": [
```

```
        {
            "Effect": "Allow",
            "Principal": {
                "Service": "appstream.amazonaws.com"
            },
            "Action": "sts:AssumeRole"
        }
    ]
}
```

To export an image

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the navigation pane, choose **Images**, **Image Registry**.
3. In the image list, select the private image you want to export.
4. Choose **Actions**, **Export**.
5. In the **Export image** dialog box, type a unique **AMI name** and optionally **AMI Description** for the AMI.
6. **IAM Role** - Select the IAM role that you have created for image export.
7. You optionally copy tags from your Image to AMI by checking the **Copy tags in export** checkbox.
8. Choose **Export Image**.

Create Your Amazon WorkSpaces Applications Image Programmatically by Using the Image Assistant CLI Operations

You can create Amazon WorkSpaces Applications images by connecting to an image builder and using the Image Assistant graphical user interface (GUI) or command line interface (CLI) operations. The Image Assistant CLI operations provide functionality that is similar to the Image Assistant GUI. With these operations, you can programmatically do the following:

- Manage the applications that are included in an image.
- Save, update, and reset default application settings.
- Enable or disable the WorkSpaces Applications dynamic application framework.
- Specify tags.
- Create an image.

You can use these operations to integrate WorkSpaces Applications image creation with your continuous integration or deployment software development process.

To work with the Image Assistant CLI operations, use the command line shell of your choice on an image builder. For example, you can use the Windows command prompt or PowerShell.

> **Note**
>
> The image builder must use a version of the WorkSpaces Applications agent that is released on or after July 26, 2019. If you don't have an image builder, you must create one. For more information, see Launch an Image Builder to Install and Configure Streaming Applications.

Contents

- Creating Default Application and Windows Settings with the Image Assistant CLI operations
- Optimizing the Launch Performance of Your Applications with the Image Assistant CLI Operations
- Process Overview for Programmatically Creating an Amazon WorkSpaces Applications Image
- Image Assistant CLI Operations for Creating and Managing Your Amazon WorkSpaces Applications Image

Creating Default Application and Windows Settings with the Image Assistant CLI operations

You can create default application and Windows settings so that your users can get started with their applications quickly. When you create these settings, WorkSpaces Applications replaces the Windows default user profile with the profile that you configure. The Windows default user profile is then used to create the initial settings for users in the fleet instance. If you create these settings by using the Image Assistant CLI operations, your application installer, or the automation, should modify the Windows default user profile directly.

To overwrite the Windows default user profile with that of another Windows user, you can also use the Image Assistant `update-default-profile` CLI operation.

For more information about configuring default application and Windows settings, see *Creating Default Application and Windows Settings for Your WorkSpaces Applications Users* in Default

Application and Windows Settings and Application Launch Performance in Amazon WorkSpaces Applications.

Optimizing the Launch Performance of Your Applications with the Image Assistant CLI Operations

WorkSpaces Applications lets you optimize the launch performance of your applications for your users' streaming sessions. When you do so by using the Image Assistant CLI operations, you can specify the files to optimize for your application launch. Adding files to the application optimization manifest reduces the time that it takes for the application to launch for the first time on a new fleet instance. However, this also increases the time that it takes for the fleet instances to be made available to users. The optimization manifest is a line-delimited text file that is per application.

> **Note**
>
> If you onboard application optimization manifests by using both the Image Assistant CLI operations and the Image Assistant GUI, the manifests are merged.

Following is an example of an application optimization manifest file:

```
C:\Program Files (x86)\Notepad++\autoCompletion
C:\Program Files (x86)\Notepad++\localization
C:\Program Files (x86)\Notepad++\plugins
C:\Program Files (x86)\Notepad++\themes
C:\Program Files (x86)\Notepad++\updater
C:\Program Files (x86)\Notepad++\userDefineLangs
C:\Program Files (x86)\Notepad++\change.log
C:\Program Files (x86)\Notepad++\config.xml
C:\Program Files (x86)\Notepad++\contextMenu.xml
C:\Program Files (x86)\Notepad++\doLocalConf.xml
C:\Program Files (x86)\Notepad++\functionList.xml
C:\Program Files (x86)\Notepad++\langs.model.xml
C:\Program Files (x86)\Notepad++\license.txt
C:\Program Files (x86)\Notepad++\notepad++.exe
C:\Program Files (x86)\Notepad++\readme.txt
C:\Program Files (x86)\Notepad++\SciLexer.dll
C:\Program Files (x86)\Notepad++\shortcuts.xml
C:\Program Files (x86)\Notepad++\stylers.model.xml
```

For more information about optimizing the launch performance of your applications, see *Optimizing the Launch Performance of Your Applications* in Default Application and Windows Settings and Application Launch Performance in Amazon WorkSpaces Applications.

Process Overview for Programmatically Creating an Amazon WorkSpaces Applications Image

You can use the Image Assistant CLI operations with your application installation automation to create a fully programmatic WorkSpaces Applications image creation workflow. After your application installation automation completes, but before the image is created, use the Image Assistant CLI operations to specify the following:

- The executable files that your users can launch
- The optimization manifests for your applications
- Other WorkSpaces Applications image metadata

The following high-level overview describes the process for programmatically creating an WorkSpaces Applications image.

1. Use your application installation automation to install the required applications on your image builder. This installation may include applications that your users will launch, any dependencies, and background applications.
2. Determine the files and folders to optimize.
3. If applicable, use the Image Assistant `add-application` CLI operation to specify the application metadata and optimization manifest for the WorkSpaces Applications image.
4. To specify additional applications for the WorkSpaces Applications image, repeat steps 1 through 3 for each application as needed.
5. If applicable, use the Image Assistant `update-default-profile` CLI operation to overwrite the default Windows profile and create default application and Windows settings for your users.
6. Use the Image Assistant `create-image` CLI operation to create the image.

Image Assistant CLI Operations for Creating and Managing Your Amazon WorkSpaces Applications Image

This section describes the Image Assistant CLI operations that you can use to create and manage your WorkSpaces Applications image.

On Windows image builders, the executable file that includes the command line interface is located at: C:\Program Files\Amazon\Photon\ConsoleImageBuilder\Image-Assistant.exe. For your convenience, this executable file is included in the Windows PATH variable. This lets you call the Image Assistant CLI operations without specifying the absolute path to the executable file. To call these operations, type the **image-assistant.exe** command.

On Linux image builders, the image assistant tool is located at /usr/local/appstream/image-assistant/AppStreamImageAssistant, with a symbolic link at /bin/AppStreamImageAssistant.

`help` operation

Retrieves a list of all Image Assistant CLI operations. For each operation in the list, a description and usage syntax is provided. To display help for a specific operation, type the name of the operation and specify the **--help** parameter. For example:

```
add-application --help
```

Synopsis

```
help
```

Output

Prints to standard out the list of available operations with a description of their function.

`add-application` operation

Adds the application to the application list for WorkSpaces Applications users. Applications in this list are included in the application catalog. The application catalog displays to users when they sign in to an WorkSpaces Applications streaming session.

> **Note**
>
> If you need to make changes to an application configuration, remove the application and add the application with the updated settings.

Synopsis

```
add-application
--name <value>
--absolute-app-path <value>
[--display-name <value>]
[--absolute-icon-path <value>]
[--working-directory <value>]
[--launch-parameters <""-escaped value>]
[--absolute-manifest-path <value>]
```

Options

`--name` (string)

A unique name for the application. The maximum length is 256 characters. You can add up to 50 applications. You cannot use whitespace characters.

`--absolute-app-path` (string)

The absolute path to the executable file, batch file, or script for the application. The path must point to a valid file.

`--display-name` (string)

The name to display for the application in the application catalog. If you don't specify a display name, WorkSpaces Applications creates a name that is derived from the executable file name. The name is created without the file extension and with underscores in place of spaces. The maximum length is 256 characters.

`--absolute-icon-path` (string)

The absolute path to the icon for the application. The path must point to a valid icon file that is one of the following types: .jpg, .png, or .bmp. The maximum dimensions are: 256 px x 256 px. If you don't specify a path, the default icon for the executable file is used, if available.
If a default icon is not available for the executable file, a default WorkSpaces Applications application icon is used.

`--working-directory` (string)

The initial working directory for the application when the application is launched.

`--absolute-manifest-path` (string)

The path to a new line-delimited text file. The file specifies the absolute paths of the files to optimize before the fleet instance is made available to a user for streaming. The path must point to a valid text file.

Message output

Exit code	Message printed to standard out	Description
0	{"status": 0, "message": "Success"}	The application was added successfully.
1	{"status": 1, "message": "Administrator privileges are required to perform this operation"}	Administrator privileges are required to complete the operation.
1	{"status": 1, "message": "Unable to add more than 50 apps to the catalog."}	The application could not be added because the maximum number of applications that can be added to the WorkSpaces Applications application catalog is 50.
1	{"status": 1, "message": "Name is not unique"}	An application with that name already exists in the WorkSpaces Applications application catalog.
1	{"status": 1, "message": "File not found (absolute-app-path)"}	The file that was specified for `absolute-app-path` could not be found.
1	{"status": 1, "message": "Unsupported file extension"}	The `Absolute-app-path` parameter only supports the following file types: .exe and .bat.

Exit code	Message printed to standard out	Description
1	{"status": 1, "message" : "Directory not found (working-directory)"	The directory that was specified for `working-directory` could not be found.
1	{"status": 1, "message": "Optimization-manifest not found: <filename>"}	The file that was specified for `optimization-manifest` could not be found.
1	{"status": 1, "message" : "File not found: <filename>"}	A file that was specified within the optimization manifest could not be found.
255	{"status": 255, "message": <error message>}	An unexpected error occurred. Try the request again. If the error persists, contact AWS Support for assistance. For more information, see AWS Support Center.

`remove-application` operation

Removes an application from the application list for the WorkSpaces Applications image. The application is not uninstalled or modified, but users will not be able to launch it from the WorkSpaces Applications application catalog.

Synopsis

```
remove-application
--name <value>
```

Options

`--name` (string)

The unique identifier of the application to remove.

Message output

Exit code	Message printed to standard out	Description
0	{"status": 0, "message": "Success"}	The application was removed successfully.
1	{"status": 1, "message": "Administrator privileges are required to perform this operation"}	Administrator privileges are required to complete the operation.
1	{"status": 1, "message": "App not found"}	The application that was specified could not be found in the WorkSpaces Applications application catalog.
255	{"status": 255, "message": <error message>}	An unexpected error occurred. Try the request again. If the error persists, contact AWS Support for assistance. For more information, see AWS Support Center.

`list-applications` operation

Lists all of the applications that are specified in the application catalog.

Synopsis

```
list-applications
```

Message output

Exit code	Message printed to standard out	Description
0	{"status": 0, "message": "Success", "applicat	List of applications in the WorkSpaces Applications application catalog.

Exit code	Message printed to standard out	Description
	ions": [{..app1.. }, { ..app2.. }]}	
255	{"status": 255, "message": <error message>}	An unexpected error occurred. Try the request again. If the error persists, contact AWS Support for assistance. For more information, see AWS Support Center.

`update-default-profile` operation

Copies the specified Windows user's profile to the Windows default user profile. New users who stream inherit the settings stored in the specified profile.

> **Note**
>
> This operation is not supported by the Linux image assistant CLI tool.

Synopsis

```
update-default-profile
[--profile <value>]
```

Options

`--profile` (string)

The name of the user whose Windows profile will be copied to the Windows default user profile. Use the following format for the name:

"<domain>\<username>"

If your image builder isn't joined to a Microsoft Active Directory domain, enter a period "." for the domain. If you don't specify a user, the WorkSpaces Applications Template User account is used.

Message output

Exit code	Message printed to standard out	Description
0	{"status": 0, "message": "Success"}	The user settings were successfully copied to the default Windows profile.
1	{"status": 1, "message": "Administrator privileges are required to perform this operation"}	Administrator privileges are required to complete the operation.
1	{"status": 1, "message": "Unable to copy file or folder: <path>. <reason>"}	The user settings could not be copied because a file or folder was unavailable.
1	{"status": 1, "message": "Cannot copy a domain user when not joined to a domain""}	A Microsoft Active Directory domain user was specified, but the image builder is not joined to an Active Directory domain.
255	{"status": 255, "message": <error message>}	An unexpected error occurred. Try the request again. If the error persists, contact AWS Support for assistance. For more information, see AWS Support Center.

`reset-user-profile` operation

Deletes the Windows user profile for the specified user.

> **Note**
>
> This operation is not supported by the Linux image assistant CLI tool.

Synopsis

```
reset-user-profile
[--profile <value>]
```

Options

`--profile` (string)

The name of the Windows user whose Windows profile will be deleted. Use the following format for the name:

"<domain>\<username>"

If your image builder isn't joined to a Microsoft Active Directory domain, enter a period "." for the domain.

Message output

Exit code	Message printed to standard out	Description
0	{"status": 0, "message": "Success"}	The specified user settings were deleted successfu lly.
1	{"status": 1, "message": "Administrator privilege s are required to perform this operation "}	Administrator privileges are required to complete the operation.
1	{"status": 1, "message" : "Unable to copy file or folder: <path>. <reason>"}	The user settings could not be reset because a file or folder was unavailable.
1	{"status": 1, "message": "Cannot copy a domain	A Microsoft Active Directory domain user was specified, but the image builder is not joined to an Active Directory domain.

Exit code	Message printed to standard out	Description
	user when not joined to a domain""}	
255	{"status": 255, "message": <error message>}	An unexpected error occurred. Try the request again. If the error persists, contact AWS Support for assistance. For more information, see AWS Support Center.

`create-image` operation

Starts the image creation workflow, resulting in an WorkSpaces Applications image that can be used for WorkSpaces Applications fleets.

Synopsis

```
create-image
--name <value>
[--description <value>]
[--display-name <value>]
[--enable-dynamic-app-catalog] | [--no-enable-dynamic-app-catalog]
[--use-latest-agent-version] | [--no-use-latest-agent-version]
[--tags <value>]
[--dry-run]
```

Options

`--name` (string)

The name for the WorkSpaces Applications image. The name must be unique within the Amazon Web Services account and AWS Region. The maximum length is 100 characters. Allowed characters are:

a-z, A-Z, 0-9, underscores (_), hyphens (-), and periods (.)

The image name cannot start with any of the following prefixes: 'aws', 'appstream', and 'amazon'. These prefixes are reserved for AWS use.

`--description` (string)

The description to display for the image. The maximum length is 256 characters.

`--display-name` (string)

The name to display for the image. The maximum length is 256 characters.

`--enable-dynamic-app-catalog` | `--no-enable-dynamic-app-catalog`

Enables or disables support for the WorkSpaces Applications dynamic application framework. If you don't specify either parameter, support for the dynamic application framework is not enabled.

The dynamic application framework provides operations within an WorkSpaces Applications streaming instance that you can use to build a dynamic app provider. Dynamic app providers can use these operations to modify the catalog of applications that your users can access in real time. For more information, see Use the WorkSpaces Applications Dynamic Application Framework to Build a Dynamic App Provider.

`--use-latest-agent-version` | `--no-use-latest-agent-version`

Specifies whether to pin the image to the version of the WorkSpaces Applications agent that is currently installed, or to always use the latest agent version. If you don't specify either parameter, the image is pinned to the version of the WorkSpaces Applications agent that is currently installed. For more information, see Manage WorkSpaces Applications Agent Versions.

`--tags` (string)

The tags to associate with the image. A tag is a key-value pair. Use the following format:

--tags "mykey" "myval" "mykey2" "myval2"

For more information about tags, see Tagging Your Amazon WorkSpaces Applications Resources.

`--dry-run` (string)

Performs validation without creating the image. Use this command to identify whether your image has any issues before you create it.

Message output

Exit code	Message printed to standard out	Description
0	{"status": 0, "message": "Success"}	The workflow to create the image was initiated successfully.
1	{"status": 1, "message": "Administrator privileges are required to perform this operation"}	Administrator privileges are required to complete the operation.
1	{"status": 1, "message": "An image with the given name already exists"}	An image with the specified name already exists in the Amazon Web Services account.
1	{"status": 1, "message": "Invalid value (tags)"}	The specified tags are not valid.
255	{"status": 255, "message": <error message>}	An unexpected error occurred. Try the request again. If the error persists, contact AWS Support for assistance. For more information, see AWS Support Center.

Create Your Linux-Based Images

You can create Linux-based Amazon WorkSpaces Applications images by connecting to a Linux image builder, installing the applications you need, creating default application settings and environment variables, and using a command line interface (CLI) tool or Image Assistant (GUI) tool to add these applications to the application catalog. To open the GUI tool, search for **Image Assistant** in the list of applications.

Contents

- Creating Default Application Settings for Your Users
- Creating Default Environment Variables for Your Linux Users

- Optimizing the Launch Performance of Your Linux Applications
- Creating Session Scripts
- Using the Image Assistant CLI Tool for Linux
- Enabling and Disabling Webcam Support
- Enabling and Disabling Heavy File Sync Mode for Home Folders
- Tutorial: Create a Custom Linux-Based WorkSpaces Applications Image
- Tutorial: Enable Japanese Support for Your Linux Images

Creating Default Application Settings for Your Users

Follow these steps to create default application settings for your users.

Contents

- Step 1: Install Linux Applications on the Image Builder
- Step 2: Create a TemplateUser Account
- Step 3: Create Default Application Settings
- Step 4: Save Default Application Settings
- Step 5: Test Default Application Settings (optional)
- Step 6: Clean Up

Step 1: Install Linux Applications on the Image Builder

In this step, you connect a Linux image builder and install your applications on the image builder.

To install applications on the image builder

1. Connect to the image builder by doing either of the following:

   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

   > **Note**
   >
   > You will be logged in as an ImageBuilderAdmin user to the Linux GNOME desktop and have root admin privileges.

2. Install the applications that you need. For example, to install a Chromium browser from a public yum repo, first open the Terminal application, then run the following command:

   [ImageBuilderAdmin]$ sudo yum update && sudo yum install chromium.x86_64

Step 2: Create a TemplateUser Account

In this step, you create a TemplateUser account, which creates default application settings for your streaming users.

To create a TemplateUser Account

1. Create a TemplateUser account that has no root permissions. For example, in a Terminal window, run the following commands to create TemplateUser on the image builder:

   [ImageBuilderAdmin]$ sudo useradd -m TemplateUser

   [ImageBuilderAdmin]$ echo -e '<*password*>\n<*password*>\n' | sudo passwd TemplateUser

2. Switch to the TemplateUser account:

   [ImageBuilderAdmin]$ su - TemplateUser

Step 3: Create Default Application Settings

In this step, you create default application settings for your WorkSpaces Applications users. Doing this enables your users to get started with applications quickly during their WorkSpaces Applications streaming sessions, without the need to create or configure these settings themselves.

To create default application settings for your users

1. Launch the application that you want to create the default settings for. For example, in a Terminal window, run following command to launch Chromium browser:

   [TemplateUser]$ chromium-browser

2. Configure the settings of the application. For example, set the home page of the Chromium browser as **`https://aws.amazon.com`**.
3. Close the applications.
4. Log out:

[TemplateUser]$ logout

Step 4: Save Default Application Settings

In this step, you copy the default application settings you added to the **/etc/skel/** directory, and make them available to your streaming users.

To save default application settings

- Run the following command in a Terminal window to copy the default application settings for your streaming users:

  [ImageBuilderAdmin]$ sudo cp -r -f /home/TemplateUser/. /etc/skel

Step 5: Test Default Application Settings (optional)

In this step, verify that the applications you've added run correctly, and the default application settings work as expected.

To test your applications and default settings on an image builder

1. Create a test user that has no root permissions. For example, in a **Terminal** window, run the following commands to create **test-user** on the image builder:

   [ImageBuilderAdmin]$ sudo useradd -m test-user

   [ImageBuilderAdmin]$ echo -e '*password*>\n<*password*>\n' | sudo passwd test-user
2. Switch to the test user:

   [ImageBuilderAdmin]$ su - test-user
3. Launch the application (e.g., Chromium) as the test user:

   [test-user]$ /usr/bin/chromium-browser
4. Verify that the default settings are available for the test user (e.g., the Chromium home page is https://aws.amazon.com/).
5. Log out:

   [test-user]$ logout

Step 6: Clean Up

Finally, your last step is to clean up.

To clean up

1. Delete TemplateUser:

   [ImageBuilderAdmin]$ sudo killall -u TemplateUser

   [ImageBuilderAdmin]$ sudo userdel -r TemplateUser
2. Delete test-user (not required if you skipped step 5):

   [ImageBuilderAdmin]$ sudo killall -u test-user

   ImageBuilderAdmin]$ sudo userdel -r test-user

Creating Default Environment Variables for Your Linux Users

You can create environment variables on a Linux Image Builder instance. Creating environment variables makes them available on streaming instances created from that image.

> **Note**
>
> On Linux fleet instances, environment variables set using the Image Assistant (GUI) tool and the default system environment variables are exported through the /etc/profile.d/appstream_system_vars.sh script. To access these environment variables, you must explicitly source the /etc/profile.d/appstream_system_vars.sh script in your applications.

To create environment variables for your users

1. If the folder `/etc/profile` doesn't exist, run the following command to create it:

   [ImageBuilderAdmin]$ sudo mkdir -p /etc/profile.d
2. To create a new shell script file (for example, my-environment.sh) in this folder, run the following command:

   [ImageBuilderAdmin]$ vim my-environment.sh
3. On first line of the script file, add the following content:

#!/bin/sh

4. For each subsequent line, add an **export** command to set the environment variables for your image. The following example adds `$HOME/bin` to the `PATH` variable:

   export PATH="$HOME/bin:$PATH"

5. Press the **Esc** key to return to command mode in vim, then run the following command to save your script and exit vim:

   :x

6. Run the following command to allow the script to run as a program:

   [ImageBuilderAdmin]$ chmod +x my-environment.sh

Optimizing the Launch Performance of Your Linux Applications

If you are using the Image Assistant GUI tool, the tool optimizes launch performance for your applications automatically.

If you are using the Image Assistant CLI, use the following steps to optimize launch performance manually. When you create and add files to an application optimization manifest, the application will launch faster when first started on a new fleet instance. However, this also increases the time that it takes for the fleet instances to be made available to users. The optimization manifest is one line-delimited text file for every application.

You can create a manifest file (such as <*your-app*>-manifest.txt) either manually or by following with the steps below.

To create a manifest file

1. Make sure that the application that you are trying to optimize is launched and running.
2. From a terminal in the Linux image builder, run the following command:

   ps -ef | grep <*application-process-name*>

3. Search for the smallest PID number from the last step's output. This is the PID for the root parent process of the application.
4. Keep the application running and make sure to use the initial components required by your users. This ensures that these components are captured by the optimization process.

5. Create a script file (e.g., ~/getfilestool.sh) with the following content:

```
#!/bin/bash
## usage getfilestool.sh $pid
lsof -p $(pstree -p $1 | grep -o '([0-9]\+)' | grep -o '[0-9]\+' | tr '\012' ,)|
grep REG | sed -n '1!p' | awk '{print $9}'|awk 'NF'
```

6. Make sure that the file can be run with the following command:

   [ImageBuilderAdmin]$ chmod u+x ~/getfilestool.sh

7. Run the following command to capture all of the running files from the root parent process found in step 3, and save it to a temporary manifest file.

   [ImageBuilderAdmin]$ sudo ~/getfilestool.sh <*root-parent-pid*> > /tmp/<y*our-app*>-manifest.txt

8. Verify the content of the optimization manifest, which is a line-delimited text file for every application.

You can specify the optimization manifest on a per-application basis by using the Image Assistant command line interface (CLI) tool. For more information, see the section called “Using the Image Assistant CLI Tool for Linux”.

Creating Session Scripts

WorkSpaces Applications provides on-instance session scripts on both Windows- and Linux-based streaming instances. For more information about session scripts, see the section called “Session Scripts to Manage Your Users' Streaming Experience”.

Session scripts are specified within an WorkSpaces Applications image. To locate the session scripts configuration file on a Linux instance, navigate to `/opt/appstream/SessionScripts/config.json`. The following code is a sample `config.json` file that specifies a session start script named “`test-session-start`” and a session end script named “`test-session-stop`” together with their runtime parameters. Make sure that the scripts referenced in `config.json` have run permissions and a command interpreter is defined (for example, #!/bin/bash).

```
{
     "SessionStart": {
          "Executables": [
               {
```

```
                    "Context": "system",
                    "Filename": "/opt/appstream/SessionScripts/test-session-start",
                    "Arguments": "arg1",
                    "S3LogEnabled": true
               }
          ],
          "WaitingTime": 30
     },
     "SessionTermination": {
          "Executables": [
               {
                    "Context": "system",
                    "Filename": "/opt/appstream/SessionScripts/test-session-stop",
                    "Arguments": "arg2",
                    "S3LogEnabled": true
               }
          ],
          "WaitingTime": 30
     }
}
```

Using the Image Assistant CLI Tool for Linux

On a Linux-based image builder, you can use the Image Assistant CLI tool **AppStreamImageAssistant** to create and manage your WorkSpaces Applications image. The tool is located at `/usr/local/appstream/image-assistant/AppStreamImageAssistant` with a symbolic link at `/bin/AppStreamImageAssistant`. This CLI tool for Linux supports many of the same operations as the Image Assistant CLI tool for Windows. For more information on these operations, see the section called “Image Assistant CLI Operations”.

Enabling and Disabling Webcam Support

WorkSpaces Applications supports real-time audio-video (AV) by redirecting local webcam video input to WorkSpaces Applications streaming sessions. This capability enables your users to use their local webcam for video and audio conferencing within an WorkSpaces Applications streaming session. With real-time AV and support for real-time audio, your users can collaborate by using familiar video and audio conferencing applications without having to leave their WorkSpaces Applications streaming session.

To use this feature, you must use a Linux WorkSpaces Applications image that uses a Linux WorkSpaces Applications agent released on or after September 21, 2022.

> **Note**
>
> Real-time AV is not supported for stream.standard.small instances powered by Rocky Linux or Red Hat Enterprise Linux. Users don't see the Camera and Mic icons on the client toolbar.

The real-time AV feature is enabled by default for Linux streaming sessions. To configure webcam permissions for your users on a Linux image builder, create `/etc/appstream/appstream.conf` and add the following contents:

> **Note**
>
> Specify **1** to enable webcam, or **0** to disable webcam.

```
[webcam]
permission = 1
```

Enabling and Disabling Heavy File Sync Mode for Home Folders

You can enable Amazon Simple Storage Service Home Folders options for your organization. When you enable Amazon S3 Home Folders for an WorkSpaces Applications stack, users of the stack can access a persistent storage folder during their application streaming sessions. No further configuration is required for your users to access their home folder. Data stored by users in their home folder is automatically backed up to an Amazon S3 bucket in your AWS account, and is made available to those users in subsequent sessions. For more information, see the section called “Administer Home Folders”.

To ensure a smooth experience and address some existing limitations, where an inconsistent file sync might be observed when users save large text files from their streaming instances to their Home Folders, WorkSpaces Applications administrators can turn on the **heavy_sync** configuration option if large file uploads to Amazon S3 are a common user scenario while using WorkSpaces Applications. Turning on this option means that it might add some latency to the home folder file sync process, but completeness of all syncs to Amazon S3 is guaranteed.

This feature is available on all Red Hat Enterprise Linux images, and Linux WorkSpaces Applications images that use a Linux WorkSpaces Applications agent released on or after September 12, 2024.

The heavy sync feature is disabled by default for Red Hat Enterprise Linux and Rocky Linux streaming sessions. To configure heavy sync permission for your users on a Red Hat Enterprise Linux or Rocky Linux image builder, create `/etc/appstream/appstream.conf` and add the following contents:

> **Note**
>
> Specify **1** to enable heavy sync, or **0** to disable heavy sync.

```
[storage]
heavy_sync = 1
```

Tutorial: Create a Custom Linux-Based WorkSpaces Applications Image

This tutorial describes how to create a custom Linux-based Amazon WorkSpaces Applications image that contains applications which you can stream to your users.

> **Important**
>
> Don't create a user named "as2-streaming-user" in your image builder. This is a reserved username for Fleet. If you create this username outside of the WorkSpaces Applications workflow, you might experience streaming issues in Fleets.

Contents

- Step 1: Install Linux Applications on the Image Builder
- Step 2: Generate Application Optimization Manifest File
- Step 3: Create an WorkSpaces Applications Application Catalog
- Step 4: Create Default Application Settings and Environment Variables
- Step 5: Test Applications and Settings
- Step 6: Finish Creating Your Image
- Step 7 (Optional): Tag and Copy an Image
- Step 8: Clean Up

Step 1: Install Linux Applications on the Image Builder

In this step, you connect a Linux image builder and install your applications on the image builder.

To install applications on the image builder

1. Connect to the image builder by doing either of the following:

   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

   > **Note**
   >
   > You will be logged in as an ImageBuilderAdmin user to the Linux GNOME desktop and have root admin privileges.

2. Install the applications that you need. For example, to install a Chromium browser from a public yum repo, first open the Terminal application, then run the following command:

   [ImageBuilderAdmin]$ sudo yum update && sudo yum install chromium.x86_64

   > **Note**
   >
   > Download and install applications only from sites that you trust.

Step 2: Generate Application Optimization Manifest File

In this step, you generate a manifest file for each application you installed in step 1.

To generate a manifest file for optimizing the launch performance of an application

1. Make sure the application (e.g., Chromium) that you are trying to optimize is launched and running.
2. In a Terminal window, run the following command to list processes related to the application:

   [ImageBuilderAdmin]$ ps -ef | grep chromium

3. Find the root parent PID from the output of the command above. The following is an example output, and the root parent PID is 16712:

Example

```
[ImageBuilderAdmin]$ ps -ef | grep chromium

ImageBu+ 16712 4128 0 Aug26 ? 00:00:44 /usr/lib64/chromium- browser/chromium-
browser --enable-plugins --enable-extensions -- enable-user- scripts --enable-
printing --enable-gpu-rasterization -- enable-sync --auto-ssl- client-auth

ImageBu+ 16726 16712 0 Aug26 ? 00:00:00 /usr/lib64/chromium- browser/chromium-
browser --type=zygote --no-zygote-sandbox ImageBu+ 16727 16712 0 Aug26 ? 00:00:00 /
usr/lib64/chromium- browser/chromium- browser --type=zygote

ImageBu+ 16731 16727 0 Aug26 ? 00:00:00 /usr/lib64/chromium- browser/chromium-
browser --type=zygot
```

4. Keep the application running and make sure to use the initial components required by your users. This ensures that these components are captured by the optimization process.
5. Create script file (e.g., ~/`getfilestool.sh`) with the following content:

```
#!/bin/bash
## usage getfilestool.sh $pid
lsof -p $(pstree -p $1 | grep -o '([0-9]\+)' | grep -o '[0-9]\+' | tr '\012' ,)|
grep REG | sed -n '1!p' | awk '{print $9}'|awk 'NF'
```

6. Verify that the file can be run by running the following command:

   [ImageBuilderAdmin]$ chmod u+x ~/getfilestool.sh

7. Run the following command to capture all running files from the root parent process found in step 3 above, and save it to a temporary manifest file:

   [ImageBuilderAdmin]$ sudo ~/getfilestool.sh 16712 > /tmp/chromium-manifest.txt

8. Verify the content of the optimization manifest, which is a line-delimited text file for each application.

Step 3: Create an WorkSpaces Applications Application Catalog

In this step, you use the CLI tool `AppStreamImageAssistant` on the image builder to create an WorkSpaces Applications application catalog by specifying applications for your image. For each

application that you plan to stream, you can specify the name, display name, executable file to launch, and icon to display.

To create an WorkSpaces Applications application catalog

1. From the image builder desktop, open **Terminal** either from the side panel or by opening the app grid.
2. Run **AppStreamImageAssistant --help** to see the list of available commands. You will use these commands to add applications and create an Image.
3. Run the following command to add an installed application (e.g., Chromium) to the application list for WorkSpaces Applications users:

   ```
   AppStreamImageAssistant add-application \
    --name Chromium \
    --absolute-app-path /usr/lib64/chromium-browser/chromium-browser \
    --display-name Chromium \
    --absolute-icon-path /usr/share/icons/hicolor/256x256/apps/chromium-browser.png \
    --absolute-manifest-path /tmp/chromium-manifest.txt
   ```

   Alternatively, run the following command:

   ```
   AppStreamImageAssistant add-application \
    --name="Chromium" \
    --absolute-app-path="/usr/lib64/chromium-browser/chromium-browser" \
    --display-name="Chromium" \
    --absolute-icon-path="/usr/share/icons/hicolor/256x256/apps/chromium-browser.png"
    \
    --absolute-manifest-path="/tmp/chromium-manifest.txt"
   ```

4. To add more applications, repeat step 3 for each additional application.
5. To see the list of applications that have been added in the catalog, along with metadata like icon paths and launch parameters, run the following command:

   AppStreamImageAssistant list-applications

6. To remove applications from the catalog, run the following command:

   AppStreamImageAssistant remove-application –-name *application_name*

Step 4: Create Default Application Settings and Environment Variables

In this step, you create default application settings and environment variables for your WorkSpaces Applications users. Doing this enables your users to get started with applications quickly during their WorkSpaces Applications streaming sessions, without the need to create or configure these settings themselves.

To create default application and environment variables for your users

1. Launch the application that you want create the default settings for. For example, in a Terminal window, run following command to launch Chromium browser:

   [ImageBuilderAdmin]$ chromium-browser
2. Configure the settings of the application. For example, set the home page of the Chromium browser as **`https://aws.amazon.com`**.
3. Make sure the Chromium application is closed, and then run the following commands to copy the configuration for Chromium to **/etc/skel**:

   [ImageBuilderAdmin]$ sudo mkdir /etc/skel/.config

   [ImageBuilderAdmin]$ sudo cp -R ~/.config/chromium /etc/skel/.config
4. Set the environment variables and add it to the script file. For example, run the following commands:

   [ImageBuilderAdmin]$ echo "export *FOO*=*BAR*" | sudo tee -a /etc/profile.d/myenvvars.sh

   [ImageBuilderAdmin]$ sudo chmod +x /etc/profile.d/myenvvars.sh

Step 5: Test Applications and Settings

In this step, verify that the applications you've added run correctly, and the default application settings and environment variables work as expected.

To test your applications and default settings on an image builder

1. Create a test user that has no root permissions. For example, in a **Terminal** window, run the following commands to create **test-user** on the image builder:

   [ImageBuilderAdmin]$ sudo useradd -m test-user

[ImageBuilderAdmin]$ echo -e 'Pa55w0rdas2!!!\nPa55w0rdas2!!!\n' | sudo passwd test-user

2. Switch to the test user:

   [ImageBuilderAdmin]$ su - test-user

3. Launch the application (e.g., Chromium) as the test user:

   [test-user]$ /usr/bin/chromium-browser

4. Verify that the default settings are available for the test user (e.g., the Chromium home page is https://aws.amazon.com/).
5. Verify that the environment variables are available for the test user. For example, run the following command:

   [test-user]$ echo $*FOO*

   This command should display the output ***BAR*** in the terminal.

6. Run the following commands to delete the test user before creating an image from this image builder:

   # logout test user

   [test-user]$ logout

   # kill test user's running processes

   [ImageBuilderAdmin]$ sudo killall -u test-user

   # delete user

   [ImageBuilderAdmin]$ sudo userdel -r test-user

Step 6: Finish Creating Your Image

In this step, choose an image name and finish creating your image.

To create the image

1. In a **Terminal** window, create an image from your Image Builder by running **AppStreamImageAssistant create-image**. This image contains your installed and registered applications, plus any session scripts and default application settings that you have configured.

   To see the list of available options, run **AppStreamImageAssistant create-image --help**. For more information, see the **create-image** operation in the section called “Create Your Image Programmatically”.
2. The remote session disconnects after a few moments. When the **Lost Connectivity** message appears, close the browser tab. While the image is created, the image builder status appears as **Snapshotting**. You cannot connect to the image builder until this process finishes.
3. Return to the console and navigate to **Images**, **Image Registry**. Verify that your new image appears in the list.

   While your image is being created, the image status in the image registry of the console appears as **Pending**. You can't connect to images that are in a **Pending** status.
4. Choose the **Refresh** icon to update the status. After your image is created, the image status changes to **Available** and the image builder is automatically stopped.

   To continue creating images, start the image builder and connect to it from the console, or create a new image builder.

Step 7 (Optional): Tag and Copy an Image

You can add one or more tags to an image during image creation or after you create an image. You can also copy the image within the same Region or to a new Region within the same Amazon Web Services account. Copying a source image results in an identical but distinct destination image. AWS does not copy any user-defined tags, however. Also, you can only copy custom images that you create, not the base images that are provided by AWS.

> **Note**
>
> You can copy up to two images at the same time to a destination. If the destination to which you are copying an image is at the image limit, you receive an error. To copy the image in this case, you must first remove images from the destination. After the destination

is below the image quota (also referred to as limit), initiate the image copy from the source Region. For more information, see Amazon WorkSpaces Applications Service Quotas.

To add tags to an existing image

1. In the navigation pane, choose **Images**, **Image Registry**.
2. In the image list, select the image to which you want to add tags.
3. Choose **Tags**, choose **Add/Edit Tags**, and then choose **Add Tag**. Specify the key and value for the tag, and then choose **Save**.

For more information, see Tagging Your Amazon WorkSpaces Applications Resources.

To copy an image

Copying an image across geographically diverse regions enables you to stream applications from multiple regions based on the same image. By streaming your applications in closer proximity to your users, you can improve your users' experience streaming applications with WorkSpaces Applications.

1. In the navigation pane, choose **Images**, **Image Registry**.
2. In the image list, select the image that you want to copy.
3. Choose **Actions**, **Copy**.
4. In the **Copy Image** dialog box, specify the following information, and then choose **Copy Image**:
   - For **Destination region**, choose the region to which to copy the new image.
   - For **Name**, specify a name that the image will have when it is copied to the destination.
   - For **Description** (optional), specify a description that the image will have when it is copied to the destination.
5. To check on the progress of the copy operation, return to the console and navigate to **Images**, **Image Registry**. Use the navigation bar to switch to the destination region (if applicable), and confirm that your new image appears in the list of images.

   The new image first appears with a status of **Copying** in the image registry of your console. After the image is successfully created, the status of the image changes to **Available**, which means that you can use the image to launch a stack and stream your applications.

Step 8: Clean Up

Finally, you can stop your running image builders to free up resources and avoid unintended charges to your account. We recommend stopping any unused, running image builders. For more information, see WorkSpaces Applications Pricing.

To stop a running image builder

1. In the navigation pane, choose **Images**, **Image Builders**, and select the running image builder instance.
2. Choose **Actions**, **Stop**.

Tutorial: Enable Japanese Support for Your Linux Images

This tutorial describes how to enable Japanese support for a Linux image. This enables applications on the image to display Japanese characters, and streaming users to use the Japanese input method in the streaming sessions from the image.

Contents

- Step 1: Install Japanese Font and Input Method
- Step 2: Set the System Time Zone
- Step 3: Set the System Locale and Display Language
- Step 4: Configure the Input Methods
- Step 5: Set the Keyboard Layout
- Step 6: Verify on Image Builder
- Step 7: Create the Image

Step 1: Install Japanese Font and Input Method

In this step, you connect a Linux image builder and install the font and input method packages of your choice.

To install Japanese font and input method

1. Connect to the image builder by doing either of the following:
   - Use the WorkSpaces Applications console (for web connections only)

- Create a streaming URL (for web or WorkSpaces Applications client connections)

> **Note**
>
> You will be logged in as an ImageBuilderAdmin user to the Linux GNOME desktop and have root admin privileges.

2. Install the font and input method that you need. To do this, open the Terminal application, then run the following commands:

   sudo yum install vlgothic-p-fonts.noarch

   sudo yum install ibus-kkc.x86_64

3. In addition to the above commands, for Rocky Linux and Red Hat Enterprise Linux, run the following command:

   sudo yum install glibc-langpack-ja

Step 2: Set the System Time Zone

To set the system time zone, run the following command:

sudo timedatectl set-timezone "Asia/Tokyo"

Step 3: Set the System Locale and Display Language

To set the system locale and display language, run the following commands.

To set the system locale and display language

1. Update the `cloud-init config` file by running the command **sudo vim /etc/cloud/cloud.cfg**, and change **locale** to **locale: ja_JP.utf8**, then save and close the file.
2. Update the system settings by running **sudo localectl set-locale LANG=ja_JP.utf8**.
3. Update the Gnome shell settings by running **sudo gsettings set org.gnome.system.locale region "ja_JP.utf8"**.

Step 4: Configure the Input Methods

Configure the input methods for the application you want to add to the image. For more information about how install an application, generate a manifest file, and create default settings, see Tutorial: Create a Custom Linux-Based WorkSpaces Applications Image. In this step, we assume that you've already installed the application Firefox, which is located at `/usr/local/firefox/firefox`.

To configure the input methods

1. Create a script by running the command **sudo vim /usr/local/bin/update-input-method.sh**, and add the following content to the script:

```
#!/bin/bash

function start_process()
{
    command=$1
    process_name=$2

    process_count=$(pgrep $process_name -c)
    echo "$(date) current $process_name count: $process_count"
    while [ $process_count -lt 1 ]
    do
        echo "$(date) starting $process_name"
        eval $command
        sleep 1
        process_count=$(pgrep $process_name -c)
    done
    echo "$(date) $process_name started"
}

start_process "ibus-daemon --xim &" "ibus-daemon"
start_process "/usr/libexec/ibus-engine-kkc --ibus &" "ibus-engine-kkc"

gsettings set org.gnome.desktop.input-sources sources "[('ibus','kkc'), ('xkb',
 'us')]"
gsettings set org.gnome.desktop.wm.keybindings switch-input-source
 "['<Control>space']"
gsettings set org.gnome.desktop.wm.keybindings switch-input-source-backward
 "['<Shift><Control>space']"
```

```
echo "$(date) updated input source and switch shortcut"
```

In the script above, the first input source ('ibus', 'kkc') is the default input method. You can change the default input method by changing the order of input sources. In addition, "Control +Space" and "Shift+Control+Space" are specified as shortcut key combinations for switching between input methods. You can specify your own key combinations that your users can use to switch input methods during streaming sessions.

2. Create the script for launching the application (Firefox) that you will add to the image. To do this, run the command **sudo vim /usr/local/bin/firefox-jp.sh**, then add the following content to the script:

```
#!/bin/bash

# Gather required environment variables from the GNOME shell session
while IFS= read -r -d $'\0' env_var; do
    case "$env_var" in
        DBUS_SESSION_BUS_ADDRESS=*|\
        GTK_IM_MODULE=*|\
        QT_IM_MODULE=*|\
        XMODIFIERS=*|\
        XAUTHORITY=*)
            echo "$env_var"
            export "$env_var"
            ;;
    esac
done < "/proc/$(pgrep -u as2-streaming-user gnome-shell | head -n1)/environ"

/usr/local/bin/update-input-method.sh > /var/tmp/update-input-method.log 2>&1 &

/usr/local/firefox/firefox &
```

3. Add run permission to both scripts by running following the commands:

   sudo chmod +x /usr/local/bin/update-input-method.sh

   sudo chmod +x /usr/local/bin/firefox-jp.sh

4. If you already created the optimization manifest file for the application run the following commands to add the application launch script to the application catalog:

```
sudo AppStreamImageAssistant add-application \
--name firefox \
--absolute-app-path /usr/local/bin/firefox-jp.sh \
--display-name firefox \
--absolute-icon-path /usr/local/firefox/browser/chrome/icons/default/default128.png
 \
--absolute-manifest-path /tmp/firefox-manifest.txt
```

Alternatively, you can also configure the input methods by adding the script update-input-method.sh as a separate application to the application catalog for the image. During streaming sessions, your users can launch this application to enable Japanese input, and switch between input methods with specified shortcut keys within the same session.

Step 5: Set the Keyboard Layout

Set the keyboard layout to match the keyboards your users will use during streaming sessions. You can use the command **localectl list-keymaps** to list all the available keymaps, and use the command **sudo localectl set-keymap jp106** to set the keymap to the Japanese keyboard of 106 keys, for example.

Step 6: Verify on Image Builder

To verify on image builder, first reboot the image builder by running the command **sudo shutdown -r now**. After reboot, connect to the image builder again, and verify that everything, including time zone, locale, language, and input method, works as expected.

Step 7: Create the Image

Create the image on the image builder. For more information, see Tutorial: Create a Custom Linux-Based WorkSpaces Applications Image. Make sure to create default application settings, including the regional settings you just configured. For more information, see "Creating Default Application Settings for Your Users" in Create Your Linux-Based Images.

All of the Linux fleet instances created from this image will have the same default time zone, locale, language, and input method settings that you configured for the image.

Use Session Scripts to Manage Your Amazon WorkSpaces Applications Users' Streaming Experience

WorkSpaces Applications provides on-instance session scripts. You can use these scripts to run your own custom scripts when specific events occur in users' streaming sessions. For example, you can use custom scripts to prepare your WorkSpaces Applications environment before your users' streaming sessions begin. You can also use custom scripts to clean up streaming instances after users complete their streaming sessions.

Session scripts are specified within an WorkSpaces Applications image. These scripts are run within the user context or the system context. If your session scripts use the standard out to write information, error, or debugging messaging, these can be optionally saved to an Amazon S3 bucket within your Amazon Web Services account.

Contents

- Run Scripts Before Streaming Sessions Begin
- Run Scripts After Streaming Sessions End
- Create and Specify Session Scripts
- Session Scripts Configuration File
- Using Windows PowerShell Files
- Logging Session Script Output
- Use Storage Connectors with Session Scripts
- Enable Amazon S3 Bucket Storage for Session Script Logs
- Use Session Scripts on Multi-Session Fleets

Run Scripts Before Streaming Sessions Begin

You can configure your scripts to run for a maximum of 60 seconds before your users' applications launch and their streaming sessions begin. Doing so enables you to customize the WorkSpaces Applications environment before users start streaming their applications. When the session scripts run, a loading spinner displays for your users. When your scripts complete successfully or the maximum waiting time elapses, your users' streaming session will begin. If your scripts don't complete successfully, an error message displays for your users. However, your users are not prevented from using their streaming session.

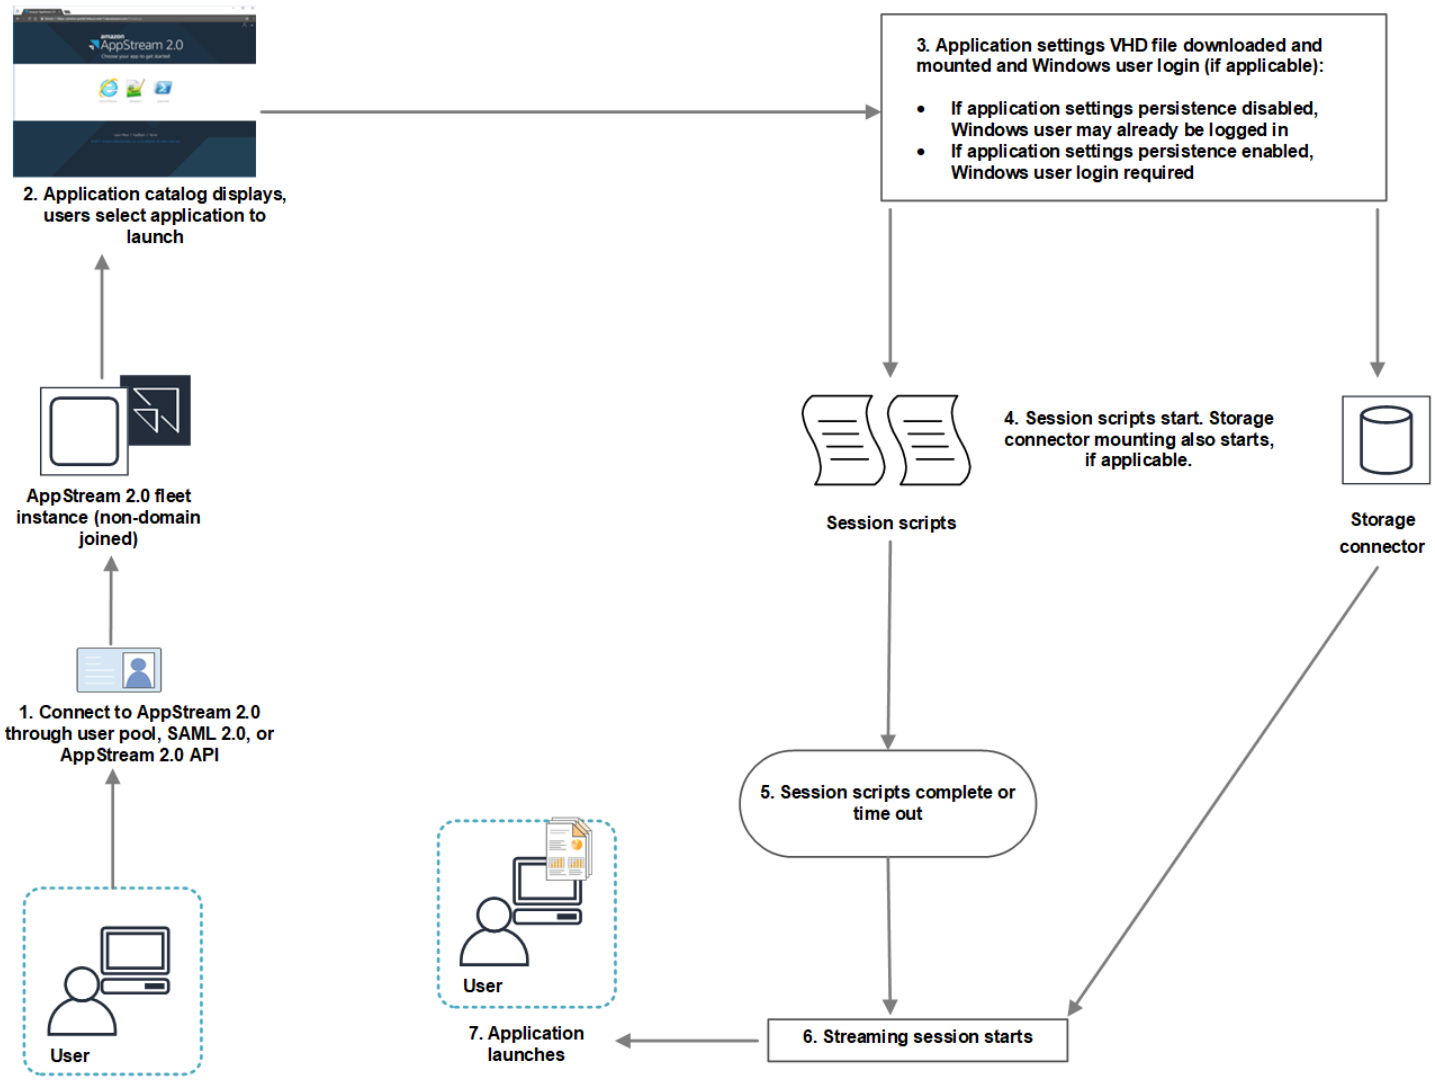
When you specify a file name on a Windows instance, you must use a double backslash. For example:

C:\\Scripts\\Myscript.bat

If you don't use a double backslash, an error displays to notify you that the .json file is incorrectly formatted.

> **Note**
>
> When your scripts complete successfully, they must return a value of 0. If your scripts return a value other than 0, WorkSpaces Applications displays the error message to the user.

When you run scripts before streaming sessions begin and the WorkSpaces Applications dynamic application framework is not enabled, the following process occurs:



1. Your users connect to an WorkSpaces Applications fleet instance that is not domain-joined. They connect by using one of the following access methods:
   - WorkSpaces Applications user pool
   - SAML 2.0
   - WorkSpaces Applications API
2. The application catalog displays in the WorkSpaces Applications portal, and your users choose an application to launch.
3. One of the following occurs:
   - If application settings persistence is enabled for your users, the application settings Virtual Hard Disk (VHD) file that stores your users' customizations and Windows settings is downloaded and mounted. Windows user login is required in this case.

For information about application settings persistence, see Enable Application Settings Persistence for Your WorkSpaces Applications Users.

- If application settings persistence is not enabled, the Windows user is already logged in.

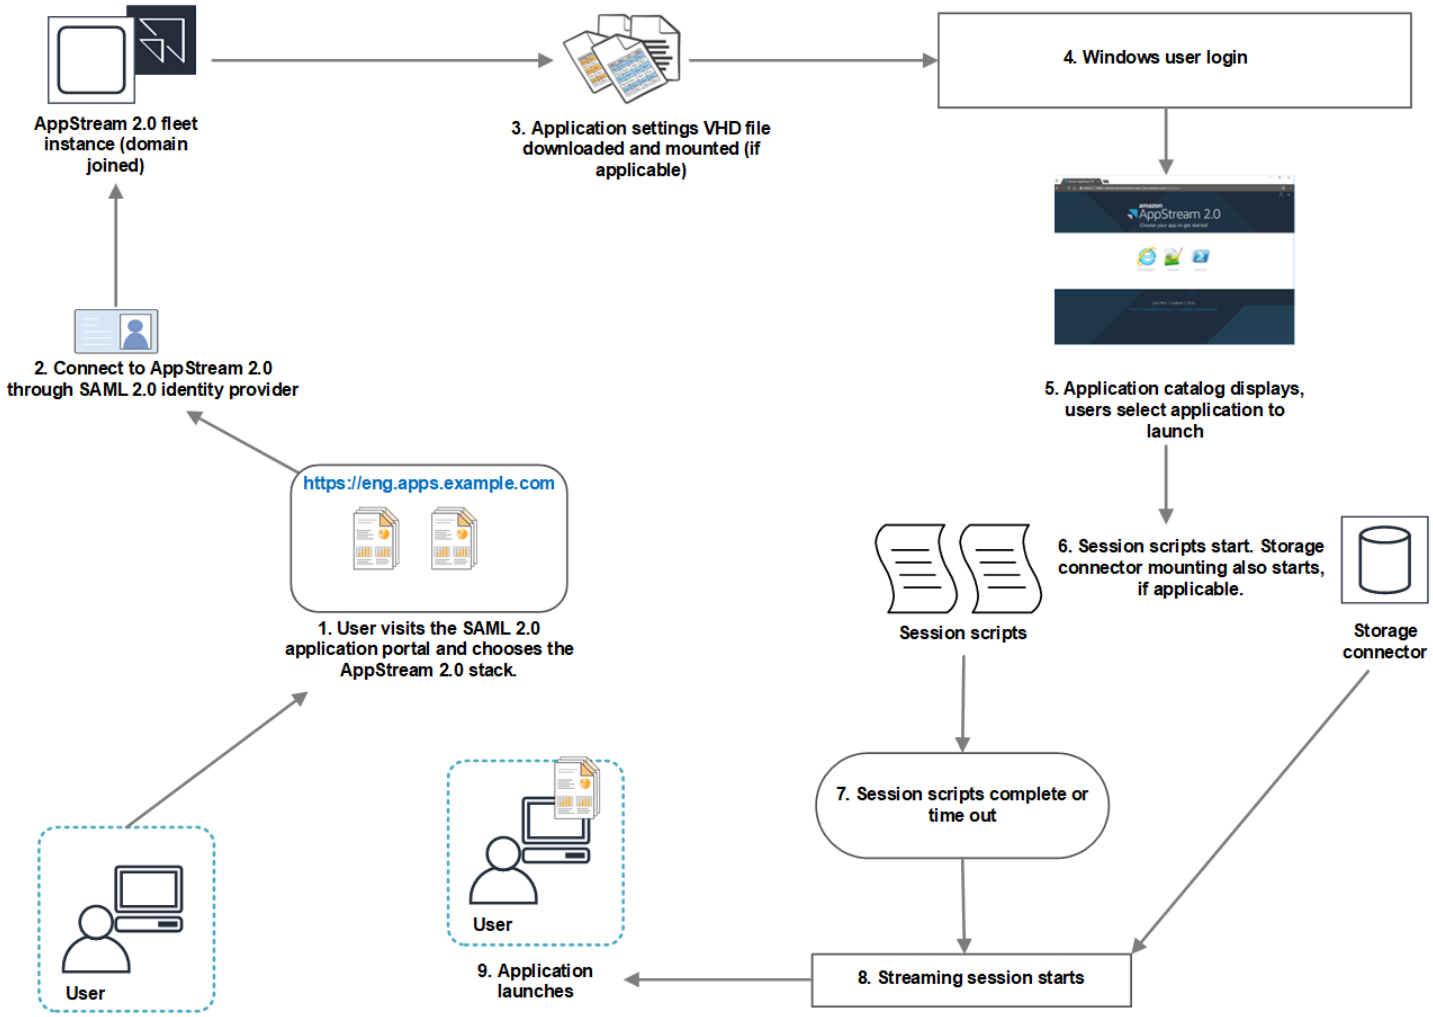
4. Your session scripts start. If persistent storage is enabled for your users, storage connector mounting also starts. For information about persistent storage, see Enable and Administer Persistent Storage for Your WorkSpaces Applications Users.

> **Note**
>
> The storage connector mount doesn't need to complete for the streaming session to start. If the session scripts complete before the storage connector mount completes, the streaming session starts.
> For information about monitoring the mount status of storage connectors, see Use Storage Connectors with Session Scripts.

5. Your session scripts complete or time out.
6. The users' streaming session starts.
7. The application that your users chose launches.

For information about the WorkSpaces Applications dynamic application framework, see Use the WorkSpaces Applications Dynamic Application Framework to Build a Dynamic App Provider.

When you run scripts before streaming sessions begin and the WorkSpaces Applications dynamic application framework is enabled, the following process occurs:



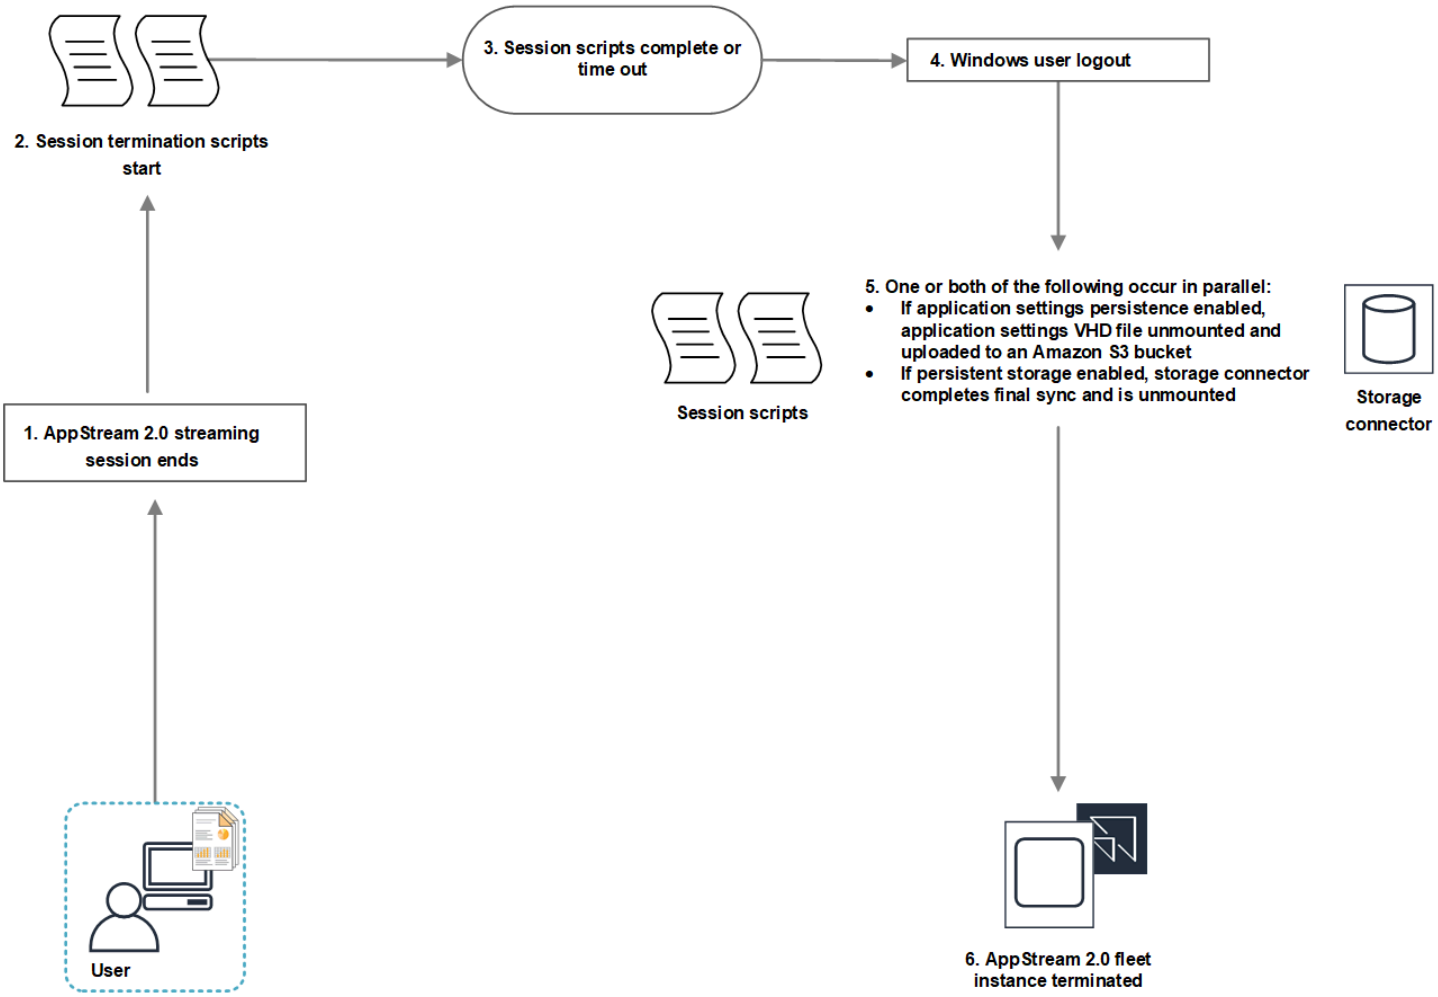
1. Your users visit the SAML 2.0 application portal for your organization, and they choose the WorkSpaces Applications stack.
2. They connect to an WorkSpaces Applications fleet instance that is domain-joined.
3. If application settings persistence is enabled for your users, the application settings VHD file that stores your users' customizations and Windows settings is downloaded and mounted.
4. Windows user logon occurs.
5. The application catalog displays in the WorkSpaces Applications portal and your users choose an application to launch.
6. Your session scripts start. If persistent storage is enabled for your users, storage connector mounting also starts.

> **Note**
>
> The storage connector mount doesn't need to complete for the streaming session to start. If the session scripts complete before the storage connector mount completes, the streaming session starts.
> For information about monitoring the mount status of storage connectors, see Use Storage Connectors with Session Scripts.

7. Your session scripts complete or time out.
8. The users' streaming session starts.
9. The application that your users chose launches.

Run Scripts After Streaming Sessions End

You can also configure your scripts to run after users' streaming sessions end. For example, you can run a script when users select **End Session** from the WorkSpaces Applications toolbar, or when they reach the maximum allowed duration for the session. You can also use these session scripts to clean up your WorkSpaces Applications environment before a streaming instance is terminated. For example, you can use scripts to release file locks or upload log files. When you run scripts after streaming sessions end, the following process occurs:



1. Your users' WorkSpaces Applications streaming session ends.
2. Your session termination scripts start.
3. The session termination scripts complete or time out.
4. Windows user logout occurs.
5. One or both of the following occur in parallel, if applicable:
   - If application settings persistence is enabled for your users, the application settings VHD file that stores your users' customizations and Windows settings is unmounted and uploaded to an Amazon S3 bucket in your account.
   - If persistent storage is enabled for your users, the storage connector completes a final synchronization and is unmounted.
6. The fleet instance is terminated.

Create and Specify Session Scripts

You can configure and specify session scripts for Always-on, On-demand, and Elastic fleets.

To configure and specify session scripts for Always-on and On-demand fleets

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the navigation pane, choose **Images**, **Image Builder**.
3. Choose an image builder that is in the **Running** state, and choose **Connect**.
4. When prompted, choose **Administrator**.
5. Navigate to `C:\AppStream\SessionScripts`, and open the `config.json` configuration file.

   For information about session script parameters, see Session Scripts Configuration File.
6. After you finish making your changes, save and close the `config.json` file.
7. On the image builder desktop, open **Image Assistant**.
8. (Optional) Specify any additional applications that you want to include in the image.
9. Follow the necessary steps in Image Assistant to finish creating your image.

   If the session scripts configuration can't be validated (for example, if the .json file is not correctly formatted), you are notified when you choose **Disconnect and create image**.

> **Note**
>
> To locate the session scripts configuration file for Linux-based image builders, navigate to `/opt/appstream/SessionScripts/config.json`.

To configure and specify session scripts for Elastic fleets

1. Create a zip file that contains the session scripts and config.json file. The scripts files will be copied to the following locations. You must use these locations for your config.json.

   - For Windows, use `C:\AppStream\SessionScripts\`*`SessionScript`*.
   - For Linux, use `/opt/appstream/SessionScripts/`*`SessionScript`*.

> **Note**
>
> In order to run the session script files, make sure that the .zip file only contains the session scripts and `config.json` files, and not the containing folder. For more information, see Session Scripts Configuration File.

2. Upload the zip file to an Amazon S3 bucket in your account.

> **Note**
>
> Your VPC must provide access to the Amazon S3 bucket. For more information, see Using Amazon S3 VPC Endpoints for WorkSpaces Applications Features.
> You must have your S3 bucket and WorkSpaces Applications fleet in the same AWS Region.
> You must have IAM permissions to perform the `S3:GetObject` action on the session scripts object in the Amazon S3 bucket. To learn more about storing the session scripts in an Amazon S3 bucket, see Store Application Icon, Setup Script, Session Script, and VHD in an S3 Bucket.

3. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
4. In the navigation pane, choose **Fleets**.
5. Choose an Elastic fleet that you want to update, and then choose **View Details**.
6. On the **Session scripts settings** tab, choose **Edit**.
7. For **Session scripts object in S3**, either enter the S3 URI that represents the session scripts object, or choose **Browse S3** to navigate to your S3 buckets and find the session scripts object.
8. After you finish making your changes, choose **Save Changes**.
9. At this point, session scripts are available for all fleet instances launched.

> **Note**
>
> You can also configure the session scripts when you create a new Elastic fleet.

Session Scripts Configuration File

To locate the session scripts configuration file in a Windows instance, navigate to C:\AppStream \SessionScripts\config.json. On a Linux instance, navigate to /opt/appstream/SessionScripts/ config.json. The file is formatted as follows.

> **Note**
>
> The configuration file is in .json format. Verify that any text you type in this file is in valid .json format.

```
{
  "SessionStart": {
    "executables": [
      {
        "context": "system",
        "filename": "",
        "arguments": "",
        "s3LogEnabled": true
      },
      {
        "context": "user",
        "filename": "",
        "arguments": "",
        "s3LogEnabled": true
      }
    ],
    "waitingTime": 30
  },
  "SessionTermination": {
    "executables": [
      {
        "context": "system",
        "filename": "",
        "arguments": "",
        "s3LogEnabled": true
      },
      {
        "context": "user",
        "filename": "",
        "arguments": "",
```

```
        "s3LogEnabled": true
      }
    ],
    "waitingTime": 30
  }
}
```

You can use the following parameters in the session scripts configuration file.

SessionStart/SessionTermination

The session scripts to run in the appropriate session event based on the name of the object.

Type: String

Required: No

Allowed values: `SessionStart`, `SessionTermination`

WaitingTime

The maximum duration of the session scripts in seconds.

Type: Integer

Required: No

Constraints: The maximum duration is 60 seconds. If the session scripts don't complete within this duration, they will be stopped. If you require a script to continue running, launch it as a separate process.

Executables

The details for the session scripts to run.

Type: String

Required: Yes

Constraints: The maximum number of scripts that can run per session event is 2 (one for the user context, one for the system context).

Context

The context in which to run the session script.

Type: String

Required: Yes

Allowed values: `user`, `system`

`Filename`

The full path to the session script to run. If this parameter is not specified, the session script is not run.

Type: String

Required: No

Constraints: The maximum length for the file name and full path is 1,000 characters.

Allowed values: `.bat`, `.exe`, `.sh`

> **Note**
>
> You can also use Windows PowerShell files. For more information, see Using Windows PowerShell Files.

`Arguments`

The arguments for your session script or executable file.

Type: String

Required: No

Length constraints: The maximum length is 1,000 characters.

`S3LogEnabled`

When the value for this parameter is set to **`True`**, an S3 bucket is created within your Amazon Web Services account to store the logs created by the session script. By default, this value is set to **`True`**. For more information, see the *Logging Session Script Output* section later in this topic.

Type: Boolean

Required: No

Allowed values: `True`, `False`

Using Windows PowerShell Files

To use Windows PowerShell files, specify the full path to the PowerShell file in the **`filename`** parameter:

```
"filename":
"C:\\Windows\\System32\\WindowsPowerShell\\v1.0\\powershell.exe",
```

Then specify your session script in the **`arguments`** parameter:

```
"arguments": "-File \"C:\\path\\to\\session\\script.ps1\"",
```

Finally, verify that the PowerShell Execution Policy allows your PowerShell file to run.

Logging Session Script Output

When this option is enabled in the configuration file, WorkSpaces Applications automatically captures the output from the session script that is written to the standard out. This output is uploaded to an Amazon S3 bucket in your account. You can review the log files for troubleshooting or debugging purposes.

> **Note**
>
> The log files are uploaded when the session script returns a value, or the value set in **`WaitingTime`** has elapsed, whichever comes first.

Use Storage Connectors with Session Scripts

When WorkSpaces Applications storage connectors are enabled, they begin mounting when the session start scripts run. If your script relies on the storage connectors being mounted, you can wait for the connectors to be available. WorkSpaces Applications maintains the mount status of the storage connectors in the Windows registry on Windows instances, at the following key:

HKEY_LOCAL_MACHINE\SOFTWARE\Amazon\AppStream\Storage\<provided user name>\<Storage connector>

The registry key values are as follows:

- Provided user name — The user ID provided through the access mode. The access modes and value for each mode are as follows:
  - User Pool — The email address for the user
  - Streaming URL — The UserID
  - SAML — The NameID. If the user name includes a slash (for example, a domain user's SAMAccountName), the slash is replaced by a "-" character.
- Storage connector — The connector for the persistent storage option that is enabled for the user. The storage connector values are as follows:
  - HomeFolder
  - GoogleDrive
  - OneDrive

Each storage connector registry key contains a **MountStatus** DWORD value. The following table lists the possible values for **MountStatus**.

> **Note**
>
> To view these registry keys, you must have Microsoft .NET Framework version 4.7.2 or later installed on your image.

Value	Description
0	Storage connector not be enabled for this user
1	Storage connector mounting is in progress
2	Storage connector mounted successfully
3	Storage connector mounting failed
4	Storage connector mounting is enabled, but not mounted yet

On Linux instances, you can check the home folder mount status by looking at the value of appstream_home_folder_mount_status in the file ~/.config/appstream-home-folder/appstream-home-folder-mount-status.

Value	Description
True	Home folder is mounted successfully
False	Home folder is not mounted yet

Enable Amazon S3 Bucket Storage for Session Script Logs

When you enable Amazon S3 logging in your session script configuration, WorkSpaces Applications captures standard output from your session script. The output is periodically uploaded to an S3 bucket within your Amazon Web Services account. For every AWS Region, WorkSpaces Applications creates a bucket in your account that is unique to your account and the Region.

You do not need to perform any configuration tasks to manage these S3 buckets. They are fully managed by the WorkSpaces Applications service. The log files that are stored in each bucket are encrypted in transit using Amazon S3's SSL endpoints and at rest using Amazon S3-managed encryption keys. The buckets are named in a specific format as follows:

```
appstream-logs-region-code-account-id-without-hyphens-random-identifier
```

region-code

This is the AWS Region code in which the stack is created with Amazon S3 bucket storage enabled for session script logs.

account-id-without-hyphens

Your Amazon Web Services account identifier. The random ID ensures that there is no conflict with other buckets in that Region. The first part of the bucket name, `appstream-logs`, does not change across accounts or Regions.

For example, if you specify session scripts in an image in the US West (Oregon) Region (us-west-2) on account number 123456789012, WorkSpaces Applications creates an Amazon S3 bucket

within your account in that Region with the name shown. Only an administrator with sufficient permissions can delete this bucket.

```
appstream-logs-us-west-2-1234567890123-abcdefg
```

Disabling session scripts does not delete any log files stored in the S3 bucket. To permanently delete log files, you or another administrator with adequate permissions must do so by using the Amazon S3 console or API. WorkSpaces Applications adds a bucket policy that prevents accidental deletion of the bucket. For more information, see *IAM Policies and the Amazon S3 Bucket for Application Settings Persistence* in Identity and Access Management for Amazon WorkSpaces Applications.

When session scripts are enabled, a unique folder is created for each streaming session that is started.

The path for the folder where the log files are stored in the S3 bucket in your account uses the following structure:

```
bucket-name/stack-name/fleet-name/access-mode/user-id-SHA-256-hash/session-id/
SessionScriptsLogs/session-event
```

bucket-name

The name of the S3 bucket in which the session scripts are stored. The name format is described earlier in this section.

stack-name

The name of the stack the session came from.

fleet-name

The name of the fleet the session script is running on.

access-mode

The identity method of the user: `custom` for the WorkSpaces Applications API or CLI, `federated` for SAML, and `userpool` for users in the user pool.

user-id-SHA-256-hash

The user-specific folder name. This name is created using a lowercase SHA-256 hash hexadecimal string generated from the user identifier.

session-id

The identifier of the user's streaming session. Each user streaming session generates a unique ID.

session-event

The event that generated the session script log. The event values are: `SessionStart` and `SessionTermination`.

The following example folder structure applies to a streaming session started from the test-stack and test-fleet. The session uses the API of user ID `testuser@mydomain.com`, from an AWS account ID of `123456789012`, and the settings group `test-stack` in the US West (Oregon) Region (us-west-2):

```
appstream-logs-us-west-2-1234567890123-abcdefg/test-stack/test-fleet/custom/
a0bcb1da11f480d9b5b3e90f91243143eac04cfccfbdc777e740fab628a1cd13/05yd1391-4805-3da6-
f498-76f5x6746016/SessionScriptsLogs/SessionStart/
```

This example folder structure contains one log file for a user context session start script, and one log file for a system context session start script, if applicable.

Use Session Scripts on Multi-Session Fleets

When using session scripts on multi-session fleets, there are additional requirements and considerations to ensure optimal performance and security.

Requirements

On a single-session fleet, for a given instance, the **SessionStart** and **SessionTermination** hooks are guaranteed to run only one time. This is because there is a 1:1 mapping of sessions to instances. When using multi-session fleets, there is an N:M mapping of sessions to instances, where each session runs its own **SessionStart** and **SessionTermination** hook. This means that the **SessionStart** and **SessionTermination** hooks can be run many times on a given instance, and in many different orderings. For the best experience, the following should be true of your session scripts when used on multi-session fleets:

- Scripts are idempotent.

When an action has already been performed, scripts should handle more than one execution on the same instance with graceful handling.

- Scripts are independent.

  Because scripts run per session, if one session is running **SessionTermination** while another is running **SessionStart**, they should not interfere with each other, or with the experience of other sessions.

- Scripts are performant.

  On multi-session instances, multiple sessions can be provisioned concurrently. This means that there can be multiple concurrent executions of the session scripts. Scripts should be efficient, not consume excessive resources, and not impact the experience of other users on the instance or the stability of the sessions.

Many of these requirements can be met by keeping session script logic focused on the specific user session for which the script is running.

Security Considerations

WorkSpaces Applications images should not be configured to allow write permission to session script files by any users. This introduces a critical attack vector for malicious users, where they could modify script files. These files could then be run as SYSTEM or another user, depending on your configuration.

> **Important**
>
> It is your responsibility to make sure that your WorkSpaces Applications images are configured securely. This is especially important for multi-session instances, where multiple users are using the same instance. If images are not configured securely, there is a security risk for all users of that instance.

The following should be true of your images and session scripts files:

- Users do not have permission to modify session script files.
- Users do not have permission to modify the session script config.json. Default behavior on the image restricts access to administrators.

Session scripts executables should be stored in a secure location where they are safe from modification at runtime.

If the service detects that a session script executable has been modified, it will fail any subsequent executions of that hook on that instance, upload log files to Amazon S3 (if Amazon S3 logging is enabled), and you will see the following message:

The session script was not executed because the executable was modified after instance provisioning. Execution was skipped for security.

If your use case requires modifying the session script executable at run time (for example, if you point to an EXE file which is modified by an automatic update process at runtime), this will fail the above checks. In this case, use a script to redirect execution to your modified executable. Leave the script unmodified at runtime when the service performs security checks.

If your session script files are excessively large (more than 100 MB), this can cause delays in instance and session provisioning, and the security checks will take additional time (depending on instance type and available resources). If your use case requires large session scripts, consider using smaller scripts to redirect execution. This will improve instance and session provisioning experiences.

Note that the service is only checking the executable defined in the session scripts config.json, and this is only a fallback/best effort mechanism. It is your responsibility to ensure that all code paths in session scripts executables are secure and cannot be modified by end users.

Applications Manager

When using an Elastic fleet, you can create app blocks and applications. *App blocks* represent a virtual hard disk (VHD) that is stored within an Amazon S3 bucket within your account that contains the application files and binaries necessary to launch the applications that your users will use. *Applications* contain the details necessary to launch your application after the VHD has been mounted. The following sections describe how to create and manage these resources.

Contents

- App Blocks
- App Block Builder
- Applications
- Store Application Icon, Setup Script, Session Script, and VHD in an S3 Bucket
- Associate Applications to Elastic Fleets
- Additional Resources

App Blocks

App blocks represent a virtual hard disk (VHD) that is stored within an Amazon S3 bucket within your account that contains the application files and binaries necessary to launch the applications your users will use. App blocks also include the setup script that informs the operating system how to handle the VHD file.

App blocks support two different types of packaging:

- Custom - Choose this option to create your application package (VHD) manually. For more information, see the section called “Custom App Blocks”.
- WorkSpaces Applications - Choose this recommended option to create your application package using app block builder. For more information, see the section called “WorkSpaces Applications App Blocks”.

Contents

- Custom App Blocks
- WorkSpaces Applications App Blocks

- Unsupported Applications

Custom App Blocks

Elastic fleet streaming instances utilize applications that are installed on virtual hard disk (VHD) files stored within an Amazon S3 bucket in your account. App blocks with custom packaging gives you the flexibility to create your own VHD file, and upload it to an Amazon S3 bucket within your account.

Contents

- Create the VHD
- Create the Setup Script for the VHD in Amazon WorkSpaces Applications
- Create a Custom App Block
- Update the App Block, VHD, and Setup Script

Create the VHD

A VHD is a single file that when mounted to the operating system is treated like a hard disk. The VHD can be mounted as a drive letter, to a folder path, or both. When the VHD is mounted, you can treat it as you would any other hard disk, including installing your application or copying files to it that your user will need.

To create the app block, you will need to create the VHD, install your applications to it, then detach it. Once detached you can test your VHD on another PC, an EC2 instance, or an WorkSpaces Applications image builder to validate the applications work as expected. Once completed, upload to an Amazon S3 bucket in your account and create the app block.

> **Note**
>
> This page describes using a VHD to deliver your application; however, the WorkSpaces Applications streaming instance will download any object from Amazon S3. The object you store in Amazon S3 can also be a zip file, application installer, or the application executable itself. You can use the setup script to configure it correctly on the streaming instance before a user launches their application.
> The WorkSpaces Applications streaming instance waits up to 120 seconds for the VHD to complete downloading before the setup script runs. If the VHD does not complete downloading within this duration, the download stops, and the setup script will not run.

We recommend a maximum size of 1.5 gigabyte for the VHD. You might be able to reduce the size of the VHD by compressing. You must use the setup script to decompress it before mounting it, because the file needs to be fully downloaded from Amazon S3 before it can be mounted and the application is launched. Larger VHDs increase the time it takes for the application to launch and the streaming session to begin.

To create a VHD for Microsoft Windows

1. From a Windows PC or Windows Amazon Elastic Compute Cloud (Amazon EC2) instance, open a command prompt with administrative privileges.
2. Launch the Microsoft **diskpart** utility by entering the following command:

   diskpart
3. Create the unformatted and uninitialized VHD file by entering the following command, where *<maximum file size>* is the size of the VHD file, in MB:

   create vdisk file=C:\path\to\new\file.vhdx maximum=*<maximum file size>* type=expandable
4. Select the newly created VHD by entering the following command:

   select vdisk file=C:\path\to\new\file.vhdx
5. Attach the newly created VHD by entering the following command:

   attach vdisk
6. Initialize the newly created VHD by entering the following command:

   convert mbr
7. Create the primary partition spanning the entire VHD by entering the following command:

   create partition primary
8. Format the newly created partition by entering the following command:

   format fs=ntfs quick
9. You can mount your newly created VHD to an unused drive letter, a folder path on the root volume, or both.

   To mount a drive letter, enter: **assign letter=*<unused drive letter>***

To mount a folder, enter: **assign mount=*C:\path\to\empty\folder\to\mount***

> **Note**
>
> To mount to a folder path, the folder must already exist and must be empty.

10. You can now install your application to the VHD, using either the drive letter or the folder mount path chosen in step 9.

After you finish installing your application(s) to the VHD, you need to detach it before you can safely upload it to an Amazon S3 bucket.

To detach a VHD for Microsoft Windows

1. Launch the Microsoft diskpart utility by entering the following command:

   diskpart

2. Select the VHD by entering the following command:

   select vdisk file=*C:\path\to\new\file.vhdx*

3. Detach the VHD by entering the following command:

   detach vdisk

4. The VHD has now been detached, and can be tested on another Windows PC, Amazon EC2 instance, or an WorkSpaces Applications image builder.

To create a VHD for Linux

1. Open a terminal:

   - For Ubuntu Pro 24.04 LTS: Use EC2 instance, or WorkSpaces powered by Ubuntu Pro

2. Create the unformatted and uninitialized VHD file:

   dd if=/dev/zero of=*<name of file>* bs=*<size of VHD>* count=1

3. Add a file system to the created VHD by entering the following command:

   sudo mkfs -t ext4 *<name of file>*

> **Note**
>
> You might see a message stating that the file is not a block special device. You can select to proceed anyway.

4. Create an empty folder to use for the mount point by entering the following command:

   sudo mkdir *`/path/to/mount/point`*

5. Mount the newly created VHD to a file system path by running the following command:

   sudo mount -t auto -o loop *`<name of file>`* *`/path/to/mount/point`*

6. You can now install your application to the VHD using the folder mount path chosen in step 4.

> **Note**
>
> The default permissions for files and folders created on the VHD can prevent non-administrator users from launching applications or reading files. Validate the permissions and change them, if necessary.

After you finish installing your application(s) to the VHD, you need to detach it before you can safely upload it to an Amazon S3 bucket.

To detach a VHD for Linux

1. Open a terminal session, and enter the following command:

   sudo umount *`/path/to/mount/point`*

2. The VHD has now been detached, and can be tested on another Ubuntu Pro 24.04 LTS Amazon EC2 instance or Ubuntu WorkSpaces Personal.

Create the Setup Script for the VHD in Amazon WorkSpaces Applications

WorkSpaces Applications uses a setup script that you provide to mount the VHD before the application launches. You can also use the setup script to complete other tasks required to make your application work. For example, you can configure registry keys, register DLLs, manage pre-requisites, or modify the user profile from the setup script. WorkSpaces Applications provides script

examples that you can use to mount your VHD. You will need to modify these scripts for your VHD and application needs.

> **Note**
>
> Setup scripts aren't required for app blocks with WorkSpaces Applications packaging. However, you can provide optional post-setup scripts to customize application installation.

Use the following links to download the example scripts:

- Ubuntu Pro 24.04 LTS bash script
- Microsoft Windows Powershell script

> **Note**
>
> WorkSpaces Applications and the Microsoft Windows operating system reserve drive letters A through E. Don't mount VHDs or network shares to these drive letters.

WorkSpaces Applications downloads the setup script and VHD to a directory on the fleet streaming instance, then runs the setup script. The setup script runs on the operating system with full administrator rights. The setup script runs in the `SYSTEM` context on Microsoft Windows, and as the `root` user on Ubuntu Pro 24.04 LTS.

File system location for the VHD and setup script:

- Ubuntu Pro 24.04 LTS:

  `/opt/appstream/AppBlocks/`*`appblock-name`*`/`

  `appblock-name`

  The name of the app block that the VHD and setup script correspond to.

- Microsoft Windows:

  `C:\AppStream\AppBlocks\`*`appblock-name`*`\`

  `appblock-name`

  The name of the app block that the VHD and setup script correspond to.

WorkSpaces Applications maintains the file name as they are on the object. For example, if your app block is named `MyApps`, with a VHD named `apps.vhd` and setup script named `mount-apps.ps1`, then the full path on a Windows streaming instance is:

- VHD

  `C:\AppStream\AppBlocks\MyApps\apps.vhd`

- Setup script

  `C:\AppStream\AppBlocks\MyApps\mount-apps.ps1`

WorkSpaces Applications captures the standard error and standard output from your setup script when it runs on a fleet streaming instance and uploads the output to an Amazon S3 bucket within your account. You can use these logs to identify and resolve issues you may have with your setup script. The buckets are named in a specific format as follows:

```
appstream-logs-region-code-account-id-without-hyphens-random-identifier
```

`region-code`

This is the AWS Region code in which the elastic fleet is created within.

`account-id-without-hyphens`

Your AWS account identifier. The random ID ensures that there is no conflict with other buckets in that Region. The first part of the bucket name, appstream-logs, does not change across accounts or Regions.

For example, if you create an elastic fleet in the US West (Oregon) Region (us-west-2) on account number 123456789012, WorkSpaces Applications creates an Amazon S3 bucket within your account in that Region with the name shown. Only an administrator with sufficient permissions can delete this bucket.

```
appstream-logs-us-west-2-1234567890123-abcdefg
```

The path for the folder where the log files are stored in the S3 bucket in your account uses the following structure:

```
bucket-name/fleet-name/instance-id/appblock-name/
```

`bucket-name`

The name of the Amazon S3 bucket in which the setup script logs are stored. The name format is described earlier in this section.

`Instance-id`

The unique identifier for the streaming instance that the setup script ran on

`appblock-name`

The name of the appblock that the setup script corresponds to.

The following example folder structure applies to a streaming session started from `test-fleet`. The session is from an AWS account ID of 123456789012, and appblock name is testappblock in the US West (Oregon) Region (us-west-2):

```
appstream-logs-us-west-2-1234567890123-abcdefg/test-fleet/
i-084427ab4a1cff7f5/testappblock/
```

This example folder structure contains one log file for the standard output, and one log file for the standard error.

Topics

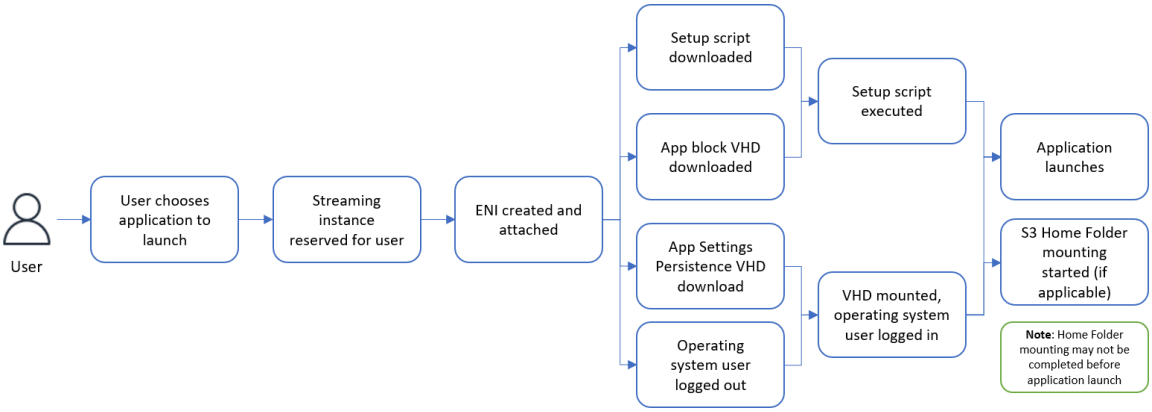
- App block setup script execution in Amazon WorkSpaces Applications

App block setup script execution in Amazon WorkSpaces Applications

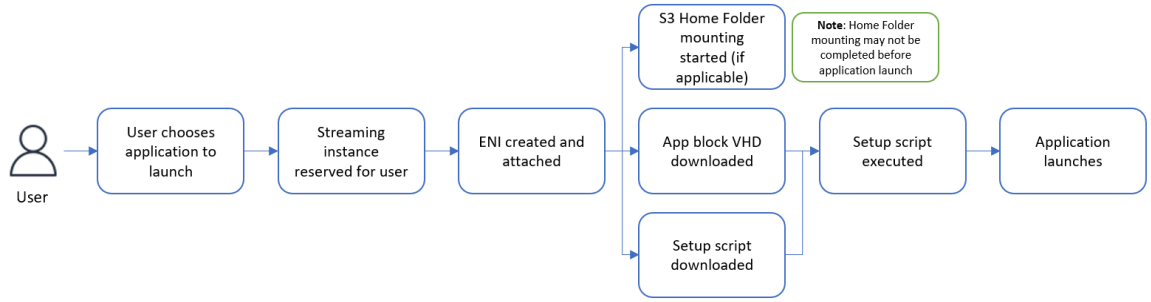
The following diagrams indicate where in the process the setup script runs. The run order is dependent upon whether Application Settings Persistence is enabled on the stack associated with the elastic fleet.

> **Note**
>
> WorkSpaces Applications uses your VPC details to download the VHD and setup script from the Amazon S3 bucket. Your VPC must provide access to the Amazon S3 bucket. For more information, see Using Amazon S3 VPC Endpoints for WorkSpaces Applications Features.

Application Settings Persistence is enabled:



Application Settings Persistence is disabled:



Create a Custom App Block

You can use the WorkSpaces Applications console to create the app block resource once you have your VHD and setup script created and uploaded to an S3 bucket in your AWS account. To learn more about storing the VHD and setup script in an Amazon S3 bucket, see the section called “Store Application Icon, Setup Script, Session Script, and VHD in an S3 Bucket”.

> **Note**
>
> You must have IAM permissions to perform the `S3:GetObject` action on the VHD and setup script objects in the Amazon S3 bucket to create the app block resource.

To create the app block resource

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. From the left-hand navigation menu, choose **Applications**, **App block**, and **Create app block**.
3. For app block packaging, select **Custom**.

4. For **App block details**, type a unique name identifier for the app block. Optionally, you can also specify the following:

   - **Display name** – A friendly name for the app block.
   - **Description** – A description for the app block.

5. For **Virtual hard disk object in S3** under **Script settings**, either enter the S3 URI that represents the VHD object, or choose **Browse S3** to navigate to your S3 buckets and find the VHD object.

6. For **Setup script object in S3** under **Script settings**, either enter the S3 URI that represents the setup script object, or choose **Browse S3** to navigate to your S3 buckets and find the setup script object.

7. For **Setup script executable** under **Script settings**, enter the executable necessary for your setup script.

   > **Note**
   >
   > If your setup script can execute directly, enter the filename of the setup script. If your setup script relies on another executable (for example, Microsoft PowerShell) to execute, enter the path to that executable.
   > Path to Microsoft PowerShell on Microsoft Windows:
   > `C:\Windows\System32\WindowsPowerShell\v1.0\powershell.exe`

8. Optionally, for **Setup script executable arguments** under **Script settings**, enter in the arguments that need to be provided to the setup script executable to execute your setup script.

   > **Note**
   >
   > If you are using a Microsoft PowerShell script, you must specify the "-File" parameter with the name of your setup script as an executable argument. Additionally, ensure that the Execution Policy allows your script to be run. To learn more, see about_Execution_Policies and What is PowerShell?.

9. For **Execution duration in seconds** under **Script settings**, enter the timeout duration for your setup script.

> **Note**
>
> The execution duration in seconds is how long WorkSpaces Applications waits for the setup script to run before continuing. If your setup script doesn't complete within this duration, an error is displayed to your user and the application will attempt to launch. The setup script is terminated after the execution duration has elapsed.

10. (Optional) For **Tags**, create tags for the app block resource
11. Review the information that you entered, and choose **Create**.
12. If your app block was created successfully, you see a success message at the top of the console. If an error occurred, you see a descriptive error message and will need to try creating the app block again.

Update the App Block, VHD, and Setup Script

App block resources are immutable and do not allow you to change them once created. If you need to make backwards compatible updates to the VHD or setup script, it is recommended that you upload a new version of the file to the Amazon S3 bucket, overwriting the current version. New Elastic fleet streaming sessions will download the latest version of the objects, and use them.

If you need to make backwards incompatible updates to the VHD or setup script, it is recommended that you upload them as new objects to the Amazon S3 bucket, and create a new app block and application resource. You can then manage the deployment to your users as part of a change window or other outage.

WorkSpaces Applications App Blocks

Elastic fleet streaming instances utilize applications that are installed on virtual hard disk (VHD) files stored within an Amazon S3 bucket in your account. When it comes to app blocks with custom packaging, you have the flexibility to create your own VHD file and upload it to an Amazon S3 bucket within your account. Alternatively, for app blocks with WorkSpaces Applications packaging, you can take advantage of the app block builder, which handles the packaging of your applications, creates a VHD file, and uploads it to your Amazon S3 bucket.

By using the WorkSpaces Applications packaged app block, you not only eliminate the need for manual steps in building a VHD file, but also remove the requirement for a setup script. It

expands application compatibility with elastic fleets, as well as reduces manual administrative steps required to create an app block. WorkSpaces Applications handles the setup of app blocks with WorkSpaces Applications packaging automatically without the need of any setup scripts. However, you can still provide optional post-setup scripts to customize the installation for your needs.

Contents

- Overview
- Unsupported Applications
- Create an WorkSpaces Applications App Block
- Activate an App Block
- Create an App Block with an Existing App Package
- Test an App Block
- Associate an App Block in Amazon WorkSpaces Applications
- Disassociate an App Block in Amazon WorkSpaces Applications

Overview

To create an app block with WorkSpaces Applications packaging, you need to initiate a streaming session with an app block builder. After the session is launched, you can download your application installers and enable the recording options. From that point onwards, WorkSpaces Applications records the file system and registry changes made on the app block builder using Application Redirection technology.

Application Redirection uses Windows filter driver framework to intercept and redirect file-system and registry changes. This redirection is seamless to the application being installed. The application will continue to interact with the original file locations on the C: drive. For example, if an installer for "TestApplication" is run on a machine with App Redirection set up, it will be installed by default to C:\Program Files\TestApplication. However, behind the scenes, all files and folders will be redirected to a mounted virtual hard disk (VHD), and a link will be created from the original file location to the actual file location. On the machine, TestApplication will still appear to be installed at C:\Program Files\TestApplication.

After all the installation changes are recorded, the VHD file is uploaded to an Amazon S3 bucket in your account.

When a user requests a session using an Elastic fleet, WorkSpaces Applications downloads the VHD file, sets up the application, runs the post-installation setup scripts (optional), and starts the application streaming.

> **Note**
>
> Application Redirection technology does not record any file system changes under %USERPROFILE%, except new directories created under %APPDATA% and %LOCALAPPDATA% directories.
> Application Redirection technology does not record any registry changes under the current user, HKEY_CURRENT_USER (HKCU).

Unsupported Applications

Applications might encounter failures when installing or running in the following scenarios:

- **Applications with location-based checks during installation**: If an application's installation process verifies the actual location of the installed files, it might result in a failure. Because WorkSpaces Applications redirects the files to the app block VHD, only links to the actual files are maintained at the original location.

If you are uncertain whether your application falls into any of these categories, you can use WorkSpaces Applications packaging to create an app block. This process involves installing your application(s) on an app block builder instance. In the event that your application(s) fail to install on the app block builder instance, you can take the following actions:

- Check the logs. The error log file for your app block builder instance can be found at C:\AppStream\AppBlocks\errorLog. This log records all installation failures, including RegKeys/File operation processing. If you see any of the following logs in the errorLog, it indicates that the packaging of your application is currently unsupported by the WorkSpaces Applications app block builder:
  - "Unable to create symbolic link"
  - "Service doesn't support file renaming"

  If there is no errorLog file, or if this file is empty, then check your application installation logs to identify the reason for failures.

- Report a problem. Select the **Report a problem** button, which is available on the application builder assistant in the app block builder. Selecting this option will gather all the WorkSpaces Applications logs from your app block builder instance, and submit them to the WorkSpaces Applications team for assistance.
- Create an app block with custom packaging: If you are unable to package your applications using the app block builder, you can try to create an app block using custom packaging methods. For more information, see the section called “Custom App Blocks”.
- If you need more help, contact AWS Support. For more information, see AWS Support Center.

It is important to consider these potential limitations, and plan accordingly when using WorkSpaces Applications packaging for your applications.

Create an WorkSpaces Applications App Block

Follow these steps to create an app block with the WorkSpaces Applications packaging type.

Step 1: Configure the app block

To configure the app block

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. From the left-hand navigation menu, choose **Applications Manager**, **App blocks**, and **Create app block**.
3. For app block packaging, select **WorkSpaces Applications**.
4. For **App block details**, type a unique name identifier for the app block. Optionally, you can also specify the following:
   - **Name** – A unique name for the app block.
   - **Display name** (optional) – A friendly name for the app block.
   - **Description** (optional) – A description for the app block.
5. (Optional) An app block with WorkSpaces Applications packaging doesn't need a setup script. You can optionally provide post-installation steps the following **Advanced Options**:
   - For **Post setup script object in S3**, either enter the Amazon S3 URI that represents the post setup script object, or choose **Browse S3** to navigate to your Amazon S3 buckets and find the setup script object.

- For **Post setup script executable**, enter the executable needed for your post setup script.

> **Note**
>
> If your setup script can execute directly, enter the filename of the setup script. If your setup script relies on another executable (for example, Microsoft PowerShell) to execute, enter the path to that executable.
> Path to Microsoft PowerShell on Microsoft Windows:
> `C:\Windows\System32\WindowsPowerShell\v1.0\powershell.exe`
> Optionally, for **Setup script executable arguments** under **Script settings**, enter in the arguments that need to be provided to the setup script executable to execute your setup script.

> **Note**
>
> If you are using a Microsoft PowerShell script, you must specify the "-File" parameter with the name of your setup script as an executable argument. Additionally, ensure that the Execution Policy allows your script to be run. To learn more, see about_Execution_Policies and What is PowerShell?.
> For **Execution duration in seconds** under **Script settings**, enter the timeout duration for your setup script.

> **Note**
>
> The execution duration in seconds is how long WorkSpaces Applications waits for the setup script to run before continuing. If your setup script doesn't complete within this duration, an error is displayed to your user and the application will attempt to launch. The setup script is terminated after the execution duration has elapsed.

6. Under **Import Settings**, choose **Create new app block application file**. For **S3 Location** under **Import settings**, either enter the Amazon S3 URI that represents the bucket, or choose **Browse S3** to navigate to your Amazon S3 buckets and select an appropriate bucket. The list of Amazon S3 buckets is global and lists all the buckets across all regions. Make sure you select the bucket in the region where you want to create your app block. For more information about

setting bucket permissions, see the section called “Store Application Icon, Setup Script, Session Script, and VHD in an S3 Bucket”.

7. Select an app block builder. Only app block builders that are not associated with other app blocks are available. If the list is empty, either create a new app block builder, or disassociate the existing ones to use. App block builder is a reusable resource that you can use to create your application package.

   > **Note**
   >
   > If you do not select an app block builder here, you can still create your app block in the **Inactive** state, and activate your app block later. For more information, see the section called “Activate an App Block”.

8. (Optional) For **Tags**, create tags for the app block resource.
9. Choose **Next**.
10. Review the information that you entered, and choose one of the following options:

    - Choose **Create app block** if you didn't select an app block builder in step 7.
    - Choose **Launch app block builder** if you chose an app block builder in step 7. Then continue to Step 2 to create your application package using the app block builder streaming session.

At this point, your app block resource is created, but it is **Inactive** and can't be used for Elastic fleets.

Step 2: Create the Application Package

Use the app block builder streaming instance to package your applications and activate your app block. The app block created using app block builder will have WorkSpaces Applications packaging, and the application package will be uploaded onto the Amazon S3 bucket in your AWS account.

To create the application package

1. After your streaming session is on, the application builder assistant automatically starts. If it doesn’t start, start it manually using the desktop icon.
2. The initial screen provides instructions for the application packaging process.
3. Bring your application installer onto your app block builder streaming session by using one of the following options:

- Download the application installers from the web.
- Use your streaming session file interface.
- Download the application installer from another AWS service using a machine role.

4. After you have all the required application installers, stop all the other apps running on the instance and choose **Start recording**. The app block builder starts recording system changes, and the screen says **Recording in progress**.
5. Start installing your applications one by one.
6. When you are done with application installation, choose **Stop recording**, and the system will stop recording changes. If you want to make any more changes to your application package, such as add more applications or remove an already installed application, choose **Start recording**, and make sure the system is in **Recording in progress** mode.

> **Note**
>
> If your application installation fails, choose **Report a problem** to collect WorkSpaces Applications related logs from the instance, and report the problem to the WorkSpaces Applications team. When you are done, end your app block builder streaming session. You can try to restart the process creating an app block by using a new app block builder instance. If the problem persists, then try to create your app block using custom packaging.

7. When you are done installing all the applications, choose **Stop recording**. You can test your application, by using the Start Menu or browsing the application using File Explorer.
8. Choose **Next** to review your app block details.

> **Note**
>
> The recommended size of an application package (VHD) file for an Elastic fleet is less than 1.5 GB. If your VHD file size is bigger than 1.5 GB, try reducing the number of applications packaged within one app block.
> Application package (VHD) file size will not shrink if you uninstall an application. Restart the application packaging process using a new app block streaming session, and install fewer applications.

9. Choose **Finish app block creation and disconnect** to create the application package and upload it to the Amazon S3 bucket. If you are successful, the streaming session will automatically disconnect, and the app block will be in an **Active** state.

> **Note**
>
> If your application installation fails, choose **Report a problem** to collect WorkSpaces Applications related logs from the instance, and report the problem to the WorkSpaces Applications team. When you are done, end your app block builder streaming session. You can try to restart the process creating an app block by using a new app block builder instance. If the problem persists, then try to create your app block using custom packaging.

Activate an App Block

If an app block with WorkSpaces Applications packaging was created, but the application package (VHD) was not attached to it, then the app block will be in an inactive state, and it can't be used to associate applications with Elastic fleets. To activate an app block, an application package (VHD) must be associated with the app block.

To create the application package

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. From the left-hand navigation menu, choose **Applications Manager**, **App blocks**.
3. Select an **Inactive** app block that you want to activate, and choose **Activate** from the **Actions** menu.
4. Select an app block builder, and choose **Launch app block builder**.

   - If the list is empty, then you either don’t have an app block builder, or all of your app block builders are associated with other app blocks. Either create a new app block builder, or disassociate an existing app block builder and test it.
   - If the app block builder is already associated with an app block, then you can continue using it for activating the app block.
   - If the selected app block builder was not associated with an app block builder, then it will be associated with the one you select, and the streaming session will launch. The app block builder remains associated with this app block after the session ends.

5. After the app block builder streaming session starts, follow the steps in the section called "Step 2: Create the Application Package" to create your application package (VHD) and activate the app block.

Create an App Block with an Existing App Package

You can use your existing application package (VHD) with WorkSpaces Applications packaging to create WorkSpaces Applications app blocks. To do this, copy your application package (VHD) file from the source Amazon S3 bucket to another destination Amazon S3 bucket. The destination bucket can be in a different region.

To create an app block with existing app package

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. From the left-hand navigation menu, choose **Applications Manager**, **App blocks**, and **Create app block**.
3. For app block packaging, select **WorkSpaces Applications**.
4. For **App block details**, type a unique name identifier for the app block. Optionally, you can also specify the following:
   - **Name** – A unique name for the app block.
   - **Display name** (optional) – A friendly name for the app block.
   - **Description** (optional) – A description for the app block.
5. (Optional) An app block with WorkSpaces Applications packaging doesn't need a setup script. You can optionally provide post-installation steps the following **Advanced Options**:
   - For **Post setup script object in S3**, either enter the Amazon S3 URI that represents the post setup script object, or choose **Browse S3** to navigate to your Amazon S3 buckets and find the setup script object.
   - For **Post setup script executable**, enter the executable needed for your post setup script.

> **Note**
>
> If your post setup script can execute directly, enter the filename of the post setup script. If your post setup script relies on another executable (for example, Microsoft PowerShell) to execute, enter the path to that executable.

Path to Microsoft PowerShell on Microsoft Windows:
`C:\Windows\System32\WindowsPowerShell\v1.0\powershell.exe`
Optionally, for **Post setup script executable arguments**, enter in the arguments that need to be provided to the setup script executable to execute your setup script.

> **Note**
>
> If you are using a Microsoft PowerShell script, you must specify the "-File" parameter with the name of your post setup script as an executable argument. Additionally, ensure that the Execution Policy allows your script to be run. To learn more, see about_Execution_Policies and What is PowerShell?.
> For **Execution duration in seconds** under **Script settings**, enter the timeout duration for your setup script.

> **Note**
>
> The execution duration in seconds is how long WorkSpaces Applications waits for the post setup script to run before continuing. If your post setup script doesn't complete within this duration, an error is displayed to your user and the application will attempt to launch. The setup script is terminated after the execution duration has elapsed.

6. Choose **Use existing app block application file** under **Import settings**. For **S3 Location**, you can enter the Amazon S3 URI for the object in an Amazon S3 bucket that represents the application package (VHD), Or, choose **Browse S3** to navigate to your Amazon S3 buckets and select the object in an Amazon S3 bucket. The list of Amazon S3 buckets is global and lists all the buckets across all regions. Make sure you select the bucket in the region where you want to create your app block.
7. Choose **Next**.
8. Review the information that you entered, and choose **Create app block**.

At this point your app block resource is created and in the **Active** state.

Test an App Block

You can use an app block builder to test your app block and verify your application functionalities. You don't need to launch an Elastic fleet for this option. You can also create multiple app block builders with different instance types or sizes, and test the performance of your application with different compute options.

> **Note**
>
> The test app block option is supported only for app blocks with WorkSpaces Applications packaging.

To test an app block

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. From the left-hand navigation menu, choose **Applications Manager**, **App blocks**.
3. Select an app block that you want to test, and choose **Test** from the **Actions** menu.
4. Select an app block builder, and choose **Launch and test app block**.
   - If the list is empty, then you either don't have an app block builder, or all of your app block builders are associated with other app blocks. Either create a new app block builder, or disassociate an existing app block builder and test it.
   - If the app block builder is already associated with an app block, then you can continue using it for activating the app block.
   - If the selected app block builder was not associated with an app block builder, then it will be associated with the one you select, and the streaming session will launch. The app block builder remains associated with this app block after the session ends.
5. App block builder launches in a separate browser window in a Desktop streaming mode. The service downloads the app block from the Amazon S3 bucket and installs it on the app block builder instance.
6. Your applications can now be streamed and tested. You can open your application by either browsing it in File Explorer or using the Start menu.
7. When you are done testing, end the streaming session.

Associate an App Block in Amazon WorkSpaces Applications

In order to create, test, or activate your app block with WorkSpaces Applications packaging, you need to associate it with an app block builder. One app block builder can only be associated with only one app block, and vice versa.

> **Note**
>
> Associating and disassociating an app block is only supported for app blocks with WorkSpaces Applications packaging.

Associate an app block with app block builder in Amazon WorkSpaces Applications

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. From the left-hand navigation menu, choose **Applications Manager**, **App blocks**.
3. Select an app block, and choose **Associate** from the **Actions** menu.
4. Select an app block builder, and choose **Associate app block builder**.

   If the list is empty, then you either don't have an app block builder, or all of your app block builders are associated with other app blocks. Either create a new app block builder, or disassociate an existing app block builder and then associate.

Disassociate an App Block in Amazon WorkSpaces Applications

If all your app block builders are associated with other app blocks, and you want to test, create, or activate another app block, then you can either create a new app block builder, or disassociate an existing app block builder from the app block and use it with the new app block.

> **Note**
>
> Associating and disassociating an app block is only supported for app blocks with WorkSpaces Applications packaging.
> Disassociation is allowed only if an app block builder is in the **STOPPED** state.

Disassociate an app block from an app block builder

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. From the left-hand navigation menu, choose **Applications Manager**, **App blocks**.
3. Select an app block, and choose **Disassociate** from the **Actions** menu.
4. Select an already associated app block builder, and choose **Disassociate app block builder**.

Unsupported Applications

Applications might encounter failures when installing or running in the following scenarios:

- **Applications requiring reboots after installation**: If an application needs to perform additional changes or configurations after installation that require a reboot, it might fail. Currently, app block builder does not support restart, which can prevent the application from completing its required post-installation steps.
- **Applications relying on user-specific details**: Applications that are intended to be installed only for the currently logged-in user on app block builder, or that rely on the logged-in user details on app block builder, such as security identifiers (SIDs) during installation, might not function correctly on Elastic fleets. This is due to the logged-in user changes within the elastic fleet environment. Additionally, application redirection does not record all directories under %USERPROFILE%. However, you have the option to configure post setup scripts to dynamically change your application configuration based on environment.
- **Applications relying on machine-specific details**: Applications that rely on machine-specific details on app block builder during installation, such as network adaptor GUID, might encounter issues on Eastic fleets. This is because the machine details, including network adaptor GUIDs, can change within the elastic fleet environment. To address this, you can configure the post setup scripts to handle the configuration of those machine-specific details.

If you are uncertain whether your application falls into any of these categories, you can use WorkSpaces Applications packaging to create an app block. This process involves installing your application(s) on an app block builder instance. In the event that your application(s) fail to install on the app block builder instance, you can take the following actions:

- Check the logs. The error log file for your app block builder instance can be found at C:\AppStream\AppBlocks\errorLog. This log records all installation failures, including registry keys and file operation processing. If you see any of the following logs in the errorLog, it indicates

that the packaging of your application is currently unsupported by the WorkSpaces Applications app block builder:

- "Unable to create symbolic link"
- "Service doesn't support file renaming"

If there is no errorLog file, or if this file is empty, then check your application installation logs to identify the reason for failures.

- Report a problem. Select the **Report a problem** button, which is available on the application builder assistant in the app block builder. Selecting this option will gather all the WorkSpaces Applications logs from your app block builder instance, and submit them to the WorkSpaces Applications team.
- Create an app block with custom packaging: If you are unable to package your applications using the app block builder, you can try to create an app block using custom packaging methods. For more information, see the section called “Custom App Blocks”.
- If you need more help, contact AWS Support. For more information, see AWS Support Center.

It is important to consider these potential limitations, and plan accordingly when using WorkSpaces Applications packaging for your applications.

App Block Builder

An app block builder is a reusable resource that you can use to package your applications (or app block). You can also use it to test your application package before associating your application to an Elastic fleet. A single app block builder can be used to create and test multiple app blocks one by one. Each time a streaming session is created for app block builder for creating or testing an app block, a new instance is created and used. After the app block builder instance is terminated, the state of the instance is not persisted.

WorkSpaces Applications Elastic fleets use Amazon EC2 instances to stream applications. You must provide your application package and associate it with your fleet. To create your own custom application packaging, connect to an app block builder instance, and then install and configure your applications for streaming. App block builder creates the packaging for your application and uploads it to an Amazon S3 bucket in your AWS account.

When you create an app block builder, you choose the following:

- An instance type — WorkSpaces Applications provides different instance sizes with various CPU and memory configurations. The instance type must align with the instance family you need.
- The VPC, subnets, and security groups to use — Make sure that the subnets and security groups provide access to the network resources that your applications require. Typical network resources required by applications might include licensing servers, database servers, file servers, and application servers. App block builder uploads the application package on to an Amazon S3 bucket in your AWS account. The VPC you choose for your fleet must provide sufficient network access to the Amazon S3 bucket. For more information, see the section called "Store Application Icon, Setup Script, Session Script, and VHD in an S3 Bucket".

Contents

- Create an App Block Builder
- Connect to an App Block Builder in Amazon WorkSpaces Applications
- App Block Builder Actions

Create an App Block Builder

You can use app block builder instance to create your application package for WorkSpaces Applications Elastic fleets.

To create an app block builder

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. Choose **Applications Manager** in the left navigation pane, then choose the **App block builders** tab and **Create app block builder**.
3. For **Step 1: Configure app block builder**, configure the app block builder by providing the following details:

   - **Name**: Type a unique name identifier for the app block builder.
   - **Display name (optional)**: Type a name to display for the app block builder (maximum of 100 characters).
   - **Operating system**: Select an operating system for your application. This must align with the operating system that you are going to select for you elastic fleet, which your end users will use to stream the application.

- **IAM role (Optional)**: When you apply an IAM role from your account to an WorkSpaces Applications app block builder, you can make AWS API requests from the app block builder instance without manually managing AWS credentials. To apply an IAM role to the app block builder, do either of the following:
  - To use an existing IAM role in your Amazon Web Services account, choose the role that you want to use from the **IAM role** list. The role must must be accessible from the image builder. For more information, see Configuring an Existing IAM Role to Use With WorkSpaces Applications Streaming Instances.
  - To create a new IAM role, choose **Create new IAM role** and follow the steps in How to Create an IAM Role to Use With WorkSpaces Applications Streaming Instances.
- **Instance Type**: Select the instance type for the app block builder. Choose a type that matches the performance requirements of the applications that you plan to install.
- **Tags (optional)**: Choose **Add Tag**, and type the key and value for the tag. To add more tags, repeat this step. For more information, see Tagging Your Amazon WorkSpaces Applications Resources.

4. Choose **Next**.
5. For **Step 2: Configure Network**, do the following:

- To add internet access for the app block builder in a VPC with a public subnet, choose **Default Internet Access**. If you are providing internet access by using a NAT gateway, leave **Default Internet Access** unselected. For more information, see Internet Access.
- For **VPC** and **Subnet 1**, choose a VPC and at least two subnets. For increased fault tolerance, we recommend that you choose three subnets in different Availability Zones. For more information, see Configure a VPC with Private Subnets and a NAT Gateway.

  If you don't have your own VPC and subnet, you can use the default VPC or create your own. To create your own, choose the **Create a new VPC** and **Create new subnet** links to create them. Choosing these links opens the Amazon VPC console. After you create your VPC and subnets, return to the WorkSpaces Applications console and choose the refresh icon to the left of the **Create a new VPC** and **Create new subnet** links to display them in the list. For more information, see Configure a VPC for WorkSpaces Applications.
- For **Security group(s)**, choose up to five security groups to associate with this image builder. If you don't have your own security group and you don't want to use the default security group, choose the **Create new security group** link to create one. After you create your subnets in the Amazon VPC console, return to the WorkSpaces Applications console and

choose the refresh icon to the left of the **Create new security group** link to display them in the list. For more information, see Security Groups in Amazon WorkSpaces Applications.

- For **VPC Endpoints (Optional)**, you can create an interface VPC endpoint (interface endpoint) in your virtual private cloud (VPC). To create the interface endpoint, choose **Create VPC Endpoint**. Selecting this link opens the VPC console. To finish creating the endpoint, follow steps 3 through 6 in the section called "Tutorial: Creating and Streaming from Interface VPC Endpoints". After you create the interface endpoint, you can use it to keep streaming traffic within your VPC.

6. Choose **Next**.
7. Choose **Review** and confirm the details for the app block builder. To change the configuration for any section, choose **Edit** and make the needed changes.
8. After you finish reviewing the configuration details, choose **Create app block builder**.

> **Note**
>
> If an error message notifies you that you don't have sufficient limits (quotas) to create the image builder, submit a limit increase request through the Service Quotas console at https://console.aws.amazon.com/servicequotas/. For more information, see Requesting a quota increase in the *Service Quotas User Guide*.

Connect to an App Block Builder in Amazon WorkSpaces Applications

You can connect to an app block builder by doing either of the following:

- Using the WorkSpaces Applications console (for browser connections only)
- Creating a streaming URL (for browser or WorkSpaces Applications client connections)

> **Note**
>
> App block builder doesn't support Active Directory domain join.

Contents

- Amazon WorkSpaces Applications Console (Browser Connection)
- Streaming URL (Amazon WorkSpaces Applications Client or Browser Connection)

Amazon WorkSpaces Applications Console (Browser Connection)

To use the WorkSpaces Applications console to connect to an app block builder through a browser, complete the following steps.

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the left navigation pane, choose **Applications Manager**, and then choose **App block builders**.
3. In the list of app block builders, choose the app block builder to which you want to connect. Verify that the status of the app block builder is **Running**, and choose **Connect**.

   For this step to work, you might need to configure your browser to allow pop-ups from https://stream.<aws-region>.amazonappstream.com/.
4. Start streaming the app block builder.

Streaming URL (Amazon WorkSpaces Applications Client or Browser Connection)

You can create a streaming URL to connect to an app block builder through a browser or the WorkSpaces Applications client. Unlike a streaming URL that you create to enable user access to a fleet instance, which is valid for a maximum of seven days, by default, a streaming URL that you create to access an image builder expires after one hour. To set a different expiration time, you must generate the streaming URL by using the CreateAppBlockBuilderStreamingURL API action.

> **Note**
>
> Streaming a URL to connect to an app block builder is not supported on the macOS client.

You can create a streaming URL in any of the following ways:

- WorkSpaces Applications console
- The CreateAppBlockBuilderStreamingURL API action
- The create-app-block-builder-streaming-url AWS CLI command

To create a streaming URL and connect to the app block builder by using the WorkSpaces Applications console, complete the steps in the following procedure.

To create a streaming URL and connect to the app block builder by using the WorkSpaces Applications console

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the navigation pane, choose **Application Manager**, **App block builders**.
3. In the list of app block builders, choose the app block builder to which you want to connect. Verify that the status of the app block builder is **Running**.
4. Choose **Actions**, **Create streaming URL**.
5. Do one of the following:
   - To save the streaming URL to connect to the app block builder later, choose **Copy Link** to copy the URL, then save it to an accessible location.
   - To connect to the app block builder through the WorkSpaces Applications client, choose **Launch in Client**. When you choose this option, the WorkSpaces Applications client sign-in page is prepopulated with the streaming URL.
   - To connect to the app block builder through a browser, choose **Launch in Browser**. When you choose this option, a browser opens with the address bar prepopulated with the streaming URL.
6. After you create the streaming URL and connect to the app block builder, start streaming the app block builder.

App Block Builder Actions

You can perform the following actions on an app block builder, depending on the current state (status) of the app block builder instance.

Delete

Permanently delete an app block builder.

The instance must be in a **Stopped** state.

Connect

Connect to a running app block builder. This action starts a desktop streaming session with the app block builder to install and add applications, and create an app block.

The instance must be in a **Running** state.

Start

Start a stopped app block builder. A running instance is billed to your account.

The instance must be in a **Stopped** state, and associated with an app block.

Stop

Stop a running app block builder. A stopped instance is not billed to your account.

The instance must be in a **Running** state.

Update

Update any of the app block builder properties, except the name.

The instance must be in a **Stopped** state.

None of these actions can be performed on an instance in any of the following intermediate states:

- **Pending**
- **Stopping**
- **Starting**
- **Deleting**

Applications

Applications contain the details necessary to launch your application after the VHD has been mounted. Applications also include the name and icon that are displayed to your user on the application catalog. Applications are associated with the app block resource that contains the files and binaries for that application.

You can use the WorkSpaces Applications console to create the application resource once you have uploaded your application icon to an Amazon S3 bucket and created the app block that contains the files and folders necessary to launch the application. To learn more about uploading the application icon to an Amazon S3 bucket, see the section called “Store Application Icon, Setup Script, Session Script, and VHD in an S3 Bucket”.

> **Note**
>
> You must have IAM permissions to perform the `S3:GetObject` action on the application icon object in the S3 bucket to create the application resource.

To create the application resource

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. From the left-hand navigation menu, choose **Applications** and **Create application**.
3. For **Name** under **Application details**, enter a unique identifier for the application.
4. (Optional) For **Display name** under **Application details**, enter a friendly name that users will see in the application catalog.
5. (Optional) For **Description** under **Application details**, enter a description for the application.
6. For **Application icon object in S3** under **Application details**, either enter the S3 URI that represents the VHD object, or choose **Browse S3** to navigate to your S3 buckets and find the application icon object.
7. For **Application executable launch path** under **Application settings**, enter the path on the streaming instance to the application's executable.
8. (Optional) For **Application working directory** in the **Application settings** section, enter the directory on the streaming instance to use for the application's working directory.
9. (Optional) For **Application launch parameters** in the **Application settings** section, enter the parameters to provide to the application executable when launching the application.
10. For **Supported operating systems (OS)** in the **Application settings** section, choose which operating systems can launch this application.
11. For **Supported instance families** in the **Application settings** section, choose which instance families can launch this application.
12. For **App block** in the **Application settings** section, choose which app block contains the files and folders necessary for this application.
13. (Optional) In the **Tags** section, create tags for the app block resource.
14. Review the information that you entered, then choose **Create**.
15. If your application was created successfully, you will see a success message at the top of the console. If an error occurred, a descriptive error message will be provided and you will need to try creating the application again.

Store Application Icon, Setup Script, Session Script, and VHD in an S3 Bucket

You must store the application icons, setup scripts, session scripts, and VHDs that you use for your applications and app blocks in an Amazon Simple Storage Service (Amazon S3) bucket in your AWS account. WorkSpaces Applications Elastic fleets download the application icon, setup script, and VHD from the S3 bucket when your user starts their streaming session. The S3 bucket must reside in the AWS Region that you intend to create WorkSpaces Applications Elastic fleets within.

We recommend that you create a new S3 bucket that is used to store only the application icons, setup scripts, session scripts, and VHDs that you intend to use with Elastic fleets. We also recommend enabling versioning on the S3 bucket. This allows reverting to previous object versions if necessary. For more information about how to create a new S3 bucket, see Creating a bucket. For more information about how to manage object versioning, see Using versioning in S3 buckets.

> **Note**
>
> WorkSpaces Applications uses your VPC to access the S3 bucket you select. The VPC you choose for your fleet must provide sufficient network access to the S3 bucket.
> Make sure that your S3 bucket content is not encrypted using keys that you manage (Customer Managed Keys).
> Currently, S3 buckets configured to use server-side encryption with customer-provided encryption keys (SSE-C) are not supported for Elastic fleets. If you require encryption at rest for your S3 objects, server-side encryption with Amazon S3-managed encryption keys (SSE-S3) is an option that will work for Elastic fleets.

Topics

- Amazon S3 Bucket Permissions

Amazon S3 Bucket Permissions

The Amazon S3 bucket that you choose must have a bucket policy that provides sufficient access to the WorkSpaces Applications service principal to access and download objects from the Amazon S3 bucket. You will need to modify the following bucket policy, then apply it to the Amazon S3 bucket you intend to use for application icons, setup scripts, and VHDs. For more information about

how to apply a policy to an Amazon S3 bucket, see Adding a bucket policy using the Amazon S3 console.

Make sure that the access control lists (ACLs) for your Amazon S3 buckets are disabled. For more information, see Disabling ACLs for all new buckets and enforcing Object Ownership.

This section presents examples of typical use cases for bucket policies. These sample policies use *bucket* as the resource value. To test these policies, replace the *user input placeholders* with your own information (such as your bucket name).

> **Note**
>
> WorkSpaces Applications doesn't support confused deputy prevention for application icons, so the statement that grants access to the application icon object can't include a condition. For more information, see the section called "Example: WorkSpaces Applications Application Amazon S3 bucket policy cross-service confused deputy prevention".

```
{
  "Version": "2012-10-17",
  "Statement": [
      {
       "Sid": "AllowAppStream20ToRetrieveObjects",
       "Effect": "Allow",
       "Principal": {
          "Service": ["appstream.amazonaws.com"]
        },
        "Action": ["s3:GetObject"],
        "Resource": [
           "arn:aws:s3:::bucket/VHD object",
           "arn:aws:s3:::bucket/Setup script object",
           "arn:aws:s3:::bucket/Session scripts zip file object"
         ],
        "Condition": {
          "StringEquals": {
            "aws:SourceAccount": "111122223333"
          }
        }
      },
      {
       "Sid": "AllowRetrievalPermissionsToS3AppIconsForAppStream",
       "Effect": "Allow",
```

```
       "Principal": {
          "Service": ["appstream.amazonaws.com"]
        },
        "Action": ["s3:GetObject"],
        "Resource": "arn:aws:s3:::bucket/Application icon object"
      }
  ]
}
```

> **Note**
>
> The bucket policy example defines specific objects in the S3 bucket that WorkSpaces Applications can access. You can also use prefixes and wildcards to simplify policy management as you increase your app blocks. For more information about bucket policies, see Using bucket policies. For more information about common bucket examples, see Bucket policy examples.

If you are using an WorkSpaces Applications app block, then WorkSpaces Applications requires additional permissions to upload the application package to your appropriate Amazon S3 bucket. For more information about WorkSpaces Applications app blocks, see the section called "WorkSpaces Applications App Blocks".

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowAppStream20ToPutAndRetrieveObjects",
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "appstream.amazonaws.com"
        ]
      },
      "Action": [
        "s3:GetObject",
        "s3:ListBucket",
        "s3:PutObject",
        "s3:GetBucketOwnershipControls"
      ],
      "Resource": [
```

```
        "arn:aws:s3:::bucket",
        "arn:aws:s3:::bucket/AppStream2/*",
        "arn:aws:s3:::bucket/Setup script object",
        "arn:aws:s3:::bucket/Session scripts zip file object"
      ],
      "Condition": {
        "StringEquals": {
          "aws:SourceAccount": "111122223333"
        }
      }
    },
    {
      "Sid": "AllowRetrievalPermissionsToS3AppIconsForAppStream",
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "appstream.amazonaws.com"
        ]
      },
      "Action": [
        "s3:GetObject"
      ],
      "Resource": "arn:aws:s3:::bucket/Application icon object"
    }
  ]
}
```

Associate Applications to Elastic Fleets

Applications must be associated to Elastic fleets before they appear to users in the application catalog to be launched. You can manage application associations using the WorkSpaces Applications console. For information about how to associate applications when creating an Elastic fleet, and how to manage application associations on existing fleets, see Update an Amazon WorkSpaces Applications Fleet.

Additional Resources

The following links provide information and other resources to help you package and deliver your applications with Elastic fleets.

Solution	Description
AWS	• Use Elastic fleets and Linux for inexpensive, secure bastion hosts in Amazon AppStream 2.0 — Describes how to package common bastion host applications and deliver them using Linux instances on Elastic fleets. • Automate AppStream 2.0 Elastic fleet application updates with AWS Systems Manager — Describes how you can automate the creation of VHDs and app blocks using an EC2 instance with AWS Systems Manager. • Stream applications at a lower cost with Amazon AppStream 2.0 Elastic fleets and Linux compatibility — Describes how you can package applications into a VHD and deliver them using Linux instances on Elastic fleets.
Liquidware FlexApp	Liquidware FlexApp integration with AppStream 2.0 Elastic fleet — Includes step-by-step instructions to package any windows-based application into FlexApp one format and deliver on WorkSpaces Applications Elastic fleets.
Numecent Cloudpaging	Introducing AppStream 2.0 Elastic Fleets: How to Maximize Application Compatibility and Portability with Cloudpaging — Numecent Cloudpaging eliminate s application compatibility issues, allowing even the most complex legacy and highly customized applications to be packaged with their dependencies and ready for deployment within a day. WorkSpaces Applications users can then stream them on-demand without the need for IT to repackage for various Windows operating systems or devices.
Turbo.Net	Introducing Turbo support for Amazon WorkSpaces Applications Elastic fleets — Turbo delivers applications instantly on major platforms and devices. Turbo's unique container technology eliminates installs and manages dependencies, conflicts, and entitlements. WorkSpaces

Solution	Description
	Applications users can immediately launch any applicati on in a Turbo Hub.

Fleets and Stacks in Fleet Type in Amazon WorkSpaces Applications

With Amazon WorkSpaces Applications, you create fleet instances and stacks as part of the process of streaming applications. A fleet consists of streaming instances that run the image that you specify. A stack consists of an associated fleet, user access policies, and storage configurations.

Contents

- Session Context in Amazon WorkSpaces Applications
- WorkSpaces Applications Fleet Types
- WorkSpaces Applications Instance Families
- Create an Amazon WorkSpaces Applications Fleet and Stack
- Customize an Amazon WorkSpaces Applications Fleet to Optimize Your Users' Application Streaming Experience
- Update an Amazon WorkSpaces Applications Fleet
- Fleet Auto Scaling for Amazon WorkSpaces Applications
- Multi-Session Recommendations

Session Context in Amazon WorkSpaces Applications

You can pass parameters to your streaming application by using either of the following methods:

- Specify session context in the CreateStreamingURL WorkSpaces Applications API operation. For more information, see CreateStreamingURL.
- Add the sts:TagSession permission to your IAM role's trust policy and specify the session context as a SAML assertion in your SAML 2.0 identity provider's authentication response. For more information, see Step 5: Create Assertions for the SAML Authentication Response and the section called “Step 5: Create Assertions for the SAML Authentication Response”.

If your image uses a version of the WorkSpaces Applications agent that was released on or after October 30, 2018, the session context is stored within the image as a Windows or Linux environment variable. For information about specific environment variables, see "User and Instance

Metadata for WorkSpaces Applications Fleets" in Customize an Amazon WorkSpaces Applications Fleet to Optimize Your Users' Application Streaming Experience.

> **Note**
>
> The session context parameter is visible to the user in the WorkSpaces Applications streaming URL. We strongly recommend that you never put confidential or sensitive information in the session context parameter. Because it is possible for users to modify the streaming URL, we recommend performing additional validation to determine that the session context is valid for the end user. For example, you can compare the session context with other session information, such as user and instance metadata for WorkSpaces Applications fleets.
> WorkSpaces Applications does not perform validation on the session context parameter.

Using Session Context to Pass Parameters to a Streaming Application

In the following steps, you'll use session context to start a web browser and automatically open a specific website. For instances running Windows, you'll use Firefox. For instances running Linux, you'll use Chromium.

To use session context to launch a website

1. In the left navigation pane, choose **Images**, **Image Builder**.
2. Choose the image builder to use, verify that it is in the **Running** state, and choose **Connect**.
3. Log in to the image builder by choosing **Administrator** on the **Local User** tab.
4. Create a child folder of `C:\`. For this example, use `C:\Scripts`.
5. Create a Windows batch file in the new folder. For this example, create `C:\Scripts\session-context-test.bat` and add a script that launches Firefox with the URL from the session context.

   Use the following script:

   ```
   CD "C:\Program Files (x86)\Mozilla Firefox"
   Start firefox.exe %APPSTREAM_SESSION_CONTEXT%
   ```

6. In Image Assistant, add `session-context-test.bat` and change the name to **`Firefox`**.

You do not need to add Firefox. This step requires that you add only the batch file.

7. Create an image, fleet, and stack. For this example, use a fleet name of **`session-context-test-fleet`** and a stack name of **`session-context-test-stack`**.
8. After the fleet is running, you can call create-streaming-url with the `session-context` parameter, as shown in this example.

```
aws appstream create-streaming-url --stack-name session-context-test-stack \
--fleet-name session-context-test-fleet \
--user-id username --validity 10000 \
--application-id firefox --session-context "www.amazon.com"
```

9. Open the streaming URL in a browser. The script file launches Firefox and loads `http://www.amazon.com`.

Similarly, you can perform the following steps to pass parameters to your Linux streaming application.

To pass parameters to your Linux streaming application

1. In the left navigation pane, choose **Images**, **Image Builder**.
2. Choose the image builder to use, verify that it is in the **Running** state, and choose **Connect**.
3. Log in to the image builder by default as **ImageBuilderAdmin**.
4. Create a script file (for example, launch-chromium.sh) by running the following command:

   sudo vim /usr/bin/launch-chromium.sh
5. Write the script and set executable permissions, such as the following:

> **Note**
>
> #!/bin/bash and source /etc/profile are always required in the script.

```
#!/bin/bash
source /etc/profile
/usr/bin/chromium-browser $APPSTREAM_SESSION_CONTEXT
```

6. Use the Image Assistant CLI to add launch-chromium.sh:

```
sudo AppStreamImageAssistant add-application \
--name chromium \
--absolute-app-path /usr/bin/launch-chromium.sh
```

7. Create an image, fleet, and stack. For this example, use a fleet name of **`session-context-test-fleet`** and a stack name of **`session-context-test-stack`**.
8. After the fleet is running, you can call create-streaming-url with the `session-context` parameter, as shown in this example.

```
aws appstream create-streaming-url --stack-name session-context-test-stack \
--fleet-name session-context-test-fleet \
--user-id username \
--application-id chromium --session-context "www.amazon.com"
```

9. Open the streaming URL in a browser. The batch file launches Chromium and loads `http://www.amazon.com`.

WorkSpaces Applications Fleet Types

The fleet type determines when your instances run and how you pay for them. You can specify a fleet type when you create a fleet. You cannot change the fleet type after you create the fleet.

The following are the possible fleet types:

Always-On

Streaming instances run all the time, even when no users are streaming applications and desktops. Streaming instances must be provisioned before a user is able to stream. The number of streaming instances provisioned is managed through auto scaling rules.

When your users choose their application or desktop, they will start streaming instantly. You are charged the running instance fee for all streaming instances, even when no users are streaming.

On-Demand

Streaming instances run only when users are streaming applications and desktops. Streaming instances not yet assigned to users are in a stopped state. Streaming instances must be provisioned before a user is able to stream. The number of streaming instances provisioned is managed through auto scaling rules.

When your users choose their application or desktop, they will start streaming after a 1-2 minute wait. You are charged a lower stopped instance fee for streaming instances that are not yet assigned to users, and the running instance fee for streaming instances that are assigned to users.

Elastic

The pool of streaming instances is managed by WorkSpaces Applications. When your users select their application or desktop to launch, they will start streaming after the app block has been downloaded and mounted to a streaming instance. For more information about creating app blocks for your Elastic fleets, see App Blocks.

You are charged the running instance fee for Elastic fleet streaming instances only for the duration of the streaming session, in seconds, with a minimum of 15 minutes.

For more information about how fleet types are charged, see Amazon WorkSpaces Applications Pricing.

Always-On and On-Demand Fleets

Always-On and On-Demand fleets represent a pool of streaming instances that you manage the capacity of using auto scaling policies. Your users use the streaming instances to stream their applications and desktops. With an Always-On fleet, your user's application launches nearly instantly, and you pay the running instance rate per instance even when a user isn't streaming. With an On-Demand fleet, your user's application launches after a 1-2 minute wait as the streaming instance is started, and you pay a lower cost stopped instance fee for instances not in use, and the running instance fee for instances that are in use.

Applications for Always-On and On-Demand fleet instances are delivered through WorkSpaces Applications images that are created by image builders. You can learn more about how to create an image builder, install your applications, and create an image by reading *Images*.

Always-On and On-Demand fleet streaming instances must be provisioned and unassigned to an existing user before a user can stream. You can use fixed or dynamic fleet autoscaling policies to manage the number of instances in your fleet, ensuring that you have sufficient available capacity to meet your user needs while controlling costs. You can learn more about scaling your fleets by reading the section called "Fleet Auto Scaling".

Elastic Fleets

Elastic fleets represent a pool of streaming instances that WorkSpaces Applications manages. You do not need to predict concurrency, or create and manage any auto scaling policies for your users to stream their applications and desktops. When your user requests a streaming instance, a streaming instance is assigned from the pool, and made available to them after configuration completes.

Elastic fleets rely on applications that are stored on app blocks. When a user chooses an application from the catalog, the app block is downloaded to the instance, mounted, and then the application launches.

AWS manages the streaming instance provisioning and availability with an Elastic fleet. You need to configure the maximum concurrency you expect when creating and updating the fleet, and ensure that you have sufficient streaming instance limits to meet your user demand.

For more information about creating app blocks for your Elastic fleets, see App Blocks.

WorkSpaces Applications Instance Families

Amazon WorkSpaces Applications users stream applications from stacks that you create. Each stack is associated with a fleet. When you create a fleet, the instance type that you specify determines the hardware of the host computers used for your fleet. Each instance type offers different compute, memory, and GPU capabilities. Instance types are grouped into *instance families* based on these capabilities. For hardware specifications and pricing information, see WorkSpaces Applications Pricing.

When you create a fleet or image builder, you must select an image that is compatible with the instance family on which you intend to run your fleet.

- When launching a new image builder, you are presented with a list of the images in your image registry. Select the appropriate base image.
- When launching a fleet, ensure that the private image you select was created from the appropriate base image.

The following table summarizes the available instance families and provides the base image naming format for each. Select an instance type from an instance family based on the

requirements of the applications that you plan to stream on your fleet, and match the base image according to the following table.

Note

If your use case involves real-time audio-video (AV) or other scenarios that require high frame rates and your display performance isn't as expected, consider scaling up to a larger instance size.

Important

Memory integrity

Memory integrity, also known as *hypervisor-protected code integrity (HVCI)* or *hypervisor enforced code integrity*, isn't supported. Before you turn on Credential Guard, you must ensure this feature is disabled. For more information, see Disable memory integrity. Enabling the feature may lead to performance or stability issues. We recommend disabling memory integrity in your Windows image. Credential Guard can remain enabled once memory integrity is disabled.

Instance Family	Description	Base Image Name
General Purpose	Basic computing resources for running web browsers and most business applications.	For stream.standard.* instance types: AppStream-*WinServer OperatingSystemVersion -MM-DD-YYYY* AppStream-RockyLinux8-*MM-DD-YYYY* AppStream-RHEL8-*MM-DD-YYYY*
		Import Image for GeneralPurpose.* instance types

Instance Family	Description	Base Image Name
Compute Optimized	Optimized for compute-bound applications that benefit from high performance processors.	For stream.compute.* instance types: AppStream-*WinServer OperatingSystemVersion -MM-DD-YYYY* AppStream-RockyLinux8-*MM-DD-YYYY* AppStream-RHEL8-*MM-DD-YYYY*
		Import Image for ComputeOptimized.* instance types
Memory Optimized	Optimized for memory-intensive applications that process large amounts of data. **Note** If you plan to use WorkSpaces Applications z1d-based instances, you must provision them from images that were created from WorkSpaces Applications base images published on or after June 12, 2018.	For stream.memory.* instance types: AppStream-*WinServer OperatingSystemVersion -MM-DD-YYYY* AppStream-RockyLinux8-*MM-DD-YYYY* AppStream-RHEL8-*MM-DD-YYYY*
		Import Image for MemoryOptimized.* instance types

Instance Family	Description	Base Image Name
Graphics G4dn	Uses NVIDIA T4 GPUs to support graphics intensive applications.	For stream.graphics.g4dn.* instance types: AppStream-Graphics-G4dn-*WinServerOperatingSystemVersion* -*MM-DD-YYYY* AppStream-Graphics-G4dn-RockyLinux8-*MM-DD-YYYY* AppStream-Graphics-G4dn-RHEL8-*MM-DD-YYYY*
		Import Image for Accelerated.g4dn.* instance types
Graphics G5	Uses NVIDIA A10G GPUs to support graphics-intensive applications such as remote workstations, video rendering, and gaming, to produce high fidelity graphics in real time.	For stream.graphics.g5.* instance types: AppStream-Graphics-G5-*WinServerOperatingSystemVersion* -*MM-DD-YYYY* AppStream-Graphics-G5-RockyLinux8-*MM-DD-YYYY* AppStream-Graphics-G5-RHEL8-*MM-DD-YYYY*
		Import Image for Accelerated.g5.* instance types

Instance Family	Description	Base Image Name
Graphics G6	Powered by NVIDIA L4 Tensor Core GPUs and third generation AMD EPYC processors. • G6 provides full GPU capabilities with 1:4 vCPU and Memory ratio. • Gr6 provides full GPU capabilities with 1:8 vCPU and Memory ratio. • G6f provides fractional GPU capabilit ies with 1:4 vCPU and Memory ratio. • Gr6f provides fractional GPU capabilities with 1:8 vCPU and Memory ratio. Note: For G6f and Gr6f, use images with date 07-28-2025 or later.	For stream.graphics.g6/gr6/ g6f/gr6f.* instance types: AppStream-Graphics-G6-*WinServerOperating SystemVersion* -*MM-DD-YYYY* AppStream-Graphics-G6-RockyLinux8-*MM-DD-YYYY* AppStream-Graphics-G6-RHEL8-*MM-DD-YYYY*
		Import Image for Accelerat ed.g6/gr6/g6f/gr6f.* instance types
Graphics G6e	Powered by NVIDIA L40S Tensor Core GPUs and third generation AMD EPYC processors.	Import Image for Accelerat ed.g6e.* instance types

Instance Family	Description	Base Image Name
Graphics G7	Powered by NVIDIA RTX PRO 4500 Blackwell Server Edition GPUs to support graphics-intensive and AI inference applications.	For stream.graphics.g7.* instance types: AppStream-Graphics-G7-*WinServerOperatingSystemVersion* -*MM-DD-YYYY* AppStream-Graphics-G7-RockyLinux8-*MM-DD-YYYY* AppStream-Graphics-G7-RHEL8-*MM-DD-YYYY*
		Import Image for Accelerated.g7.* instance types
Graphics NV (unified)	A single graphics base image compatible with all supported graphics families: Graphics G4dn, G5, G6, Gr6, G6f, Gr6f, and G7. Use this image to create a fleet on any family, or to change a fleet's instance type across families. See Unified graphics images.	For stream.graphics.g4dn/g5/g6/gr6/g6f/gr6f/g7.* instance types: AppStream-Graphics-NV-*WinServerOperatingSystemVersion* -*MM-DD-YYYY* AppStream-Graphics-NV-RockyLinux8-*MM-DD-YYYY* AppStream-Graphics-NV-RHEL8-*MM-DD-YYYY*

The graphics instance families that are available for a graphics image depend on the installed graphics driver version. If the installed driver version doesn't support an instance type, that

instance type isn't available for selection. For supported driver versions, see Install NVIDIA GRID drivers.

WorkSpaces Applications instances have one 200 GB fixed-size volume, which is used for the C drive. Because WorkSpaces Applications is non-persistent, each instance's volume is immediately deleted after each user session.

For more information, see the following:

- WorkSpaces Applications Base Image and Managed Image Update Release Notes
- Amazon WorkSpaces Applications Service Quotas
- WorkSpaces Applications Pricing

Create an Amazon WorkSpaces Applications Fleet and Stack

To stream your applications, Amazon WorkSpaces Applications requires an environment that includes a fleet that is associated with a stack, as well as at least one application image. This tutorial describes the steps to set up a fleet and stack, and how to give users access to the stack. If you haven't already done so, we recommend that you try the procedures in Get Started with Amazon WorkSpaces Applications: Set Up With Sample Applications first.

If you want to create an image to use, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

If you plan to join a fleet to an Active Directory domain, configure your Active Directory domain before completing the following steps. For more information, see Using Active Directory with WorkSpaces Applications.

Tasks

- Create a Fleet in Amazon WorkSpaces Applications
- Create a Stack in Amazon WorkSpaces Applications
- Provide Access to Users in Amazon WorkSpaces Applications
- Clean Up Resources in Amazon WorkSpaces Applications

Create a Fleet in Amazon WorkSpaces Applications

Set up and create a fleet from which user applications are launched and streamed.

> **Note**
>
> To create an Always-On or On-Demand fleet, you must have an image that has applications installed to create an Always-On or On-Demand fleet that your users can stream from. To create an image, see the section called "Tutorial: Create a Custom Image by Using the Console". To create an Elastic fleet, you must have applications associated to app blocks. To create applications and app blocks for an Elastic fleet, see Applications Manager.

To set up and create a fleet

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. Choose **Get Started** if you are new to the console, or **Fleets** from the left navigation pane. Choose **Create Fleet**.
3. For **Step 1: Select fleet type**, review the details of the fleet types, choose the type of fleet to create based on your use case, and select **Next**.

   > **Note**
   >
   > The fleet type determines its immediate availability and how you pay for it. For more information, see WorkSpaces Applications Fleet Types.

4. If you chose to create an Always-On or On-Demand fleet, for **Step 2: Choose an Image**, choose an image that meets your needs and then choose **Next**.
5. If you chose to create an Elastic fleet, for **Step 2: Assign applications**, choose the applications that users can launch from this fleet.
6. For **Step 3: Configure fleet**, enter the following **details**:

   - For **Name**, enter a unique name identifier for the fleet. Special characters aren't allowed.
   - For **Display Name**, enter a name to display for the fleet (maximum of 100 characters). Special characters aren't allowed.
   - For **Description**, enter a description for the fleet (maximum of 256 characters).
   - For **Choose instance type**, choose the instance type that meets the performance requirements of your applications. All streaming instances in your fleet launch with the

instance type that you select. For more information, see WorkSpaces Applications Instance Families.

> **Note**
>
> If the image that you selected is a unified graphics image, the console lists the graphics instance families that the image supports (the families allowed by the image's graphics driver version). You can select any of these instance types, and the fleet launches on that graphics family. For more information, see Unified graphics images.

- You can use stream.* instance types for images with `type = "native"` and `type = "BYOL"`. To use any of the following instance types, you must use the Import Image option to import an EC2 AMI and create an image with `type = "custom"`.
  - GeneralPurpose.*
  - MemoryOptimized.*
  - ComputeOptimized.*
  - Accelerated.*
- Configure **storage** volumes for Always-On or On-Demand fleet instances. By default, the storage volume matches your image volume size, with a service default of 200 GB included in the hourly instance rate. You can customize your storage capacity from 200 GB up to 500 GB based on your requirements.

> **Note**
>
> Note: Storage volume size cannot be set below the image volume size. Storage capacity can be increased up to 500 GB, with additional charges applying to any storage beyond the included 200 GB. These charges apply to fleet instances regardless of their running state (both running and stopped instances).

- For Elastic fleets, for **Choose platform type**, choose the operating system that matches the requirements of your users' applications.
- For **Maximum session duration in minutes**, choose the maximum amount of time that a streaming session can remain active. If users are still connected to a streaming instance five minutes before this limit is reached, they are prompted to save any open documents

before being disconnected. After this time elapses, the instance is terminated and replaced by a new instance. The maximum session duration that you can set in the WorkSpaces Applications console is 5760 minutes (96 hours). The maximum session duration that you can set using the WorkSpaces Applications API and CLI is 432000 seconds (120 hours).

- For **Disconnect timeout in minutes**, choose the amount of time that a streaming session remains active after users disconnect. If users try to reconnect to the streaming session after a disconnection or network interruption within this time interval, they are connected to their previous session. Otherwise, they are connected to a new session with a new streaming instance. If you associate a stack with a fleet for which a redirect URL is specified, after users' streaming sessions end, the users are redirected to that URL.

  If a user ends the session by choosing **End Session** or **Logout** on the WorkSpaces Applications toolbar, the disconnect timeout does not apply. Instead, the user is prompted to save any open documents, and then immediately disconnected from the streaming instance. The instance the user was using is then terminated.

- For **Idle disconnect timeout in minutes**, choose the amount of time that users can be idle (inactive) before they are disconnected from their streaming session and the **Disconnect timeout in minutes** time interval begins. Users are notified before they are disconnected due to inactivity. If they try to reconnect to the streaming session before the time interval specified in **Disconnect timeout in minutes** has elapsed, they are connected to their previous session. Otherwise, they are connected to a new session with a new streaming instance. Setting this value to 0 disables it. When this value is disabled, users are not disconnected due to inactivity.

  > **Note**
  >
  > Users are considered idle when they stop providing keyboard or mouse input during their streaming session. For domain-joined fleets, the countdown for the idle disconnect timeout doesn't begin until users log in with their Active Directory domain password or with a smart card. File uploads and downloads, audio in, audio out, and pixels changing do not qualify as user activity. If users continue to be idle after the time interval in **Idle disconnect timeout in minutes** elapses, they are disconnected.

- For Elastic fleets, for **Max concurrent sessions**, specify the maximum number of concurrent sessions this fleet should have.

> **Note**
>
> If you get an error message that says “The maximum number of concurrent sessions for your account was exceeded," you can submit a limit increase, through the Service Quotas console at https://console.aws.amazon.com/servicequotas/. For more information, see Requesting a quota increase in the *Service Quotas User Guide*.

- **Multiple user sessions** — Choose this option if you want to provision multiple user sessions on a single instance. By default, every unique user session is served by an instance (single-session).

> **Note**
>
> Multi-session is available only on Always-on and On-demand fleets powered by a Windows operating system. Multi-session is not available on Elastic fleets or the Linux operating system.
> Only base images and managed image updates released on or after May 15, 2023 support multi-session fleets. For more details, see the section called “Base Image and Managed Image Update Release Notes”.

- **Maximum sessions per instance** — Maximum number of user sessions on an instance. You must choose this value based on your end users' application performance needs. You can also adjust the maximum sessions per instance for a fleet after it is provisioned. In that case, the existing user sessions and instances will not be impacted, but the fleet will become consistent with the new value of maximum sessions per instance. The value must be between 2 and 50. Before setting this value for your fleet, see the section called “Multi-Session Recommendations”.
- For Always-On and On-Demand fleets, for **Minimum capacity**, choose a minimum number of instances (for single-session fleets) or user sessions (for multi-session fleets) for your fleet based on the minimum number of expected concurrent users.
- For Always-On and On-Demand fleets, for **Maximum capacity**, choose a maximum number of instances (for single-session fleets) or user sessions (for multi-session fleets) for your fleet based on the maximum number of expected concurrent users.

> **Note**
>
> For multi-session, you must specify the capacity based on the number of user sessions. The service will calculate the required number of instances to be launched, based on your fleet configuration and the value of maximum sessions per instance.

- For **Stream view**, choose the WorkSpaces Applications view that is displayed to your users during their streaming sessions. Choose **Application** to display only the windows of applications opened by users. Choose **Desktop** to display the standard desktop that is provided by the operating system.

  > **Note**
  >
  > By default, WorkSpaces Applications displays only the windows of applications opened by users during their streaming sessions. To enable **Desktop** view for your users, configure your fleet to use an WorkSpaces Applications image that uses a version of the WorkSpaces Applications agent released on or after February 19, 2020.

- For **Scaling details (Advanced)**, specify the scaling policies that WorkSpaces Applications uses to increase and decrease the capacity of your fleet. Note that the size of your fleet is limited by the minimum and maximum capacity that you specified. For more information, see Fleet Auto Scaling for Amazon WorkSpaces Applications.
- For **IAM role (Advanced)**, when you apply an IAM role from your account to an WorkSpaces Applications fleet instance, you can make AWS API requests from the fleet instance without manually managing AWS credentials. To apply an IAM role, do either of the following:
  - To use an existing IAM role in your AWS account, choose the role that you want to use from the **IAM role** list. The role must be accessible from the fleet instance. For more information, see Configuring an Existing IAM Role to Use With WorkSpaces Applications Streaming Instances.
  - To create a new IAM role, choose **Create new IAM role** and follow the steps in How to Create an IAM Role to Use With WorkSpaces Applications Streaming Instances.
- For Elastic fleets, for **USB Redirection (advanced)**, you can specify up to 10 strings that specify what types of USB devices that are attached to the local device can be redirected

into the streaming session when using the Windows native client. For more information, see the section called “Qualify USB Devices for Use with Streaming Applications”.

7. Choose **Next**.
8. If you chose to create an Always-On or On-Demand fleet, for **Step 3: Choose an Image**, choose an image that meets your needs and then choose **Next**.
9. If you chose to create an Elastic fleet, for **Step 3: Assign applications**, choose the applications that users can launch from this fleet.
10. For **Step 4: Configure Network**, do the following:

    - To add internet access for fleet instances in a VPC with a public subnet, choose **Default Internet Access**. If you are providing internet access by using a NAT gateway, leave **Default Internet Access** unselected. For more information, see Internet Access.

      > **Note**
      >
      > Your VPC must provide access to Amazon Simple Storage Service (S3) if you enable features that rely on saving to an S3 bucket. For more information, see the section called “Amazon S3 VPC Endpoints”.

    - For **VPC** and **Subnet 1**, choose a VPC and subnets that have access to the network resources that your application needs. Fleets require that you specify at least two subnets that are in different Availability Zones. For more information, see Configure a VPC with Private Subnets and a NAT Gateway.

      If you don't have your own VPC and subnet, you can use the default VPC or create your own. To create your own, choose the **Create a new VPC** and **Create new subnet** links to create them. Choosing these links opens the Amazon VPC console. After you create your VPC and subnets, return to the WorkSpaces Applications console and choose the refresh icon to the left of the **Create a new VPC** and **Create new subnet** links to display them in the list. For more information, see Configure a VPC for WorkSpaces Applications.

    - For **Security group(s)**, choose up to five security groups to associate with this fleet. If you don't have your own security group and you don't want to use the default security group, choose the **Create new security group** link to create one. After you create your subnets in the Amazon VPC console, return to the WorkSpaces Applications console and choose the refresh icon to the left of the **Create new security group** link to display them in the list. For more information, see Security Groups in Amazon WorkSpaces Applications.

- For Always-On and On-Demand fleets, for **Active Directory Domain (Optional)**, choose the Active Directory and organizational unit (OU) for your streaming instance computer objects. Ensure that the network access settings you selected enable DNS resolvability and communication with your directory. For more information, see Using Active Directory with WorkSpaces Applications.

11. Choose **Next**.
12. For **Step 5: Review**, confirm the details for the fleet. To change the configuration for any section, choose **Edit** and make the needed changes. After you finish reviewing the configuration details, choose **Create**.
13. In the pricing acknowledgement dialog box, select the acknowledgement check box, and choose **Create**.

> **Note**
>
> If an error message notifies you that you don't have sufficient limits (quotas) to create the fleet, submit a limit increase request through the Service Quotas console at https://console.aws.amazon.com/servicequotas/. For more information, see Requesting a quota increase in the *Service Quotas User Guide*.

14. While your fleet is being created, the status of your fleets displays as **Starting** in the **Fleets** list. Choose the **Refresh** icon periodically to update the fleet status until the status is **Running**. You cannot associate the fleet with a stack and use it for streaming sessions until the status of the fleet is **Running**.

Create a Stack in Amazon WorkSpaces Applications

Set up and create a stack to control access to your fleet.

> **Note**
>
> You can enable Google Drive, OneDrive, and Application Settings Persistence only for stacks associated with a Windows fleet. Before you associate an existing stack with a Linux fleet, please make sure these settings are disabled.

To set up and create a stack

1. In the left navigation pane, choose **Stacks**, and then choose **Create Stack**.
2. For **Step 1: Stack Details**, Under **Stack details**, enter a unique name identifier for the stack. Optionally, you can do the following:

   - **Display name** — Enter a name to display for the stack (maximum of 100 characters).
   - **Description**— Enter a description for the stack (maximum of 256 characters).
   - **Redirect URL** — Specify a URL to which users are redirected after their streaming sessions end.
   - **Feedback URL** — Specify a URL to which users are redirected after they click the **Send Feedback** link to submit feedback about their application streaming experience. If you do not specify a URL, this link is not displayed.
   - **Fleet** — Select an existing fleet or create a new one to associate with your stack.
   - **AI agent access** — Select **Enable AI agent access** to allow AI agents to connect to streaming sessions and interact with desktop applications through the managed MCP server. When you enable agent access, Steps 2 and 3 of the stack creation wizard are configured for agent-specific settings. For more information, see Overview.
   - **Streaming Protocol Preference** — Specify the streaming protocol you'd like your stack to prefer, UDP or TCP. UDP is currently only supported in the Windows native client. For more information, see System Requirements and Feature Support (WorkSpaces Applications Client).
   - **Tags** — Choose **Add Tag**, and type the key and value for the tag. To add more tags, repeat this step. For more information, see Tagging Your Amazon WorkSpaces Applications Resources.
   - **VPC Endpoints (Advanced)** — You can create a private link, which is an interface VPC endpoint (interface endpoint), in your virtual private cloud (VPC). To start creating the interface endpoint, select **Create VPC Endpoint**. Selecting this link opens the VPC console. To finish creating the endpoint, follow steps 3 through 6 in *To create an interface endpoint*, in Tutorial: Creating and Streaming from Interface VPC Endpoints.

     After you create the interface endpoint, you can use it to keep streaming traffic within your VPC.

> **Note**
>
> Streaming from Interface VPC Endpoints is currently not supported for fleets powered by Windows Desktop operating system (such as Windows 11).

- **Embed WorkSpaces Applications (Optional)** — To embed an WorkSpaces Applications streaming session in a webpage, specify the domain to host the embedded streaming session. Embedded streaming sessions are only supported over HTTPS [TCP port 443].

> **Note**
>
> You must meet prerequisites and perform additional steps to configure embedded WorkSpaces Applications streaming sessions. For more information, see Embed Amazon WorkSpaces Applications Streaming Sessions.

3. Choose **Next.**
4. For **Step 2: Enable Storage**, you can provide persistent storage for your users by choosing one or more of the following:

   - **Home Folders** — Users can save their files to their home folder and access existing files in their home folder during application streaming sessions. For information about requirements for enabling home folders, see Enable Home Folders for Your WorkSpaces Applications Users.
   - **Google Drive for Google Workspace** — Users can link their Google Drive for Google Workspace account to WorkSpaces Applications. During application streaming sessions, they can sign in to their Google Drive account, save files to Google Drive, and access their existing files in Google Drive. You can enable Google Drive for accounts in Google Workspace domains only, not for personal Gmail accounts.

   > **Note**
   >
   > Enabling Google Drive is not supported for Linux-based stacks.
   > For multi-session fleets, Google Drive requires an image that uses an WorkSpaces Applications agent released on or after June 29, 2026. Images using Update an Image by Using Managed WorkSpaces Applications Image Updates released on or after June 29, 2026 are also supported.

> **Note**
>
> After you select **Enable Google Drive**, type the name of at least one organizational domain that is associated with your Google Workspace account. Access to Google Drive during application streaming sessions is limited to users that are in the domains that you specify. You can specify up to 10 domains. For more information about requirements for enabling Google Drive, see Enable Google Drive for Your WorkSpaces Applications Users.

- **OneDrive for Business** — Users can link their OneDrive for Business account to WorkSpaces Applications. During application streaming sessions, they can sign in to their OneDrive account, save files to OneDrive, and access their existing files in OneDrive. You can enable OneDrive for accounts in OneDrive domains only, not for personal accounts.

> **Note**
>
> Enabling OneDrive is not supported for Linux-based stacks.
> For multi-session fleets, OneDrive requires an image that uses an WorkSpaces Applications agent released on or after June 29, 2026. Images using Update an Image by Using Managed WorkSpaces Applications Image Updates released on or after June 29, 2026 are also supported.

> **Note**
>
> After you select **Enable OneDrive**, enter the name of least one organizational domain that is associated with your OneDrive account. Access to OneDrive during application streaming sessions is limited to users that are in the domains that you specify. You can specify up to 10 domains. For more information about requirements for enabling OneDrive, see Enable OneDrive for Your WorkSpaces Applications Users.

5. Choose **Next**.
6. For **Step 3: User Settings**, configure the following settings. When you're done, choose **Review**.

> **Note**
>
> If you enabled AI agent access in Step 1, Step 3 displays **Edit agent settings** instead of **User Settings**. You configure computer input, computer vision, screenshot storage, screen resolution, screen image type, and application settings persistence for agents. For more information, see Overview.

Clipboard, file transfer, print to local device, and authentication permissions:

> **Note**
>
> **Smart card sign in for Active Directory** is currently not available for multi-session fleets.

- **Clipboard** — By default, users can copy and paste data between their local device and streaming applications. You can limit Clipboard options so that users can paste data to their remote streaming session only or copy data to their local device only. You can also disable Clipboard options entirely. Users can still copy and paste between applications in their streaming session. You can choose **Copy to local device character limit** or **Paste to remote session character limit** or both to limit the amount of data that users can copy or paste when using the clipboard, either in or out of their WorkSpaces Applications streaming session. The value can be between 1 and 20,971,520 (20 MB), and defaults to the maximum value when unspecified.
- **File transfer** — By default, users can upload and download files between their local device and streaming session. You can limit file transfer options so that users can upload files to their streaming session only or download files to their local device only. You can also disable file transfer entirely.

  > **Important**
  >
  > If your users require WorkSpaces Applications file system redirection to access local drives and folders during their streaming sessions, you must enable both file upload and download. To use file system redirection, your users must have WorkSpaces

Applications client version 1.0.480 or later installed. For more information, see Enable File System Redirection for Your WorkSpaces Applications Users.

- **Print to local device** — By default, users can print to their local device from within a streaming application. When they choose **Print** in the application, they can download a .pdf file that they can print to a local printer. You can disable this option to prevent users from printing to a local device.
- **Password sign in for Active Directory** — Users can enter their Active Directory domain password to sign in to an WorkSpaces Applications streaming instance that is joined to an Active Directory domain.

  You can also enable **Smart card sign in for Active Directory**. At least one authentication must be enabled.
- **Smart card sign in for Active Directory** — Users can use a smart card reader and smart card connected to their local computer to sign in to an WorkSpaces Applications streaming instance that is joined to an Active Directory domain.

  You can also enable **Password sign in for Active Directory**. At least one authentication method must be enabled.

> **Note**
>
> **Clipboard, file transfer, and print to local device settings** — These settings control only whether users can use WorkSpaces Applications data transfer features. If your image provides access to a browser, network printer, or other remote resource, your users might be able to transfer data to or from their streaming session in other ways. **Authentication settings** — These settings control only the authentication method that can be used for Windows sign in to an WorkSpaces Applications streaming instance (fleet or image builder). They do not control the authentication method that can be used for in-session authentication, after a user signs in to a streaming instance. For information about configuration requirements for using smart cards for Windows sign in and in-session authentication, see Smart Cards. These settings are not supported for Linux-based stacks.

Time zone:

- **Set time zone automatically for remote session** — This setting syncs the time zone used for streaming to match the time zone set on the user's device. Users can override this and set their own preferred time zone.

Application settings persistence:

- **Enable Application Settings Persistence** — Users' application customizations and Windows settings are automatically saved after each streaming session and applied during the next session. These settings are saved to an Amazon Simple Storage Service (Amazon S3) bucket in your account, within the AWS Region in which application settings persistence is enabled.
- **Settings Group** — The settings group determines which saved application settings are used for a streaming session from this stack. If the same settings group is applied to another stack, both stacks use the same application settings. By default, the settings group value is the name of the stack.

> **Note**
>
> For information about requirements for enabling and administering application settings persistence, see Enable Application Settings Persistence for Your WorkSpaces Applications Users.

7. For **Step 4: Review**, confirm the details for the stack. To change the configuration for any section, choose **Edit** and make the needed changes. After you finish reviewing the configuration details, choose **Create**.

After the service sets up resources, the **Stacks** page appears. The status of your new stack appears as **Active** when it is ready to use.

Provide Access to Users in Amazon WorkSpaces Applications

After you create a stack with an associated fleet, you can provide access to users through the WorkSpaces Applications user pool, SAML 2.0 [single sign-on (SSO)], or the WorkSpaces Applications API. For more information, see User Pool Administration in Amazon WorkSpaces Applications and Amazon WorkSpaces Applications Integration with SAML 2.0.

> **Note**
>
> Users in the WorkSpaces Applications user pool can't be assigned to stacks with fleets that are joined to an Active Directory domain.

After you provide your users with access to WorkSpaces Applications, they can start WorkSpaces Applications streaming sessions by using a web browser or by using the WorkSpaces Applications client application for a supported device. If you provide access to users through the WorkSpaces Applications user pool, they must use a web browser for streaming sessions. If you use SAML 2.0 or the WorkSpaces Applications API, you can make the WorkSpaces Applications client available to them. The WorkSpaces Applications client is a native application that is designed for users who require additional functionality during their WorkSpaces Applications streaming sessions. For more information, see Provide Access Through the WorkSpaces Applications Client.

Clean Up Resources in Amazon WorkSpaces Applications

You can stop your running fleet and delete your active stack to free up resources and to avoid unintended charges to your account. We recommend stopping any unused, running fleets.

Note that you cannot delete a stack with an associated fleet.

> **Important**
>
> To delete a fleet that has an associated auto scaling policy, the IAM role used to perform the deletion must include the `application-autoscaling:DeregisterScalableTarget` permission. Without this permission, the fleet deletion will not fail, however, the service cannot deregister the auto scaling target associated with the fleet and it will remain as an orphan resource. Ensure this permission is included in the user's IAM policy before attempting to delete a fleet with active scaling policies.

To clean up your resources

1. In the navigation pane, choose **Stacks**.
2. Select the stack and choose **Actions**, **Disassociate Fleet**. In the confirmation dialog box, choose **Disassociate.**

3. In the navigation pane, choose **Fleets**.
4. Select the fleet that you want to stop, choose **Actions**, and then choose **Stop**. It takes about 5 minutes to stop a fleet.
5. When the status of the fleet is **Stopped**, choose **Actions**, **Delete**.
6. In the navigation pane, choose **Stacks**.
7. Select the stack and choose **Actions**, **Delete**.

Customize an Amazon WorkSpaces Applications Fleet to Optimize Your Users' Application Streaming Experience

By customizing WorkSpaces Applications fleet instances, you can define specific aspects of your WorkSpaces Applications environment to optimize your users' application streaming experience. For example, you can persist environment variables to dynamically pass settings across applications and set default file associations that are applied to all of your users. At a high level, customizing a fleet instance includes the following tasks:

- Connecting to an image builder and customizing it as needed.
- On the image builder, using Image Assistant to create a new image that includes your customizations.
- Creating a new fleet instance or modifying an existing one. When you configure the fleet instance, select the new customized image that you created.
- Creating a new stack or modifying an existing one and associating it with your fleet instance.

> **Note**
>
> For certain fleet customizations, in Active Directory environments, you might need to use the Group Policy Management Console (GPMC) to update Group Policy object (GPO) settings on a domain-joined computer.

Contents

- Persist Environment Variables in Amazon WorkSpaces Applications
- Set Default File Associations for Your Users in Amazon WorkSpaces Applications

- Disable Internet Explorer Enhanced Security Configuration in Amazon WorkSpaces Applications
- Change the Default Internet Explorer Home Page for Users' Streaming Sessions in Amazon WorkSpaces Applications
- User and Instance Metadata for Amazon WorkSpaces Applications Fleets

Persist Environment Variables in Amazon WorkSpaces Applications

Environment variables enable you to dynamically pass settings across applications. For example, many engineering applications rely on environment variables to specify the IP address or host name of a license server to locate and check out a license from that server.

Follow the steps in these procedures to make environment variables available across your fleet instances.

> **Note**
>
> The following instructions apply to Windows fleets only.

Contents

- Change System Environment Variables
- Change User Environment Variables
- Create an Environment Variable That is Limited in Scope

> **Note**
>
> If you are using Active Directory and Group Policy with WorkSpaces Applications, keep in mind that streaming instances must be joined to an Active Directory domain to use Group Policy for environment variables. For information about how to configure the Group Policy **Environment Variable** preference item, see Configure an Environment Variable Item in the Microsoft documentation.

Change System Environment Variables

Follow these steps to change system environment variables across your fleet instances.

To change system environment variables on an image builder

This procedure applies only to system environment variables, not user environment variables. To change user environment variables that persist across your fleet instances, perform the steps in the next procedure.

1. Connect to the image builder on which to change system environment variables and sign in with an account that has local administrator permissions. To do so, do either of the following:

   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

   > **Note**
   >
   > If the image builder that you want to connect to is joined to an Active Directory domain and your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection. For information about smart card sign in, see Smart Cards.

2. Choose the Windows **Start** button, open the context (right-click) menu for **Computer**, and then choose **Properties**.
3. In the navigation pane, choose **Advanced system settings**.
4. In **System variables**, change the environment variables that you want to persist across your fleet instances, and then choose **OK**.
5. On the image builder desktop, open Image Assistant.
6. Follow the necessary steps in Image Assistant to finish creating your image. For more information, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

   The changes to the system environment variables persist across your fleet instances and are available to streaming sessions launched from those instances.

   > **Note**
   >
   > Setting AWS CLI credentials as system environment variables might prevent WorkSpaces Applications from creating the image.

Change User Environment Variables

Follow these steps to change user environment variables across your fleet instances.

To change user environment variables

1. Connect to the image builder on which to change system environment variables and sign in as a **Template User**. To do so, do either of the following:

   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

   > **Note**
   >
   > If the image builder that you want to connect to is joined to an Active Directory domain and your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection. For information about smart card sign in, see Smart Cards.

   Template User lets you create default application and Windows settings for your users. For more information, see "Creating Default Application and Windows Settings for Your WorkSpaces Applications Users" in Default Application and Windows Settings and Application Launch Performance in Amazon WorkSpaces Applications.
2. On the image builder, choose the Windows **Start** button, **Control Panel**, **User Accounts**.
3. Choose **User Accounts** again. In the left navigation pane, choose **Change my environment variables.**
4. Under **User environment variables** for **DefaultProfileUser**, set or create the user environment variables as needed, then choose **OK**.
5. This disconnects your current session and opens the login menu. Log in to the image builder by doing either of the following:

   - If your image builder is not joined to an Active Directory domain, on the **Local User** tab, choose **Administrator**.
   - If your image builder is joined to an Active Directory domain, choose the **Directory User** tab, and log in as a domain user who has local administrator permissions on the image builder.
6. On the image builder desktop, open Image Assistant.

7. Follow the necessary steps in Image Assistant to finish creating your image. For more information, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

Create an Environment Variable That is Limited in Scope

Follow these steps to create an environment variable that is limited in scope to the processes that are spawned off the script. This approach is useful when you need to use the same environment variable name with different values for different applications. For example, if you have two different applications that use the environment variable "LIC_SERVER", but each application has a different value for "LIC_SERVER".

To create an environment variable that is limited in scope

1. Connect to the image builder on which to create an environment variable that is limited in scope and sign in with an account that has local administrator permissions. To do so, do either of the following:

   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

   > **Note**
   >
   > If the image builder that you want to connect to is joined to an Active Directory domain and your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection. For information about smart card sign in, see Smart Cards.

2. Create a child folder of C:\ drive for the script (for example, C:\Scripts).
3. Open Notepad to create the new script, and enter the following lines:

   ```
   set variable=value

   start " " "C:\path\to\application.exe"
   ```

   Where:

   variable is the variable name to be used

value is the value for the given variable name

> **Note**
>
> If the application path includes spaces, the entire string must be encapsulated within quotation marks. For example:
> `start " " "C:\Program Files\application.exe"`

4. Choose **File**, **Save**. Name the file and save it with the .bat extension to C:\Scripts. For example, name the file LaunchApp.bat.
5. If needed, repeat steps 4 and 5 to create a script for each additional application that requires its own environment variable and values.
6. On the image builder desktop, start Image Assistant.
7. Choose **Add App**, navigate to C:\Scripts, and select one of the scripts that you created in step 5. Choose **Open**.
8. In the **App Launch Settings** dialog box, keep or change the settings as needed. When you're done, choose **Save**.
9. If you created multiple scripts, repeat steps 8 and 9 for each script.
10. Follow the necessary steps in Image Assistant to finish creating your image. For more information, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

    The environment variable and specific value are now available for processes that are run from the script. Other processes cannot access this variable and value.

Set Default File Associations for Your Users in Amazon WorkSpaces Applications

The associations for application file extensions are set on a per-user basis and so are not automatically applied to all users who launch WorkSpaces Applications streaming sessions. For example, if you set Adobe Reader as the default application for .pdf files on your image builder, this change is not applied to your users.

> **Note**
>
> The following steps apply to Windows fleets only.

> **Note**
>
> The following steps must be performed on an image builder that is joined to an Active Directory domain. In addition, your fleet must be joined to an Active Directory domain. Otherwise, the default file associations that you set are not applied.

To set default file associations for your users

1. Connect to the image builder on which to set default file associations and sign in with a domain account that has local administrator permissions on the image builder. To do so, do either of the following:

   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

   > **Note**
   >
   > If your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection. For information about smart card sign in, see Smart Cards.

2. Set default file associations as needed.
3. Open the Windows command prompt as an administrator.
4. At the command prompt, type the following command to export the image builder file associations as an XML file, and then press ENTER:

   ```
   dism.exe /online /export-DefaultAppAssociations:c:
   \default_associations.xml
   ```

   If you receive an error message stating that you cannot service a running 64-bit operating system with a 32-bit version of DISM, close the command prompt window. Open File Explorer,

browse to C:\Windows\System32, right-click cmd.exe, choose **Run as Administrator**, and run the command again.

5. You can use either Local Group Policy Editor or the GPMC to set a default associations configuration file:

   - Local Group Policy Editor:

     On your image builder, open the command prompt as an administrator, type `gpedit.msc`, and then press ENTER.

     In the console tree, under **Computer Configuration**, expand **Administrative Templates**, **Windows Components**, and then choose **File Explorer**.

   - GPMC:

     In your directory or on a domain controller, open the command prompt as an administrator, type `gpmc.msc`, and then press ENTER.

     In the left console tree, select the OU in which you want to create a new GPO, or use an existing GPO, and then do either of the following:

     - Create a new GPO by opening the context (right-click) menu and choosing **Create a GPO in this domain, Link it here**. For **Name**, provide a descriptive name for this GPO.
     - Select an existing GPO.

     Open the context menu for the GPO, and choose **Edit**.

     Under **User Configuration**, expand **Policies**, **Administrative Templates**, **Windows Components**, and then choose **File Explorer**.

6. Double-click **Set a default associations configuration file**.
7. In the **Set a default associations configuration file properties** dialog box, choose **Enabled**, and do one of the following:

   - If you are using Local Group Policy Editor, enter this path: `c:\default_associations.xml`.
   - If you are using the GPMC, enter a network path. For example, `\\networkshare\default_associations.xml`.

8. Choose **Apply**, **OK**.
9. Close Local Group Policy Editor or the GPMC.

10. On the image builder desktop, open Image Assistant.
11. Follow the necessary steps in Image Assistant to finish creating your image. For more information, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

    The file associations that you configured are applied to the fleet instances and user streaming sessions that are launched from those instances.

Disable Internet Explorer Enhanced Security Configuration in Amazon WorkSpaces Applications

Internet Explorer Enhanced Security Configuration (ESC) places servers and Internet Explorer in a configuration that limits exposure to the internet. However, this configuration can impact the WorkSpaces Applications end user experience. Users who are connected to WorkSpaces Applications streaming sessions may find that websites do not display or perform as expected when:

- Internet Explorer ESC is enabled on fleet instances from which users' streaming sessions are launched
- Users run Internet Explorer during their streaming sessions
- Applications use Internet Explorer to load data

> **Note**
>
> The following steps apply to Windows fleets only.

To disable Internet Explorer Enhanced Security Configuration

1. Connect to the image builder on which to disable Internet Explorer ESC and sign in with an account that has local administrator permissions. To do so, do either of the following:
   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

> **Note**
>
> If the image builder that you want to connect to is joined to an Active Directory domain and your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection. For information about smart card sign in, see Smart Cards.

2. On the image builder, disable Internet Explorer ESC by doing the following:

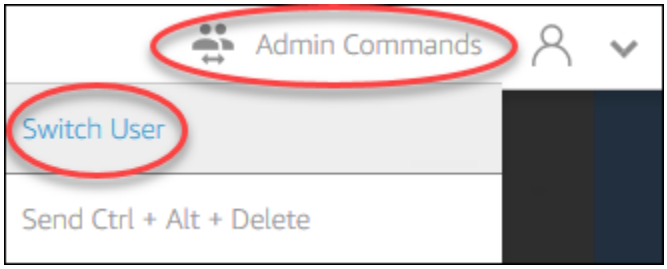
   a. Open Server Manager. Choose the Windows **Start** button, and then choose **Server Manager**.

   b. In the left navigation pane, choose **Local Server**.

   c. In the right properties pane, choose the **On** link next to IE Enhanced Security Configuration.

   d. In the **Internet Explorer Enhanced Configuration** dialog box, choose the **Off** option under **Administrators** and **Users**, then choose **OK**.

3. In the upper right area of the image builder desktop, choose **Admin Commands**, **Switch User**.



4. This disconnects your current session and opens the login menu. Log in to the image builder by doing either of the following:

   - If your image builder is not joined to an Active Directory domain, on the **Local User** tab, choose **Template User**.
   - If your image builder is joined to an Active Directory domain, choose the **Directory User** tab, and log in as a domain user who does not have local administrator permissions on the image builder.

5. Open Internet Explorer and reset your settings by doing the following:

   a. In the upper right area of the Internet Explorer browser window, choose the **Tools** icon, then choose **Internet options**.

b. Choose the **Advanced** tab, then choose **Reset**.

c. When prompted to confirm your choice, choose **Reset** again.

d. When the **Reset Internet Explorer Settings** message displays, choose **Close**.

6. Reboot image builder.
7. Choose **Admin Commands**, **Switch User**, and then do either of the following:

   - If your image builder is not joined to an Active Directory domain, on the **Local User** tab, choose **Administrator**.
   - If your image builder is joined to an Active Directory domain, choose the **Directory User** tab, and log in with the same domain account that you used in step 4.

8. On the image builder desktop, open Image Assistant.
9. In **Step 2. Configure Apps**, choose **Save settings**.
10. Follow the necessary steps in Image Assistant to finish creating your image. For more information, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

Change the Default Internet Explorer Home Page for Users' Streaming Sessions in Amazon WorkSpaces Applications

You can use Group Policy to change the default Internet Explorer home page for users' streaming sessions. Alternatively, if you do not have Group Policy in your environment or prefer not to use Group Policy, you can use the WorkSpaces Applications Template User account instead.

> **Note**
>
> The following steps apply to Windows fleets only.

Contents

- Use Group Policy to Change the Default Internet Explorer Home Page
- Use the WorkSpaces Applications Template User Account to Change the Default Internet Explorer Home Page

Use Group Policy to Change the Default Internet Explorer Home Page

In Active Directory environments, you use the Group Policy Management (GPMC) MMC-snap-in to set a default home page that users can't change. If Active Directory is not in your environment, you can use Local Group Policy Editor to perform this task. To set a home page that users can change, you must use the GPMC.

To use the GPMC, do the following first:

- Obtain access to a computer or an EC2 instance that is joined to your domain.
- Install the GPMC. For more information, see Installing or Removing Remote Server Administration Tools for Windows 7 in the Microsoft documentation.
- Log in as a domain user with permissions to create GPOs. Link GPOs to the appropriate organizational units (OUs).

To change the default Internet Explorer home page by using a Group Policy administrative template

You can use a Group Policy administrative template to set a default home page that users can't change. For more information about administrative templates, see Edit Administrative Template Policy Settings in the Microsoft documentation.

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. If you are not using Active Directory in your environment, open Local Group Policy Editor. If you are using Active Directory, open the GPMC. Locate the **Scripts (Logon\Logoff)** policy setting:

   - Local Group Policy Editor:

     On your image builder, open the command prompt as an administrator, type `gpedit.msc`, and then press ENTER.

     Under **User Configuration**, expand **Administrative Templates**, **Windows Components**, and then choose **Internet Explorer**.
   - GPMC:

     In your directory or on a domain controller, open the command prompt as an administrator, type `gpmc.msc`, and then press ENTER.

In the left console tree, select the OU in which you want to create a new GPO, or use an existing GPO, and then do either of the following: :

- Create a new GPO by opening the context (right-click) menu and choosing **Create a GPO in this domain, Link it here**. For **Name**, provide a descriptive name for this GPO.
- Select an existing GPO.

Open the context menu for the GPO, and choose **Edit**.

Under **User Configuration**, expand **Policies**, **Administrative Templates**, **Windows Components**, and then choose **Internet Explorer**.

3. Double-click **Disable changing home page settings**, choose **Enabled**, and in **Home Page**, enter a URL.
4. Choose **Apply**, **OK**.
5. Close Local Group Policy Editor or the GPMC.

To change the default Internet Explorer home page by using Group Policy preferences

You can use Group Policy preferences to set a default home page that users can change. For more information about working with Group Policy preferences, see Configure a Registry Item and Group Policy Preferences Getting Started Guide in the Microsoft documentation.

1. In your directory or on a domain controller, open the command prompt as an administrator, type `gpmc.msc`, and then press ENTER.
2. In the left console tree, select the OU in which you want to create a new GPO, or use an existing GPO, and then do either of the following:

   - Create a new GPO by opening the context (right-click) menu and choosing **Create a GPO in this domain, Link it here**. For **Name**, provide a descriptive name for this GPO.
   - Select an existing GPO.
3. Open the context menu for the GPO, and choose **Edit**.
4. Under **User Configuration**, expand **Preferences**, and then choose **Windows Settings**.
5. Open the context (right-click) menu for **Registry** and choose **New**, **Registry Item**.
6. In the **New Registry Properties** dialog box, specify the following registry settings for Group Policy to configure:

- For **Action**, choose **Update**.
- For **Hive**, choose **HKEY_CURRENT_USER**.
- For **Key Path**, browse to and select HKEY_CURRENT_USER\SOFWARE\Microsoft\Internet Explorer\Main.
- For **Value Name**, enter **Start Page**.
- For **Value Data**, enter your home page URL.

7. On the **Common** tab, choose **Apply Once**, **Do not Re-Apply**.

> **Note**
>
> To enable your users to choose the **Use Default** button in their Internet Explorer browser settings and reset their default home page to your company home page, you can also set a value for Default_Page_URL without choosing **Apply Once** and **Do not Re-Apply**.

8. Choose **OK** and close the GPMC.

Use the WorkSpaces Applications Template User Account to Change the Default Internet Explorer Home Page

Follow these steps to use the **Template User** account to change the default Internet Explorer home page.

To change the default Internet Explorer Home page by using the Template User account

1. Connect to the image builder on which to change the default Internet Explorer home page and sign in with the **Template User** account. To do so, do either of the following:

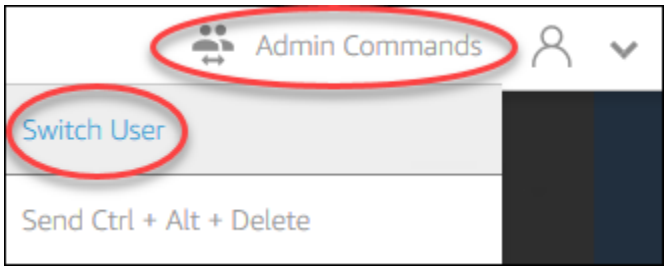
   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

   > **Note**
   >
   > If the image builder that you want to connect to is joined to an Active Directory domain and your organization requires smart card sign in, you must create a

streaming URL and use the WorkSpaces Applications client for the connection. For information about smart card sign in, see Smart Cards.

Template User lets you create default application and Windows settings for your users. For more information, see "Creating Default Application and Windows Settings for Your WorkSpaces Applications Users" in Default Application and Windows Settings and Application Launch Performance in Amazon WorkSpaces Applications.

2. Open Internet Explorer and complete the necessary steps to change the default home page.
3. In the upper right area of the image builder desktop, choose **Admin Commands**, **Switch User**.



4. This disconnects your current session and opens the login menu. Log in to the image builder by doing either of the following:
   - If your image builder is not joined to an Active Directory domain, on the **Local User** tab, choose **Administrator**.
   - If your image builder is joined to an Active Directory domain, choose the **Directory User** tab, and log in as a domain user who has local administrator permissions on the image builder.
5. On the image builder desktop, open Image Assistant.
6. Follow the necessary steps in Image Assistant to finish creating your image. For more information, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

User and Instance Metadata for Amazon WorkSpaces Applications Fleets

WorkSpaces Applications fleet instances have user and instance metadata available through Windows environment variables. You can use the following environment variables in your applications and scripts to modify your environment based on the fleet instance details.

Environment Variable	Context	Description
AppStream_Stack_Name	User	The name of the stack from which the streaming session started.
AppStream_User_Access_Mode	User	The access mode used to manage user access to the stream. The available values are **`custom`**, **`userpool`**, or **`saml`**.
AppStream_Session_Reservation_DateTime	User	The date and time when the user's streaming session started.
AppStream_UserName	User	The user name associated with the user.
AppStream_Session_ID	User	The session identifier for the user's streaming session.
APPSTREAM_SESSION_CONTEXT	Machine	Contains the parameters passed to your streaming application when a session is started. For more information, see Session Context in Amazon WorkSpaces Applications. **Note** This environment variable is only available after the first application launch.
AppStream_Image_Arn	Machine	The ARN of the image that was used to create the streaming instance.
AppStream_Instance_Type	Machine	The streaming instance's type. For example, **`stream.standard.medium`**.
AppStream_Resource_Type	Machine	The type of WorkSpaces Applications resource. The value is either **`fleet`** or **`image-builder`**.

Environment Variable	Context	Description
AppStream _Resource _Name	Machine	The fleet's name.

On Linux fleet instances, these environment variables are exported through the following profile.d scripts:

- **User environment variables** in /etc/profile.d/appstream_user_vars.sh
- **System environment variables** in /etc/profile.d/appstream_system_vars.sh

To access the environment variables, you must explicitly source these files in your applications.

Update an Amazon WorkSpaces Applications Fleet

You can make updates to an existing WorkSpaces Applications fleet.

When you create a new WorkSpaces Applications image, you must update your Always-On and On-Demand fleets to make the applications and data on the new image available to users. If your update is minor (for example, patching applications or the operating system), you can update your running fleet. When new streaming instances are created, they are created from the updated image. Changing the image on a running fleet does not disrupt users who have active streaming sessions. Unused streaming instances are replaced periodically, while streaming instances that users are connected to are terminated after the streaming sessions are finished.

You can update a fleet with a new image that runs the same operating system when the fleet is in a **Running** or **Stopped** state. However, you can update a fleet with a new image that runs a different operating system only when the fleet is in a **Stopped** state.

> **Note**
>
> The application catalog that WorkSpaces Applications displays to users is based on the current image that is associated with the fleet. If the updated image contains applications

that are not specified in the older image, the applications may not launch if the user streams from an instance that is based on the older image.

Topics

- Update a Fleet with a New Image in Amazon WorkSpaces Applications
- Manage Applications Associated to an Elastic Fleet in Amazon WorkSpaces Applications

Update a Fleet with a New Image in Amazon WorkSpaces Applications

To apply operating system updates or make new applications available to users, create a new image that has these changes. Then, update the fleet with the new image.

To update an WorkSpaces Applications fleet with a new image

1. Connect to the image builder that you want to use and sign in with an account that has local administrator permissions on the image builder. To do so, do either of the following:

   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

   > **Note**
   >
   > If your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection. For information about smart card sign in, see Smart Cards.

2. Do either or both of the following as required:

   - Install updates to the operating system.
   - Install applications.

   If an application requires the Windows operating system to restart, let it do so. Before the operating system restarts, you are disconnected from your image builder. After the restart is complete, connect to the image builder again, then finish installing the application.

3. On the image builder desktop, open Image Assistant.

4. Follow the necessary steps in Image Assistant to finish creating your image. For more information, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

   After the image status changes to **Available**, you can update the fleet with your new image.
5. In the left navigation pane, choose **Fleets**.
6. Select the fleet that you want to update with the new image.
7. On the fleet detail page, choose the **Image** tab.
8. In the **Image details** panel, choose **Edit**. This opens the **Edit: Fleet image** page.
9. Select the new image from the **Images** list. You can use the **Filter by attribute or keyword** box to search; the list is paginated.
10. Choose **Save** to apply the change.

> **Note**
>
> The fleet does not need to be in the **Stopped** state to change its image. Optionally, stop the fleet first if you prefer.

> **Note**
>
> All existing instances and user sessions will continue to run with the old image, but all new instance launches will spin up from the new image. For multi-session fleets, it is possible that instances keep running with the older image for a longer duration of time because the service will not terminate an instance if there is an active session on the instance, and if user sessions keep getting provisioned on these instances, it is possible the instance will continue to run with the old image. To get rid of long-running instances on multi-session fleets, evaluate the option of putting them in drain mode. To learn more, refer to the section called “Manage Multi-Session Fleet Instances”.

You can also change the instance type of an existing fleet. By default, you can change a fleet's instance type only within the same instance family. If the fleet uses a unified graphics image, you can change the instance type across supported graphics families by using the `UpdateFleet` operation when the fleet is stopped. For more information, see Unified graphics images.

To change the instance type of a fleet that uses a unified graphics image

1. In the left navigation pane, choose **Fleets**.
2. Select the fleet that you want to change.
3. Choose **Actions**, **Stop** to stop the fleet. You must stop the fleet before you can change its instance type across graphics families.
4. After the fleet is in the **Stopped** state, choose **Actions**, **Edit**.
5. For **Instance Type**, select the new instance type (for example, stream.graphics.g7.4xlarge).
6. Choose **Save**.
7. Choose **Actions**, **Start** to start the fleet. New instances launch on the selected graphics family.

For example, suppose that the fleet my-cad-fleet uses a unified graphics image and runs on stream.graphics.g6.4xlarge. To move the fleet to the Graphics G7 family, stop my-cad-fleet, edit the fleet, select stream.graphics.g7.4xlarge for the instance type, save the change, and start the fleet. New instances then launch on the Graphics G7 family. You can also change the instance type by using the UpdateFleet API operation or the update-fleet AWS CLI command.

Manage Applications Associated to an Elastic Fleet in Amazon WorkSpaces Applications

You can associate and disassociate applications from an Elastic fleet at any time. Changes to the applications associated to an Elastic fleet are visible to users currently streaming from the fleet, but may not take effect. For example, if you disassociate an application from a fleet, it will be removed from the application catalog, but the virtual hard disk will remain mounted to existing streaming sessions.

To manage applications associated to an Elastic fleet

1. Open the WorkSpaces Applications console.
2. In the left navigation pane, choose **Fleets**, then either select the name of the fleet, or select the fleet radio button, then choose **View details**.
3. To associate a new application to the fleet, choose **Associate** in **Assigned applications**, select the application to be associated, and choose **Associate**.
4. To disassociate an existing application from the fleet, select the application to disassociate, choose **Disassociate**, and confirm that you want to disassociate the selected application by choosing **Disassociate**.

Fleet Auto Scaling for Amazon WorkSpaces Applications

Fleet Auto Scaling lets you change the size of your WorkSpaces Applications Always-On or On-Demand fleet automatically to match the supply of available instances to user demand. The size of your fleet determines the number of users who can stream concurrently. For a multi-session fleet, more than one user can use a single instance. For a non multi-session fleet, one instance is required for each user session. You can specify your fleet capacity in terms of instances (for single-session fleets) and user sessions (for multi-session fleets). Based on your fleet configurations and auto scaling policies, the required number of instances will be made available. You can define scaling policies that adjust the size of your fleet automatically based on a variety of utilization metrics, and optimize the number of available instances to match user demand. You can also choose to turn off automatic scaling and make the fleet run at a fixed size.

Note

Elastic fleet capacity is automatically managed by WorkSpaces Applications for you. You do not need to create auto scaling rules to manage the number of fleet streaming instances that are available for Elastic fleets.

Note

As you develop your plan for WorkSpaces Applications fleet scaling, make sure that your network configuration meets your requirements.

Before you can use Fleet Auto Scaling, Application Auto Scaling needs permissions to access Amazon CloudWatch alarms and WorkSpaces Applications fleets. For more information, see Using AWS Managed Policies and Linked Roles to Manage Administrator Access to WorkSpaces Applications Resources and Using IAM Policies to Manage Administrator Access to Application Auto Scaling.

Note

When you use scaling, you work with the Application Auto Scaling API. For Fleet Auto Scaling to work correctly for WorkSpaces Applications, Application Auto Scaling requires permission to describe and update your WorkSpaces Applications fleets and describe

your Amazon CloudWatch alarms, and permissions to modify your fleet capacity on your behalf. For more information, see Roles Required for WorkSpaces Applications, Application Auto Scaling, and AWS Certificate Manager Private CA and Using IAM Policies to Manage Administrator Access to Application Auto Scaling.

The following topics provide information to help you understand and use WorkSpaces Applications Fleet Auto Scaling.

Contents

- Scaling Concepts for Amazon WorkSpaces Applications
- Managing Fleet Scaling Using the Amazon WorkSpaces Applications Console
- Managing Fleet Scaling Using the AWS CLI for Amazon WorkSpaces Applications
- Additional Resources for Auto Scaling Amazon WorkSpaces Applications

Scaling Concepts for Amazon WorkSpaces Applications

WorkSpaces Applications scaling is provided by Application Auto Scaling. For more information, see the Application Auto Scaling API Reference.

For step-by-step guidance for working with WorkSpaces Applications Fleet Auto Scaling, see Scaling Your Desktop Application Streams with Amazon AppStream 2.0 in the *AWS Compute Blog*.

To use Fleet Auto Scaling effectively, you must understand the following terms and concepts.

Multi-session vs. Single-session

In a single-session scenario, each user session has its own dedicated instance. In a multi-session mode, more than one user session can be provisioned on an instance. Fleet capacity and auto scaling policies must be configured in terms of user sessions, and the service will calculate and launch the required number of instances.

Minimum Capacity/Minimum User Sessions for fleet

The minimum number of instances (for single session fleets) or user sessions (for multi-session fleets). The number of instances (for single session fleets) or user sessions (for multi-session fleets) can't be below this value, and scaling policies will not scale your fleet below this value. For example, in a single-session scenario, if you set the minimum capacity for a fleet to 2,

your fleet will never have less than 2 instances. Similarly, in a multi-session scenario, with the maximum number of sessions on an instance set to 5, if you set the minimum capacity for a fleet to 12, your fleet will never have less than roundup (12/5) = 3 instances.

Maximum Capacity/Maximum User Sessions for fleet

The maximum number of instances (for single session fleets) or user sessions (for multi-session fleets). The number of instances (for single session fleets) or user sessions (for multi-session fleets) can't be above this value, and scaling policies will not scale your fleet above this value. For example, in a single-session scenario, if you set the maximum capacity for a fleet to 10, your fleet will never have more than 10 instances. Similarly, in a multi-session scenario, with maximum number of sessions on an instance set to 5, if you set the maximum capacity for a fleet to 52, your fleet will never have more than roundup (52/5) = 11 instances.

Desired Capacity

The total number of instances (for single session fleets) or user sessions (for multi-session fleets) that are either running or pending. This value represents the total number of concurrent streaming sessions that your fleet can support in a steady state. To set the value for **Desired Capacity**, edit **Fleet Details**. We do not recommend changing the **Desired Capacity** value manually when you use **Scaling Policies**.

If the value of **Desired Capacity** is set below the value of **Minimum Capacity** and a scale-out activity is triggered, Application Auto Scaling scales the **Desired Capacity** value up to the value of **Minimum Capacity** and then continues to scale out as required, based on the scaling policy. However, in this case, a scale-in activity does not adjust **Desired Capacity**, because it is already below the **Minimum Capacity** value.

If the value of **Desired Capacity** is set above the value of **Maximum Capacity** and a scale-in activity is triggered, Application Auto Scaling scales the **Desired Capacity** value down to the value of **Maximum Capacity** and then continues to scale in as required, based on the scaling policy. However, in this case, a scale-out activity does not adjust **Desired Capacity**, because it is already above the **Maximum Capacity** value.

Scaling Policy Action

The action that scaling policies perform on your fleet when the **Scaling Policy Condition** is met. You can choose an action based on **% capacity** or **number of instance(s)** (for single session fleets) or **user sessions** (for multi-session fleets). For example, if **Current Capacity** is 4 and **Scaling Policy Action** is set to "Add 25% capacity", **Desired Capacity** is increased will be set to 5 when **Scaling Policy Condition** is met.

Scaling Policy Condition

The condition that triggers the action set in **Scaling Policy Action**. This condition includes a scaling policy metric, a comparison operator, and a threshold. For example, to scale a fleet if the utilization of the fleet is greater than 50%, your scaling policy condition should be "If Capacity Utilization > 50%".

Scaling Policy Metric

Your scaling policy is based on this metric. The following metrics are available for scaling policies:

Capacity Utilization

The percentage of instances in a fleet that are being used. You can use this metric to scale your fleet based on usage of the fleet. For example, **Scaling Policy Condition**: "If Capacity Utilization < 25%" perform **Scaling Policy Action**: "Remove 25 % capacity".

Available Capacity

The number of instances (for single session fleets) or user sessions (for multi-session fleets) in your fleet that are available for users. You can use this metric to maintain a buffer in your capacity available for users to start streaming sessions. For example, **Scaling Policy Condition**: "If Available Capacity < 5" perform **Scaling Policy Action**: "Add 5 instance(s) (for single session fleets) or user session(s) (for multi-session fleets)".

Insufficient Capacity Error

The number of session requests rejected due to lack of capacity. You can use this metric to provision new instances for users who can't start streaming sessions due to lack of capacity. For example, **Scaling Policy Condition**: "If Insufficient Capacity Error > 0" perform **Scaling Policy Action**: "Add 1 instance(s) (fr single session fleets) or user session(s) (for multi-session fleets)".

Managing Fleet Scaling Using the Amazon WorkSpaces Applications Console

You can set up and manage fleet scaling by using the WorkSpaces Applications console in either of the following two ways: During fleet creation, or any time, by using the **Fleets** tab. Two default scaling policies are associated with newly created fleets after launch. You can edit these policies on

the **Scaling Policies** tab in the WorkSpaces Applications console. For more information, see Create a Fleet in Amazon WorkSpaces Applications.

For user environments that vary in number, define scaling policies to control how scaling responds to demand. If you expect a fixed number of users or have other reasons for disabling scaling, you can set your fleet with a fixed number of instances ro user sessions.

To set a fleet scaling policy using the console

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the navigation pane, choose **Fleets**.
3. Select the fleet and then choose **Scaling Policies**.
4. Edit existing policies by choosing the edit icon next to each value. Set the desired values in the edit field and choose **Update**. The policy changes go into effect within a few minutes.
5. Add (create) new policies using the **Add Policy** link. Set the desired values in the edit field and choose **Create**. The new policy goes into effect within a few minutes.

You can use the **Fleet Usage** tab to monitor the effects of your scaling policy changes. The following is an example usage graph of scaling activity when five users connect to the fleet and then disconnect. This example is from a fleet using the following scaling policy values:

- Minimum Capacity = 10
- Maximum Capacity = 50
- Scale Out = Add 5 instances (for single session fleets) or user sessions (for multi-session fleets) if Capacity Utilization > 75%
- Scale In = Remove 6 instances (for single session fleets) or user sessions (for multi-session fleets) if Capacity Utilization < 25%

> **Note**
>
> The above policy is applicable in both single-session and multi-session scenarios. In a single session scenario, 5 new instances will be launched during a scale out event, and 4 instances will be reclaimed during the scale down event. In a multi-session scenario, with the maximum sessions per instance = 4, the scale out event will trigger a launch of roundup (add 5 user sessions/ maximum sessions per instance 4) = 2 instances. During a

scale in event, services will reclaim roundup (remove 6 user sessions/maximum sessions per instance 4) = 2 instances. Instances with running user sessions will not be reclaimed. Only instances with no user sessions running will be reclaimed.

To set a fixed capacity fleet using the console

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the navigation pane, choose **Fleets**.
3. Select the fleet.
4. For **Scaling Policies**, remove all policies associated with the fleet.
5. For **Fleet Details**, edit the fleet to set **Desired Capacity**.

The fixed fleet has constant capacity based on the value that you specified as **Desired Capacity**. Note that a fixed fleet has the desired number of instances available at all times and the fleet must be stopped to stop billing costs for that fleet.

Important

To delete a fleet that has an associated auto scaling policy, the IAM role used to perform the deletion must include the `application-autoscaling:DeregisterScalableTarget` permission. Without this permission, the fleet deletion will not fail, however, the service cannot deregister the auto scaling target associated with the fleet and it will remain as an orphan resource. Ensure this permission is included in the user's IAM policy before attempting to delete a fleet with active scaling policies.

Managing Fleet Scaling Using the AWS CLI for Amazon WorkSpaces Applications

You can set up and manage fleet scaling by using the AWS Command Line Interface (AWS CLI). For more advanced features such as setting up multiple scaling policies or setting scale-in and scale-out cooldown times, use the AWS CLI. Before running scaling policy commands, you must register your fleet as a scalable target. To do so, use the following register-scalable-target command:

```
aws application-autoscaling register-scalable-target
  --service-namespace appstream \
  --resource-id fleet/fleetname \
  --scalable-dimension appstream:fleet:DesiredCapacity \
  --min-capacity 1 --max-capacity 5
```

Examples

- Example 1: Applying a Scaling Policy Based on Capacity Utilization
- Example 2: Applying a Scaling Policy Based on Insufficient Capacity Errors
- Example 3: Applying a Scaling Policy Based on Low Capacity Utilization
- Example 4: Change the Fleet Capacity Based on a Schedule
- Example 5: Applying a Target Tracking Scaling Policy

Example 1: Applying a Scaling Policy Based on Capacity Utilization

This AWS CLI example sets up a scaling policy that scales out a fleet by 25% if Utilization >= 75%.

The following put-scaling-policy command defines a utilization-based scaling policy:

```
aws application-autoscaling put-scaling-policy --cli-input-json file://scale-out-
utilization.json
```

The contents of the file `scale-out-utilization.json` are as follows:

```
{
    "PolicyName": "policyname",
    "ServiceNamespace": "appstream",
    "ResourceId": "fleet/fleetname",
    "ScalableDimension": "appstream:fleet:DesiredCapacity",
    "PolicyType": "StepScaling",
    "StepScalingPolicyConfiguration": {
        "AdjustmentType": "PercentChangeInCapacity",
        "StepAdjustments": [
            {
                "MetricIntervalLowerBound": 0,
                "ScalingAdjustment": 25
            }
        ],
        "Cooldown": 120
```

```
    }
}
```

If the command is successful, the output is similar to the following, although some details are unique to your account and Region. In this example, the policy identifier is `e3425d21-16f0-d701-89fb-12f98dac64af`.

```
{"PolicyARN": "arn:aws:autoscaling:us-west-2:123456789012:scalingPolicy:e3425d21-16f0-
d701-89fb-12f98dac64af:resource/appstream/fleet/SampleFleetName:policyName/scale-out-
utilization-policy"}
```

Now, set up a CloudWatch alarm for this policy. Use the names, Region, account number, and policy identifier that apply to you. You can use the policy ARN returned by the previous command for the `--alarm-actions` parameter.

```
aws cloudwatch put-metric-alarm
--alarm-name alarmname \
--alarm-description "Alarm when Capacity Utilization exceeds 75 percent" \
--metric-name CapacityUtilization \
--namespace AWS/AppStream \
--statistic Average \
--period 300 \
--threshold 75 \
--comparison-operator GreaterThanOrEqualToThreshold \
--dimensions "Name=Fleet,Value=fleetname" \
--evaluation-periods 1 --unit Percent \
--alarm-actions "arn:aws:autoscaling:your-region-code:account-
number-without-hyphens:scalingPolicy:policyid:resource/appstream/
fleet/fleetname:policyName/policyname"
```

Example 2: Applying a Scaling Policy Based on Insufficient Capacity Errors

This AWS CLI example sets up a scaling policy that scales out the fleet by 1 if the fleet returns an `InsufficientCapacityError` error.

The following command defines a insufficient capacity-based scaling policy:

```
aws application-autoscaling put-scaling-policy --cli-input-json file://scale-out-
capacity.json
```

The contents of the file `scale-out-capacity.json` are as follows:

```
{
    "PolicyName": "policyname",
    "ServiceNamespace": "appstream",
    "ResourceId": "fleet/fleetname",
    "ScalableDimension": "appstream:fleet:DesiredCapacity",
    "PolicyType": "StepScaling",
    "StepScalingPolicyConfiguration": {
        "AdjustmentType": "ChangeInCapacity",
        "StepAdjustments": [
            {
                "MetricIntervalLowerBound": 0,
                "ScalingAdjustment": 1
            }
        ],
        "Cooldown": 120
    }
}
```

If the command is successful, the output is similar to the following, although some details are unique to your account and Region. In this example, the policy identifier is `f4495f21-0650-470c-88e6-0f393adb64fc`.

```
{"PolicyARN": "arn:aws:autoscaling:us-
west-2:123456789012:scalingPolicy:f4495f21-0650-470c-88e6-0f393adb64fc:resource/
appstream/fleet/SampleFleetName:policyName/scale-out-insufficient-capacity-policy"}
```

Now, set up a CloudWatch alarm for this policy. Use the names, Region, account number, and policy identifier that apply to you. You can use the policy ARN returned by the previous command for the `--alarm-actions` parameter.

```
aws cloudwatch put-metric-alarm
--alarm-name alarmname \
--alarm-description "Alarm when out of capacity is > 0" \
--metric-name InsufficientCapacityError \
--namespace AWS/AppStream \
--statistic Maximum \
--period 300 \
--threshold 0 \
--comparison-operator GreaterThanThreshold \
--dimensions "Name=Fleet,Value=fleetname" \
--evaluation-periods 1 --unit Count \
```

```
--alarm-actions "arn:aws:autoscaling:your-region-code:account-
number-without-hyphens:scalingPolicy:policyid:resource/appstream/
fleet/fleetname:policyName/policyname"
```

Example 3: Applying a Scaling Policy Based on Low Capacity Utilization

This AWS CLI example sets up a scaling policy that scales in the fleet to reduce actual capacity when `CapacityUtilization` is low.

The following command defines an excess capacity-based scaling policy:

```
aws application-autoscaling put-scaling-policy --cli-input-json file://scale-in-
capacity.json
```

The contents of the file `scale-in-capacity.json` are as follows:

```
{
    "PolicyName": "policyname",
    "ServiceNamespace": "appstream",
    "ResourceId": "fleet/fleetname",
    "ScalableDimension": "appstream:fleet:DesiredCapacity",
    "PolicyType": "StepScaling",
    "StepScalingPolicyConfiguration": {
        "AdjustmentType": "PercentChangeInCapacity",
        "StepAdjustments": [
            {
                "MetricIntervalUpperBound": 0,
                "ScalingAdjustment": -25
            }
        ],
        "Cooldown": 360
    }
}
```

If the command is successful, the output is similar to the following, although some details are unique to your account and Region. In this example, the policy identifier is `12ab3c4d-56789-0ef1-2345-6ghi7jk8lm90`.

```
{"PolicyARN": "arn:aws:autoscaling:us-
west-2:123456789012:scalingPolicy:12ab3c4d-56789-0ef1-2345-6ghi7jk8lm90:resource/
appstream/fleet/SampleFleetName:policyName/scale-in-utilization-policy"}
```

Now, set up a CloudWatch alarm for this policy. Use the names, Region, account number, and policy identifier that apply to you. You can use the policy ARN returned by the previous command for the `--alarm-actions` parameter.

```
aws cloudwatch put-metric-alarm
--alarm-name alarmname \
--alarm-description "Alarm when Capacity Utilization is less than or equal to 25
 percent" \
--metric-name CapacityUtilization \
--namespace AWS/AppStream \
--statistic Average \
--period 120 \
--threshold 25 \
--comparison-operator LessThanOrEqualToThreshold \
--dimensions "Name=Fleet,Value=fleetname" \
--evaluation-periods 10 --unit Percent \
--alarm-actions "arn:aws:autoscaling:your-region-code:account-
number-without-hyphens:scalingPolicy:policyid:resource/appstream/
fleet/fleetname:policyName/policyname"
```

Example 4: Change the Fleet Capacity Based on a Schedule

Changing your fleet capacity based on a schedule lets you scale your fleet capacity in response to predictable changes in demand. For example, at the start of a work day, you might expect a certain number of users to request streaming connections at one time. To change your fleet capacity based on a schedule, you can use the Application Auto Scaling PutScheduledAction API action or the put-scheduled-action AWS CLI command.

Before changing your fleet capacity, you can list your current fleet capacity by using the WorkSpaces Applications describe-fleets AWS CLI command.

```
aws appstream describe-fleets --name fleetname
```

The current fleet capacity will appear similar to the following output (shown in JSON format):

```
{
    {
            "ComputeCapacityStatus": {
                "Available": 1,
                "Desired": 1,
                "Running": 1,
```

```
                "InUse": 0
            },
}
```

Then, use the `put-scheduled-action` command to create a scheduled action to change your fleet capacity. For example, the following command changes the minimum capacity to 3 and the maximum capacity to 5 every day at 9:00 AM UTC.

> **Note**
>
> For cron expressions, specify when to perform the action in UTC. For more information, see Cron Expressions.

```
aws application-autoscaling put-scheduled-action --service-namespace appstream \
--resource-id fleet/fleetname \
--schedule="cron(0 9 * * ? *)" \
--scalable-target-action MinCapacity=3,MaxCapacity=5 \
--scheduled-action-name ExampleScheduledAction \
--scalable-dimension appstream:fleet:DesiredCapacity
```

To confirm that the scheduled action to change your fleet capacity was successfully created, run the describe-scheduled-actions command.

```
aws application-autoscaling describe-scheduled-actions --service-namespace appstream --
resource-id fleet/fleetname
```

If the scheduled action was successfully created, the output appears similar to the following.

```
{
    "ScheduledActions": [
        {
            "ScalableDimension": "appstream:fleet:DesiredCapacity",
            "Schedule": "cron(0 9 * * ? *)",
            "ResourceId": "fleet/ExampleFleet",
            "CreationTime": 1518651232.886,
            "ScheduledActionARN": "<arn>",
            "ScalableTargetAction": {
                "MinCapacity": 3,
                "MaxCapacity": 5
```

```
            },
            "ScheduledActionName": "ExampleScheduledAction",
            "ServiceNamespace": "appstream"
        }
    ]
}
```

For more information, see Scheduled Scaling in the *Application Auto Scaling User Guide*.

Example 5: Applying a Target Tracking Scaling Policy

With target tracking scaling, you can specify a capacity utilization level for your fleet.

When you create a target tracking scaling policy, Application Auto Scaling automatically creates and manages CloudWatch alarms that trigger the scaling policy. The scaling policy adds or removes capacity as required to keep capacity utilization at, or close to, the specified target value. To ensure application availability, your fleet scales out proportionally to the metric as fast as it can but scales in more gradually.

The following put-scaling-policy command defines a target tracking scaling policy that attempts to maintain 75% capacity utilization for an WorkSpaces Applications fleet.

```
aws application-autoscaling put-scaling-policy --cli-input-json file://config.json
```

The contents of the file `config.json` are as follows:

```
{
  "PolicyName":"target-tracking-scaling-policy",
  "ServiceNamespace":"appstream",
  "ResourceId":"fleet/fleetname",
  "ScalableDimension":"appstream:fleet:DesiredCapacity",
  "PolicyType":"TargetTrackingScaling",
  "TargetTrackingScalingPolicyConfiguration":{
    "TargetValue":75.0,
    "PredefinedMetricSpecification":{
      "PredefinedMetricType":"AppStreamAverageCapacityUtilization"
    },
    "ScaleOutCooldown":300,
    "ScaleInCooldown":300
  }
}
```

If the command is successful, the output is similar to the following, although some details are unique to your account and Region. In this example, the policy identifier is 6d8972f3-efc8-437c-92d1-6270f29a66e7.

```
{
    "PolicyARN": "arn:aws:autoscaling:us-west-2:123456789012:scalingPolicy:6d8972f3-
efc8-437c-92d1-6270f29a66e7:resource/appstream/fleet/fleetname:policyName/target-
tracking-scaling-policy",
    "Alarms": [
        {
            "AlarmARN": "arn:aws:cloudwatch:us-
west-2:123456789012:alarm:TargetTracking-fleet/fleetname-AlarmHigh-d4f0770c-b46e-434a-
a60f-3b36d653feca",
            "AlarmName": "TargetTracking-fleet/fleetname-AlarmHigh-d4f0770c-b46e-434a-
a60f-3b36d653feca"
        },
        {
            "AlarmARN": "arn:aws:cloudwatch:us-
west-2:123456789012:alarm:TargetTracking-fleet/fleetname-AlarmLow-1b437334-d19b-4a63-
a812-6c67aaf2910d",
            "AlarmName": "TargetTracking-fleet/fleetname-AlarmLow-1b437334-d19b-4a63-
a812-6c67aaf2910d"
        }
    ]
}
```

For more information, see Target Tracking Scaling Policies in the *Application Auto Scaling User Guide*.

Additional Resources for Auto Scaling Amazon WorkSpaces Applications

For step-by-step guidance for working with WorkSpaces Applications Fleet Auto Scaling, see Scaling Your Desktop Application Streams with Amazon AppStream 2.0 in the *AWS Compute Blog*.

To learn more about using the Application Auto Scaling AWS CLI commands or API actions, see the following resources:

- application-autoscaling section of the *AWS CLI Command Reference*
- Application Auto Scaling API Reference
- Application Auto Scaling User Guide

Multi-Session Recommendations

When operating a multi-session environment, consider the following two key areas to ensure optimal performance and streaming experience:

- Determining the maximum number of user sessions per instance
- Managing your multi-session image effectively

Session Density

When deciding the maximum number of user sessions on an instance, consider the following factors to determine the optimum session count:

- *Evaluate resource requirements*: Understand the resource requirements of the applications being used within the sessions. Consider factors such as CPU, memory, disk I/O, and network bandwidth. This evaluation will help determine the amount of resources each user session typically requires.
- *Consider instance specifications*: Take into account the specifications of the instance, including the number of CPUs, available memory, and GPU specifications. Instances with higher specifications can handle a larger number of user sessions. For more information on different instance types supported by WorkSpaces Applications and pricing, see WorkSpaces Applications pricing.
- *Performance testing*: Conduct performance testing on the applications and workload expected to run within the user sessions. Measure resource utilization, response times, and overall system performance. Use this data to assess the impact of concurrent user sessions on performance, and determine the optimal session-to-instance ratio. You can run these assessments across different instance types offered by WorkSpaces Applications to find the optimal instance type or size for your end users. For more information on different instance types offered by WorkSpaces Applications, see the section called “Instance Families”.
- *Monitor resource utilization*: Continuously monitor the resource utilization of the instance during normal usage. Observe CPU, memory, and disk utilization. Ensure that the resource utilization remains within acceptable limits to avoid performance degradation. For a multi-session environment, you can view these metrics on WorkSpaces Applications and the CloudWatch console. For more information, see the section called “Monitoring Resources”.
- *Consider user behavior patterns*: Analyze user behavior patterns to understand peak usage periods and potential concurrent usage. Some users might have intermittent or sporadic usage

patterns, while others might have consistent usage throughout the day. Account for these patterns when determining the maximum number of user sessions to avoid resource contention during peak periods.

With WorkSpaces Applications, you can configure a maximum of 50 user sessions per instance, regardless of the instance type or size that you choose. However, this is only an upper limit, and not a recommended limit.

The following is an example table to help you determine the maximum number of user sessions on an instance in a multi-session fleet. The recommended maximum number of users listed in the table is based on general guidelines and assumptions. Testing with the real-life workload is crucial, since actual performance can vary, depending on the workload's individual characteristics, the application's resource requirements, and user behavior.

Recommendations based on workload types

End User Category	Workload Type	Example Users	Use Cases	Recommended Configuration(s)
End users who conduct a single task and use minimal applications	Light	Task workers, Front desk users	Data entry applications, Text editing, Bastion host	4 users per vCPU on Stream.standard.xlarge/2xlarge or Stream.compute.xlarge+ or Stream.memory.xlarge+
End users who conduct a single task and use minimal applications	Light to Medium	Task workers, Front desk users, Contact center employees	Data entry applications, Text editing, Bastion host, Chat, Email, Messaging apps	2 users per vCPU on Stream.standard.xlarge/2xlarge or Stream.compute.xlarge+ or Stream.memory.xlarge+

End User Category	Workload Type	Example Users	Use Cases	Recommended Configuration(s)
End users who create complex spreadsheets, presentations, and large documents	Medium	Task workers, Contact center employees, Business analysts	Data entry applications, Chat, Email, Messaging apps, Productivity apps	2 users per vCPU on Stream.memory.xlarge+ or Stream.compute.xlarge+
End users with high performance workloads	Medium to Heavy	Knowledge workers, Software developers, Business intelligence analysts	Software Scripting	1 user per vCPU on Stream.memory.xlarge+ or Stream.compute.xlarge+
End users with high performance workloads	Heavy	Knowledge workers, Software developers, Data scientists	Screen sharing, Data analytics, Audio conferencing	1 user per 2 vCPUs on Stream.memory.xlarge+ or Stream.compute.xlarge+
End users with workloads that require graphics and heavy compute/memory resources	Heavy to Accelerated	Graphics/ Architecture designers, CAD/ CAM users	Audio conferencing, Graphics-intensive applications, such as remote graphics workstations	1 user per 2 vCPUs Graphics.g4dn.*

End User Category	Workload Type	Example Users	Use Cases	Recommended Configuration(s)
End users with workloads that require graphics and heavy compute/memory resources	Accelerated	Video editors, Gamers and game developers, Data miners, GIS data engineers, AI scientists	Audio conferencing, Video transcoding and 3D rendering, Photo-realistic design, Graphics workstations, ML model training, ML inference	1 user per 2 vCPUs Graphics.G5.*

Image Management

- *Image considerations*: Remote Desktop Session Host and Remote Desktop Licensing must not be installed on multi-session fleet images. WorkSpaces Applications configures those components automatically during instance provisioning.
- *Image updates*: If you update your multi-session fleet image through Image Builder or Managed Image Update, test it on a pre-production fleet before updating your production fleets or sharing the image with other AWS accounts.

Multi-session image optimization

This section describes Windows image optimizations that reduce CPU and memory usage on WorkSpaces Applications multi-session fleets. These optimizations let you support more concurrent sessions per instance at the same performance level.

Before you begin

- *If you are using Image Builder*: Launch an Image Builder session from the WorkSpaces Applications console and connect to it. All steps in Part 1 are performed inside this session.
- *If you are importing an existing image*: Launch a new Image Builder session based on your imported image first, then follow the steps in Part 1. Do not apply optimizations directly to a fleet instance.

Part 1: Image-level optimizations

These steps are performed once inside an Image Builder session and baked into your image.

1. *Step 1: Download the Windows Desktop Optimization Tool (WDOT).* Inside the Image Builder session, open a browser and see the Windows Desktop Optimization Tool repository on the GitHub website. Choose **Code**, choose **Download ZIP**, and extract it to a local folder, for example `C:\WDOT`. Alternatively, if Git is available, run the following command.

   ```
   git clone https://github.com/The-Virtual-Desktop-Team/Windows-Desktop-Optimization-
   Tool.git C:\WDOT\Windows-Desktop-Optimization-Tool-main
   ```

2. *Step 2: Open an elevated PowerShell session.* Choose **Start**, search for **PowerShell**, right-click **Windows PowerShell**, and then choose **Run as administrator**. Navigate to the WDOT folder and allow the scripts to run for the current session.

   ```
   cd C:\WDOT\Windows-Desktop-Optimization-Tool-main
   Set-ExecutionPolicy -ExecutionPolicy Bypass -Scope Process -Force
   ```

3. *Step 3: Create and customize a configuration profile.*

   ```
   # Create a configuration profile
   .\New-WVDConfigurationFiles.ps1 -FolderName "WorkSpaces-Production"

   # Apply all recommended settings for each optimization category
   .\Set-WVDConfigurations.ps1 -ConfigurationFile "Services" -ConfigFolderName
    "WorkSpaces-Production" -ApplyAll
   .\Set-WVDConfigurations.ps1 -ConfigurationFile "ScheduledTasks" -ConfigFolderName
    "WorkSpaces-Production" -ApplyAll
   .\Set-WVDConfigurations.ps1 -ConfigurationFile "AppxPackages" -ConfigFolderName
    "WorkSpaces-Production" -ApplyAll
   ```

   > **Note**
   >
   > Review the generated JSON files under `C:\WDOT\Configurations\WorkSpaces-Production\` before proceeding. Set any item to `Skip` if it conflicts with your applications.

4. *Step 4: Apply optimizations.*

```
.\Windows_Optimization.ps1 -ConfigProfile "WorkSpaces-Production" -Optimizations
 @("Services", "ScheduledTasks", "AppxPackages") -AcceptEULA
```

This does the following:

- Disables unnecessary Windows services.
- Disables scheduled tasks that run in the background.
- Removes unused AppX packages.

> **Important**
>
> If you are optimizing AppX packages, ensure the AppX service is not in a STOPPED state before creating your image. Always test in a non-production environment first to verify application compatibility.

5. *Step 5: Create your image*. Once optimizations are applied, proceed with your normal image creation process from the WorkSpaces Applications console.

Part 2: User session optimizations

These optimizations apply per-user registry settings (HKCU). They must run after a user logs on, and they cannot be baked into the image directly. There are two ways to apply them.

Option A: Session scripts (recommended)

WorkSpaces Applications session scripts allow you to automatically run a script in the user's context at session start. This is the recommended approach because it requires no user action.

1. *Step 1: Save the optimization script to your image*. Inside your Image Builder session, create the folder and save the script.

```
New-Item -ItemType Directory -Path "C:\Scripts" -Force
```

Save the PowerShell script shown in Step 3 as `C:\Scripts\UserOptimizations.ps1`.

2. *Step 2: Configure the session scripts configuration file*. The session scripts configuration file is located at `C:\AppStream\SessionScripts\config.json` and is present by default on the image. Update it with the following content.

```
{
  "SessionStart": {
    "executables": [
      {
        "context": "system",
        "filename": "",
        "arguments": "",
        "s3LogEnabled": true
      },
      {
        "context": "user",
        "filename": "C:\\Windows\\System32\\WindowsPowerShell\\v1.0\
\powershell.exe",
        "arguments": "-ExecutionPolicy Bypass -File C:\\Scripts\
\UserOptimizations.ps1",
        "s3LogEnabled": true
      }
    ],
    "waitingTime": 30
  },
  "SessionTermination": {
    "executables": [
      {
        "context": "system",
        "filename": "",
        "arguments": "",
        "s3LogEnabled": true
      },
      {
        "context": "user",
        "filename": "",
        "arguments": "",
        "s3LogEnabled": true
      }
    ],
    "waitingTime": 30
  }
}
```

For full session scripts documentation, see the section called "Use Session Scripts on Multi-Session Fleets".

3. *Step 3: Review the optimization script*, saved as `C:\Scripts\UserOptimizations.ps1`.

```
# ------------------------------------------------
# WorkSpaces Application User Session Optimization Script
# Run as: User context, at session start
# ------------------------------------------------

# 1. Disable DWM visual effects
# Disabling transparency, animations, and blur reduces DWM compositing work —
# the single largest CPU and memory optimization in multi-session environments.
$visualFxPath = "HKCU:\Software\Microsoft\Windows\CurrentVersion\Explorer
\VisualEffects"
If (!(Test-Path $visualFxPath)) { New-Item -Path $visualFxPath -Force | Out-Null }
Set-ItemProperty -Path $visualFxPath -Name VisualFXSetting -Value 2 -Type DWord

$desktopPath = "HKCU:\Control Panel\Desktop"
Set-ItemProperty -Path $desktopPath -Name UserPreferencesMask -Value ([byte[]]
(0x90,0x12,0x03,0x80,0x10,0x00,0x00,0x00)) -Type Binary

# Apply visual effects changes immediately without disrupting the session
rundll32.exe user32.dll, UpdatePerUserSystemParameters

# 2. Disable Search UI web background extensions
$searchGlobal = "HKCU:\SOFTWARE\Microsoft\Windows\CurrentVersion\Search"
If (!(Test-Path $searchGlobal)) { New-Item -Path $searchGlobal -Force | Out-Null }
Set-ItemProperty -Path $searchGlobal -Name BingSearchEnabled -Value 0 -Type DWord
Set-ItemProperty -Path $searchGlobal -Name CortanaConsent -Value 0 -Type DWord

# 3. Restrict background apps
$backgroundPath = "HKCU:\SOFTWARE\Microsoft\Windows\CurrentVersion
\BackgroundAccessApplications"
If (!(Test-Path $backgroundPath)) { New-Item -Path $backgroundPath -Force | Out-
Null }
Set-ItemProperty -Path $backgroundPath -Name GlobalUserDisabled -Value 1 -Type
 DWord

$privacyPath = "HKCU:\SOFTWARE\Microsoft\Windows\CurrentVersion\Privacy"
If (!(Test-Path $privacyPath)) { New-Item -Path $privacyPath -Force | Out-Null }
Set-ItemProperty -Path $privacyPath -Name LetAppsRunInBackground -Value 2 -Type
 DWord
```

```
# 4. Disable News & Interests (Taskbar widgets)
$feedPath = "HKCU:\SOFTWARE\Microsoft\Windows\CurrentVersion\Feeds"
If (!(Test-Path $feedPath)) { New-Item -Path $feedPath -Force | Out-Null }
Set-ItemProperty -Path $feedPath -Name ShellFeedsTaskbarViewMode -Value 2 -Type
 DWord

# 5. Disable Tips & Suggestions
$contentPath = "HKCU:\SOFTWARE\Microsoft\Windows\CurrentVersion
\ContentDeliveryManager"
If (!(Test-Path $contentPath)) { New-Item -Path $contentPath -Force | Out-Null }
Set-ItemProperty -Path $contentPath -Name "SubscribedContent-338389Enabled" -Value
 0 -Type DWord
Set-ItemProperty -Path $contentPath -Name SoftLandingEnabled -Value 0 -Type DWord
```

Option B: Default user profile (no session scripts required)

If you prefer not to use session scripts, you can apply these settings to the Windows default user profile during image creation. Any new user logging in to the instance inherits these settings automatically. In your elevated PowerShell session on the Image Builder, run the following commands.

```
# Load the default user hive
reg load "HKU\DefaultUser" "C:\Users\Default\NTUSER.DAT"

# DWM visual effects
reg add "HKU\DefaultUser\Software\Microsoft\Windows\CurrentVersion\Explorer
\VisualEffects" /v VisualFXSetting /t REG_DWORD /d 2 /f
reg add "HKU\DefaultUser\Control Panel\Desktop" /v UserPreferencesMask /t REG_BINARY /d
 9012038010000000 /f

# Search UI web background extensions
reg add "HKU\DefaultUser\SOFTWARE\Microsoft\Windows\CurrentVersion\Search" /v
 BingSearchEnabled /t REG_DWORD /d 0 /f
reg add "HKU\DefaultUser\SOFTWARE\Microsoft\Windows\CurrentVersion\Search" /v
 CortanaConsent /t REG_DWORD /d 0 /f

# Background apps
reg add "HKU\DefaultUser\SOFTWARE\Microsoft\Windows\CurrentVersion
\BackgroundAccessApplications" /v GlobalUserDisabled /t REG_DWORD /d 1 /f
reg add "HKU\DefaultUser\SOFTWARE\Microsoft\Windows\CurrentVersion\Privacy" /v
 LetAppsRunInBackground /t REG_DWORD /d 2 /f
```

```
# News & Interests
reg add "HKU\DefaultUser\SOFTWARE\Microsoft\Windows\CurrentVersion\Feeds" /v
 ShellFeedsTaskbarViewMode /t REG_DWORD /d 2 /f

# Tips & Suggestions
reg add "HKU\DefaultUser\SOFTWARE\Microsoft\Windows\CurrentVersion
\ContentDeliveryManager" /v "SubscribedContent-338389Enabled" /t REG_DWORD /d 0 /f
reg add "HKU\DefaultUser\SOFTWARE\Microsoft\Windows\CurrentVersion
\ContentDeliveryManager" /v SoftLandingEnabled /t REG_DWORD /d 0 /f

# Unload the hive
reg unload "HKU\DefaultUser"
```

Note

Always unload the hive with `reg unload` before creating your image. Leaving it loaded will cause issues with the image.

Part 3: Antivirus configuration

If you run antivirus software on your fleet instances, improper configuration can significantly reduce multi-session performance. Antivirus scans running concurrently across many sessions compound CPU and I/O overhead rapidly.

Follow the guidance in the section called “Antivirus Software” to ensure your antivirus is correctly configured for multi-session environments. Key practices include the following:

- Exclude WorkSpaces Applications process paths and user session directories from real-time scanning.
- Schedule full scans outside of active session hours.
- Use a VDI-optimized antivirus agent where available.

Bring Your Own Windows Desktop Licenses for Amazon WorkSpaces Applications

Amazon WorkSpaces Applications now supports Windows Desktop operating systems (such as Windows 11) through Bring Your Own License (BYOL). You can use your existing Windows Desktop OS licenses and custom images with Amazon WorkSpaces Applications. This lets you stream business-critical applications that require Microsoft Windows Desktop OS while using your current licensing investments.

Overview

With BYOL for Amazon WorkSpaces Applications, you can:

- **Migrate Windows 11 applications** that are non-compliant with Windows Server editions to AWS and stream them to your users
- **Reduce costs** by leveraging existing Microsoft VDA E3/E5 licenses, eliminating duplicate licensing and saving RDS SAL fee per user per month compared to non-BYOL options
- **Maintain full compliance** with Microsoft licensing requirements — AWS provisions dedicated hardware for your BYOL instances

How it works

BYOL for Amazon WorkSpaces Applications reuses the existing Amazon WorkSpaces BYOL infrastructure. This means:

- If you are **already using BYOL on Amazon WorkSpaces**, you can simply import your existing BYOL images into Amazon WorkSpaces Applications
- If you are **new to Amazon WorkSpaces BYOL**, you first set up your dedicated BYOL account in Amazon WorkSpaces, create an image, and then import that image into Amazon WorkSpaces Applications
- **Shared dedicated infrastructure** — Customers use the same BYOL dedicated infrastructure for both WorkSpaces and Amazon WorkSpaces Applications

For detailed information on setting up BYOL in Amazon WorkSpaces, see Bring Your Own Windows Desktop Licenses in WorkSpaces.

Prerequisites

Before you begin, verify the following:

- **Licensing requirement** — Your Microsoft licensing agreement allows Windows to run in a virtual hosted environment. VDA E3/E5 user licenses (purchased under subscription from Microsoft) are required for Windows Desktop Client BYOL on AWS. See Amazon Web Services and Microsoft.
- **Minimum WorkSpaces instances commitment** — You must use a minimum of 50 WorkSpaces per Region. These 50 WorkSpaces can be any mix of AlwaysOn/On-Demand Amazon WorkSpaces Applications unique users streaming sessions, and/or WorkSpaces Personal instances in the same region per month.
- **Windows versions supported** — You will need a Windows virtual machine image or Windows ISO image file that uses a supported Windows OS version.
- **Network connectivity** — Ensure required HTTPS endpoints are accessible for the image import process.

Supported Windows Versions

The following Windows 11 versions are supported for Amazon WorkSpaces Applications BYOL:

- Windows 11 Version 24H2 (October 2024 release)
- Windows 11 Version 25H2 (September 2025 release)
- Windows 11 Version 26H2 (September 2026 release)

> **Note**
>
> Windows 10 images imported through WorkSpaces are not supported for Amazon WorkSpaces Applications. Only Windows 11 images can be used with Amazon WorkSpaces Applications.

Getting Started

If your account is already enabled for BYOL, you do not need to provision new dedicated infrastructure to run Amazon WorkSpaces Applications fleets powered by Windows Desktop OS. You can skip **Step 1**. Amazon WorkSpaces Applications will use the same dedicated BYOL account for provisioning the fleets. The same dedicated hardware that supports your existing WorkSpaces Personal or Pools BYOL workloads can be used for Amazon WorkSpaces Applications.

Step 1: Enable your account for BYOL (First time only)

1. Log in to the AWS Console
2. Choose the Region where you want to enable BYOL
3. Navigate to **Services** → **Amazon WorkSpaces** → **Account Settings** (in the left navigation)
4. In the **Bring Your Own License (BYOL)** section, choose **Get Started with BYOL**
5. Choose **Enable account for BYOL** and confirm the minimum requirements
6. **Choose IP range (First time only)** — Select an IP/CIDR range for your BYOL management network interface. Once chosen, it cannot be altered.

For detailed steps on image import, see Bring Your Own Windows Desktop Licenses in WorkSpaces.

> **Note**
>
> Graphics BYOL enablement is not yet supported through the WorkSpaces Console. Please create an AWS Support ticket for Graphics BYOL.

> **Note**
>
> BYOL enablement is shared between WorkSpaces and Amazon WorkSpaces Applications through the unified "Account Settings" interface.

Step 2: Create a BYOL image in Amazon WorkSpaces

Option A: For existing WorkSpaces BYOL customers

If you already have BYOL images in Amazon WorkSpaces, you can skip directly to **Step 3** to import your existing images into Amazon WorkSpaces Applications.

Option B: For new BYOL Customers

- **(Optional) Validate your image before importing** — If you are importing a customized virtual machine image, run the WorkSpaces Image Checker tool to ensure compatibility. If importing a Windows ISO, you can skip this step.
- **Import the image** — After enabling BYOL, choose **Import Image**. You have three options:
  - **VM import** — Imports a virtual machine image (VHDX, VMDK, or OVF file) that has already been customized
  - **ISO import** — Imports a Windows ISO image downloaded from Microsoft that has not been customized
  - **AMI import** — Imports an existing Amazon EC2 AMI to use as your BYOL image

For detailed steps on image import, see Bring Your Own Windows Desktop Licenses in WorkSpaces.

Step 3: Import BYOL Image into Amazon WorkSpaces Applications

Once your BYOL image is successfully created in Amazon WorkSpaces:

1. Navigate to the **Amazon WorkSpaces Applications Console**. Note: Your console role will also need to include the `workspaces:DescribeWorkspaceImages` permission.
2. Go to **Images** → **Import Image**
3. Select **Amazon WorkSpaces Image** as the image source. Enter the WorkSpaces Image ID (starts with "wsi-").
4. Provide image details (name, display name, description)
5. Complete the import process

> **Note**
>
> At the time of launch, you can import your Windows 11 and Windows Server 2022/Server 2025 images into Amazon WorkSpaces Applications.

Step 4: Customize your image (Optional)

You can optionally use Image Builder to configure custom applications, optimize launch performance, and create a snapshot to produce a custom image:

1. Launch an Image Builder instance from your imported BYOL image
2. Connect to the Image Builder and install/configure your applications
3. Create a snapshot to generate your customized image

Step 5: Create fleet and stream to users

1. Navigate to the **Fleets** console in Amazon WorkSpaces Applications
2. Create a fleet using your BYOL image
3. Configure fleet settings (instance type, scaling policies, etc.)
4. Once the fleet is ready, your end users can start streaming from BYOL instances

Supported instance types

The following instance types are available for BYOL Amazon WorkSpaces Applications:

- **Standard:** stream.standard.medium, stream.standard.large, stream.standard.xlarge, stream.standard.2xlarge
- **Compute-optimized:** stream.compute.large, stream.compute.xlarge, stream.compute.2xlarge, stream.compute.4xlarge, stream.compute.8xlarge
- **Memory-optimized:** stream.memory.large, stream.memory.xlarge, stream.memory.2xlarge, stream.memory.4xlarge, stream.memory.8xlarge
- **Graphics (requires separate approval):** stream.graphics.g6 and stream.graphics.g7 families

Note

BYOL is not supported with unified graphics images. BYOL graphics deployments continue to use per-family graphics images (for example, the stream.graphics.g6 and stream.graphics.g7 families), each of which requires separate approval. For more information about unified graphics images, see Unified graphics images.

Region availability

Windows BYOL for Amazon WorkSpaces Applications is available in all regions where both Amazon WorkSpaces Applications and WorkSpaces Personal are offered:

- US East (N. Virginia)
- US East (Ohio)
- US West (Oregon)
- Asia Pacific (Malaysia)
- Asia Pacific (Mumbai)
- Asia Pacific (Seoul)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)
- Canada (Central)
- Europe (Frankfurt)
- Europe (Ireland)
- Europe (London)
- Europe (Paris)
- Israel (Tel Aviv)
- South America (São Paulo)
- AWS GovCloud (US-West)
- AWS GovCloud (US-East)

Important

Windows BYOL for Amazon WorkSpaces Applications is only available in a subset of the Availability Zones for each supported Region. For the list of AZ IDs that you can use in each Region, see Availability Zones for WorkSpaces Personal in the *Amazon WorkSpaces Administration Guide*. To see the mapping of AZ IDs to Availability Zones in your account, see AZ IDs for Your Resources in the *AWS RAM User Guide*.

Important considerations

- **Dedicated infrastructure** — BYOL images powered by Windows Desktop OS (Windows 11) run on dedicated infrastructure only. Windows Server OS images continue to run on shared infrastructure.
- **No multi-session support** — Multi-session Windows 11 is not available due to Microsoft licensing restrictions.
- **Image sharing** — You can share BYOL images only with AWS accounts that are enabled for BYOL and are part of your organization (under the same payer account).
- **Separate BYOL account not required** — You can use your existing AWS account for Amazon WorkSpaces Applications BYOL workloads. The same dedicated infrastructure account (DP account) can host WorkSpaces Personal, WorkSpaces Pools, and Amazon WorkSpaces Applications.
- **Image operations** — Import creates images in both WorkSpaces and Amazon WorkSpaces Applications (if selected). Copy, update, and delete operations are service-specific and do not reflect across services.

Note

Streaming from Interface VPC Endpoints is currently not supported for fleets powered by Windows Desktop operating system (such as Windows 11).

Frequently asked questions

Can I use my existing BYOL WorkSpaces images with Amazon WorkSpaces Applications?

Yes. You can import your existing WorkSpaces Personal/Pools images into Amazon WorkSpaces Applications. The image OS and version must be supported by Amazon WorkSpaces Applications (Windows 11 and Windows Server 2022/Server 2025 at launch).

Can I use existing Amazon WorkSpaces Applications images for BYOL?

No. Existing Amazon WorkSpaces Applications images are for Windows Server operating systems, whereas BYOL is for Windows Desktop operating systems only.

How can I provide software updates to my end users?

You can use the managed image updates procedure to get an automated way to update your image with the latest Windows OS updates, driver updates, and Amazon WorkSpaces Applications agent software.

Can I delete a BYOL image?

Yes, you can delete a BYOL image if there is no active fleet or Image Builder referencing it. Deletion removes the image from the respective service only.

Application onboarding guidelines

Before you make an application available to end users through WorkSpaces Applications, validate that the application works correctly in the WorkSpaces Applications cloud environment, verify that the application can be streamed without rendering artifacts, and size your fleet appropriately for the application's resource profile. This page provides a structured onboarding checklist that you can follow for each application you plan to stream. If you plan to use multi-session, you need to pay specific attention to the ability to run multiple versions of the application on the same host. If you plan to use native application mode, you need to verify that the application does not present any compatibility problems with this mode.

These guidelines apply whether you are onboarding a new application or migrating an existing application from another delivery model.

Overview of the onboarding process

Onboarding an application to WorkSpaces Applications has two parallel tracks:

- **Application compatibility validation:** Confirm the application behaves correctly in the streaming environment across the features your users rely on.
- **Instance sizing and capacity planning:** Choose an instance type and fleet scaling policy that match the application's CPU, memory, and GPU profile and your expected concurrent user count.

We recommend completing compatibility validation first (on an Image Builder and a pilot fleet), then using the measurements from that validation to drive instance sizing decisions.

Part 1: Application compatibility validation

General compatibility

WorkSpaces Applications establishes a streaming session in this order:

1. The user's client connects to a streaming instance.
2. The user is logged into the server operating system session on the instance.
3. The application is started.

4. The streaming session begins, at which point environment settings that depend on the client (for example, client display resolution, DPI, client time zone, and client-redirected devices like printers) are applied.

Because client-dependent environment settings are applied *after* the application starts, an application that reads these values only once during startup will not react to the user's actual client configuration. Validate the following:

- The application reads or subscribes to display resolution, DPI, and time zone changes after startup, **or** you configure the streaming instance with defaults that match your user population before the client connects.
- The application tolerates the user's local time zone being different from the streaming instance's time zone.
- The application does not fail or hang when the client disconnects and reconnects mid-session.

Multi-session compatibility

If you plan to run the application on a multi-session fleet, validate the following on an Image Builder and a pilot fleet with two or more concurrent users:

- The application's license permits concurrent multi-user sessions on a single server instance. Some applications' per-device or per-user licenses explicitly forbid this configuration.
- User profiles, application settings, and user data are isolated between sessions. Test that a change one user makes is not visible to another concurrent user.
- User-specific data is written to per-user locations such as `%APPDATA%` and `%LOCALAPPDATA%`, not to shared directories such as `C:\Program Files` or `C:\ProgramData`.
- The application does not rely on system services shared across all sessions, or such services can correctly handle multiple concurrent sessions.
- The application does not contain hardcoded paths that assume a single-user environment.
- System device dialogs enumerate only the current user's redirected devices like scanners and printers, not all devices visible to the server.
- File open and save dialogs correctly resolve mapped client drives and redirected folders for the current user.
- The application's installer, updater, and any background processes do not require interactive administrator access while users are signed in.

> **Note**
>
> Multi-session fleets currently do not support webcam, Dynamic Application Framework, and Smart Card Authentication. If you need any of these features, we recommend that you use single-session fleets.

Native application mode compatibility

Native application mode streams each remote application as a separate window on the user's local device, with its own taskbar icon. Applications that work correctly in classic mode can behave differently in native application mode because window management, focus, and rendering are handled differently. For an overview of the feature, see Native Application Mode.

Validate the following in native application mode on an Image Builder and a pilot fleet:

- **Window display:** All application windows, including dialogs, splash screens, and modal prompts, display correctly and can be interacted with. Pay particular attention to windows that are drawn manually rather than with the standard Windows UI toolkits (for example, custom-rendered splash screens or startup dialogs that appear before the main window). These are more likely to have compatibility issues.
- **Transparent or non-rectangular windows:** Check for any windows that have transparent areas or that have rounded or non-rectangular shapes (for example, balloon tips, custom tooltips, or skinned windows). These may not render correctly in native application mode.
- **System tray:** Applications that require the Windows notification area (system tray) are not currently supported in native application mode. If your application uses a tray icon only for secondary features, confirm that the core workflow still functions when the tray is unavailable.
- **DPI awareness:** Streaming sessions may run at resolutions and DPI settings that differ from the local client. If the application is not DPI-aware, Windows itself scales the output, which results in blurred rendering. Test on at least one client with a DPI scaling other than 100% (for example, a high-DPI laptop at 125% or 150%).
- **Multi-window workflows:** Test workflows that span multiple application windows (for example, switching between a main window and a modal dialog, or between two documents opened in separate windows). Confirm that focus transitions and taskbar click-to-activate behave as expected.

- **Alt+Tab and taskbar behavior:** Switch between the application and other local applications by using Alt+Tab and by clicking the taskbar icon. The remote application should come to the foreground without bringing unrelated remote windows with it.
- **Modal dialogs:** When a modal dialog is open in the remote application, the underlying remote window should correctly indicate it is disabled, and clicking it should flash or activate the modal.
- **Print workflows:** Print from within the application and verify that the print dialog is visible (not hidden behind the main window) and that redirected printers are enumerated. Print dialogs displayed by some printer drivers may attach to the wrong window. If you observe this, consider using the **DCV PDF Printer** instead of **Microsoft Print to PDF**.
- **Browser tab docking:** When users try to dock or undock tabs in one browser window into separate windows during a streaming session in native application mode, the remote streaming browser does not work the same way as a local browser. Users must press the Alt key until the tabs dock into separate browser windows. If your users rely on frequent tab undocking, plan user training for this behavior.
- **Mode switching:** Test that the application continues to function if the user switches between native application mode and classic mode during a session.

We recommend starting native application mode validation with a pilot user group and documenting any application-specific limitations before full deployment. Application behavior and performance can vary between streaming modes, so testing in classic mode is not a substitute for testing in native application mode.

File handling and redirection

- Open and save files from and to mapped client drives and redirected folders.
- If the application uses temporary files, confirm they are written to per-user temp directories.
- Test large-file operations over the session's file transfer mechanism if your users work with files larger than the typical size handled by the streaming session.

Printing

- Test printing to each redirected printer type your users will use (network printers, Microsoft Print to PDF, PDF printer drivers such as the DCV PDF Printer, and third-party redirected printers).
- Test printing from each application that has a print workflow, including applications that embed web content (such as Chromium-based views).

- Validate print dialog behavior in native application mode (see the previous section).

Local device interactions

- Audio input and output (microphone, speakers, and headsets).
- Webcam, if used by the application.
- USB device redirection, if used by the application. See USB device redirection for the list of supported devices.
- Smart card authentication, if required by the application.

Network performance

- Measure application responsiveness over a network connection representative of your worst-case user (for example, a remote user on a consumer broadband connection with 100 ms round-trip time). Streaming sessions are sensitive to round-trip time and packet loss.
- Validate that the application tolerates brief network interruptions and session reconnects.

Multi-monitor support

- Test workflows that span multiple monitors on the client side.
- If the application reads monitor geometry, confirm it reads the client's monitor layout, not the streaming instance's layout.

Audio and video features

Real-time audio-video scenarios (voice, video conferencing, and collaboration tools embedded in the application) require higher frame rates and may need a larger instance type. See Part 2: Instance sizing and capacity planning.

Validation environment

Perform the checks in this section on an Image Builder first, then on a pilot fleet with a small group of representative users, before rolling the image out to your full user base. Do not rely on classic-mode testing as a substitute for native-application-mode testing.

Part 2: Instance sizing and capacity planning

Choose an instance family

Select an instance family based on the resource profile of the application. For hardware specifications and pricing, see WorkSpaces Applications Instance Families and WorkSpaces Applications Pricing.

Application profile	Recommended instance family
Office, web browsers, most line-of-business applications	General Purpose
Compute-bound applications (heavy client-side computation, local analytics)	Compute Optimized
Memory-intensive applications (large datasets in memory, in-memory databases)	Memory Optimized
Graphics applications using DirectX, OpenGL, or OpenCL	Graphics G4dn, G5, or G6 family
Real-time audio-video for high-frame-rate scenarios	Scale up the instance size within your chosen family; consider a Graphics family instance if the application also uses GPU acceleration

Each WorkSpaces Applications instance has a 200 GB fixed-size C drive that is deleted after every user session. Do not rely on instance-local storage for user data; use home folders, file shares, or profile management for persistence.

Size an instance for a single user

Before sizing for concurrent users, measure the application's resource use for a single session:

- Provision an Image Builder of the smallest instance size in your chosen family that meets the application's stated minimum requirements.

- Sign in as a single user and run the representative workload end to end. Include all dependencies that will run in the streaming session (for example, background sync clients, security agents, and profile management clients).
- Measure, using Windows Performance Monitor or an equivalent tool: peak and sustained CPU use, peak and sustained working-set (private bytes) memory use, disk I/O rate, and if applicable, GPU use and video memory.
- If peak CPU exceeds approximately 80% or peak memory exceeds approximately 75% of instance capacity during normal workflows, move up to the next instance size.
- Leave headroom for the Windows Server operating system, the WorkSpaces Applications agent, Amazon DCV, anti-malware, and any other management agents. A rule of thumb is to reserve approximately 1 vCPU and 1 GB of memory for base system overhead on a single-session instance.

Size for multi-session fleets

Multi-session fleets run multiple users concurrently on a single Windows Server instance. The supported maximum users per instance depends on the instance size and the resource profile of the application.

- Start from your single-user measurement (from the previous section).
- Apply a concurrency multiplier based on the application's behavior. For applications that have a predominantly idle resource profile (for example, office applications used interactively), plan for aggregate CPU and memory use to scale roughly linearly with user count, but with a 20–30% reduction due to shared OS overhead. For applications that have a consistently active resource profile (for example, data processing tools or browsers running heavy web applications), plan for nearly linear scaling with no reduction.
- Calculate a candidate maximum users per instance as:

```
max_users_per_instance = min(
    (instance_vcpus - 1) / peak_single_user_vcpus_under_concurrency,
    (instance_memory_gb - 1) / peak_single_user_memory_gb_under_concurrency
)
```

- Validate the candidate value in an actual multi-session pilot. Run a test with the candidate number of concurrent users (for example, with a load-generation tool or with actual pilot users).

Monitor CPU use (target less than 80% peak, less than 70% sustained), available memory (target greater than 15% of total at peak), disk queue length (target less than 2 sustained), and DCV session responsiveness (subjective — does the session feel interactive?).

- If the pilot fails any of these targets, reduce the maximum users per instance by one and retest. If the pilot passes comfortably, you may increase the maximum users per instance by one and retest, or leave headroom for workload spikes.

Configure fleet scaling

After you know the maximum users per instance, configure fleet scaling based on your expected concurrent user count over time. See Fleet Auto Scaling for WorkSpaces Applications for the full mechanics. This section summarizes the decisions you should make as part of onboarding.

- **Minimum capacity.** Set based on the lowest expected concurrent user count during business hours, divided by the users per instance. Provisioning takes several minutes per instance, so a minimum capacity of zero or a value that is too low can result in users waiting for an instance to launch at the start of the workday. For predictable morning ramp-ups, use a **scheduled scaling policy** to increase the minimum capacity before the workday starts and decrease it before the workday ends.
- **Maximum capacity.** Set to an upper bound that accounts for peak concurrent users plus a safety margin. The peak is typically 1.2–1.5× the average during business hours, but measure your own traffic to set this accurately.
- **Target utilization.** For fleets with unpredictable demand, use a **target tracking scaling policy**. Choose a target utilization such that `100% - target utilization` exceeds your expected user turnover rate (churn) within a 15-minute window. For example, if 10% of users start and end sessions within any 15-minute window, set the target to 90% or lower. See Best Practices for Scaling Policy Design in the whitepaper for more detail.
- **InsufficientCapacityError alarm.** Create an Amazon CloudWatch alarm on the `InsufficientCapacityError` metric for each fleet, so that administrators are alerted when automatic scaling cannot keep up with demand.

Validate end-to-end with a pilot

Before rolling an application out to all users, run a pilot with 10–50 users for at least one full business week. During the pilot:

- Confirm that application compatibility validation results from Part 1 hold under real user workloads.
- Confirm that the chosen instance size supports the observed peak concurrent users per instance.
- Confirm that the fleet scaling policy handles the start-of-day ramp and end-of-day ramp-down without `InsufficientCapacityError` events.
- Collect feedback on session responsiveness and application behavior from pilot users.

Onboarding checklist

Use this checklist to track the status of each application you are onboarding.

Application compatibility

- The application license permits the intended deployment (single-session or multi-session, concurrent users).
- General compatibility checks pass (client-dependent settings applied after application start).
- Multi-session compatibility checks pass (if targeting multi-session fleets).
- Native application mode checks pass (windows, dialogs, tray, DPI, focus, print, and mode switching).
- File handling and redirection checks pass.
- Printing workflows validated for all redirected printer types users will use.
- Local device interactions (audio, webcam, USB, and smart card) validated for all devices users will use.
- Multi-monitor workflows validated.

Instance sizing and capacity

- Instance family chosen based on application resource profile.
- Single-user resource use measured on an Image Builder.
- Maximum users per instance calculated and validated in a multi-session pilot (if applicable).
- Fleet scaling policy configured (minimum capacity, maximum capacity, target utilization or scheduled scaling).
- `InsufficientCapacityError` alarm configured.

Pilot and rollout

- Pilot run with 10–50 users for at least one business week.
- Pilot feedback reviewed and any blocking issues resolved.
- Application documented for end-user training (including any native application mode known behaviors such as browser tab docking via the Alt key).
- Rollout plan agreed with stakeholders.

User Authentication

The following topics provide information about Amazon WorkSpaces Applications user authentication and authorization.

Contents

- Amazon WorkSpaces Applications User Pools
- Amazon WorkSpaces Applications Integration with SAML 2.0

Amazon WorkSpaces Applications User Pools

The WorkSpaces Applications user pool provides a simplified way to manage access to applications for your users through a persistent portal for each AWS Region. This feature is a built-in alternative to user management through Active Directory and SAML 2.0 federation. Stacks can't be assigned to users in the user pool if the stacks are associated with a fleet that is joined to an Active Directory domain.

The WorkSpaces Applications user pool provides the following key features:

- Users can access application stacks through a persistent URL and login credentials by using their email address and a password that they choose.
- Users' email addresses are case-sensitive. During login, if they specify an email address that doesn't use the same capitalization as the email address specified when their user pool account was created, a "user does not exist" error message displays.
- You can assign multiple stacks to users. Doing so enables WorkSpaces Applications to display multiple application catalogs to users when they log in.
- When you create new users, a welcome email is automatically sent to them. The email includes instructions, a login portal link, and a temporary password for connecting to the login portal.
- After you create users, they are enabled unless you specifically disable them.
- You can control which users have access to which application stacks, or disable access completely.

Topics

- User Pool End User Experience for Amazon WorkSpaces Applications
- Resetting a Forgotten Password in Amazon WorkSpaces Applications

- User Pool Administration in Amazon WorkSpaces Applications

User Pool End User Experience for Amazon WorkSpaces Applications

The following steps summarize the initial connection experience for users in the user pool.

1. You create new users in the Region you want by specifying their email addresses.
2. WorkSpaces Applications sends them a welcome email.
3. You assign one or more stacks to the users.
4. WorkSpaces Applications sends them an optional notification email. This email includes information about how to access the stacks that are newly assigned to them.
5. The users connect to the login portal by entering the information included in the welcome email, and they set a permanent password. The login portal link never expires and can be used any time.
6. They sign in to WorkSpaces Applications by entering their email address and permanent password.
7. After they sign in, the users can view their application catalogs.

The login portal link provided in the welcome email should be saved for future use, as it does not change and is valid for all users in the user pool. The login portal URL and users in the user pool are managed on a per-Region basis.

Resetting a Forgotten Password in Amazon WorkSpaces Applications

If users forget their password, follow these steps to connect to the login portal link (provided in the welcome email) and choose a new password.

To choose a new password

1. Open the WorkSpaces Applications login portal by using the login link provided in the welcome email.
2. Choose **Forgot Password?**.
3. Type the email address that you used to create the user in the user pool, and choose **Next**.

Your email address is case-sensitive. During login, if your email address doesn't use the same capitalization as the email address specified when your user pool account was created, a "user does not exist" error message displays.

4. Check your email for the password reset request message. If you are having difficulty finding the email, check your spam email folder. Type the verification code from the email in **Verification Code**.

> **Note**
>
> The verification code is valid for 24 hours. If a new password is not chosen within this time, request a new verification code.

5. Following the password rules shown, type and confirm your new password. Choose **Reset Password**.

User Pool Administration in Amazon WorkSpaces Applications

To create and manage users in the user pool, sign in to the WorkSpaces Applications console for the AWS Region you want and choose **User Pool** in the left navigation pane. The User Pool dashboard supports bulk operations on a list of users for some actions. You can select multiple users on which to perform the same action from the **Actions** list. Users in the user pool are created and managed on a per-Region basis.

WorkSpaces Applications does not support bulk user creation or disable. However, you can use Amazon Cognito with the CreateStreamingURL API action to manage access efficiently for multiple users. Amazon Cognito user pools let you quickly create your own directory to sign up and sign in users. In addition, you can use Amazon Cognito user pools to store user profiles. For information about how to integrate WorkSpaces Applications with your Cognito User Pool, see the Create a SaaS Portal with Amazon WorkSpaces Applications tutorial.

> **Note**
>
> WorkSpaces Applications sends email to users on your behalf when you create a new user created or assign a user to a stack. To ensure the email is delivered, add `no-reply@accounts.`*`aws-region-code`*`.amazonappstream.com` to your allow list, where

aws-region-code is a valid AWS Region code in which you are working. If users are having difficulty finding the emails, ask them to check their "spam" email folder.

Tasks

- Creating a User in Amazon WorkSpaces Applications
- Deleting a User in Amazon WorkSpaces Applications
- Assigning Stacks to Users in Amazon WorkSpaces Applications
- Unassigning Stacks from Users in Amazon WorkSpaces Applications
- Disabling Users in Amazon WorkSpaces Applications
- Enabling Users in Amazon WorkSpaces Applications
- Re-Sending Welcome Email in Amazon WorkSpaces Applications

Creating a User in Amazon WorkSpaces Applications

You must enter a valid and unique email address for each new user within a Region. However, you can reuse an email address for a new user in another Region.

When you create a new user, be aware of the following:

- You cannot change the email address, first name, or last name for a user that you have already created. To change this information for a user, disable the user. Then, recreate the user (as a new user) and specify the updated information as needed.
- Users' email addresses are case-sensitive. During login, if they specify an email address that doesn't use the same capitalization as the email address specified when their user pool account was created, a "user does not exist" error message displays.
- You can assign one or more stacks to the user after the user is created.

To create a new user

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the left navigation pane, choose **User Pool**, **Create User**.
3. For **Email**, type the unique email address for the user.
4. Type the user's first name and last name in the corresponding fields. These fields need not be unique.

5. Choose **Create User**.

After users are created, WorkSpaces Applications sends them a welcome email. This email includes the login portal link, the login email address to be used, and a temporary password. By browsing to the login portal and typing their temporary password, users can set a permanent password to access their applications.

By default, the new user's status is **Enabled**, meaning you can assign one or more stacks to the user or perform other administrative actions.

Deleting a User in Amazon WorkSpaces Applications

You can enable or disable a user, but you cannot delete a user by using the WorkSpaces Applications console. To delete a user, you must use the DeleteUser API action.

Assigning Stacks to Users in Amazon WorkSpaces Applications

You can assign one or more stacks to one or more users in the user pool. After they are assigned to at least one stack, users can log in to WorkSpaces Applications and launch applications. If users are assigned to more than one stack, they are presented with a list of stacks as catalogs to choose from before launching applications.

> **Note**
>
> Stacks can't be assigned to users if the stacks are associated with a fleet that is joined to an Active Directory domain.

To assign a stack to users

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the left navigation pane, choose **User Pool** and select the users you want.
3. Choose **Actions**, **Assign stack**. For more information, see Using Active Directory with WorkSpaces Applications.
4. Review the list to confirm that the correct users are specified. For **Stack**, choose the stack you want to assign.
5. By default, **Send email notification to user** is enabled. Clear this option if you do not want to send the notification email to users now.

6. Choose **Assign stack**.

Unassigning Stacks from Users in Amazon WorkSpaces Applications

You can unassign a stack from one or more users in the user pool. After a stack is unassigned from users, they can't launch applications from the stack. If users are connected when you unassign the stack, their sessions remain active until the session cookie expires (about one hour).

To unassign a stack from users

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the left navigation pane, choose **User Pool** and select the users you want.
3. Choose **Actions**, **Unassign stack**.
4. Review the list to confirm that the correct users are specified. For **Stack**, choose the stack you want to unassign. The list includes all stacks, assigned or unassigned.
5. Choose **Unassign stack**.

Disabling Users in Amazon WorkSpaces Applications

You can disable one or more users in the user pool, one at a time. After they are disabled, users can no longer log in to WorkSpaces Applications until they are re-enabled. This action does not delete users. If users are connected when you disable them, their sessions remain active until the session cookie expires (about one hour). Stack assignments for the users are retained. If the users are re-enabled, their stack assignments become active again.

To disable a user

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the left navigation pane, choose **User Pool** and select the user you want.
3. Choose **Actions**, **Disable user**.
4. Confirm that the correct user is specified, and choose **Disable User**.

Enabling Users in Amazon WorkSpaces Applications

You can enable one or more users in the user pool, one at a time. After they are enabled, users can log in to WorkSpaces Applications and launch applications from the stacks to which they are assigned. If the users were disabled, these assignments are retained.

To enable users

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the left navigation pane, choose **User Pool** and select the user you want.
3. Choose **Actions**, **Enable user**.
4. Confirm that the correct user is specified, and choose **Enable User**.

Re-Sending Welcome Email in Amazon WorkSpaces Applications

You can re-send the welcome email with connection instructions to users in the user pool. Unused passwords expire after seven days. To provide a new temporary password, you must re-send the welcome email. This option is only available until users set their permanent password. If they've already set their password and forgotten it, they can set a new one. For more information, see Resetting a Forgotten Password in Amazon WorkSpaces Applications.

To resend the welcome email for a user

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the left navigation pane, choose **User Pool** and select the user you want.
3. For **User Details**, choose **Resend welcome email**.
4. Confirm that the success message displays at the top of the User Pool dashboard.

Amazon WorkSpaces Applications Integration with SAML 2.0

Amazon WorkSpaces Applications supports identity federation to WorkSpaces Applications stacks through Security Assertion Markup Language 2.0 (SAML 2.0). You can use an identity provider (IdP) that supports SAML 2.0—such as Active Directory Federation Services (AD FS) in Windows Server, Ping One Federation Server, or Okta—to provide an onboarding flow for your WorkSpaces Applications users.

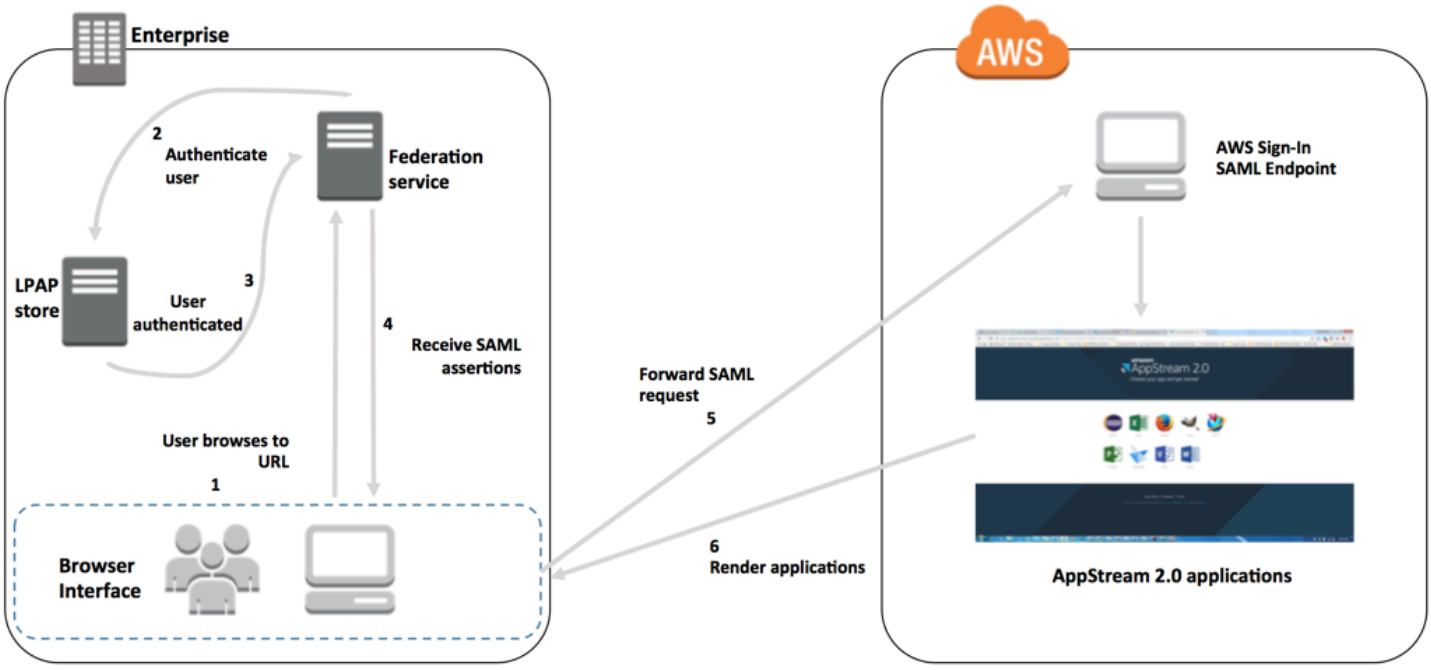
This feature offers your users the convenience of one-click access to their WorkSpaces Applications applications using their existing identity credentials. You also have the security benefit of identity authentication by your IdP. By using your IdP, you can control which users have access to a particular WorkSpaces Applications stack.

Topics

- Example Authentication Workflow
- Setting Up SAML
- WorkSpaces Applications Integration with SAML 2.0

Example Authentication Workflow

The following diagram illustrates the authentication flow between WorkSpaces Applications and a third-party identity provider (IdP). In this example, the administrator has set up a sign-in page to access WorkSpaces Applications, called `applications.exampleco.com`. The webpage uses a SAML 2.0–compliant federation service to trigger a sign-on request. The administrator has also set up a user to allow access to WorkSpaces Applications.



1. The user browses to `https://applications.exampleco.com`. The sign-on page requests authentication for the user.
2. The federation service requests authentication from the organization's identity store.

3. The identity store authenticates the user and returns the authentication response to the federation service.
4. On successful authentication, the federation service posts the SAML assertion to the user's browser.
5. The user's browser posts the SAML assertion to the AWS Sign-In SAML endpoint (`https://signin.aws.amazon.com/saml`). AWS Sign-In receives the SAML request, processes the request, authenticates the user, and forwards the authentication token to WorkSpaces Applications.

   For information about working with SAML in the AWS GovCloud (US) Regions, see AWS Identity and Access Management in the *AWS GovCloud (US) User Guide*.
6. Using the authentication token from AWS, WorkSpaces Applications authorizes the user and presents applications to the browser.

From the user's perspective, this process happens transparently. The user starts at your organization's internal portal and is automatically redirected to an WorkSpaces Applications application portal without being required to enter AWS credentials.

Setting Up SAML

To enable users to sign in to WorkSpaces Applications by using their existing credentials, and start streaming applications, you can set up identity federation using SAML 2.0. To do this, use an IAM role and a relay state URL to configure your SAML 2.0-compliant identity provider (IdP) and enable AWS to permit your federated users to access an WorkSpaces Applications stack. The IAM role grants users the permissions to access the stack. The relay state is the stack portal to which users are forwarded after successful authentication by AWS.

Contents

- Prerequisites
- Step 1: Create a SAML Identity Provider in AWS IAM
- Step 2: Create a SAML 2.0 Federation IAM Role
- Step 3: Embed an Inline Policy for the IAM Role
- Step 4: Configure Your SAML-Based IdP
- Step 5: Create Assertions for the SAML Authentication Response
- Step 6: Configure the Relay State of Your Federation

- Table 1: WorkSpaces Applications relay state region endpoints (Recommended)
- Table 2: Old WorkSpaces Applications relay state region endpoints (Not recommended)
- Table 3: Relay state URL parameters

Prerequisites

Complete the following prerequisites before configuring your SAML 2.0 connection.

1. Configure your SAML-based IdP to establish a trust relationship with AWS.
   - Inside your organization's network, configure your identity store to work with a SAML-based IdP. For configuration resources, see WorkSpaces Applications Integration with SAML 2.0.
   - Use your SAML-based IdP to generate and download a federation metadata document that describes your organization as an IdP. This signed XML document is used to establish the relying party trust. Save this file to a location that you can access from the IAM console later.
2. Use the WorkSpaces Applications management console to create an WorkSpaces Applications stack. You need the stack name to create the IAM policy and to configure your IdP integration with WorkSpaces Applications, as described later in this topic.

   You can create an WorkSpaces Applications stack by using the WorkSpaces Applications management console, AWS CLI, or WorkSpaces Applications API. For more information, see Create an Amazon WorkSpaces Applications Fleet and Stack.

Step 1: Create a SAML Identity Provider in AWS IAM

First, create a SAML IdP in AWS IAM. This IdP defines your organization's IdP-to-AWS trust relationship using the metadata document generated by the IdP software in your organization. For more information, see Creating and Managing a SAML Identity Provider (Amazon Web Services Management Console) in the *IAM User Guide*. For information about working with SAML IdPs in the AWS GovCloud (US) Regions, see AWS Identity and Access Management in the *AWS GovCloud (US) User Guide*.

Step 2: Create a SAML 2.0 Federation IAM Role

Next, create a SAML 2.0 federation IAM role. This step establishes a trust relationship between IAM and your organization's IdP, which identifies your IdP as a trusted entity for federation.

To create an IAM role for the SAML IdP

1. Open the IAM console at https://console.aws.amazon.com/iam/.
2. In the navigation pane, choose **Roles**, **Create role**.
3. For **Role type**, choose **SAML 2.0 federation**.
4. For **SAML Provider**, select the SAML IdP that you created.

> **Important**
>
> Do not choose either of the two SAML 2.0 access methods (**Allow programmatic access only** or **Allow programmatic and Amazon Web Services Management Console access**).

5. For **Attribute**, choose **SAML:sub_type**.
6. For **Value**, enter **`https://signin.aws.amazon.com/saml`**. This value restricts role access to SAML user streaming requests that include a SAML subject type assertion with a value of persistent. If the SAML:sub_type is persistent, your IdP sends the same unique value for the NameID element in all SAML requests from a particular user. For more information about the SAML:sub_type assertion, see the *Uniquely Identifying Users in SAML-Based Federation* section in Using SAML-Based Federation for API Access to AWS.

> **Note**
>
> Even though the **`https://signin.aws.amazon.com/saml`** endpoint is highly available, it is only hosted in the us-east-1 region of AWS. To prevent service interruptions in the unlikely event that one of the regional endpoint's availability is affected, use regional endpoints and set up additional SAML sign-in endpoints for failover. For more information, please see How to use regional SAML endpoints for failover.

7. Review your SAML 2.0 trust information, confirming the correct trusted entity and condition, and then choose **Next: Permissions**.
8. On the **Attach permissions policies** page, choose **Next: Tags**.
9. (Optional) Enter a key and value for each tag that you want to add. For more information, see Tagging IAM Users and Roles.

10. When you're done, choose **Next: Review**. Later, you'll create and embed an inline policy for this role.
11. For **Role name**, enter a name that identifies the purpose of this role. Because multiple entities might reference the role, you can't edit the role's name once it is created.
12. (Optional) For **Role description**, enter a description for the new role.
13. Review the role details and choose **Create role**.
14. (Optional) If you plan to use session context or attribute-based application entitlements using a third-party SAML 2.0 identity provider or certificate-based authentication, you must add the sts:TagSession permission to your new IAM role's trust policy. For more information, see Attribute-Based Application Entitlements Using a Third-Party SAML 2.0 Identity Provider and Passing session tags in AWS STS.

    In your new IAM role's details, choose the **Trust relationships** tab, and then choose **Edit trust relationship**. The Edit Trust Relationship policy editor starts. Add the **sts:TagSession** permission, as follows:

    JSON

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "Federated": "arn:aws:iam::111122223333:saml-
provider/IDENTITY-PROVIDER"
            },
            "Action": [
                "sts:AssumeRoleWithSAML",
                "sts:TagSession"
            ],
            "Condition": {
                "StringEquals": {
                    "SAML:sub_type": "persistent"
                }
            }
        }
    ]
}
```

Replace *IDENTITY-PROVIDER* with the name of the SAML IdP you created in Step 1. Then choose **Update Trust Policy**.

Step 3: Embed an Inline Policy for the IAM Role

Next, embed an inline IAM policy for the role that you created. When you embed an inline policy, the permissions in that policy can't be accidentally attached to the wrong principal entity. The inline policy provides federated users with access to the WorkSpaces Applications stack that you created.

1. In the details for the IAM role that you created, choose the **Permissions** tab, and then choose **Add inline policy**. The Create policy wizard starts.
2. In **Create policy**, choose the **JSON** tab.
3. Copy and paste the following JSON policy into the JSON window. Then, modify the resource by entering your AWS Region Code, account ID, and stack name. In the following policy, `"Action": "appstream:Stream"` is the action that provides your WorkSpaces Applications users with permissions to connect to streaming sessions on the stack that you created.

   JSON

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "appstream:Stream",
            "Resource": "arn:aws:appstream:us-
east-1:111122223333:stack/STACK-NAME",
            "Condition": {
                "StringEquals": {
                    "appstream:userId": "${saml:sub}"
                }
            }
        }
    ]
}
```

Replace *REGION-CODE* with the AWS Region where your WorkSpaces Applications stack exists. Replace *STACK-NAME* with the name of the stack. *STACK-NAME* is case-sensitive, and must match the exact case and spelling as the stack name shown in the the **Stacks** dashboard of the WorkSpaces Applications management console.

For resources in the AWS GovCloud (US) Regions, use the following format for the ARN:

`arn:aws-us-gov:appstream:`*REGION-CODE*`:`*ACCOUNT-ID-WITHOUT-HYPHENS*`:stack/`*STACK-NAME*

4. (Optional) If you are planning to use attribute-based application entitlements using a third-party SAML 2.0 identity provider with **SAML 2.0 Multi-stack Application Catalogs**, the Resource in your IAM role inline policy must be **`"Resource": "arn:aws:appstream:`*REGION-CODE*`:`*ACCOUNT-ID-WITHOUT-HYPHENS*`:stack/*"`** to allow application entitlements to control streaming access to stacks. To enforce additional protection on a stack resource, you can add an explicit deny in the policy. For more information, see Attribute-Based Application Entitlements Using a Third-Party SAML 2.0 Identity Provider and Policy evaluation logic.
5. When you're done, choose **Review policy**. The Policy Validator reports any syntax errors.

Step 4: Configure Your SAML-Based IdP

Next, depending on your SAML-based IdP, you may need to manually update your IdP to trust AWS as a service provider by uploading the `saml-metadata.xml` file at https://signin.aws.amazon.com/static/saml-metadata.xml to your IdP. This step updates your IdP's metadata. For some IdPs, the update may already be configured. If this is the case, proceed to the next step.

If this update is not already configured in your IdP, review the documentation provided by your IdP for information about how to update the metadata. Some providers give you the option to type the URL, and the IdP obtains and installs the file for you. Others require you to download the file from the URL and then provide it as a local file.

Step 5: Create Assertions for the SAML Authentication Response

Next, you may need to configure the information that your IdP sends to AWS as SAML attributes in its authentication response. Depending on your IdP, this information may already be pre-configured. If this is the case, skip this step and continue to Step 6,

If this information is not already configured in your IdP, provide the following:

- **SAML Subject NameID** – The unique identifier for the user who is signing in.

> **Note**
>
> For stacks with domain-joined fleets, the NameID value for the user must be provided in the format of "*domain*\username" using the sAMAccountName or "username@domain.com" using userPrincipalName. If you are using the sAMAccountName format, you can specify the *domain* by using either the NetBIOS name or the fully qualified domain name (FQDN). The sAMAccountName format is required for Active Directory one-way trust scenarios. For more information, see Using Active Directory with WorkSpaces Applications.

- **SAML Subject Type** (with a value set to `persistent`) – Setting the value to `persistent` ensures that your IdP sends the same unique value for the `NameID` element in all SAML requests from a particular user. Make sure that your IAM policy includes a condition to only allow SAML requests with a SAML sub_type set to `persistent`, as described in the section called "Step 2: Create a SAML 2.0 Federation IAM Role".
- **`Attribute` element with the `Name` attribute set to https://aws.amazon.com/SAML/Attributes/Role** – This element contains one or more `AttributeValue` elements that list the IAM role and SAML IdP to which the user is mapped by your IdP. The role and IdP are specified as a comma-delimited pair of ARNs.
- **`Attribute` element with the `Name` attribute set to https://aws.amazon.com/SAML/Attributes/RoleSessionName** – This element contains one `AttributeValue` element that provides an identifier for the AWS temporary credentials that are issued for SSO. The value in the `AttributeValue` element must be between 2 and 64 characters long, can contain only alphanumeric characters, underscores, and the following characters: + (plus sign), = (equals sign), , (comma), . (period), @ (at symbol), and - (hyphen). It cannot contain spaces. The value is typically a user ID (bobsmith) or an email address (bobsmith@example.com). It should not be a value that includes a space, such as a user's display name (Bob Smith).
- **`Attribute` element with the `Name` attribute set to https://aws.amazon.com/SAML/Attributes/PrincipalTag:SessionContext (optional)** – This element contains one `AttributeValue` element that provides parameters that can be used to pass session context parameters to your streaming applications. For more information, see Session Context in Amazon WorkSpaces Applications.

- **`Attribute` element with the `Name` attribute set to https://aws.amazon.com/SAML/Attributes/PrincipalTag:ObjectSid (optional)** – This element contains one `AttributeValue` element that provides the Active Directory security identifier (SID) for the user who is signing in. This parameter is used with certificate-based authentication to enable strong mapping to the Active Directory user.
- **`Attribute` element with the `Name` attribute set to https://aws.amazon.com/SAML/Attributes/PrincipalTag:Domain (optional)** – This element contains one `AttributeValue` element that provides the Active Directory DNS fully qualified domain name (FQDN) for the user who is signing in. This parameter is used with certificate-based authentication when the Active Directory `userPrincipalName` for the user contains an alternative suffix. The value must be provided in the format of **`domain.com`**, including any subdomains.
- **`Attribute` element with the `SessionDuration` attribute set to https://aws.amazon.com/SAML/Attributes/SessionDuration (optional)** – This element contains one `AttributeValue` element that specifies the maximum amount of time that a federated streaming session for a user can remain active before reauthentication is required. The default value is 60 minutes, or 3,600 seconds. For more information, see the *An optional Attribute element with the SessionDuration attribute set to https://aws.amazon.com/SAML/Attributes/SessionDuration* section in Configuring SAML Assertions for the Authentication Response.

  > **Note**
  >
  > Although `SessionDuration` is an optional attribute, we recommend that you include it in the SAML response. If you do not specify this attribute, the session duration is set to a default value of 60 minutes.
  > If your users access their streaming applications in WorkSpaces Applications by using the WorkSpaces Applications native client or using the web browser on the new experience, their sessions are disconnected after their session duration expires. If your users access their streaming applications in WorkSpaces Applications by using a web browser on the old/classic experience, after the users' session duration expires and they refresh their browser page, their sessions are disconnected.

For more information about how to configure these elements, see Configuring SAML Assertions for the Authentication Response in the *IAM User Guide*. For information about specific configuration requirements for your IdP, see the documentation for your IdP.

Step 6: Configure the Relay State of Your Federation

Finally, use your IdP to configure the relay state of your federation to point to the WorkSpaces Applications stack relay state URL. After successful authentication by AWS, the user is directed to the WorkSpaces Applications stack portal, defined as the relay state in the SAML authentication response.

The format of the relay state URL is as follows:

```
https://relay-state-region-endpoint?stack=stackname&accountId=aws-account-id-without-hyphens
```

Construct your relay state URL from your Amazon Web Services account ID, stack name, and the relay state endpoint associated with the Region in which your stack is located.

Optionally, you can specify the name of the application that you want to launch automatically. To find the application name, select the image in the WorkSpaces Applications console, choose the **Applications** tab, and note the name that displays in the **Application Name** column. Alternatively, if you haven't yet created the image, connect to the image builder where you installed the application, and open Image Assistant. The names of applications display in the **Add Apps** tab.

If your fleet is enabled for the **Desktop** stream view, you can also choose to launch directly to the operating system desktop. To do so, specify **`Desktop`** at end of the relay state URL, after **`&app=`**.

With an identity provider (IdP) -initiated flow, in the system default browser, after users sign into the IdP and select the WorkSpaces Applications application from the IdP user portal, they are redirected to an WorkSpaces Applications sign-in page in the system default browser with the following options:

- **Continue with Browser**
- **Open WorkSpaces Applications client**

On the page, users can choose to start the session either in the browser, or with the WorkSpaces Applications client application. Optionally, you can also specify which client should be used for a SAML 2.0 federation. To do this, specify either `native` or `web` at end of the relay state URL, after `&client=`. When the parameter is present in a relay state URL, the corresponding sessions will start in the specified client automatically, without users making the choice.

> **Note**
>
> This feature is only available if you use the new relay state region endpoints (in Table 1 below) to construct the relay state URL, and use the WorkSpaces Applications client version 1.1.1300 and later. Also, users should always use their system default browser to sign into the IdP. The feature will not work if they use a non-default browser.

With attribute-based application entitlements using a third-party SAML 2.0 identity provider, you can enable access to multiple stacks from a single relay state URL. Remove the stack and app (if present) parameters from the relay state URL, as follows:

```
https://relay-state-region-endpoint?accountId=aws-account-id-without-hyphens
```

When users federate to the WorkSpaces Applications application catalog, they will be presented with all of the stacks where application entitlements have matched one or more applications to the user for the account ID and relay state endpoint associated with the Region in which your stacks are located. When a user selects a catalog, application entitlements will only display the applications the user is entitled to.

> **Note**
>
> Users cannot stream from multiple stacks at the same time.

For more information, see Attribute-Based Application Entitlements Using a Third-Party SAML 2.0 Identity Provider.

Table 1: WorkSpaces Applications relay state region endpoints (Recommended)

Table 1 below lists the relay state endpoints for the Regions where WorkSpaces Applications is available. The relay state endpoints in Table 1 are compatible with the section called "WorkSpaces Applications Web Browser Access (Version 2)" and the Windows client application version 1.1.1300 and later. If you are using older versions of the Windows client, you should use the old relay state endpoints listed in Table 2 to configure your SAML 2.0 federation. If you want your users to stream using a FIPS-compliant connection, you must use a FIPS-compliant endpoint. For more information about FIPS endpoints, see the section called "FIPS Endpoints".

Table 1: WorkSpaces Applications relay state region endpoints (Recommended)

Region	Relay state endpoint
US East (N. Virginia)	`https://appstream2.euc-sso.us-east-1.aws.amazon.com/saml` (FIPS) `https://appstream2.euc-sso-fips.us-east-1.aws.amazon.com/saml`
US East (Ohio)	`https://appstream2.euc-sso.us-east-2.aws.amazon.com/saml`
US West (Oregon)	`https://appstream2.euc-sso.us-west-2.aws.amazon.com/saml` (FIPS) `https://appstream2.euc-sso-fips.us-west-2.aws.amazon.com/saml`
Asia Pacific (Malaysia)	`https://appstream2.euc-sso.ap-southeast-5.aws.amazon.com/saml`
Asia Pacific (Osaka)	`https://appstream2.euc-sso.ap-northeast-3.aws.amazon.com/saml`
Asia Pacific (Mumbai)	`https://appstream2.euc-sso.ap-south-1.aws.amazon.com/saml`
Asia Pacific (Seoul)	`https://appstream2.euc-sso.ap-northeast-2.aws.amazon.com/saml`
Asia Pacific (Singapore)	`https://appstream2.euc-sso.ap-southeast-1.aws.amazon.com/saml`
Asia Pacific (Sydney)	`https://appstream2.euc-sso.ap-southeast-2.aws.amazon.com/saml`
Asia Pacific (Tokyo)	`https://appstream2.euc-sso.ap-northeast-1.aws.amazon.com/saml`

Region	Relay state endpoint
Canada (Central)	`https://appstream2.euc-sso.ca-central-1.aws.amazon.com/saml`
Canada West (Calgary)	`https://appstream2.euc-sso.ca-west-1.aws.amazon.com/saml`
Europe (Frankfurt)	`https://appstream2.euc-sso.eu-central-1.aws.amazon.com/saml`
Europe (Ireland)	`https://appstream2.euc-sso.eu-west-1.aws.amazon.com/saml`
Europe (London)	`https://appstream2.euc-sso.eu-west-2.aws.amazon.com/saml`
Europe (Milan)	`https://appstream2.euc-sso.eu-south-1.aws.amazon.com/saml`
Europe (Paris)	`https://appstream2.euc-sso.eu-west-3.aws.amazon.com/saml`
Europe (Spain)	`https://appstream2.euc-sso.eu-south-2.aws.amazon.com/saml`
Europe (Zurich)	`https://appstream2.euc-sso.eu-central-2.aws.amazon.com/saml`

Region	Relay state endpoint
AWS GovCloud (US-East)	`https://appstream2.euc-sso.us-gov-east-1.amazonaws-us-gov.com/saml` (FIPS) `https://appstream2.euc-sso-fips.us-gov-east-1.amazonaws-us-gov.com/saml` **Note** For more information about using WorkSpaces Applications in AWS GovCloud (US) Regions, see Amazon WorkSpaces Applications in the *AWS GovCloud (US) User Guide*.
AWS GovCloud (US-West)	`https://appstream2.euc-sso.us-gov-west-1.amazonaws-us-gov.com/saml` (FIPS) `https://appstream2.euc-sso-fips.us-gov-west-1.amazonaws-us-gov.com/saml` **Note** For more information about using WorkSpaces Applications in AWS GovCloud (US) Regions, see Amazon WorkSpaces Applications in the *AWS GovCloud (US) User Guide*.
South America (São Paulo)	`https://appstream2.euc-sso.sa-east-1.aws.amazon.com/saml`
Israel (Tel Aviv)	`https://appstream2.euc-sso.il-central-1.aws.amazon.com/saml`

Table 2: Old WorkSpaces Applications relay state region endpoints (Not recommended)

Table 2 below lists the old relay state endpoints that are still available. However, it is recommended that you use the new relay state endpoints listed in Table 1 to configure your SAML 2.0 federations. In particular, with the new relay state endpoints, you can enable your users to launch the WorkSpaces Applications client application (version 1.1.1300 and later) from IdP-initiated streaming sessions. The new relay state endpoints in Table 1 also allow users to sign into other AWS applications in different tabs of the same web browser, without impacting the ongoing WorkSpaces Applications streaming session. The old relay state endpoints in Table 2 do not support this. For more information, see the section called “My WorkSpaces Applications client users are getting disconnected from their WorkSpaces Applications session after every 60 minutes.”

Table 2: Old WorkSpaces Applications relay state region endpoints (Not recommended)

Region	Relay state endpoint
US East (N. Virginia)	`https://appstream2.us-east-1.aws.amazon.com/saml` (FIPS) `https://appstream2-fips.us-east-1.aws.amazon.com/saml`
US East (Ohio)	`https://appstream2.us-east-2.aws.amazon.com/saml`
US West (Oregon)	`https://appstream2.us-west-2.aws.amazon.com/saml` (FIPS) `https://appstream2-fips.us-west-2.aws.amazon.com/saml`
Asia Pacific (Mumbai)	`https://appstream2.ap-south-1.aws.amazon.com/saml`
Asia Pacific (Seoul)	`https://appstream2.ap-northeast-2.aws.amazon.com/saml`

Region	Relay state endpoint
Asia Pacific (Singapore)	`https://appstream2.ap-southeast-1.aws.amazon.com/saml`
Asia Pacific (Sydney)	`https://appstream2.ap-southeast-2.aws.amazon.com/saml`
Asia Pacific (Tokyo)	`https://appstream2.ap-northeast-1.aws.amazon.com/saml`
Canada (Central)	`https://appstream2.ca-central-1.aws.amazon.com/saml`
Europe (Frankfurt)	`https://appstream2.eu-central-1.aws.amazon.com/saml`
Europe (Ireland)	`https://appstream2.eu-west-1.aws.amazon.com/saml`
Europe (London)	`https://appstream2.eu-west-2.aws.amazon.com/saml`
AWS GovCloud (US-East)	`https://appstream2.us-gov-east-1.amazonaws-us-gov.com/saml` (FIPS) `https://appstream2-fips.us-gov-east-1.amazonaws-us-gov.com/saml` **Note** For more information about using WorkSpaces Applications in AWS GovCloud (US) Regions, see Amazon WorkSpaces Applications in the *AWS GovCloud (US) User Guide*.

Region	Relay state endpoint
AWS GovCloud (US-West)	`https://appstream2.us-gov-west-1.amazonaws-us-gov.com/saml` (FIPS) `https://appstream2-fips.us-gov-west-1.amazonaws-us-gov.com/saml` **Note** For more information about using WorkSpaces Applications in AWS GovCloud (US) Regions, see Amazon WorkSpaces Applications in the *AWS GovCloud (US) User Guide*.
South America (São Paulo)	`https://appstream2.sa-east-1.aws.amazon.com/saml`

Table 3: Relay state URL parameters

Table 3 below lists all of the available parameters that you can use to construct a relay state URL.

Table 3: Relay state URL parameters

Parameter	Required	Format	Supported by
accountId	Required	12-character AWS account ID	New and old endpoints in Table 1 and 2
stack	Optional	Stack name	New and old endpoints in Table 1 and 2
app	Optional	App name or "Desktop"	New and old endpoints in Table 1 and 2

Parameter	Required	Format	Supported by
client	Optional	“native" or "web"	New endpoints in Table 1 only

WorkSpaces Applications Integration with SAML 2.0

The following links help you configure third-party SAML 2.0 identity provider solutions to work with WorkSpaces Applications.

IdP solution	More information
AWS IAM Identity Center	Enable federation with IAM Identity Center and Amazon WorkSpaces Applications — Describes how to use IAM Identity Center to federate user access to your WorkSpaces Applications applications with their existing enterprise credentials.
Active Directory Federation Services (AD FS) for Windows Server	Enabling Identity Federation with AD FS 3.0 and Amazon AppStream 2.0 on the AWS Compute Blog — Describes how to provide users with SSO access to WorkSpaces Applications by using their existing enterprise credentials. You can configure federated identities for WorkSpaces Applications by using AD FS 3.0.
Azure Active Directory (Azure AD)	Enabling Federation with Azure AD Single Sign-On and Amazon WorkSpaces Applications — Describes how to configure federated user access for Amazon WorkSpaces Applications by using Azure AD SSO for enterprise applications.
GG4L School Passport™	Enabling Identity Federation with GG4L’s School Passport™ and Amazon WorkSpaces Applications — Describes how to configure GG4L’s School Passport™ to federate login to WorkSpaces Applications.

IdP solution	More information
Google	Setting up G Suite SAML 2.0 federation with Amazon WorkSpaces Applications — Describes how to use the G Suite Admin console to set up SAML federation to WorkSpaces Applications for users in G Suite domains.
Okta	How to Configure SAML 2.0 for Amazon WorkSpaces Applications — Describes how to use Okta to set up SAML federation to WorkSpaces Applications. For stacks that are joined to a domain, the "Application username format" must be set to "AD user principal name".
Ping Identity	Configuring an SSO connection to Amazon WorkSpaces Applications — Describes how to set up single sign-on (SSO) to WorkSpaces Applications.

IdP solution	More information
Shibboleth	Single Sign-On: Integrating AWS, OpenLDAP, and Shibboleth — Describes how to set up the initial federatio n between the Shibboleth IdP and the AWS Managemen t Console. You must complete the following additional steps to enable federation to WorkSpaces Applications. Step 4 of the AWS Security whitepaper describes how to create IAM roles that define the permissions that federated users have to the AWS Management Console. After you create these roles and embed the inline policy as described in the whitepaper, modify this policy so that it provides federated users with permissions to access only an WorkSpaces Applications stack. To do this, replace the existing policy with the policy noted in *Step 3: Embed an Inline Policy for the IAM Role*, in Setting Up SAML. When you add the stack relay state URL as described in *Step 6: Configure the Relay State of Your Federation*, in Setting Up SAML, add the relay state parameter to the federation URL as a target request attribute. The URL must be encoded. For information about configuring relay state parameters, see the SAML 2.0 section in the Shibboleth documentation. For more information, see Enabling Identity Federation with Shibboleth and Amazon WorkSpaces Applications.
VMware WorkSpace ONE	Federating Access to Amazon WorkSpaces Applications from VMware Workspace ONE — Describes how to use the VMware Workspace ONE platform to federate user access to your WorkSpaces Applications applications.
SimpleSAMLphp	Enabling Federation with SimpleSAMLphp and Amazon WorkSpaces Applications — Describes how to configure SAML 2.0 federation for WorkSpaces Applications using SimpleSAMLphp.

IdP solution	More information
OneLogin Single Sign-On (SSO)	OneLogin SSO with Amazon WorkSpaces Applications — Describes how to configure federated user access for WorkSpaces Applications using OneLogin SSO.
JumpCloud Single Sign-On (SSO)	Enable federation with JumpCloud SSO and Amazon WorkSpaces Applications — Describes how to configure federated user access for WorkSpaces Applications using JumpCloud SSO.
BIO-key PortalGuard	Enable federation with Bio-key PortalGuard and Amazon AppStream 2.0 — Describes how to configure BIO-key PortalGuard for federated logins to WorkSpaces Applicati ons.

For solutions to common problems you may encounter, see Troubleshooting.

For more information about additional supported SAML providers, see Integrating Third-Party SAML Solution Providers with AWS in the *IAM User Guide*.

Using Active Directory with WorkSpaces Applications

You can join your Amazon WorkSpaces Applications Always-On and On-Demand Windows fleets and image builders to domains in Microsoft Active Directory and use your existing Active Directory domains, either cloud-based or on-premises, to launch domain-joined streaming instances. You can also use AWS Directory Service for Microsoft Active Directory, also known as AWS Managed Microsoft AD, to create an Active Directory domain and use that to support your WorkSpaces Applications resources. For more information about using AWS Managed Microsoft AD, see Microsoft Active Directory in the *AWS Directory Service Administration Guide*.

Note

Amazon Linux 2 fleets, image builders, elastic fleets, and app block builders currently do not support domain join.

Active Directory trust requirements

WorkSpaces Applications supports only two-way (bidirectional) trusts between the user domain and the WorkSpaces Applications computer object domain. Selective authentication may be used if WorkSpaces Applications computer objects are allowed to authenticate to the user domain, gather the user's security identifier (SID), and translate the user principal name (UPN) to the Security Account Manager account name (SAMAccountName).

By joining WorkSpaces Applications to your Active Directory domain, you can:

- Allow your users and applications to access Active Directory resources such as printers and file shares from streaming sessions.
- Use Group Policy settings that are available in the Group Policy Management Console (GPMC) to define the end user experience.
- Stream applications that require users to be authenticated using their Active Directory login credentials.
- Apply your enterprise compliance and security policies to your WorkSpaces Applications streaming instances.

Contents

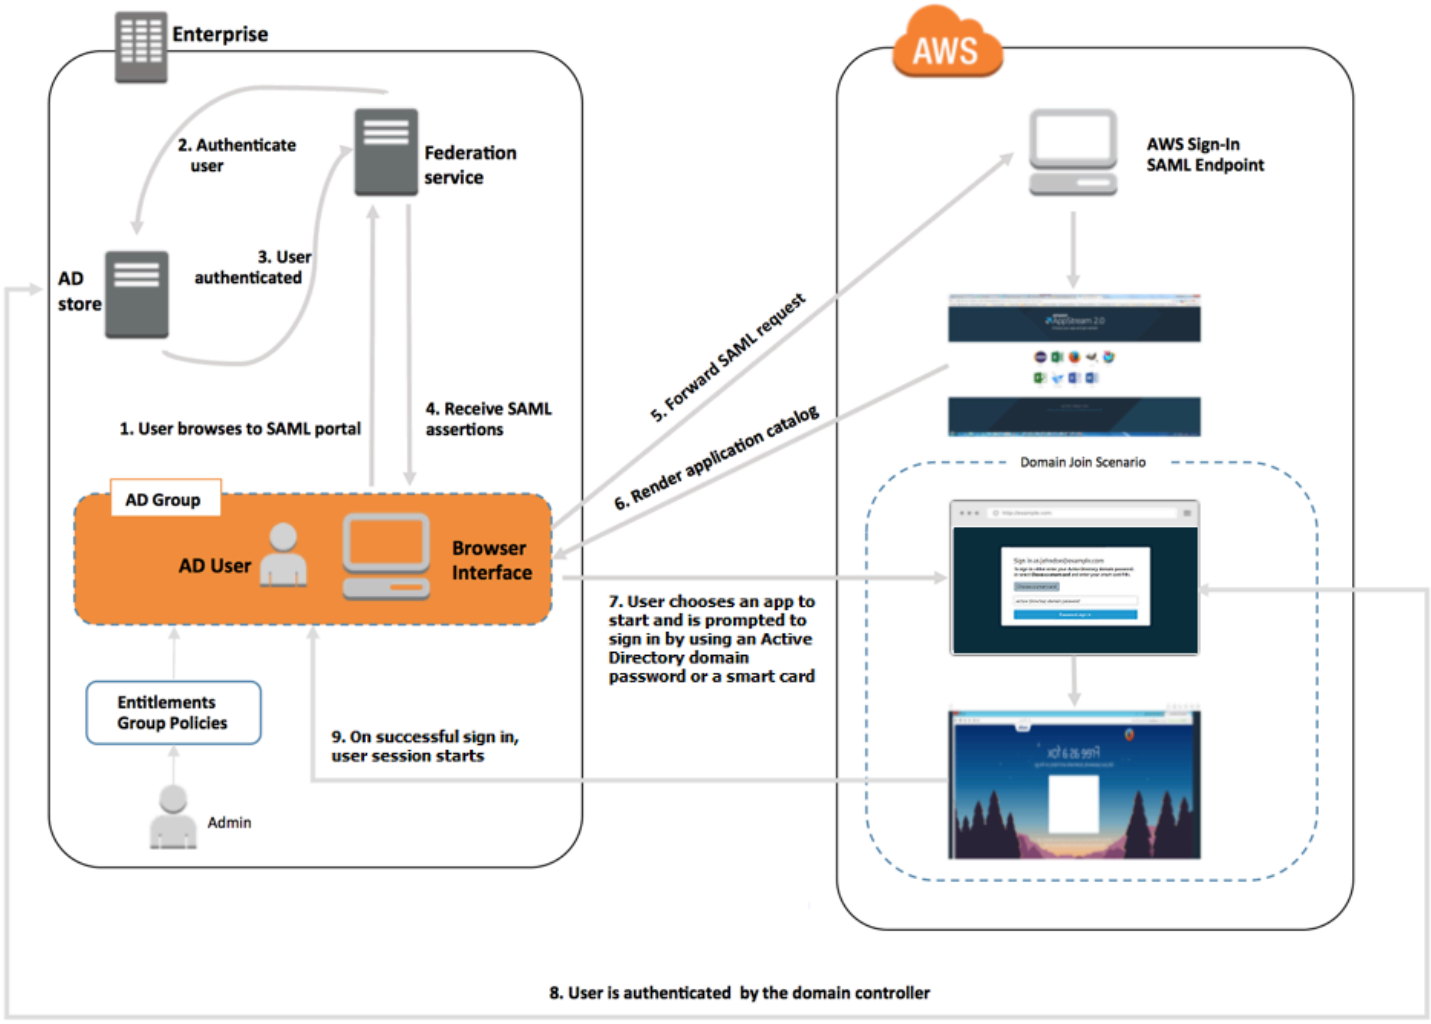
- Overview of Active Directory Domains
- Before You Begin Using Active Directory with Amazon WorkSpaces Applications
- Tutorial: Setting Up Active Directory
- Certificate-Based Authentication
- WorkSpaces Applications Active Directory Administration
- More Info

Overview of Active Directory Domains

Using Active Directory domains with WorkSpaces Applications requires an understanding of how they work together and the configuration tasks that you'll need to complete. You'll need to complete the following tasks:

1. Configure Group Policy settings as needed to define the end user experience and security requirements for applications.
2. Create the domain-joined application stack in WorkSpaces Applications.
3. Create the WorkSpaces Applications application in the SAML 2.0 identity provider and assign it to end users either directly or through Active Directory groups.

For your users to be authenticated to a domain, several steps must occur when these users initiate an WorkSpaces Applications streaming session. The following diagram illustrates the end-to-end user authentication flow from the initial browser request through SAML and Active Directory authentication.



User Authentication Flow

1. The user browses to `https://applications.exampleco.com`. The sign-on page requests authentication for the user.
2. The federation service requests authentication from the organization's identity store.
3. The identity store authenticates the user and returns the authentication response to the federation service.
4. On successful authentication, the federation service posts the SAML assertion to the user's browser.
5. The user's browser posts the SAML assertion to the AWS Sign-In SAML endpoint (`https://signin.aws.amazon.com/saml`). AWS Sign-In receives the SAML request, processes the request, authenticates the user, and forwards the authentication token to the WorkSpaces Applications service.

6. Using the authentication token from AWS, WorkSpaces Applications authorizes the user and presents applications to the browser.
7. The user chooses an application and, depending on the Windows login authentication method that is enabled on the WorkSpaces Applications stack, they're prompted to enter their Active Directory domain password or choose a smart card. If both authentication methods are enabled, the user can choose whether to enter their domain password or use their smart card. Certificate-based authentication can also be used to authenticate users, removing the prompt.
8. The domain controller is contacted for user authentication.
9. After being authenticated with the domain, the user's session starts with domain connectivity.

From the user's perspective, this process is transparent. The user starts by navigating to your organization's internal portal and is redirected to an WorkSpaces Applications application portal, without having to enter AWS credentials. Only an Active Directory domain password or smart card credentials are required.

Before a user can initiate this process, you must configure Active Directory with the required entitlements and Group Policy settings and create a domain-joined application stack.

Before You Begin Using Active Directory with Amazon WorkSpaces Applications

Before you use Microsoft Active Directory domains with WorkSpaces Applications, be aware of the following requirements and considerations.

Contents

- Active Directory Domain Environment
- Domain-Joined WorkSpaces Applications Streaming Instances
- Group Policy Settings
- Smart Card Authentication

Active Directory Domain Environment

Your active directory domain environment must meet the following requirements.

- You must have a Microsoft Active Directory domain to which to join your streaming instances. If you don't have an Active Directory domain or you want to use your on-premises Active Directory environment, see Active Directory Domain Services on AWS Partner Solution Deployment Guide.
- You must have a domain service account with permissions to create and manage computer objects in the domain that you intend to use with WorkSpaces Applications. For information, see How to Create a Domain Account in Active Directory in the Microsoft documentation.

  When you associate this Active Directory domain with WorkSpaces Applications, provide the service account name and password. WorkSpaces Applications uses this account to create and manage computer objects in the directory. For more information, see Granting Permissions to Create and Manage Active Directory Computer Objects.
- When you register your Active Directory domain with WorkSpaces Applications, you must provide an organizational unit (OU) distinguished name. Create an OU for this purpose. The default Computers container is not an OU and cannot be used by WorkSpaces Applications. For more information, see Finding the Organizational Unit Distinguished Name.
- The directories that you plan to use with WorkSpaces Applications must be accessible through their fully qualified domain names (FQDNs) through the virtual private cloud (VPC) in which your streaming instances are launched. For more information, see Active Directory and Active Directory Domain Services Port Requirements in the Microsoft documentation.
- Domain controller access can also be supported over IPv6, and require DHCP options set updates.

Domain-Joined WorkSpaces Applications Streaming Instances

SAML 2.0-based user federation is required for application streaming from domain-joined Always-On and On-Demand fleets. You cannot launch sessions to domain-joined instances by using CreateStreamingURL or the WorkSpaces Applications user pool.

Also, you must use an image that supports joining image builders and fleets to an Active Directory domain. All public images published on or after July 24, 2017 support joining an Active Directory domain. For more information, see WorkSpaces Applications Base Image and Managed Image Update Release Notes and Tutorial: Setting Up Active Directory.

> **Note**
>
> You can join Always-On and On-Demand fleet streaming instances to an Active Directory domain. Supported operating systems include Windows, Red Hat Enterprise Linux, and Rocky Linux.

Group Policy Settings

Verify your configuration for the following Group Policy settings. If required, update the settings as described in this section so that they don't block WorkSpaces Applications from authenticating and logging in your domain users. Otherwise, when your users try to log in to WorkSpaces Applications the login may not succeed. Instead, a message displays, notifying users that "An unknown error occurred."

- **Computer Configuration > Administrative Templates > Windows Components > Windows Logon Options > Disable or Enable software Secure Attention Sequence** — Set this to **Enabled** for **Services**.
- **Computer Configuration > Administrative Templates > System > Logon > Exclude credential providers** — Ensure that the following CLSID are *not* listed: `e7c1bab5-4b49-4e64-a966-8d99686f8c7c` and `f148bAed-5f7f-40c9-8D48-51e24e571825`
- **Computer Configuration > Policies > Windows Settings > Security Settings > Local Policies > Security Options > Interactive Logon > Interactive Logon: Message text for users attempting to log on** — Set this to **Not defined**.
- **Computer Configuration > Policies > Windows Settings > Security Settings > Local Policies > Security Options > Interactive Logon > Interactive Logon: Message title for users attempting to log on** — Set this to **Not defined**.
- **Computer Configuration > Policies > Windows Settings > Security Settings > Local Policies > User Rights Assignment > Allow log on locally** — Set this to **Not defined** or add the domain user/group to this list.
- **Computer Configuration > Policies > Windows Settings > Security Settings > Local Policies > User Rights Assignment > Deny log on locally** — Set this to **Not defined** or make sure that domain users/groups are not included in the list.

If you are using multi-session fleets, you also need the following Group Policy settings, in addition to the settings specified above.

- **Computer Configuration > Policies > Windows Settings > Security Settings > Local Policies > User Rights Assignment > Allow log on through Remote Desktop Services** — Set this to **Not defined** or add the domain user/group to this list.
- **Computer Configuration > Policies > Windows Settings > Security Settings > Local Policies > User Rights Assignment > Deny log on through Remote Desktop Services** — Set this to **Not defined** or make sure that domain users/groups are not included in the list.

Smart Card Authentication

WorkSpaces Applications supports the use of Active Directory domain passwords or smart cards such as Common Access Card (CAC) and Personal Identity Verification (PIV) smart cards for Windows sign in to WorkSpaces Applications streaming instances. For information about how to configure your Active Directory environment to enable smart card sign in by using third-party certification authorities (CAs), see Guidelines for enabling smart card logon with third-party certification authorities in the Microsoft documentation.

> **Note**
>
> WorkSpaces Applications also supports the use of smart cards for in-session authentication after a user signs in to a streaming instance. This feature is supported only for users who have WorkSpaces Applications client for Windows version 1.1.257 or later installed. For information about additional requirements, see Smart Cards.

Tutorial: Setting Up Active Directory

To use Active Directory with WorkSpaces Applications, you must first register your directory configuration by creating a Directory Config object in WorkSpaces Applications. This object includes the information required to join streaming instances to an Active Directory domain. You create a Directory Config object by using the WorkSpaces Applications management console, AWS SDK, or AWS CLI. You can then use your directory configuration to launch domain-joined Always-On and On-Demand fleets and image builders.

> **Note**
>
> You can only join Always-On and On-Demand fleet streaming instances to an Active Directory domain.

Tasks

- Step 1: Create a Directory Config Object
- Step 2: Create an Image by Using a Domain-Joined Image Builder
- Step 3: Create a Domain-Joined Fleet
- Step 4: Configure SAML 2.0

Step 1: Create a Directory Config Object

The Directory Config object that you create in WorkSpaces Applications will be used in later steps.

If you are using the AWS SDK, you can use the CreateDirectoryConfig operation. If you are using the AWS CLI, you can use the create-directory-config command.

To create a Directory Config object by using the WorkSpaces Applications console

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the navigation pane, choose **Directory Configs**, **Create Directory Config**.
3. For **Directory Name**, provide the fully qualified domain name (FQDN) of the Active Directory domain (for example, `corp.example.com`). Each region can have only one **Directory Config** value with a specific directory name.
4. For **Service Account Name**, enter the name of an account that can create computer objects and that has permissions to join the domain. For more information, see Granting Permissions to Create and Manage Active Directory Computer Objects. The account name must be in the format `DOMAIN\username`.
5. For **Password** and **Confirm Password**, type the directory password for the specified account.
6. For **Organizational Unit (OU)**, type the distinguished name of at least one OU for streaming instance computer objects.

> **Note**
>
> The OU name can't contain spaces. If you specify an OU name that contains spaces, when a fleet or image builder attempts to rejoin the Active Directory domain, WorkSpaces Applications cannot cycle the computer objects correctly and the domain rejoin does not succeed. For information about how to troubleshoot this issue, see the *DOMAIN_JOIN_INTERNAL_SERVICE_ERROR* topic for "The account already exists" message in Active Directory Domain Join.
> In addition, the default Computers container is not an OU and cannot be used by WorkSpaces Applications. For more information, see Finding the Organizational Unit Distinguished Name.

7. To add more than one OU, select the plus sign (**+**) next to **Organizational Unit (OU)**. To remove OUs, choose the **x** icon.
8. Choose **Next**.
9. Review the configuration information and choose **Create**.

Step 2: Create an Image by Using a Domain-Joined Image Builder

Next, using the WorkSpaces Applications image builder, create a new image with Active Directory domain-join capabilities. Note that the fleet and image can be members of different domains. You join the image builder to a domain to enable domain join and to install applications. Fleet domain join is discussed in the next section.

To create an image for launching domain-joined fleets

1. Follow the procedures in Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.
2. For the base image selection step, use an AWS base image released on or after July 24, 2017. For a current list of released AWS images, see WorkSpaces Applications Base Image and Managed Image Update Release Notes.
3. For **Step 3: Configure Network**, select a VPC and subnets with network connectivity to your Active Directory environment. Select the security groups that are set up to allow access to your directory through your VPC subnets.

4. Also in **Step 3: Configure Network**, expand the **Active Directory Domain (Optional)** section, and select values for the **Directory Name** and **Directory OU** to which the image builder should be joined.
5. Review the image builder configuration and choose **Create**.
6. Wait for the new image builder to reach a **Running** state, and choose **Connect**.
7. Log in to the image builder in Administrator mode or as a directory user with local administrator permissions. For more information, see Granting Local Administrator Rights on Image Builders.
8. Complete the steps in Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console to install applications and create a new image.

Step 3: Create a Domain-Joined Fleet

Using the private image created in the previous step, create an Active Directory domain-joined Always-On or On-Demand fleet for streaming applications. The domain can be different than the one that you used for the image builder to create the image.

To create a domain-joined Always-On or On-Demand fleet

1. Follow the procedures in Create a Fleet in Amazon WorkSpaces Applications.
2. For the image selection step, use the image that was created in the previous step, Step 2: Create an Image by Using a Domain-Joined Image Builder.
3. For **Step 4: Configure Network**, select a VPC and subnets with network connectivity to your Active Directory environment. Select the security groups that are set up to allow communication to your domain.
4. Also in **Step 4: Configure Network**, expand the **Active Directory Domain (Optional)** section and select the values for the **Directory Name** and **Directory OU** to which the fleet should be joined.
5. Review the fleet configuration and choose **Create**.
6. Complete the remaining steps in Create an Amazon WorkSpaces Applications Fleet and Stack so that your fleet is associated with a stack and running.

Step 4: Configure SAML 2.0

Your users must use your SAML 2.0-based identity federation environment to launch streaming sessions from your domain-joined fleet.

To configure SAML 2.0 for single sign-on access

1. Follow the procedures in Setting Up SAML.
2. WorkSpaces Applications requires that the SAML_Subject `NameID` value for the user who is logging in be provided in either of the following formats:

   - *`domain`*`\username` using the sAMAccountName
   - `username@domain.com` using the userPrincipalName

   If you are using the sAMAccountName format, you can specify the *`domain`* by using either the NetBIOS name or the fully qualified domain name (FQDN).
3. Provide access to your Active Directory users or groups to enable access to the WorkSpaces Applications stack from your identity provider application portal.
4. Complete the remaining steps in Setting Up SAML.

To log in a user with SAML 2.0

1. Log in to your SAML 2.0 provider's application catalog and open the WorkSpaces Applications SAML application that you created in the previous procedure.
2. When the WorkSpaces Applications application catalog is displayed, select an application to launch.
3. When a loading icon is displayed, you are prompted to provide a password. The domain user name provided by your SAML 2.0 identity provider appears above the password field. Enter your password, and choose **log in**.

The streaming instance performs the Windows login procedure, and the selected application opens.

Certificate-Based Authentication

You can use certificate-based authentication with WorkSpaces Applications fleets joined to Microsoft Active Directory. This removes the user prompt for the Active Directory domain password when a user logs in. By using certificate-based authentication with your Active Directory domain, you can:

- Rely on your SAML 2.0 identity provider to authenticate the user and provide SAML assertions to match the user in Active Directory.
- Create a single sign-on logon experience with fewer user prompts.
- Enable passwordless authentication flows using your SAML 2.0 identity provider.

Certificate-based authentication uses AWS Private Certificate Authority (AWS Private CA) resources in your AWS account. With AWS Private CA, you can create private certificate authority (CA) hierarchies, including root and subordinate CAs. You can also create your own CA hierarchy and issue certificates from it that authenticate internal users. For more information, see What is AWS Private CA.

When you use AWS Private CA for certificate-based authentication, WorkSpaces Applications requests certificates for your users automatically at session reservation for each WorkSpaces Applications fleet instance. It authenticates users to Active Directory with a virtual smart card provisioned with the certificates.

Certificate-based authentication (CBA) is supported on WorkSpaces Applications domain-joined fleets (both single-session and multi-session fleets) that run Windows instances. To enable CBA on multi-session fleets, you must use an WorkSpaces Applications image that uses an WorkSpaces Applications agent released on or after 02-07-2025. Or, your image must use managed WorkSpaces Applications image updates released on or after 02-11-2025.

Contents

- Prerequisites
- Enable Certificate-based Authentication
- Manage Certificate-based Authentication
- Enable Cross-account PCA Sharing

Prerequisites

Complete the following steps before you use certificate-based authentication.

1. Set up a domain-joined fleet and configure SAML 2.0. Ensure that you use the `username@domain.com` `userPrincipalName` format for the SAML_Subject `NameID`. For more information, see the section called "Step 5: Create Assertions for the SAML Authentication Response".

   > **Note**
   >
   > Don't enable **Smart card sign in for Active Directory** in your stack if you want to use certificate-based authentication. For more information, see the section called "Smart Cards".

2. Use WorkSpaces Applications agent version 10-13-2022 or later with your image. For more information, see the section called "Keep Your Image Up-to-Date".
3. Configure the `ObjectSid` attribute in your SAML assertion. You can use this attribute to perform strong mapping with the Active Directory user. Certificate-based authentication fails if the `ObjectSid` attribute doesn't match the Active Directory security identifier (SID) for the user specified in the SAML_Subject `NameID`. For more information, see the section called "Step 5: Create Assertions for the SAML Authentication Response". The `ObjectSid` is mandatory for certificate-based authentication after September 10, 2025. For more information, see KB5014754: Certificate-based authentication changes on Windows domain controllers.
4. Add the `sts:TagSession` permission to the IAM role trust policy that you use with your SAML 2.0 configuration. For more information, see Passing session tags in AWS STS. This permission is required to use certificate-based authentication. For more information, see the section called "Step 2: Create a SAML 2.0 Federation IAM Role".
5. Create a private certificate authority (CA) using AWS Private CA, if you don't have one configured with your Active Directory. AWS Private CA is required to use certificate-based authentication. For more information, see Planning your AWS Private CA deployment. The following AWS Private CA settings are common for many certificate-based authentication use cases:
   - **CA type options**
     - **Short-lived certificate CA usage mode** – Recommended if the CA only issues end user certificates for certificate-based authentication.

- **Single level hierarchy with a Root CA** – Choose a subordinate CA to integrate it with an existing CA hierarchy.
- **Key algorithm options** – RSA 2048
- **Subject distinguished name options** – Use the most appropriate options to identify this CA in your Active Directory Trusted Root Certification Authorities store.
- **Certificate revocation options** – CRL distribution

> **Note**
>
> Certificate-based authentication requires an online CRL distribution point accessible from both the WorkSpaces Applications fleet instance and the domain controller. This requires unauthenticated access to the Amazon S3 bucket configured for AWS Private CA CRL entries, or a CloudFront distribution with access to the Amazon S3 bucket if it blocks public access. For more information about these options, see Planning a certificate revocation list (CRL).

6. Tag your private CA with a key entitled `euc-private-ca` to designate the CA for use with WorkSpaces Applications certificate-based authentication. This key doesn't require a value. For more information, see Managing tags for your private CA. For more information about the AWS managed policies used with WorkSpaces Applications to grant permissions to resources in your AWS account, see the section called "AWS Managed Policies Required to Access WorkSpaces Applications Resources".
7. Certificate-based authentication uses virtual smart cards to log on. For more information, see Guidelines for enabling smart card logon with third-party certification authorities. Follow these steps:
   a. Configure domain controllers with a domain controller certificate to authenticate smart card users. If you have an Active Directory Certificate Services enterprise CA configured in your Active Directory, it automatically enrolls domain controllers with certificates that enable smart card logon. If you don't have Active Directory Certificate Services, see Requirements for domain controller certificates from a third-party CA. AWS recommends Active Directory enterprise certificate authorities to automatically manage enrollment for domain controller certificates.

> **Note**
>
> If you use AWS Managed Microsoft AD, you can configure Certificate Services on an Amazon EC2 instance that satisfies the requirement for domain controller certificates. See Deploy Active Directory to a new Amazon Virtual Private Cloud for example deployments of AWS Managed Microsoft AD configured with Active Directory Certificate Services.
> With AWS Managed Microsoft AD and Active Directory Certificate Services, you must also create outbound rules from the controller's VPC security group to the Amazon EC2 instance running Certificate Services. You must provide the security group access to TCP port 135, and ports 49152 through 65535 to enable certificate auto-enrollment. The Amazon EC2 instance must also allow inbound access on these same ports from domain instances, including domain controllers. For more information on locating the security group for AWS Managed Microsoft AD, see Configure your VPC subnets and security groups.

b. On the AWS Private CA console, or with the SDK or CLI, export the private CA certificate. For more information, see Exporting a private certificate.

c. Publish the private CA to Active Directory. Log on to a domain controller or a domain-joined machine. Copy the private CA certificate to any *<path>**<file>* and run the following commands as a domain administrator. You can also use Group Policy and the Microsoft PKI Health Tool (PKIView) to publish the CA. For more information, see Configuration instructions.

```
certutil -dspublish -f <path>\<file> RootCA
```

```
certutil -dspublish -f <path>\<file> NTAuthCA
```

Make sure that the commands complete successfully, then remove the private CA certificate file. Depending on your Active Directory replication settings, it can take several minutes for the CA to publish to your domain controllers and WorkSpaces Applications fleet instances.

> **Note**
>
> Active Directory must distribute the CA to the Trusted Root Certification Authorities and Enterprise NTAuth stores automatically for WorkSpaces Applications fleet instances when they join the domain.

For Windows operating systems, the distribution of the CA (Certificate Authority) happens automatically. However, for Rocky Linux and Red Hat Enterprise Linux, you must download the root CA certificate(s) from the CA used by your WorkSpaces Applications Directory Config. If your KDC root CA certificate(s) are different, you must also download those. Before using certificate-based authentication, it's necessary to import these certificates onto an image or snapshot.

On the image, there should be a file named `/etc/sssd/pki/sssd_auth_ca_db.pem`. It should look like the following:

```
-----BEGIN CERTIFICATE-----
Base64-encoded certificate chain from ACM Private CA
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
Base64-encoded certificate body from ACM private CA
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
Base64-encoded root CA KDC certificate chain
-----END CERTIFICATE-----
```

> **Note**
>
> When copying an image across regions or accounts, or re-associating an image with a new Active Directory, this file will need to be reconfigured with the relevant certificates on an image builder and snapshotted again before use.

Below are instructions for downloading the root CA certificates:

1. On the image builder, create a file named `/etc/sssd/pki/sssd_auth_ca_db.pem`.
2. Open the AWS Private CA console.
3. Choose the private certificate used with your WorkSpaces Applications Directory Config.

4. Choose the **CA certificate** tab.
5. Copy the certificate chain and certificate body to `/etc/sssd/pki/sssd_auth_ca_db.pem` on the image builder.

If the root CA certificates used by the KDCs are different from the root CA certificate used by your WorkSpaces Applications Directory Config, follow these example steps to download them:

1. Connect to a Windows instance joined to the same domain as your image builder.
2. Open `certlm.msc`.
3. In the left pane, choose **Trusted Root Certificate Authorities**, and then choose **Certificates**..
4. For each root CA certificate, open the context (right-click) menu.
5. Choose **All Tasks**, choose **Export** to open the Certificate Export Wizard, and then choose **Next**.
6. Choose **Base64-encoded X.509 (.CER)**, and choose **Next**.
7. Choose **Browse**, enter a file name, and choose **Next**.
8. Choose **Finish**.
9. Open the exported certificate in a text editor.
10. Copy the contents of the file to `/etc/sssd/pki/sssd_auth_ca_db.pem` on the image builder.

Enable Certificate-based Authentication

Complete the following steps to enable certificate-based authentication.

To enable certificate-based authentication

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the navigation pane, choose **Directory Configs**. Select the directory config you want to configure, and choose **Edit**.
3. Choose **Enable Certificate-Based Authentication.**
4. Confirm that your private CA ARN is associated in the list. To appear in the list, you should store the private CA in the same AWS account and AWS Region. You must also tag the private CA with a key named `euc-private-ca`.
5. Configure directory log in fallback. With Fallback, users can log in with their AD domain password if certificate-based authentication is unsuccessful. This is recommended only in

cases where users know their domain passwords. When fallback is turned off, a session can disconnect the user if a lock screen or Windows log off occurs. If fallback is turned on, the session prompts the user for their AD domain password.

6. Choose **Save Changes**.
7. Certificate-based authentication is now enabled. When users authenticate with SAML 2.0 to an WorkSpaces Applications stack using the domain-joined fleet from the WorkSpaces Applications web client or the client for Windows (version 1.1.1099 and later), they will no longer receive a prompt for the domain password. Users will see a “Connecting with certificate-based authentication...” message when connecting to a session enabled for certificate-based authentication.

Manage Certificate-based Authentication

After you enable certificate-based authentication, review the following tasks.

Topics

- Private CA Certificate
- End User Certificates
- Audit Reports
- Logging and Monitoring

Private CA Certificate

In a typical configuration, the private CA certificate has a validity period of 10 years. For more information about replacing a private CA with an expired certificate, or reissuing the private CA with a new validity period, see Managing the private CA lifecycle

End User Certificates

End user certificates issued by AWS Private CA for WorkSpaces Applications certificate-based authentication don't require renewal or revocation. These certificates are short-lived. WorkSpaces Applications automatically issues a new certificate for each new session, or every 24 hours for sessions with a long duration. The WorkSpaces Applications session governs the use of these end user certificates. If you end a session, WorkSpaces Applications stops using that certificate. These end user certificates have a shorter validity period than a typical AWS Private CA CRL distribution. As a result, end user certificates don't need to be revoked and won't appear in a CRL.

Audit Reports

You can create an audit report to list all of the certificates that your private CA has issued or revoked. For more information, see Using audit reports with your private CA.

Logging and Monitoring

You can use CloudTrail to record API calls to a private CA by WorkSpaces Applications. For more information see What Is AWS CloudTrail? and Using CloudTrail. In CloudTrail Event history you can view **GetCertificate** and **IssueCertificate** event names from **acm-pca.amazonaws.com** event source made by the WorkSpaces Applications**EcmAssumeRoleSession** user name. These events will be recorded for every WorkSpaces Applications certificate-based authentication request. For more information, see Viewing events with CloudTrail Event history.

Enable Cross-account PCA Sharing

Private CA (PCA) cross-account sharing offers the ability to grant permissions for other accounts to use a centralized CA. The CA can generate and issue certificates by using AWS Resource Access Manager (RAM) to manage the permissions. This removes the need for a Private CA in every account. Private CA cross-account sharing can be used with WorkSpaces Applications certificate-based Authentication (CBA) within the same AWS Region.

To use a shared Private CA resource with WorkSpaces Applications CBA, complete the following steps:

1. Configure the Private CA for CBA in a centralized AWS account. For more information, see the section called “Certificate-Based Authentication”.
2. Share the Private CA with the resource AWS accounts where WorkSpaces Applications resources utilize CBA. To do this, follow the steps in How to use AWS RAM to share your ACM Private CA cross-account. You do not need to complete step 3 to create a certificate. You can either share the Private CA with individual AWS accounts, or share through AWS Organizations. If you share with individual accounts, you need to accept the shared Private CA in your resource account by using the AWS Resource Access Manager console or APIs.

   When configuring the share, confirm that the AWS Resource Access Manager resource share for the Private CA in the resource account is using the `AWSRAMBlankEndEntityCertificateAPICSRPassthroughIssuanceCertificateAuthority` managed permission template. This template aligns with the PCA template used by the WorkSpaces Applications service role when issuing CBA certificates.

3. After the share is successful, view the shared Private CA by using the Private CA console in the resource account.
4. Use the API or CLI to associate the Private CA ARN with CBA in your WorkSpaces Applications Directory Config. At this time, the WorkSpaces Applications console does not support selection of shared Private CA ARNs. The following are example CLI commands:

```
aws appstream update-directory-config --directory-
name <value> --certificate-based-auth-properties
Status=<value>,CertificateAuthorityArn=<value>
```

WorkSpaces Applications Active Directory Administration

Setting up and using Active Directory with WorkSpaces Applications involves the following administrative tasks.

Tasks

- Granting Permissions to Create and Manage Active Directory Computer Objects
- Finding the Organizational Unit Distinguished Name
- Granting Local Administrator Rights on Image Builders
- Updating the Service Account Used for Joining the Domain
- Locking the Streaming Session When the User is Idle
- Editing the Directory Configuration
- Deleting a Directory Configuration
- Configuring WorkSpaces Applications to Use Domain Trusts
- Managing WorkSpaces Applications Computer Objects in Active Directory
- Configuring Active Directory Sites and Services for WorkSpaces Applications

Granting Permissions to Create and Manage Active Directory Computer Objects

To allow WorkSpaces Applications to perform Active Directory computer object operations, you need an account with sufficient permissions. As a best practice, use an account that has only the minimum privileges necessary. The minimum Active Directory organizational unit (OU) permissions are as follows:

- Create Computer Object
- Change Password
- Reset Password
- Write Description

Before setting up permissions, you'll need to do the following first:

- Obtain access to a computer or an EC2 instance that is joined to your domain.
- Install the Active Directory User and Computers MMC snap-in. For more information, see Installing or Removing Remote Server Administration Tools for Windows 7 in the Microsoft documentation.
- Log in as a domain user with appropriate permissions to modify the OU security settings.
- Create or identify the user, service account, or group for which to delegate permissions.

To set up minimum permissions

1. Open **Active Directory Users and Computers** in your domain or on your domain controller.
2. In the left navigation pane, select the first OU on which to provide domain join privileges, open the context (right-click) menu , and then choose **Delegate Control**.
3. On the **Delegation of Control Wizard** page, choose **Next**, **Add**.
4. For **Select Users, Computers, or Groups**, select the pre-created user, service account, or group, and then choose **OK**.
5. On the **Tasks to Delegate** page, choose **Create a custom task to delegate**, and then choose **Next**.
6. Choose **Only the following objects in the folder**, **Computer objects**.
7. Choose **Create selected objects in this folder**, **Next**.
8. For **Permissions**, choose **Read**, **Write**, **Change Password**, **Reset Password**, **Next**.
9. On the **Completing the Delegation of Control Wizard** page, verify the information and choose **Finish**.
10. Repeat steps 2-9 for any additional OUs that require these permissions.

If you delegated permissions to a group, create a user or service account with a strong password and add that account to the group. This account will then have sufficient privileges to connect your

streaming instances to the directory. Use this account when creating your WorkSpaces Applications directory configuration.

Finding the Organizational Unit Distinguished Name

When you register your Active Directory domain with WorkSpaces Applications, you must provide an organizational unit (OU) distinguished name. Create an OU for this purpose. The default Computers container is not an OU and cannot be used by WorkSpaces Applications. The following procedure shows how to obtain this name.

> **Note**
>
> The distinguished name must start with **OU=** or it cannot be used for computer objects.

Before you complete this procedure, you'll need to do the following first:

- Obtain access to a computer or an EC2 instance that is joined to your domain.
- Install the Active Directory User and Computers MMC snap-in. For more information, see Installing or Removing Remote Server Administration Tools for Windows 7 in the Microsoft documentation.
- Log in as a domain user with appropriate permissions to read the OU security properties.

To find the distinguished name of an OU

1. Open **Active Directory Users and Computers** in your domain or on your domain controller.
2. Under **View**, ensure that **Advanced Features** is enabled.
3. In the left navigation pane, select the first OU to use for WorkSpaces Applications streaming instance computer objects, open the context (right-click) menu, and then choose **Properties**.
4. Choose **Attribute Editor**.
5. Under **Attributes**, for **distinguishedName**, choose **View**.
6. For **Value**, select the distinguished name, open the context menu, and then choose **Copy**.

Granting Local Administrator Rights on Image Builders

By default, Active Directory domain users do not have local administrator rights on image builder instances. You can grant these rights by using Group Policy preferences in your directory, or manually, by using the local administrator account on an image builder. Granting local administrator rights to a domain user allows that user to install applications on and create images in an WorkSpaces Applications image builder.

Contents

- Using Group Policy preferences
- Using the local Administrators group on the image builder

Using Group Policy preferences

You can use Group Policy preferences to grant local administrator rights to Active Directory users or groups and to all computer objects in the specified OU. The Active Directory users or groups to which you want to grant local administrator permissions must already exist. To use Group Policy preferences, you'll need to do the following first:

- Obtain access to a computer or an EC2 instance that is joined to your domain.
- Install the Group Policy Management Console (GPMC) MMC snap-in. For more information, see Installing or Removing Remote Server Administration Tools for Windows 7 in the Microsoft documentation.
- Log in as a domain user with permissions to create Group Policy objects (GPOs). Link GPOs to the appropriate OUs.

To use Group Policy preferences to grant local administrator permissions

1. In your directory or on a domain controller, open the command prompt as an administrator, type `gpmc.msc`, and then press ENTER.
2. In the left console tree, select the OU where you will create a new GPO or use an existing GPO, and then do either of the following:
   - Create a new GPO by opening the context (right-click) menu and choosing **Create a GPO in this domain, Link it here**. For **Name**, provide a descriptive name for this GPO.
   - Select an existing GPO.

3. Open the context menu for the GPO, and choose **Edit**.
4. In the console tree, choose **Computer Configuration**, **Preferences**, **Windows Settings**, **Control Panel Settings**, and **Local Users and Groups**.
5. Select **Local Users and Groups** selected, open the context menu , and choose **New**, **Local Group**.
6. For **Action**, choose **Update**.
7. For **Group name**, choose **Administrators (built-in)**.
8. Under **Members**, choose **Add...** and specify the Active Directory users or groups to which to assign local administrator rights on the streaming instance. For **Action**, choose **Add to this group**, and choose **OK**.
9. To apply this GPO to other OUs, select the additional OU, open the context menu and choose **Link an Existing GPO**.
10. Using the new or existing GPO name that you specified in step 2, scroll to find the GPO, and then choose **OK**.
11. Repeat steps 9 and 10 for additional OUs that should have this preference.
12. Choose **OK** to close the **New Local Group Properties** dialog box.
13. Choose **OK** again to close the GPMC.

To apply the new preference to the GPO, you must stop and restart any running image builders or fleets. The Active Directory users and groups that you specified in step 8 are automatically granted local administrator rights on the image builders and fleets in the OU to which the GPO is linked.

Using the local Administrators group on the image builder

To grant Active Directory users or groups local administrator rights on your image builder, you can manually add these users or groups to the local Administrators group on the image builder. Image builders that are created from images with these rights maintain the same rights.

The Active Directory users or groups to which to grant local administrator rights must already exist.

To add Active Directory users or groups to the local Administrators group on the image builder

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. Connect to the image builder in Administrator mode. The image builder must be running and domain-joined. For more information, see Tutorial: Setting Up Active Directory.
3. Choose **Start**, **Administrative Tools**, and then double-click **Computer Management**.

4. In the left navigation pane, choose **Local Users and Groups** and open the **Groups** folder.
5. Open the **Administrators** group and choose **Add...**.
6. Select all Active Directory users or groups to which to assign local administrator rights and choose **OK**. Choose **OK** again to close the **Administrator Properties** dialog box.
7. Close Computer Management.
8. To log in as an Active Directory user and test whether that user has local administrator rights on the image builder, choose **Admin Commands**, **Switch user**, and then enter the credentials of the relevant user.

Updating the Service Account Used for Joining the Domain

To update the service account that WorkSpaces Applications uses for joining the domain, we recommend using two separate service accounts for joining image builders and fleets to your Active Directory domain. Using two separate service accounts ensures that there is no disruption in service when a service account needs to be updated (for example, when a password expires).

To update a service account

1. Create an Active Directory group and delegate the correct permissions to the group.
2. Add your service accounts to the new Active Directory group.
3. When needed, edit your WorkSpaces Applications Directory Config object by entering the sign-in credentials for the new service account.

After you've set up the Active Directory group with the new service account, any new streaming instance operations will use the new service account, while in-process streaming instance operations continue to use the old account without interruption.

The service account overlap time while the in-process streaming instance operations complete is very short, no more than a day. The overlap time is needed because you shouldn't delete or change the password for the old service account during the overlap period, or existing operations can fail.

Locking the Streaming Session When the User is Idle

WorkSpaces Applications relies on a setting that you configure in the GPMC to lock the streaming session after your user is idle for specified amount of time. To use the GPMC, you'll need to do the following first:

- Obtain access to a computer or an EC2 instance that is joined to your domain.
- Install the GPMC. For more information, see Installing or Removing Remote Server Administration Tools for Windows 7 in the Microsoft documentation.
- Log in as a domain user with permissions to create GPOs. Link GPOs to the appropriate OUs.

To automatically lock the streaming instance when your user is idle

1. In your directory or on a domain controller, open the command prompt as an administrator, type `gpmc.msc`, and then press ENTER.
2. In the left console tree, select the OU where you will create a new GPO or use an existing GPO, and then do either of the following:
   - Create a new GPO by opening the context (right-click) menu and choosing **Create a GPO in this domain, Link it here**. For **Name**, provide a descriptive name for this GPO.
   - Select an existing GPO.
3. Open the context menu for the GPO, and choose **Edit**.
4. Under **User Configuration**, expand **Policies**, **Administrative Templates**, **Control Panel**, and then choose **Personalization**.
5. Double-click **Enable screen saver**.
6. In the **Enable screen saver** policy setting, choose **Enabled**.
7. Choose **Apply**, and then choose **OK**.
8. Double-click **Force specific screen saver**.
9. In the **Force specific screen saver** policy setting, choose **Enabled**.
10. Under **Screen saver executable name**, enter **`scrnsave.scr`**. When this setting is enabled, the system displays a black screen saver on the user's desktop.
11. Choose **Apply**, and then choose **OK**.
12. Double-click **Password protect the screen saver**.
13. In the **Password protect the screen saver** policy setting, choose **Enabled**.
14. Choose **Apply**, and then choose **OK**.
15. Double-click **Screen saver timeout**.
16. In the **Screen saver timeout** policy setting, choose **Enabled**.
17. For **Seconds**, specify the length of time that users must be idle before the screen saver is applied. To set the idle time to 10 minutes, specify 600 seconds.

18. Choose **Apply**, and then choose **OK**.
19. In the console tree, under **User Configuration**, expand **Policies**, **Administrative Templates**, **System**, and then choose **Ctrl+Alt+Del Options**.
20. Double-click **Remove Lock Computer**.
21. In the **Remove Lock Computer** policy setting, choose **Disabled**.
22. Choose **Apply**, and then choose **OK**.

Editing the Directory Configuration

After a WorkSpaces Applications directory configuration has been created, you can edit it to add, remove, or modify organizational units, update the service account username, or update the service account password.

To update a directory configuration

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the left navigation pane, choose **Directory Configs** and select the directory configuration to edit.
3. Choose **Actions**, **Edit**.
4. Update the fields to be changed. To add additional OUs, select the plus sign (**+**) next to the topmost OU field. To remove an OU field, select the **x** next to the field.

   > **Note**
   >
   > At least one OU is required. OUs that are currently in use cannot be removed.

5. To save changes, choose **Update Directory Config**.
6. The information in the **Details** tab should now update to reflect the changes.

Changes to the service account sign-in credentials do not impact in-process streaming instance operations. New streaming instance operations use the updated credentials. For more information, see Updating the Service Account Used for Joining the Domain.

Deleting a Directory Configuration

You can delete an WorkSpaces Applications directory configuration that is no longer needed. Directory configurations that are associated with any image builders or fleets cannot be deleted.

To delete a directory configuration

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the left navigation pane, choose **Directory Configs** and select the directory configuration to delete.
3. Choose **Actions**, **Delete**.
4. Verify the name in the pop-up message, and choose **Delete**.
5. Choose **Update Directory Config**.

Configuring WorkSpaces Applications to Use Domain Trusts

WorkSpaces Applications supports Active Directory domain environments where network resources such as file servers, applications, and computer objects reside in one domain, and the user objects reside in another. The domain service account used for computer object operations does not need to be in the same domain as the WorkSpaces Applications computer objects.

When creating the directory configuration, specify a service account that has the appropriate permissions to manage computer objects in the Active Directory domain where the file servers, applications, computer objects and other network resources reside.

Your end user Active Directory accounts must have the "Allowed to Authenticate" permissions for the following:

- WorkSpaces Applications computer objects
- Domain controllers for the domain

For more information, see Granting Permissions to Create and Manage Active Directory Computer Objects.

Managing WorkSpaces Applications Computer Objects in Active Directory

WorkSpaces Applications does not delete computer objects from Active Directory. These computer objects can be easily identified in your directory. Each computer object in the directory is created with the `Description` attribute, which specifies a fleet or an image builder instance and the name.

Computer Object Description Examples

Type	Name	Description Attribute
Fleet	ExampleFleet	`WorkSpaces Applications - fleet:ExampleFleet`
Image builder	ExampleImageBuilder	`WorkSpaces Applications - image-builder:ExampleImageBuilder`

You can identify and delete inactive computer objects created by WorkSpaces Applications by using the following `dsquery computer` and `dsrm` commands. For more information, see Dsquery computer and Dsrm in the Microsoft documentation.

The `dsquery` command identifies inactive computer objects over a certain period of time and uses the following format. The `dsquery` command should also be run with the parameter `-desc "WorkSpaces Applications*"` to display only WorkSpaces Applications objects.

```
dsquery computer "OU-distinguished-name" -desc "WorkSpaces Applications*" -
inactive number-of-weeks-since-last-login
```

- *`OU-distinguished-name`* is the distinguished name of the organizational unit. For more information, see Finding the Organizational Unit Distinguished Name. If you don't provide the *`OU-distinguished-name`* parameter, the command searches the entire directory.
- *`number-of-weeks-since-last-log-in`* is the desired value based on how you want to define inactivity.

For example, the following command displays all computer objects in the `OU=ExampleOU,DC=EXAMPLECO,DC=COM` organizational unit that have not been logged into within the past two weeks.

```
dsquery computer OU=ExampleOU,DC=EXAMPLECO,DC=COM -desc "WorkSpaces Applications*" -
inactive 2
```

If any matches are found, the result is one or more object names. The `dsrm` command deletes the specified object and uses the following format:

```
dsrm objectname
```

Where *`objectname`* is the full object name from the output of the `dsquery` command. For example, if the `dsquery` command above results in a computer object named "ExampleComputer", the `dsrm` command to delete it would be as follows:

```
dsrm "CN=ExampleComputer,OU=ExampleOU,DC=EXAMPLECO,DC=COM"
```

You can chain these commands together by using the pipe (`|`) operator. For example, to delete all WorkSpaces Applications computer objects, prompting for confirmation for each, use the following format. Add the `-noprompt` parameter to `dsrm` to disable confirmation.

```
dsquery computer OU-distinguished-name -desc "WorkSpaces Applications*" –
inactive number-of-weeks-since-last-log-in | dsrm
```

Configuring Active Directory Sites and Services for WorkSpaces Applications

As with all multi-site Active Directory deployments, Windows authentication within WorkSpaces Applications is most reliable when Active Directory Sites and Services includes the WorkSpaces Applications fleet subnets. Assign those subnets to a site that includes a domain controller in the same AWS Region as WorkSpaces Applications to prevent logon and replication delays.

Because WorkSpaces Applications streaming instances are ephemeral, they are frequently terminated, re-provisioned, and joined to the domain with new computer object credentials. If you configure Sites and Services incorrectly, the Windows domain join mechanism might target one site for the domain join and another site for authentication. In larger Active Directory

environments where the sites are not directly linked, syncing the latest domain join information to the randomly selected authentication site can take multiple hours. This delay occurs because site-to-site replication is three hours by default. The delay can create database inconsistencies for the computer account and cause the user's logon to fail.

At a high level, a basic configuration uses two existing domain controllers. Domain controller A is in the site `"us-east-1-site"`, and domain controller B is in the site `"us-east-2-site"`. To configure Sites and Services for this setup, complete the following steps:

1. Create WorkSpaces Applications fleet A in `us-east-1` in the summary subnet `10.0.0.0/24`.
2. Create WorkSpaces Applications fleet B in `us-east-2` in the summary subnet `172.16.0.0/24`.
3. Create a new subnet (`10.0.0.0/24`) in Active Directory Sites and Services and assign it to `"us-east-1-site"`.
4. Create a new subnet (`172.16.0.0/24`) in Active Directory Sites and Services and assign it to `"us-east-2-site"`.

With this configuration, WorkSpaces Applications fleet A targets domain controller A as the primary, with failover to domain controller B. WorkSpaces Applications fleet B targets domain controller B as the primary, with failover to domain controller A.

For more information about Microsoft Active Directory Sites and Services and the domain controller locator process, see the following Microsoft documentation: Designing the site topology and How Domain Controllers Are Located in Windows.

More Info

For more information related to this topic, see the following resources:

- Troubleshooting Notification Codes—Resolutions to notification code errors.
- Troubleshooting Active Directory—Help with common difficulties.
- Microsoft Active Directory—Information about using Directory Service.

Add Your Custom Branding to Amazon WorkSpaces Applications

To create a familiar experience for your users when they stream applications, you can customize the appearance of WorkSpaces Applications with your own branding images, text, and website links, and you can choose from one of several color palettes. When you customize WorkSpaces Applications, your branding is displayed to users during application streaming sessions rather than the default WorkSpaces Applications branding.

Topics

- Custom Branding Options for Amazon WorkSpaces Applications
- Adding Your Custom Branding to Amazon WorkSpaces Applications
- Specifying a Custom Redirect URL and Feedback URL in Amazon WorkSpaces Applications
- Previewing Your Custom Branding Changes in Amazon WorkSpaces Applications
- Color Theme Palettes in Amazon WorkSpaces Applications

Custom Branding Options for Amazon WorkSpaces Applications

You can customize the appearance of the streaming application catalog page by using the following branding options.

> **Note**
>
> Custom branding is not available for the user pool sign-in portal or for the email notifications that WorkSpaces Applications sends to user pool users.

Branding element	Description	Requirements and recommendations
Organization logo	Enables you to display an image that is familiar to your users. The image appears	File type: .png, jpg, .jpeg, or .gif Maximum dimensions: 1000 px x 500 px Maximum file size: 300 KB

Branding element	Description	Requirements and recommendations
	in the header of the streaming application catalog page, which is displayed to users after they sign in to WorkSpaces Applicati ons.	
Organization website links	Enables you to display links to helpful resources for your users, such as your organizat ion's IT support and product marketing sites. The links are displayed in the footer of the streaming application catalog page.	Maximum number of links: 3 Format (URL): https://example.com or http://ex ample.com Maximum length (display name): 100 letters, spaces, and numbers Special characters allowed (display name): @ . / # & + $
Color theme	Applied to website links, text, and buttons. These colors are also applied as accents in the background for the streaming application catalog page.	Predefined themes from which to choose: 4 For information about each color theme, see Color Theme Palettes in Amazon WorkSpaces Applicati ons later in this topic.
Page title	Displayed at the top of the browser tab during users' applicati on streaming sessions.	Maximum length: 200 letters, spaces, and numbers. Special characters allowed: @ . / # & + $

Branding element	Description	Requirements and recommendations
Favicon	Enables your users to recognize their application streaming site in a browser full of tabs or bookmarks. The favicon icon is displayed at the top of the browser tab for the application streaming site during users' streaming sessions.	File type: .png, .jpg, .jpeg, .gif, or .ico Maximum dimensions: 128 px x 128 px Maximum file size: 50 KB
Redirect URL	Enables you to specify a URL to which users are redirected when they end a streaming session.	Format: https://example.com or http://example.com This URL is configured in the **Details** page for a stack when you create or edit a stack, not in the **Branding** page.
Feedback URL	Enables you to specify a URL for a **Send Feedback** link, so that your users can submit feedback to the organization. If you do not specify a URL, the **Send Feedback** link is not displayed. Your users can still submit new portal feedback by selecting **Provide New Portal Feedback**, which is submitted to AWS.	Format: https://example.com or http://example.com This URL is configured in the **Details** page for a stack when you create or edit a stack, not in the **Branding** page.

Adding Your Custom Branding to Amazon WorkSpaces Applications

To customize WorkSpaces Applications with your organizational branding, use the WorkSpaces Applications console to select the stack to customize, and then add your branding.

If you want to choose your organization logo or favicon from your Amazon S3 buckets, make sure that your Amazon S3 bucket content is not encrypted using keys that you manage (Customer Managed Keys). Amazon S3 buckets configured to use server-side encryption with customer-provided encryption keys (SSE-C) are not supported for organization logo and favicon. If you require encryption at rest for your Amazon S3 objects, server-side encryption with Amazon S3-managed encryption keys (SSE-S3) is an option for organization logo and favicon.

To add your custom branding to WorkSpaces Applications

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the left pane, choose **Stacks**.
3. In the stack list, select the stack to customize with your branding.
4. Choose **Branding**, **Custom**.
5. For **Application catalog page**, customize how the streaming application catalog page appears to users after they sign in to WorkSpaces Applications.

   a. For **Organization logo**, do either of the following:

      - Either enter the Amazon S3 URI that represents the organization logo, or choose **Browse S3** to navigate to your Amazon S3 buckets and find the organization logo.
      - If you already uploaded a organization logo and want to view it, choose **View**. You can change it by entering another Amazon S3 URI that represents the organization logo, or choose **Browse S3** to find another organization logo.

   b. For **Organization website links**, specify up to three website links to display in the page footer. For each link, choose the **Add Link** button, and then enter a display name and URL. To add more links, repeat these steps for each link to add. To remove a link, choose the **Remove** button under the link URL.

c. For **Color theme**, choose the colors to use for your website links, body text, and buttons, and as an accent for the page background. For information about each color theme, see Color Theme Palettes in Amazon WorkSpaces Applications later in this topic.

6. For **Browser tab**, customize the page title and icon to display to users at the top of their browser tab during streaming sessions.

   a. For **Page title**, enter the title to display at the top of the browser tab.

   b. For **Favicon**, do either of the following:

      - Enter the Amazon S3 URI that represents the favicon, or choose **Browse S3** to navigate to your Amazon S3 buckets and find the favicon.
      - If you already uploaded a favicon and want to view it, **choose View**. Or, you can change it by entering another Amazon S3 URI that represents the favicon, or choose **Browse S3** to find another favicon.

7. Do either of the following:

   - To apply your branding changes, choose **Save**. When users connect to new streaming sessions that are launched for the stack, your branding changes are displayed.

     > **Note**
     >
     > WorkSpaces Applications retains the custom branding changes that you save. If you save your custom branding changes, but then choose to restore the WorkSpaces Applications default branding, your custom branding changes are saved for later use. If you restore the WorkSpaces Applications default branding and decide later to reapply your custom branding, choose **Custom**, **Save**. In this case, the most recently saved custom branding is displayed to your users.

   - To discard your branding changes, choose **Cancel**. When prompted to confirm your choice, choose **Confirm**. If you cancel your changes, the most recently saved branding is displayed to your users.

Specifying a Custom Redirect URL and Feedback URL in Amazon WorkSpaces Applications

You can specify a URL to which your users are redirected when they end their streaming session, as well as a URL where your users can submit feedback. By default, WorkSpaces Applications displays a **Send Feedback** link that enables users to submit feedback to AWS about the quality of their application streaming session. To enable your users to submit feedback to a site that you specify, you can provide a custom feedback URL. You can specify the redirect URL and feedback URL when you create a new stack or edit the details for an existing stack. For more information, see Create a Stack in Amazon WorkSpaces Applications.

Previewing Your Custom Branding Changes in Amazon WorkSpaces Applications

You can preview how your branding changes will appear to your users by applying your branding changes to a test stack before you apply them to a production stack, and then creating a streaming URL for the test stack. After you validate your branding changes, you can them deploy them to your production stack. For information, see Step 2: Provide Access to Users in *Getting Started with Amazon WorkSpaces Applications*.

Color Theme Palettes in Amazon WorkSpaces Applications

When you choose a color theme, the colors for that theme are applied to the website links, text, and buttons in your streaming application catalog page. A color is also applied as an accent in the background for your streaming application catalog page. For each color in a color theme palette, the hex value is also noted.

Color Themes

- Red
- Light Blue
- Blue
- Pink

Red

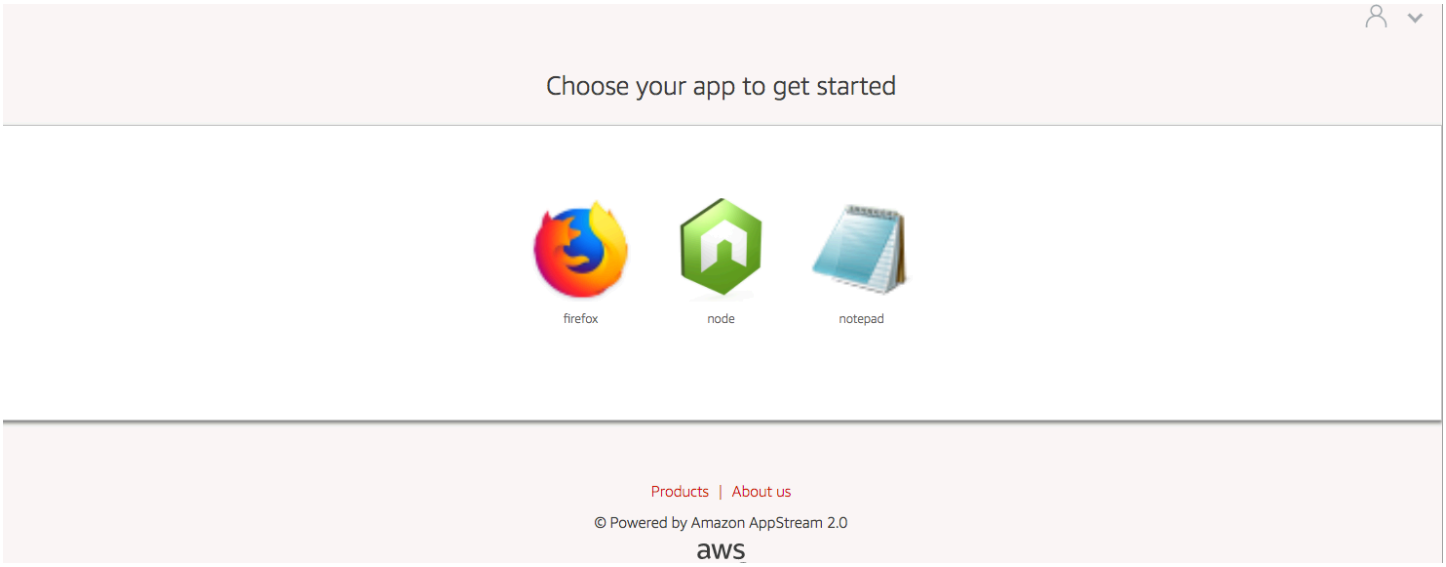
The following colors are applied when you select the red color theme.

Red (#d51900) – Used for buttons and website links.

White (#faf9f7) – Used as a background accent.

Dark grey (#404040) – Used for the body text and in the progress spinner.

When you choose the red color theme, the website links, body text, and background accent appear in your streaming application catalog page as follows.



Light Blue

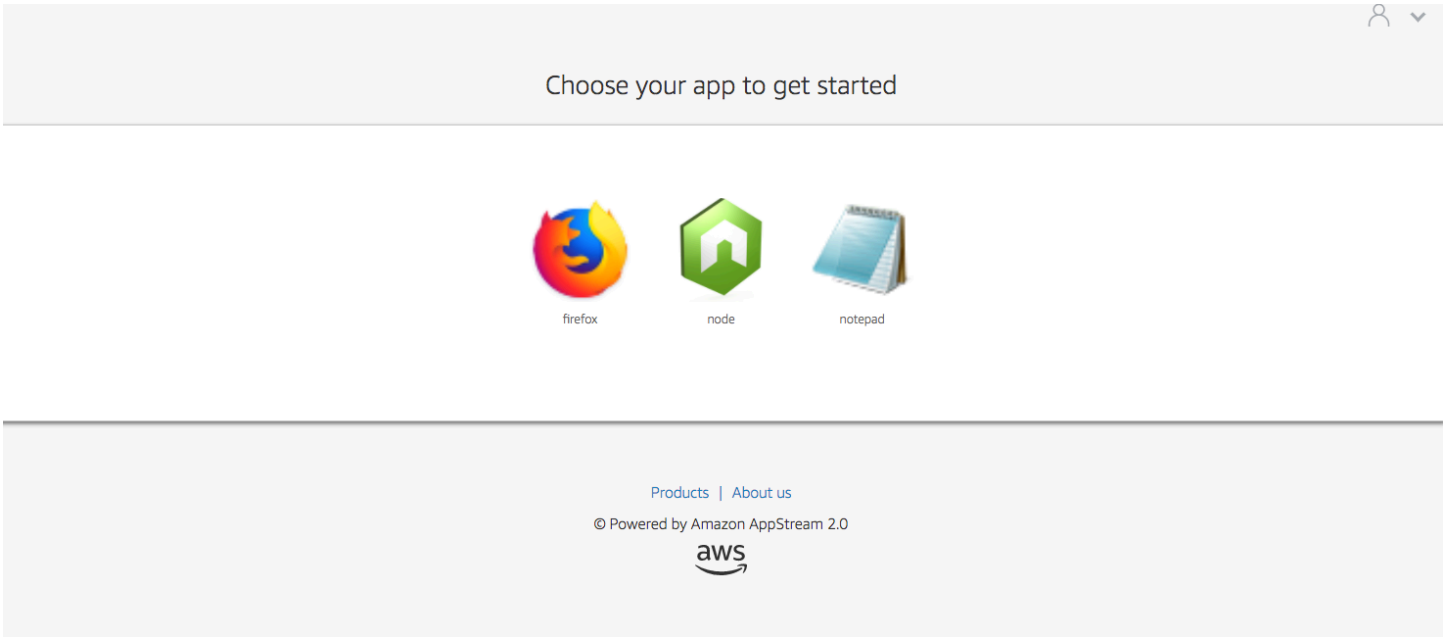
The following colors are applied when you select the light blue color theme:

Light blue (#1d83c2) – Used for buttons and website links.

White (#f6f6f6) – Used as a background accent.

Dark grey (#333333) – Used for the body text and in the progress spinner.

When you choose the light blue color theme, the website links, body text, and background accent appear in your streaming application catalog page as follows.



Blue

The following colors are applied when you select the blue color theme:

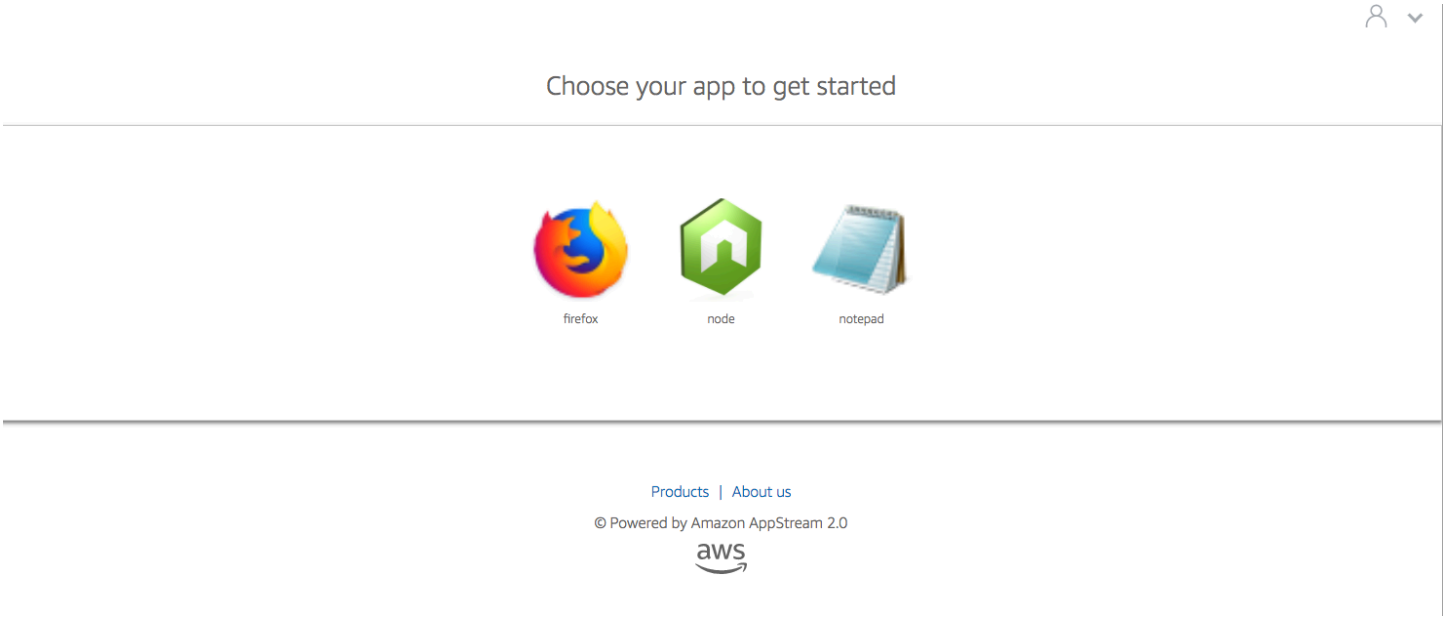
Blue (#0070ba) – Used for website links.

White (#ffffff) – Used as a background accent.

Light green (#8ac53e) – Used for buttons.

Grey (#666666) – Used for the body text and in the progress spinner.

When you choose the blue color theme, the website links, body text, and background accent appear in your streaming application catalog page as follows.



Pink

The following colors are applied when you select the pink color theme:

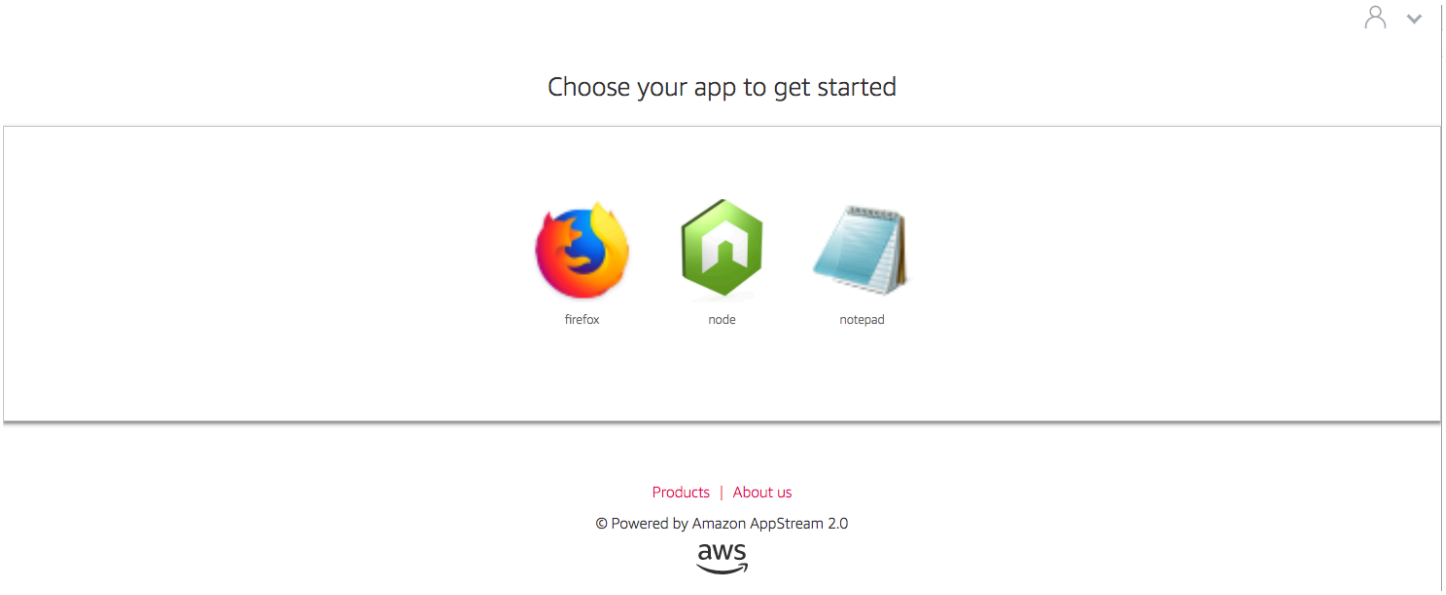
Pink (#ec0069) – Used for website links.

White (#ffffff) – Used as a background accent.

Blue (#3159a2) – Used for buttons.

Dark grey (#333333) – Used for the body text and in the progress spinner.

When you choose the pink color theme, the website links, body text, and background accent appear in your streaming application catalog page as follows.



Embed Amazon WorkSpaces Applications Streaming Sessions

You can create a dynamic, interactive, and customized experience for your users by embedding an WorkSpaces Applications streaming session within your website. Embedded WorkSpaces Applications streaming sessions let your users interact with 3D models, maps, and datasets directly from your website. For example, users can view training instructions or educational materials alongside their WorkSpaces Applications streaming session.

Contents

- Prerequisites for Embedding Amazon WorkSpaces Applications Streaming Sessions
- Recommendations and Usage Considerations for Embedding Amazon WorkSpaces Applications Streaming Sessions
- Step 1: Specify a Host Domain to Embedded Amazon WorkSpaces Applications Streaming Sessions
- Step 2: Create a Streaming URL for User Authentication
- Step 3: Download the Embedded Amazon WorkSpaces Applications Files
- Step 4. Configure Your Website for Integration with Amazon WorkSpaces Applications
- Constants, Functions, and Events for Embedded Amazon WorkSpaces Applications Streaming Sessions

Prerequisites for Embedding Amazon WorkSpaces Applications Streaming Sessions

To embed an WorkSpaces Applications streaming session in a website, you must have the following:

- A configured WorkSpaces Applications environment that includes an WorkSpaces Applications image, fleet, and stack. For information about how to create these resources, see the following topics in the *WorkSpaces Applications Administration Guide*:
  - Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console or Create Your Amazon WorkSpaces Applications Image Programmatically by Using the Image Assistant CLI Operations

  - Create a Fleet in Amazon WorkSpaces Applications
  - Create a Stack in Amazon WorkSpaces Applications
- A streaming URL for user authentication. SAML 2.0 and WorkSpaces Applications user pools are currently not supported as authentication methods for embedded WorkSpaces Applications streaming sessions.
- Optionally, you can use custom domains for embedded WorkSpaces Applications streaming sessions. You can use custom domains so that your own company URL displays for users rather than an WorkSpaces Applications URL. Custom domains are required if your users have web browsers that block third-party cookies.

  > **Note**
  >
  > You can configure custom domains by using Amazon CloudFront. For information, see Using Custom Domains with WorkSpaces Applications.

  When you use a custom domain, you must:

  - Create a streaming URL that uses the same domain.
  - Add **`appstream-custom-url-domain`** to the header of the webpage that will host the embedded WorkSpaces Applications streaming sessions. For the header value, use the domain that your reverse proxy displays to users. For more information, see Configuration Requirements for Using Custom Domains.

Recommendations and Usage Considerations for Embedding Amazon WorkSpaces Applications Streaming Sessions

Consider the following recommendations and usage notes for embedded WorkSpaces Applications streaming sessions.

- To maintain maximum control over the embedded WorkSpaces Applications streaming experience for your users, we recommend that you configure short-lived streaming URLs that last approximately 5 seconds. Any user can inspect the contents of a webpage and view its source. This includes the document object model (DOM) and the src (source) URL of the iframe. If the URL is still valid when a user copies it, that user can paste the URL in a separate browser tab and

stream the session with the standard WorkSpaces Applications portal user interface, without the embed options.

- Concurrent sessions are not supported when custom domains are used for embedded WorkSpaces Applications streaming sessions. Concurrent sessions occur when users start two embedded WorkSpaces Applications streaming sessions either on the same webpage or across two different browser tabs. You can't have a single user with concurrent sessions, but you can have multiple users. For example, a user logs into your app, your app generates a streaming URL to give to the customer (which counts as a unique user for billing), a customer loads the streaming URL, and the customer is assigned to an appstream instance within your specified pool.

Step 1: Specify a Host Domain to Embedded Amazon WorkSpaces Applications Streaming Sessions

To embed an WorkSpaces Applications streaming session in a webpage, first update your stack to specify the domain to host the embedded streaming session. This a security measure to ensure that only authorized website domains can embed WorkSpaces Applications streaming sessions. WorkSpaces Applications adds the domain or domains that you specify to the **`Content-Security-Policy`** (CSP) header. For more information, see Content Security Policy (CSP) in the Mozilla MDN Web Docs documentation.

To update your stack to specify the domain to host the embedded streaming session, use any of the following methods:

- The WorkSpaces Applications console
- The `EmbedHostDomains` API action
- The `embed-host-domains` AWS command line interface (AWS CLI) command

To specify a host domain by using the WorkSpaces Applications console, perform the following steps.

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the left navigation pane, choose **Stacks**, and select the stack that you want.
3. Choose **Edit**.

4. Expand **Embed WorkSpaces Applications (Optional)**.
5. In **Host Domains**, specify a valid domain. For example: **`training.example.com`**.

> **Note**
>
> Embedded streaming sessions are only supported over HTTPS [TCP port 443].

6. Choose **Update**.

Step 2: Create a Streaming URL for User Authentication

You must create a streaming URL to authenticate users for embedded WorkSpaces Applications streaming sessions. SAML 2.0 and user pools are currently not supported for embedded streaming sessions. To create a streaming URL, use one of the following methods:

- WorkSpaces Applications console
- The CreateStreamingURL API action
- The create-streaming-url AWS CLI command

Configuration Requirements for Using Custom Domains

Whether you use custom domains to apply your company branding or to ensure that embedded WorkSpaces Applications streaming sessions work with browsers that block third-party cookies, the configuration requirements are the same.

For web browsers that block third-party cookies, custom domains are required. WorkSpaces Applications uses browser cookies to authenticate streaming sessions and lets users reconnect to an active session without being prompted to provide their sign-in credentials every time. By default, WorkSpaces Applications streaming URLs include **`appstream.com`** as the domain. When you embed a streaming session within your website, **`appstream.com`** is treated as a third-party domain. As a result, streaming sessions may be blocked when modern browsers are used that block third-party cookies by default.

To avoid embedded WorkSpaces Applications streaming sessions from being blocked in this scenario, follow these steps:

1. Specify a custom domain to host your embedded WorkSpaces Applications streaming sessions.

When you configure your custom domain, make sure that the domain is a subdomain of the webpage in which you plan to embed WorkSpaces Applications. For example, if you update your stack to specify **`training.example.com`** as the host domain, you can create a subdomain called **`content.training.example.com`** for your embedded streaming sessions.

2. Create a streaming URL for embedded WorkSpaces Applications streaming sessions that uses the same custom subdomain. To create the streaming URL, use the CreateStreamingURL API action or the create-streaming-url AWS CLI command. You cannot use the WorkSpaces Applications console to create a streaming URL in this scenario.

   To create a streaming URL for embedded WorkSpaces Applications streaming sessions, in the URL, replace **`appstream2.`***`region`***`.aws.amazon.com`** with your own domain.

   By default, WorkSpaces Applications streaming URLs are formatted as follows:

   ```
   https://appstream2.region.aws.amazon.com/authenticate?parameters=authenticationcode
   ```

   If your subdomain is **`content.training.example.com`**, your new streaming URL follows this format:

   ```
   https://content.training.example.com/authenticate?parameters=authenticationcode
   ```

   > **Note**
   >
   > When you create a custom domain, you can use the domain for embedded WorkSpaces Applications streaming sessions only in the AWS Region for which it was configured. If you plan to support custom domains in multiple Regions, create a custom domain for each applicable Region. Also, embedded streaming sessions are only supported over HTTPS [TCP port 443].

3. Add **`appstream-custom-url-domain`** to the header of the webpage that will host the embedded streaming sessions. For the header value, use the domain that your reverse proxy displays to users. For example:

   ```
   Header name: appstream-custom-url-domain
   Header value: training.example.com
   ```

Setting a custom domain and creating a streaming URL that specifies the same domain lets the cookies be saved as first-party cookies. For information about how to configure custom domains by using Amazon CloudFront, see Using Custom Domains with WorkSpaces Applications.

After you set up a custom domain for your embedded WorkSpaces Applications streaming sessions, if your streaming URLs don't redirect to your custom domain, or if your custom domain doesn't display correctly for your users, see the following troubleshooting topics:

- I set up a custom domain for my embedded WorkSpaces Applications streaming sessions, but my WorkSpaces Applications streaming URLs aren't redirecting to my custom domain.

Step 3: Download the Embedded Amazon WorkSpaces Applications Files

To host embedded WorkSpaces Applications streaming sessions, you must download and configure the provided WorkSpaces Applications API JavaScript file.

1. On the Embedding WorkSpaces Applications in Your Website webpage, choose the link in step 1 to download the WorkSpaces Applications Embed Kit .zip file, **appstream_embed_<version>.zip**.
2. Navigate to the location where you downloaded the .zip file, and extract the contents of the file.
3. The extracted contents of the file comprise one folder, **appstream-embed**. In addition to the **COPYRIGHT.txt** and **THIRD_PARTY_NOTICES.txt** file, this folder contains the following two files:

   - **appstream-embed.js** — Provides the embedded WorkSpaces Applications API. This JavaScript file includes the functions and API actions for configuring and controlling your embedded WorkSpaces Applications streaming session.
   - **embed-sample.html** — Describes how to use the embedded WorkSpaces Applications API to initialize a streaming session, call functions, and listen for events. This sample file expands on the information in this topic, to provide an example use case for developers.

Step 4. Configure Your Website for Integration with Amazon WorkSpaces Applications

The following sections provide information about how to configure your webpage to host embedded WorkSpaces Applications streaming sessions.

Contents

- Import the appstream-embed JavaScript File
- Initialize and Configure the AppStream.Embed Interface Object
- Examples for Hiding Items in the WorkSpaces Applications User Interface

Import the appstream-embed JavaScript File

1. On the webpage where you plan to embed the WorkSpaces Applications streaming session, import the **appstream-embed.js** file into the webpage by adding the following code:

   ```
   <script type="text/javascript" src="./appstream_embed.js"> </script>
   ```

2. Next, create an empty container div. The ID of the div that you set is passed into the WorkSpaces Applications embed constructor. It's then used to inject an iframe for the streaming session. To create the div, add the following code:

   ```
   <div id="appstream-container"> </div>
   ```

Initialize and Configure the `AppStream.Embed` Interface Object

To initialize the `AppStream.Embed` interface object in JavaScript, you must add code that creates an `AppStream.Embed` object with options for the streaming URL and user interface configuration. These options, and the div ID that you created, are stored in an object called `appstreamOptions`.

The following example code shows how to initialize the `AppStream.Embed` interface object.

```
var appstreamOptions = {
     sessionURL: 'https://appstream2.region.aws.amazon.com/authenticate?
parameters=authenticationcode...',
     userInterfaceConfig:{[AppStream.Embed.Options.HIDDEN_ELEMENTS]:
[AppStream.Embed.Elements.TOOLBAR]}
```

```
};
appstreamEmbed = new AppStream.Embed("appstream-container", appstreamOptions);
```

In the code, replace *sessionURL* and *userInterfaceConfig* with your own values.

> **Note**
>
> The value specified for *userInterfaceConfig* hides the entire WorkSpaces Applications toolbar. This value, which is included as an example, is optional.

sessionUrl

The streaming URL that you created by using the WorkSpaces Applications console, the CreateStreamingURL API action, or the create-streaming-url AWS CLI command. This parameter is case-sensitive.

Type: String

Required: Yes

userInterfaceConfig

The configuration that generates the initial state of the user interface elements. The configuration is a key-value pair.

The key, `AppStream.Embed.Options.HIDDEN_ELEMENTS`, specifies the user interface objects that are initially hidden when the embedded WorkSpaces Applications streaming session is initialized. Later, you can return both hidden and visible objects by using the `getInterfaceState` parameter.

The value is an array of constants (toolbar buttons). For a list of constants that you can use, see Working with `HIDDEN_ELEMENTS`.

Type: Map (*key*:*value*)

Required: No

Examples for Hiding Items in the WorkSpaces Applications User Interface

The examples in this section show how to hide items in the WorkSpaces Applications user interface from users during their embedded WorkSpaces Applications streaming sessions.

Examples

- Example 1: Hide the entire WorkSpaces Applications toolbar
- Example 2: Hide a specific button on the WorkSpaces Applications toolbar
- Example 3: Hide multiple buttons on the WorkSpaces Applications toolbar

Example 1: Hide the entire WorkSpaces Applications toolbar

To prevent users from accessing any button on the WorkSpaces Applications toolbar during embedded streaming sessions, use the `AppStream.Embed.Elements.TOOLBAR` constant. This constant lets you hide all WorkSpaces Applications toolbar buttons.

```
var appstreamOptions = {
     sessionURL: 'https://appstream2.region.aws.amazon.com/authenticate?
parameters=authenticationcode...',
     userInterfaceConfig:{[AppStream.Embed.Options.HIDDEN_ELEMENTS]:
[AppStream.Embed.Elements.TOOLBAR]}
 };
```

Example 2: Hide a specific button on the WorkSpaces Applications toolbar

You can display the WorkSpaces Applications toolbar, while preventing users from accessing a specific toolbar button during embedded streaming sessions. To do so, specify the constant for the button that you want to hide. The following code uses the `AppStream.Embed.Elements.FILES_BUTTON` constant to hide the **My Files** button. This prevents users from accessing persistent storage options during embedded streaming sessions.

```
var appstreamOptions = {
     sessionURL: 'https://appstream2.region.aws.amazon.com/authenticate?
parameters=authenticationcode...',
     userInterfaceConfig:{[AppStream.Embed.Options.HIDDEN_ELEMENTS]:
[AppStream.Embed.Elements.FILES_BUTTON]}
```

```
};
```

Example 3: Hide multiple buttons on the WorkSpaces Applications toolbar

You can display the WorkSpaces Applications toolbar, while preventing users from accessing more than one toolbar button during embedded streaming sessions. To do so, specify the constants for the buttons that you want to hide. The following code uses the `AppStream.Embed.Elements.END_SESSION_BUTTON` and `AppStream.Embed.Elements.FULLSCREEN_BUTTON` constants to hide the **End Session** and **Fullscreen** buttons.

> **Note**
>
> Separate each constant with a comma, with no preceding or following space.

```
var appstreamOptions = {
     sessionURL: 'https://appstream2.region.aws.amazon.com/authenticate?
parameters=authenticationcode... (https://appstream2.region.aws.amazon.com/#/)',
     userInterfaceConfig:{[AppStream.Embed.Options.HIDDEN_ELEMENTS]:
[AppStream.Embed.Elements.END_SESSION_BUTTON,AppStream.Embed.Elements.FULLSCREEN_BUTTON]}
 };
```

Constants, Functions, and Events for Embedded Amazon WorkSpaces Applications Streaming Sessions

The following topics provide reference information for constants, functions, and events that you can use to configure embedded WorkSpaces Applications streaming sessions.

Contents

- Working with HIDDEN_ELEMENTS
- Functions for the AppStream.Embed Object
- Events for Embedded WorkSpaces Applications Streaming Sessions
- Examples for Adding Event Listeners and Ending an Embedded WorkSpaces Applications Streaming Session

The following WorkSpaces Applications user interface elements can be passed into the HIDDEN_ELEMENTS configuration option when an embedded WorkSpaces Applications streaming session is initialized.

Working with HIDDEN_ELEMENTS

The following WorkSpaces Applications user interface elements can be passed as constants into the HIDDEN_ELEMENTS configuration option when an embedded WorkSpaces Applications streaming session is initialized.

```
AppStream.Embed.Elements.TOOLBAR
AppStream.Embed.Elements.FULLSCREEN_BUTTON
AppStream.Embed.Elements.END_SESSION_BUTTON
AppStream.Embed.Elements.TOOLBAR
AppStream.Embed.Elements.CATALOG_BUTTON
AppStream.Embed.Elements.WINDOW_SWITCHER_BUTTON
AppStream.Embed.Elements.FILES_BUTTON
AppStream.Embed.Elements.CLIPBOARD_BUTTON
AppStream.Embed.Elements.COPY_LOCAL_BUTTON
AppStream.Embed.Elements.PASTE_REMOTE_BUTTON
AppStream.Embed.Elements.SETTINGS_BUTTON
AppStream.Embed.Elements.STREAMING_MODE_BUTTON
AppStream.Embed.Elements.SCREEN_RESOLUTION_BUTTON
AppStream.Embed.Elements.REGIONAL_SETTINGS_BUTTON
AppStream.Embed.Elements.FULLSCREEN_BUTTON
AppStream.Embed.Elements.END_SESSION_BUTTON
```

The following three elements can be passed as strings into HIDDEN_ELEMENTS, rather than as constants.

String	Description
'adminCommandsButton'	When you are connected to an WorkSpaces Applications image builder, the **Admin Commands** button displays on the top right corner of the WorkSpaces Applications toolbar. Passing this string into HIDDEN_ELEMENTS hides the **Admin Commands** button.
'softKeyboardButton'	During WorkSpaces Applications streaming sessions on touch-enabled devices, users can tap the keyboard icon

String	Description
	on the WorkSpaces Applications toolbar to display the on-screen keyboard. Passing this string into `HIDDEN_ELEMENTS` hides the keyboard icon.
`'keyboardShortcutsButton'`	During WorkSpaces Applications streaming sessions on touch-enabled devices, users can tap the Fn icon on the WorkSpaces Applications toolbar to display keyboard shortcuts. Passing this string into `HIDDEN_ELEMENTS` hides the Fn icon.

Functions for the `AppStream.Embed` Object

The following table lists the functions that can be performed on the `AppStream.Embed` object.

Function	Description
`AppStream.Embed(containerId:string, options:object)`	The `AppStream.Embed` object constructor. This constructor initializes and communicates with the `AppStream.Embed` object, and it uses a div container ID. The ID is used to inject the iframe. It also injects an object that includes the configuration options for `appstreamOptions` (*`sessionURL`* and `HIDDEN_ELEMENTS`).
`endSession()`	This function ends the streaming session, but does not destroy the iframe. If you specify a redirect URL, the iframe attempts to load the URL. Depending on the CORS headers of the page, the URL may not load.
`launchApp(appId:string)`	This function programmatically launches an application with the application ID that was specified during image creation.

Function	Description
`launchAppSwitcher()`	This function sends the **AppSwitcher** command to the WorkSpaces Applications portal. This triggers the **AppSwitcher** command on the instance.
`getSessionState()`	This function returns an object for `sessionStatus` . For more information, see Events for Embedded WorkSpaces Applications Streaming Sessions.
`getUserInterfaceState()`	This function returns an object for `UserInterfaceState` . The object contains the key-value pairs for the following: `sessionStatus` : State enumeration `sessionTerminationReason` : String `sessionDisconnectionReason` : String For more information, see Events for Embedded WorkSpaces Applications Streaming Sessions.
`addEventListener(name, callback)`	This function adds a callback function to call when the specified event is triggered. For a list of the events that can be triggered, see Events for Embedded WorkSpaces Applications Streaming Sessions.
`removeEventListener(name, callback)`	This function removes the callback for the specified events.
`destroy()`	This function deletes the iframe and cleans up resources. This function does not affect streaming sessions that are in progress.

Events for Embedded WorkSpaces Applications Streaming Sessions

The following table lists the events that can be triggered during embedded WorkSpaces Applications streaming sessions.

Event	Data	Description
`AppStream.Embed.Events.SESSION_STATE_CHANGE`	`sessionStatus` : `State` enumeration `sessionTerminationReason` : String `sessionDisconnectionReason` : String	This event is triggered when any session state change occurs. The event includes a map of the states that changed. To retrieve the full session state, use the `getSessionState()` function. Following are the session states: `AppStream.Embed.SessionStatus.Unknown` — The session has not started and is not reserved `AppStream.Embed.SessionStat`

Event	Data	Description
		us.Reserved — The session is reserved but has not started. AppStream.Embed.SessionStatus.Started — The user connected to the session and started streaming. AppStream.Embed.SessionStatus Disconnected — The user disconnected from the session. AppStream.Embed.SessionStatus.Ended — The session was marked as

Event	Data	Description
		ended or expired.
`AppStream.Embed.Events.SESSION_INTERFACE_STATE_CHANGE`	`hiddenElements` : Array of strings `isFullscreen` : Boolean `isSoftKeyboardVisible` : Boolean	This event is triggered when any session state change occurs. The event includes a map of the states that changed. To retrieve the full session state, use the `getSessionState()` function.
`AppStream.Embed.Events.SESSION_ERROR`	`errorCode` : Number `errorMessage` : String	This event is triggered when any errors occur during a session.

Examples for Adding Event Listeners and Ending an Embedded WorkSpaces Applications Streaming Session

The examples in this section show how to do the following:

- Add event listeners for embedded WorkSpaces Applications streaming sessions.
- Programmatically end an embedded WorkSpaces Applications streaming session.

Example 1: Add event listeners for embedded WorkSpaces Applications streaming sessions

To add event listeners for session state changes, session interface state changes, and session errors during embedded streaming sessions, use the following code:

```
appstreamEmbed.addEventListener(AppStream.Embed.Events.SESSION_STATE_CHANGE,
 updateSessionStateCallback);

appstreamEmbed.addEventListener(AppStream.Embed.Events.SESSION_INTERFACE_STATE_CHANGE,
 updateUserInterfaceStateCallback);

appstreamEmbed.addEventListener(AppStream.Embed.Events.SESSION_ERROR, errorCallback);
```

In this example, `AppStream.Embed.Events.SESSION_STATE_CHANGE`, `AppStream.Embed.Events.SESSION_INTERFACE_STATE_CHANGE`, and `AppStream.Embed.Events.SESSION_ERROR` are event names.

The `updateSessionStateCallback`, `updateUserInterfaceStateCallback`, and `errorCallback` functions are ones that you implement. These functions are passed into the `addEventListener` function and called when an event is triggered.

Example 2: Programmatically end an embedded WorkSpaces Applications streaming session

To end an embedded WorkSpaces Applications streaming sessions, use the following function:

```
appstreamEmbed.endSession();
```

Enable and Administer Persistent Storage for Your WorkSpaces Applications Users

Amazon WorkSpaces Applications supports the following persistent storage options for users in your organization:

- Home folders
- Google Drive for Google Workspace
- OneDrive for Business
- Custom shared folders (Server Message Block (SMB) network drives)

You can enable one or more options for your organization. As an WorkSpaces Applications administrator, you must understand how to perform the following tasks to enable and administer persistent storage for your users.

Contents

- Enable and Administer Home Folders for Your WorkSpaces Applications Users
- Enable and Administer Google Drive for Your WorkSpaces Applications Users
- Enable and Administer OneDrive for Business for Your WorkSpaces Applications Users
- Enable and Administer Custom Shared Folders (Server Message Block (SMB) Network Drives) for Your WorkSpaces Applications Users

For troubleshooting information, see Troubleshooting Persistent Storage Issues.

Enable and Administer Home Folders for Your WorkSpaces Applications Users

WorkSpaces Applications supports the following persistent storage options for users in your organization:

- Home folders
- Google Drive for Google Workspace
- OneDrive for Business

- Custom shared folders (Server Message Block (SMB) network drives)

You can enable one or more options for your organization. When you enable home folders for an WorkSpaces Applications stack, users of the stack can access a persistent storage folder during their application streaming sessions. No further configuration is required for your users to access their home folder. Data stored by users in their home folder is automatically backed up to an Amazon Simple Storage Service bucket in your Amazon Web Services account and is made available to those users in subsequent sessions.

Files and folders are encrypted in transit using Amazon S3's SSL endpoints. Files and folders are encrypted at rest using Amazon S3-managed encryption keys.

Home folders are stored on fleet instances in the following default locations:

- For single-session, non-domain-joined Windows instances: C:\Users\PhotonUser\My Files\Home Folder
- For multi-session, non-domain-joined Windows instances: C:\Users\as2-xxxxxxxx\My Files\Home Folder , where as2-xxxxxxxxx is a random user name assigned to each user session. You can determine your local user name through env variable $USERNAME.
- Domain-joined Windows instances: C:\Users\%username%\My Files\Home Folder
- Linux instances: ~/MyFiles/HomeFolder

As an administrator, use the applicable path if you configure your applications to save to the home folder. In some cases, your users may not be able to find their home folder because some applications do not recognize the redirect that displays the home folder as a top-level folder in File Explorer. If this is the case, your users can access their home folder by browsing to the same directory in File Explorer.

Contents

- Files and Directories Associated with Compute-Intensive Applications
- Enable Home Folders for Your WorkSpaces Applications Users
- Administer Your Home Folders

Files and Directories Associated with Compute-Intensive Applications

During WorkSpaces Applications streaming sessions, saving large files and directories associated with compute-intensive applications to persistent storage can take longer than saving files and directories required for basic productivity applications. For example, it might take longer for applications to save a large amount of data or frequently modify the same files than it would to save files created by applications that perform a single write action. It might also take longer to save many small files.

If your users save files and directories associated with compute-intensive applications and WorkSpaces Applications persistent storage options aren't performing as expected, we recommend that you use a Server Message Block (SMB) solution such as Amazon FSx for Windows File Server or an AWS Storage Gateway file gateway. Following are examples of files and directories associated with compute-intensive applications that are more suitable for use with these SMB solutions:

- Workspace folders for integrated development environments (IDEs)
- Local database files
- Scratch space folders created by graphics simulation applications

For more information, see:

- *Amazon FSx for Windows File Server Windows User Guide*
- Using Amazon FSx with Amazon WorkSpaces Applications
- File gateways in the *AWS Storage Gateway User Guide*

Enable Home Folders for Your WorkSpaces Applications Users

Before enabling home folders, you must do the following:

- Check that you have the correct AWS Identity and Access Management (IAM) permissions for Amazon S3 actions. For more information, see Using IAM Policies to Manage Administrator Access to the Amazon S3 Bucket for Home Folders and Application Settings Persistence.
- Use an image that was created from an AWS base image released on or after May 18, 2017. For a current list of released AWS images, see WorkSpaces Applications Base Image and Managed Image Update Release Notes.

- Enable network connectivity to Amazon S3 from your virtual private cloud (VPC) by configuring internet access or a VPC endpoint for Amazon S3. For more information, see Networking and Access for Amazon WorkSpaces Applications and Using Amazon S3 VPC Endpoints for WorkSpaces Applications Features.

You can enable or disable home folders while creating a stack (see Create a Stack in Amazon WorkSpaces Applications), or after the stack is created by using the AWS Management Console for WorkSpaces Applications, AWS SDK, or AWS CLI. For each AWS Region, home folders are backed up by an Amazon S3 bucket.

The first time you enable home folders for an WorkSpaces Applications stack in an AWS Region, the service creates an Amazon S3 bucket in your account in that same Region. The same bucket is used to store the content of home folders for all users and all stacks in that Region. For more information, see Amazon S3 Bucket Storage.

> **Note**
>
> For guidance that you can provide your users to help them get started with using home folders during WorkSpaces Applications streaming sessions, see Use Home Folders.

To enable home folders while creating a stack

- Follow the steps in Create a Stack in Amazon WorkSpaces Applications, and make sure that **Enable Home Folders** is selected.

To enable home folders for an existing stack

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the left navigation pane, choose **Stacks**, and select the stack for which to enable home folders.
3. Below the stacks list, choose **Storage** and select **Enable Home Folders**.
4. In the **Enable Home Folders** dialog box, choose **Enable**.

Administer Your Home Folders

Review the following topics to learn how to administer your home folders.

Contents

- Disable Home Folders
- Amazon S3 Bucket Storage
- Home Folder Content Synchronization
- Home Folder Formats
- Using the AWS Command Line Interface or AWS SDKs
- Additional Resources

Disable Home Folders

You can disable home folders for a stack without losing user content already stored in home folders. Disabling home folders for a stack has the following effects:

- Users who are connected to active streaming sessions for the stack receive an error message. They are informed that they can no longer store content in their home folder.
- Home folders do not appear for any new sessions that use the stack with home folders disabled.
- Disabling home folders for one stack does not disable it for other stacks.
- Even if home folders are disabled for all stacks, WorkSpaces Applications does not delete the user content.

To restore access to home folders for the stack, enable home folders again by following the steps described earlier in this topic.

To disable home folders while creating a stack

- Follow the steps in Create a Stack in Amazon WorkSpaces Applications and make sure that the **Enable Home Folders** option is cleared.

To disable home folders for an existing stack

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the left navigation pane, choose **Stacks**, and select the stack.
3. Below the stacks list, choose **Storage** and clear **Enable Home Folders**.
4. In the **Disable Home Folders** dialog box, type `CONFIRM` (case-sensitive) to confirm your choice, then choose **Disable**.

Amazon S3 Bucket Storage

WorkSpaces Applications manages user content stored in home folders by using Amazon S3 buckets created in your account. For every AWS Region, WorkSpaces Applications creates a bucket in your account. All user content generated from streaming sessions of stacks in that Region is stored in that bucket. The buckets are fully managed by the service without any input or configuration from an administrator.

The new enhanced buckets are named in a specific format (Version 2) as follows:

```
appstream2-36fb080bb8-region-code-account-id-without-hyphens-random-identifier
```

Where *region-code* is the AWS Region code in which the stack is created and *account-id-without-hyphens* is your Amazon Web Services account ID. The first part of the bucket name, appstream2-36fb080bb8-, does not change across accounts or Regions.

For example, if you enable home folders for stacks in the US West (Oregon) Region (us-west-2) on account number 123456789012, the service creates an Amazon S3 bucket in that Region with the name shown. Only an administrator with sufficient permissions can delete this bucket.

```
appstream2-36fb080bb8-us-west-2-123456789012-abcdefg
```

The old version buckets are named as follows. Accounts created before WorkSpaces Applications introduced the new enhanced bucket naming (Version 2) will follow the old naming format.

```
appstream2-36fb080bb8-region-code-account-id-without-hyphens
```

Where *region-code* is the AWS Region code in which the stack is created and *account-id-without-hyphens* is your Amazon Web Services account ID. The first part of the bucket name, appstream2-36fb080bb8-, does not change across accounts or Regions.

For example, if you enable home folders for stacks in the US West (Oregon) Region (us-west-2) on account number 123456789012, the service creates an Amazon S3 bucket in that Region with the name shown. Only an administrator with sufficient permissions can delete this bucket.

```
appstream2-36fb080bb8-us-west-2-123456789012
```

As mentioned earlier, disabling home folders for stacks does not delete any user content stored in the Amazon S3 bucket. To permanently delete user content, an administrator with adequate

access must do so from the Amazon S3 console. WorkSpaces Applications adds a bucket policy that prevents accidental deletion of the bucket. For more information, see Using IAM Policies to Manage Administrator Access to the Amazon S3 Bucket for Home Folders and Application Settings Persistence.

Home Folder Content Synchronization

When home folders are enabled, WorkSpaces Applications creates a unique folder for each user in which to store their content. The folder is created as a unique Amazon S3 prefix that uses a hash of the user name within an S3 bucket for your Amazon Web Services account and Region. After WorkSpaces Applications creates the home folder in Amazon S3, it copies the accessed content in that folder from the S3 bucket to the fleet instance. This enables the user to access their home folder content quickly, from the fleet instance, during their streaming session. Changes that you make to a user's home folder content in an S3 bucket and that the user makes to their home folder content on a fleet instance are synchronized between Amazon S3 and WorkSpaces Applications as follows.

1. At the beginning of a user's WorkSpaces Applications streaming session, WorkSpaces Applications catalogs the home folder files that are stored for that user in the Amazon S3 bucket for your Amazon Web Services account and Region.
2. A user's home folder content is also stored on the WorkSpaces Applications fleet instance from which they stream. When a user accesses their home folder on the WorkSpaces Applications fleet instance, the list of cataloged files is displayed.
3. WorkSpaces Applications downloads a file from the S3 bucket to the fleet instance only after the user uses a streaming application to open the file during their streaming session.
4. After WorkSpaces Applications downloads the file to the fleet instance, synchronization occurs after the file is accessed
5. If the user changes the file during their streaming session, WorkSpaces Applications uploads the new version of the file from the fleet instance to the S3 bucket periodically or at the end of the streaming session. However, the file is not downloaded from the S3 bucket again during the streaming session.

The following sections describe synchronization behavior when you add, replace, or remove a user's home folder file in Amazon S3.

Contents

- Synchronization of files that you add to a user's home folder in Amazon S3

- Synchronization of files that you replace in a user's home folder in Amazon S3
- Synchronization of files that you remove from a user's home folder in Amazon S3

Synchronization of files that you add to a user's home folder in Amazon S3

If you add a new file to a user's home folder in an S3 bucket, WorkSpaces Applications catalogs the file and displays it in the list of files in the user's home folder within a few minutes. However, the file isn't downloaded from the S3 bucket to the fleet instance until the user opens the file with an application during their streaming session.

Synchronization of files that you replace in a user's home folder in Amazon S3

If a user opens a file in their home folder on the fleet instance during their streaming session, and you replace the same file in their home folder in an S3 bucket with a new version during that user's active streaming session, the new version of the file is not immediately downloaded to the fleet instance. The new version is downloaded from the S3 bucket to the fleet instance only after the user starts a new streaming session and opens the file again.

Synchronization of files that you remove from a user's home folder in Amazon S3

If a user opens a file in their home folder on the fleet instance during their streaming session, and you remove the file from their home folder in an S3 bucket during that user's active streaming session, there may be up to a 60 second delay before the deletion is reflected on the fleet instance. After this delay, the file is removed from the fleet instance when the user does either of the following:

- Opens the home folder again
- Refreshes the home folder

Home Folder Formats

The hierarchy of a user folder depends on how a user launches a streaming session, as described in the following sections.

AWS SDKs and AWS CLI

For sessions launched using `CreateStreamingURL` or `create-streaming-url` the user folder structure is as follows:

```
bucket-name/user/custom/user-id-SHA-256-hash/
```

Where *bucket-name* is in the format shown in Amazon S3 Bucket Storage and *user-id-SHA-256-hash* is the user-specific folder name created using a lowercase SHA-256 hash hexadecimal string generated from the `UserId` value passed to the CreateStreamingURL API operation or `create-streaming-url` command. For more information, see CreateStreamingURL in the *Amazon WorkSpaces Applications API Reference* and create-streaming-url in the *AWS CLI Command Reference*.

The following example folder structure applies to session access using the API or AWS CLI with a `UserId` testuser@mydomain.com, account id 123456789012 in the US West (Oregon) Region (us-west-2):

```
appstream2-36fb080bb8-us-west-2-123456789012/user/custom/
a0bcb1da11f480d9b5b3e90f91243143eac04cfccfbdc777e740fab628a1cd13/
```

You can identify the folder for a user by generating the lowercase SHA-256 hash value of the `UserId` using websites or open source coding libraries available online.

SAML 2.0

For sessions created using SAML federation, the user folder structure is as follows:

```
bucket-name/user/federated/user-id-SHA-256-hash/
```

In this case, *user-id-SHA-256-hash* is the folder name created using a lowercase SHA-256 hash hexadecimal string generated from the `NameID` SAML attribute value passed in the SAML federation request. To differentiate users who have the same name but belong to two different domains, send the SAML request with `NameID` in the format `domainname\username`. For more information, see Amazon WorkSpaces Applications Integration with SAML 2.0.

The following example folder structure applies to session access using SAML federation with `NameID` SAMPLEDOMAIN\testuser, account ID 123456789012 in the US West (Oregon) Region:

```
appstream2-36fb080bb8-us-west-2-123456789012/user/
federated/8dd9a642f511609454d344d53cb861a71190e44fed2B8aF9fde0C507012a9901
```

When part or all of the NameID string is capitalized (as the domain name *SAMPLEDOMAIN* is in the example), WorkSpaces Applications generates the hash value based on the capitalization

used in the string. Using this example, the hash value for SAMPLEDOMAIN\testuser is 8DD9A642F511609454D344D53CB861A71190E44FED2B8AF9FDE0C507012A9901.
In the folder for that user, this value is displayed in lowercase, as follows: 8dd9a642f511609454d344d53cb861a71190e44fed2B8aF9fde0C507012a9901.

You can identify the folder for a user by generating the SHA-256 hash value of the `NameID` using websites or open source coding libraries available online.

Using the AWS Command Line Interface or AWS SDKs

You can enable and disable home folders for a stack by using the AWS CLI or AWS SDKs.

Use the following create-stack command to enable home folders while creating a new stack:

```
aws appstream create-stack --name ExampleStack --storage-connectors
 ConnectorType=HOMEFOLDERS
```

Use the following update-stack command to enable home folders for an existing stack:

```
aws appstream update-stack --name ExistingStack --storage-connectors
 ConnectorType=HOMEFOLDERS
```

Use the following command to disable home folders for an existing stack. This command does not delete any user data.

```
aws appstream update-stack --name ExistingStack --delete-storage-connectors
```

Additional Resources

For more information about managing Amazon S3 buckets and best practices, see the following topics in the *Amazon Simple Storage Service User Guide*:

- You can provide offline access to user data for your users with Amazon S3 policies. For more information, see Amazon S3: Allows IAM Users Access to Their S3 Home Directory, Programmatically and In the Console in the *IAM User Guide*.
- You can enable file versioning for content stored in Amazon S3 buckets used by WorkSpaces Applications. For more information, see Using Versioning.

Enable and Administer Google Drive for Your WorkSpaces Applications Users

Note

Amazon WorkSpaces Applications's use and transfer to any other app of information received from Google APIs will adhere to Google API Services User Data Policy, including the Limited Use requirements.

Amazon WorkSpaces Applications supports the following persistent storage options for users in your organization:

- Google Drive for Google Workspace
- OneDrive for Business
- Home folders

You can enable one or more options for your organization. When you enable Google Drive for Google Workspace for an WorkSpaces Applications stack, users of the stack can link their Google Drive for Google Workspace account to WorkSpaces Applications. Then they can sign into their Google Drive for Google Workspace account and access their Google Drive folder during application streaming sessions. Any changes that they make to files or folders in Google Drive during those sessions are automatically backed up and synchronized, so that they are available to users outside of their streaming sessions.

Important

You can enable Google Drive for Google Workspace for accounts in your Google Workspace domains only, but not for personal Gmail accounts.

Note

You can enable Google Drive for Windows stacks, but not for Linux stacks.
To enable Google Drive for stacks associated with multi-session fleets, the image must use WorkSpaces Applications Agent Release Notes released on or after June 29, 2026 or

your image is using Update an Image by Using Managed WorkSpaces Applications Image Updates released on or after June 29, 2026.

Contents

- Enable Google Drive for Your WorkSpaces Applications Users
- Disable Google Drive for Your WorkSpaces Applications Users
- Controlling Google Drive Re-authentication

Enable Google Drive for Your WorkSpaces Applications Users

Before enabling Google Drive, you must do the following:

- Have an active Google Workspace account with a valid organizational domain and users in the domain to use with WorkSpaces Applications.
- Configure an WorkSpaces Applications stack with an associated fleet.

  The fleet must use an image that uses a version of the WorkSpaces Applications agent released on or after May 31, 2018. For more information, see WorkSpaces Applications Agent Release Notes. The fleet must also have access to the internet.
- Add Amazon WorkSpaces Applications as a trusted app in one or more domains associated with your Google Workspace account. You can enable Google Drive for up to 10 domains.
- Have a Windows-based stack. (Linux-based stacks are not supported).

Follow these steps to add Amazon WorkSpaces Applications as a trusted app in your Google Workspace domains.

To add Amazon WorkSpaces Applications as a trusted app in your Google Workspace domains

1. Sign in to the Google Workspace Admin console at https://admin.google.com/.
2. In the left navigation sidebar, choose **Security**, **Access and data control**, **API controls**.
3. At the top of the page, in the **App access control** section, choose **MANAGE THIRD-PARTY APP ACCESS**.
4. Choose **Add app**, and then choose **OAuth App Name Or Client ID**.

5. Enter the Amazon WorkSpaces Applications OAuth client ID for your AWS Region, and then choose **SEARCH**. For a list of client IDs, see the table that follows this procedure.
6. In the search results, choose Amazon WorkSpaces Applications, and then choose **Select**.
7. In the **Client ID** page, under **OAuth Client ID**, verify that the correct ID appears in the list, and then select the check box to the left of the ID.
8. On the lower right of the page, choose **SELECT**.
9. Configure which organizational units in your Google Workspace organization should gain access.
10. Under **Access to Google Data**, choose **Trusted: Can access all Google services**, and then choose **CONTINUE**.
11. Review that the selections made are correct, then when you are satisfied, choose **FINISH**.
12. Verify that the Amazon WorkSpaces Applications app, with the correct OAuth ID, appears in the list of connected apps.

Amazon WorkSpaces Applications OAuth2 client IDs

Region	Amazon WorkSpaces Applications OAuth client ID
Asia Pacific (Malaysia)	`526025990430-7u0f5r0rg4caj0impcs3atv atjtmdeld.apps.googleusercontent.com`
Asia Pacific (Mumbai)	`325827353178-coqs1c374mf388ctllrlls3 74dc1bmb2.apps.googleusercontent.com`
Asia Pacific (Osaka)	`725129131438-7tsavvqnqitn9u7nggji294 3npo9rndm.apps.googleusercontent.com`
Asia Pacific (Seoul)	`562383781419-am1i2dnvt050tmdltsvr36i 8l2js40dj.apps.googleusercontent.com`
Asia Pacific (Singapore)	`856871139998-4eia2n1db5j6gtv4c1rdte1 fh1gec8vs.apps.googleusercontent.com`
Asia Pacific (Sydney)	`151535156524-b889372osskprm4dt1clpm5 3mo3m9omp.apps.googleusercontent.com`

Region	Amazon WorkSpaces Applications OAuth client ID
Asia Pacific (Tokyo)	`922579247628-qpl9kpihg3hu5dul2lphbjs4qbg6mjm2.apps.googleusercontent.com`
AWS GovCloud (US-East)	`20306576244-gqqkappmhhv9fj06sdk7as60he89e7ce.apps.googleusercontent.com` **Note** For more information about using WorkSpaces Applications in the AWS GovCloud (US) Regions, see Amazon WorkSpaces Applications in the *AWS GovCloud (US) User Guide*.
AWS GovCloud (US-West)	`996065833880-litfkb2vfd7c65nt7s24r7t81e5bc9bl.apps.googleusercontent.com` **Note** For more information about using WorkSpaces Applications in the AWS GovCloud (US) Regions, see Amazon WorkSpaces Applications in the *AWS GovCloud (US) User Guide*.
Canada (Central)	`872792838542-t39aqh72jv895c89thtk6v83sl6jugm2.apps.googleusercontent.com`
Canada West (Calgary)	`971152077170-30hu17p9bc2bklct219i6ovob3m31olk.apps.googleusercontent.com`
Europe (Frankfurt)	`643727794574-1se5360a77i84je9j3ap12obov1ib76q.apps.googleusercontent.com`
Europe (Ireland)	`599492309098-098muc7ofjfo9vua5rm5u9q2k3mlok3j.apps.googleusercontent.com`

Region	Amazon WorkSpaces Applications OAuth client ID
Europe (London)	682555519925-usbn2sk1ffgo8odgf23nj66ri71na0k5.apps.googleusercontent.com
Europe (Milan)	357829681601-84bcnr1ve68rthka1st5dboj0jgsrki7.apps.googleusercontent.com
Europe (Paris)	1018958878172-sgg1u1hlqiq8v53gdpq9aiu6ta48q1de.apps.googleusercontent.com
Europe (Spain)	258457153543-0qbtkg4a99stlj12pi8sdqaetfb96b1s.apps.googleusercontent.com
Europe (Zurich)	551239834406-k9ucmqo8l63t53ipmll44smu6g5ngerk.apps.googleusercontent.com
Israel (Tel Aviv)	1089892007469-bkqfmc1sm3ahqrssjjp4ees1mmiuium0.apps.googleusercontent.com
South America (São Paulo)	891888628791-1ltbtedva29esqvqadiatlj4htcgcjfo.apps.googleusercontent.com
US East (N. Virginia)	266080779488-15n5q5nkiclp6m524qibnmhmbsg0hk92.apps.googleusercontent.com
US East (Ohio)	723951369598-6tvdlf52g2qh0qa141o4k1avasvnj51i.apps.googleusercontent.com
US West (Oregon)	1026466167591-i4jmemrggsjomp9tnkkcs5tniggfiujb.apps.googleusercontent.com

Follow these steps to enable Google Drive for your WorkSpaces Applications users.

To enable Google Drive while creating a stack

- Follow the steps in Create a Stack in Amazon WorkSpaces Applications, make sure that **Enable Google Drive** is selected, and that you have specified at least one organizational domain associated with your Google Workspace account.

To enable Google Drive for an existing stack

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the left navigation pane, choose **Stacks**, and select the stack for which to enable Google Drive.
3. Below the stacks list, choose **Storage** and select **Enable Google Drive for Google Workspace**.
4. In the **Enable Google Drive for Google Workspace** dialog box, in **Google Workspace domain name**, type the name of at least one organizational domain that is associated with your Google Workspace account. To specify another domain, choose **Add another domain**, and type the name of the domain.
5. After you add domain names, choose **Enable**.

> **Note**
>
> For guidance that you can provide your users to help them get started with using Google Drive during WorkSpaces Applications streaming sessions, see Use Google Drive.

Disable Google Drive for Your WorkSpaces Applications Users

You can disable Google Drive for a stack without losing user content that is already stored on Google Drive. Disabling Google Drive for a stack has the following effects:

- Users who are connected to active streaming sessions for the stack receive an error message. They are informed that they do not have permissions to access their Google Drive.
- Any new sessions that use the stack with Google Drive disabled do not display Google Drive.
- Only the specific stack for which Google Drive is disabled is affected.
- Even if Google Drive is disabled for all stacks, WorkSpaces Applications does not delete the user content stored in their Google Drive.

Follow these steps to disable Google Drive for an existing stack.

To disable Google Drive for an existing stack

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.

2. In the left navigation pane, choose **Stacks**, and select the stack for which to disable Google Drive.
3. Below the stacks list, choose **Storage**, and clear the **Enable Google Drive for Google Workspace** option.
4. In the **Disable Google Drive for Google Workspace** dialog box, type `CONFIRM` (case-sensitive) to confirm your choice, then choose **Disable**.

   When users of the stack start their next WorkSpaces Applications streaming session, they can no longer access their Google Drive folder from within that session and future sessions.

Controlling Google Drive Re-authentication

When you enable Google Drive as a storage connector on a stack, users authenticate with their Google account during their first streaming session. By default, this authentication persists across subsequent sessions, so users do not need to sign in again.

You can change this behavior by configuring the storage connector authentication setting at the stack level. When configured to force re-authentication, stored Google identity provider tokens are cleared when a session ends, requiring users to re-authenticate with Google Drive at the start of each new session.

When to use forced re-authentication:

- Shared workstation environments where multiple users may access the same streaming URL
- Security and compliance requirements that mandate per-session authentication for Google Drive access
- Environments where users should not retain access to previously linked Google Drive storage across sessions

You can configure this setting when creating or updating a stack using the WorkSpaces Applications API. For details on the available values and API usage, see CreateStack and UpdateStack in the *WorkSpaces Applications API Reference*.

Important

- This setting applies at the stack level and affects all cloud storage connectors (OneDrive for Business and Google Drive) configured on the stack. You cannot set different re-authentication behaviors for individual connectors.
- This setting has no effect on Home Folders.
- Even with forced re-authentication disabled (authentication set to persist across sessions), Google may expire or revoke credentials based on your Google Workspace admin policies, OAuth token lifetime settings, or security events. Persistent authentication does not guarantee a permanent sign-in — users may still be prompted to re-authenticate based on Google's credential expiration policies.

Enable and Administer OneDrive for Business for Your WorkSpaces Applications Users

WorkSpaces Applications supports the following persistent storage options for users in your organization.

- OneDrive for Business
- Google Drive for Google Workspace
- Home folders

You can enable one or more options for your organization. When you enable OneDrive for Business for an WorkSpaces Applications stack, users of the stack can link their OneDrive for Business account to WorkSpaces Applications. Then they can sign into their OneDrive for Business account and access their OneDrive folder during application streaming sessions. Any changes that they make to files or folders in OneDrive during those sessions are automatically backed up and synchronized, so that they are available to users outside of their streaming sessions.

Important

You can enable OneDrive for Business for accounts in your OneDrive domains only, but not for personal accounts. Unless an administrator grants consent for the whole tenant and you select **Reuse existing Microsoft consent** for the domain, you must configure your

Microsoft Azure Active Directory environment to allow end-user consent to applications. For more information, see Configure how end-users consent to applications in the Azure Active Directory Application management documentation.
Your Azure Active Directory environment, not WorkSpaces Applications, decides whether users can consent on their own or need administrator approval. If your environment requires administrator consent, an administrator must grant consent to WorkSpaces Applications for the whole tenant. Then select **Reuse existing Microsoft consent** for each domain, as described in Enable OneDrive for Your WorkSpaces Applications Users. If that option is cleared, users are prompted for consent every time they link their account, and users who can't consent on their own are blocked even after an administrator grants consent.

Note

You can enable OneDrive for Business for Windows stacks, but not for Linux stacks.
To enable OneDrive for Business for stacks associated with multi-session fleets, the image must use WorkSpaces Applications Agent Release Notes released on or after June 29, 2026 or your image is using Update an Image by Using Managed WorkSpaces Applications Image Updates released on or after June 29, 2026.

Contents

- Enable OneDrive for Your WorkSpaces Applications Users
- Disable OneDrive for Your WorkSpaces Applications Users
- Controlling OneDrive Re-authentication

Enable OneDrive for Your WorkSpaces Applications Users

Before enabling OneDrive, you must do the following:

- Have an active Microsoft Office 365 or OneDrive for Business account with a valid organizational domain and users in the domain to use with WorkSpaces Applications.
- Configure an WorkSpaces Applications stack with an associated fleet.

The fleet must use an image that uses a version of the WorkSpaces Applications agent released on or after July 26, 2018. For more information, see WorkSpaces Applications Agent Release Notes. The fleet must also have access to the internet.

- Have a Windows-based stack. (Linux-based stacks are not supported).

Follow these steps to enable OneDrive for your WorkSpaces Applications users.

To enable OneDrive while creating a stack

- Follow the steps in Create a Stack in Amazon WorkSpaces Applications, make sure that **Enable OneDrive** is selected, and that you have specified at least one organizational domain that is associated with your OneDrive for Business account.

To enable OneDrive for an existing stack

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the left navigation pane, choose **Stacks**, and select the stack for which to enable OneDrive.
3. Below the stacks list, choose **Storage**, and select **Enable OneDrive for Business**.
4. In the **Enable OneDrive for Business** dialog box, in **OneDrive domain name**, type the name of at least one organizational domain that is associated with your OneDrive account. To specify another domain, choose **Add another domain**, and type the name of the domain.
5. For each domain, choose whether users reuse existing Microsoft consent. **Reuse existing Microsoft consent** is selected by default. When it is selected, users are prompted for consent only if they or an administrator have not already granted it. When it is cleared, users are prompted for consent every time they link their OneDrive for Business account. Your Azure Active Directory environment, not WorkSpaces Applications, decides whether a user needs administrator approval.
6. After you add OneDrive domain names, choose **Enable**.

Before your users can use OneDrive with WorkSpaces Applications, you must provide them with permissions to link their OneDrive account with third-party web applications. To do so, follow the steps in the next section.

Important

Unless an administrator grants consent for the whole tenant and you select **Reuse existing Microsoft consent** for the domain, you must configure your Microsoft Azure Active Directory environment to allow end-user consent to applications. For more information, see Configure how end-users consent to applications in the Azure Active Directory Application management documentation.

Provide Your Users with Permissions to Link OneDrive with WorkSpaces Applications

You must enable Integrated Apps in your Office 365 or OneDrive for Business admin console before users can link their OneDrive for Business account to WorkSpaces Applications.

1. Sign in to Office 365 or the OneDrive for Business admin console.
2. In the left navigation pane of the console, choose **Settings**, **Services & add-ins**.
3. From the list of services and add-ins, choose **Integrated Apps**.
4. On the **Integrated apps** page, turn on the option to allow users in your organization to let third party web apps access their Office 365 information.

Note

For guidance that you can provide your users to help them get started with using OneDrive during WorkSpaces Applications streaming sessions, see Use OneDrive for Business.

Disable OneDrive for Your WorkSpaces Applications Users

You can disable OneDrive for a stack without losing user content that is already stored on OneDrive. Disabling OneDrive for a stack has the following effects:

- Users who are connected to active streaming sessions for the stack receive an error message. They are informed that they do not have permissions to access their OneDrive.
- Any new sessions that use the stack with OneDrive disabled do not display OneDrive.
- Only the specific stack for which OneDrive is disabled is affected.

- Even if OneDrive is disabled for all stacks, WorkSpaces Applications does not delete the user content stored in their OneDrive.

Follow these steps to disable OneDrive for an existing stack.

To disable OneDrive for an existing stack

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the left navigation pane, choose **Stacks**, and select the stack for which to disable OneDrive.
3. Below the stacks list, choose **Storage**, and clear **Enable OneDrive for Business** option.
4. In the **Disable OneDrive for Business** dialog box, type `CONFIRM` (case-sensitive) to confirm your choice, then choose **Disable**.

   When users of the stack start their next WorkSpaces Applications streaming session, they can no longer access their OneDrive folder from within that session and future sessions.

Controlling OneDrive Re-authentication

When you enable OneDrive for Business as a storage connector on a stack, users authenticate with their Microsoft account during their first streaming session. By default, this authentication persists across subsequent sessions, so users do not need to sign in again.

You can change this behavior by configuring the storage connector authentication setting at the stack level. When configured to force re-authentication, stored Microsoft identity provider tokens are cleared when a session ends, requiring users to re-authenticate with OneDrive for Business at the start of each new session.

When to use forced re-authentication:

- Shared workstation environments where multiple users may access the same streaming URL
- Security and compliance requirements that mandate per-session authentication for OneDrive access
- Environments where users should not retain access to previously linked OneDrive storage across sessions

You can configure this setting when creating or updating a stack using the WorkSpaces Applications API. For details on the available values and API usage, see CreateStack and UpdateStack in the *WorkSpaces Applications API Reference*.

Important

- This setting applies at the stack level and affects all cloud storage connectors (OneDrive for Business and Google Drive) configured on the stack. You cannot set different re-authentication behaviors for individual connectors.
- This setting has no effect on Home Folders.
- Even with forced re-authentication disabled (authentication set to persist across sessions), Microsoft may expire or revoke credentials based on your Entra ID tenant policies, token lifetime configurations, or security events. Persistent authentication does not guarantee a permanent sign-in — users may still be prompted to re-authenticate based on Microsoft's credential expiration policies.

Enable and Administer Custom Shared Folders (Server Message Block (SMB) Network Drives) for Your WorkSpaces Applications Users

You can enable one or more options for your organization. When you enable and map the Server Message Block (SMB) network drives, multiple users can access the same data from Windows WorkSpaces Applications sessions. Any changes that users make to SMB network drives during those sessions are automatically backed up and synchronized.

Note

- Server Message Block (SMB) network drives mapping are supported only on domain-joined fleets
- To use this feature, you must use an WorkSpaces Applications image that uses the WorkSpaces Applications agent released after September 18, 2024. For more information, see Manage WorkSpaces Applications Agent Versions and WorkSpaces Applications Base Image and Managed Image Update Release Notes.

Before you map Server Message Block (SMB) network drives, ensure that for inbound rules, the security group that your users use to connect to fleets exposes TCP port 445 (SMB protocol) to the domain controller and the security group.

Contents

- Map Server Message Block (SMB) Network Drives

Map Server Message Block (SMB) Network Drives

You can use any machine that is under the targeted network of the SMBs. If you prefer to configure the setup through session scripts, you need to first create a script that gets invoked when user is logged on, as session script is configured per image.

To map Server Message Block (SMB) network drives, do the following steps.

Step 1: Ensure services are running

From the Start Menu, open **services.msc** and make sure the following services are all running:

- DNS Client
- Function Discovery Resource Publication
- SSDP Discovery
- UPnP Device Host

Step 2: Create a SMB folder

You can create an SMB with File Explorer.

To use File Explorer to configure your SMB shared folders

1. Right-click the SMB folder and choose **Properties**, **Sharing**.
2. Choose **Advanced Sharing**.
3. For **Advanced Sharing**, check **Share this folder**, and then choose **Permissions**.
4. If you want to provide permissions for all your users, leave it as the default setting.

   If you want to add specific users, under **Share Permissions**, choose **Everyone**, **Remove**. Then choose **Add** and enter the users or groups you want to access the file share.

For each user or group you add, choose **Allow** to assign **Full Control**, **Change**, or **Read permissions**.

5. Choose **Apply**, **OK**, **OK**, **Close**.

Step 3: Verify that the SMB is accessible in the domain

Open the file explorer from another server that uses the same security group and joins to the same domain. Access the network share through the provided network path by navigating to the network path folder. Choose **Properties**, **Sharing**, **Network Path**.

Step 4: Enable users to create symbolic links from local/domain Group Policy

Enable creating symbolic links from local/domain Group Policy for your users to ensure the session script or logon script defined in group policy. This allows you to create a script in Step 5 with user permissions.

To enable users to create symbolic links from local/domain Group Policy

1. In the GPO, which will be used to define this policy, choose **Computer Configuration**, **Windows Settings**, **Security Settings**, **User Rights Assignment**, **Policy**, **Create symbolic links**. Then, update the permission for users to include. For more information about creating symbolic links, see Create symbolic links.
2. By default, remote-to-remote (for example, a symlink mapping to a network share within another similar symlink) and remote-to-local (for example, a symlink mapping to a local share within a symlink mapping to a network share) accesses are disabled. If symlink mapping is needed, run the commands below:

   - For enabling remote-to-remote access - `fsutil behavior set SymlinkEvaluation R2R:1`
   - For enabling remote-to-local access - `fsutil behavior set SymlinkEvaluation R2L:1`

Step 5: Create a script that gets invoked when user is logged on

Create a script that gets invoked when user is logged on by either using an WorkSpaces Applications session script or GPO logon script. If you choose to use the WorkSpaces Applications session script, the session script will only get applied to that specific WorkSpaces Applications

image. If you use the GPO logon script, the GPOs will be applied to the domain / OU, which can be configured to your fleets. That way you don't need configure scripts for every image that you own.

Option 1: Use a session script to mount the SMB shared folder under My Files (using Powershell)

To use a session script to mount the SMB shared folder under My Files (using Powershell)

1. After you've successfully defined user permissions, configure the following example script using user context or system context.

   The following is the example config.json script that uses user context.

```
"SessionStart": {
    "executables": [
    {
        "context": "system",
        "filename": "",
        "arguments": "",
        "s3LogEnabled": true
    },
    {
        "context": "user",
        "filename": "C:\\Windows\\System32\\WindowsPowerShell\\v1.0\
\powershell.exe",
        "arguments": "-File \"C:\\AppStream\\SessionScripts\\userStart.ps1\"",
        "s3LogEnabled": true
    }
 ],
"waitingTime": 30
```

   The following is the example script that uses system context.

```
"SessionStart": {
    "executables": [
    {
        "context": "system",
        "filename": "C:\\Windows\\System32\\WindowsPowerShell\\v1.0\
\powershell.exe",
        "arguments": "-File \"C:\\AppStream\\SessionScripts\\systemStart.ps1\"",
        "s3LogEnabled": true
    },
```

```
    {
        "context": "user",
        "filename": "",
        "arguments": "",
        "s3LogEnabled": true
    }
 ],
"waitingTime": 30
```

2. If you're using multi-session fleets, you can use the system environment variable `$env:AppStream_Session_UserName` to navigate to the your user's My Files folder. This allows mapping to `Admin` instead of the user name when using the system context `$env:USERNAME`.

```
# Define the target application path
$targetPaths = "<SMB-PATH>"

# Define the shortcut location
$symlinkLocation = "C:\Users\$Env:AppStream_Session_UserName\My Files\Custom
 Folder"

# Create the junction for Custom Home Folder under MyFiles
New-Item -ItemType SymbolicLink -Path $symlinkLocation -Target $targetPaths
```

Option 2: Use GPO Logon Script to mount SMB shared folders to be under My Files

1. Mount SMB shared folders by creating a symbolic link to a file or folder. For more information, see Example 7: Create a symbolic link to a file or folder
2. Assign user logon scripts.
3. Add the following script to create a junction for Custom Home Folders, under My Files.

```
# Define the target application path
$targetPaths = "<SMB-PATH>"

# Define the shortcut location
$symlinkLocation = "C:\Users\$env:Username\My Files\Custom Folder"

# Create the junction for Custom Home Folder under MyFiles
New-Item -ItemType SymbolicLink -Path $symlinkLocation -Target $targetPaths
```

If you are using Windows Server 2022 images, you might experience an issue where your My Files folder doesn't get created until the Logon Script is completed successfully. This might can cause a timeout when your SMB mounting operation is done through Logon Script. To resolve this issue, while also mounting your SMB, trigger an independent process (`Start-Process`) using your Logon Script by doing the following:

a. Create a Logon Script.

```
# Define the log file path
$logFilePath = "<This-is-where-your-log-files-are-saved>"

# Function to write log messages
function Write-Log {
    param (
        [string]$message
    )
    $timestamp = get-date -format "yyyy-MM-dd HH:mm:ss"
    $logMessage = "$timestamp - $message"
    $logMessage | Out-File -FilePath $logFilePath -Append -Encoding UTF8
}

try {
    Write-Log "Setting execution policy..."
    Set-ExecutionPolicy -ExecutionPolicy Bypass -Scope Process -Force
    Write-Log "Unblocking logon script file..."
    $filePath = "<This-is-where-your-actual-logon-script-is-linked>"
    Unblock-File -Path $filePath
    Write-Log "Running actual logon script..."
    Start-Process -FilePath 'Powershell.exe' -ArgumentList "-File
 `"$filePath`""
} catch {
    Write-Log "An error occurred: $_" "ERROR"
}
```

b. Update this Logon Script delay configuration using Group Policy, if needed. For more information, see Configure Logon Script Delay on the Microsoft Learn website. Logon Script delay will be the amount of time it will delay before triggering your async Logon Script. The default delay is 5 minutes.

c. Restart your fleet to apply the Logon Script delay.

Enable Application Settings Persistence for Your WorkSpaces Applications Users

WorkSpaces Applications supports persistent application settings for Windows-based stacks. This means that your users' application customizations and Windows settings are automatically saved after each streaming session and applied during the next session. Examples of persistent application settings that your users can configure include, but are not limited to, browser favorites, settings, webpage sessions, application connection profiles, plugins, and UI customizations. These settings are saved to an Amazon Simple Storage Service (Amazon S3) bucket in your account, within the AWS Region in which application settings persistence is enabled. They are available in each WorkSpaces Applications streaming session.

Note

Enabling application settings persistence is currently not supported for Linux-based stacks.

Note

Standard Amazon S3 charges may apply to data that is stored in your S3 bucket. For more information, see Amazon S3 Pricing.

Contents

- How Application Settings Persistence Works
- Enabling Application Settings Persistence
- Administer the VHDs for Your Users' Application Settings

How Application Settings Persistence Works

Persistent application settings are saved to a Virtual Hard Disk (VHD) file. This file is created the first time a user streams an application from a stack on which application settings persistence is enabled. If the fleet associated with the stack is based on an image that contains default application and Windows settings, the default settings are used for the user's first streaming

session. For more information about default settings, see *Step 3: Create Default Application and Windows Settings* in Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

When the streaming session ends, the VHD is unmounted and uploaded to an Amazon S3 bucket within your account. The bucket is created when you enable persistent application settings for the first time for a stack in an AWS Region. The bucket is unique to your AWS account and the Region. The VHD is encrypted in transit using Amazon S3 SSL endpoints, and at rest using AWS Managed CMKs.

The VHD is mounted to the streaming instance in both C:\Users\%username% and D:\%username%. If your instance is not joined to an Active Directory domain, the Windows user name is PhotonUser. If your instance is joined to an Active Directory domain, the Windows user name is that of the logged in user.

Application settings persistence does not work across different operating system versions. For example, if you enable application settings persistence on a stack and the stack is associated with a fleet that uses a Windows Server 2012 R2 image, if you update the fleet to use an image that runs a different operating system (such as Windows Server 2016), settings from previous streaming sessions are not saved for users of the stack. Instead, after you update the fleet to use the new image, when users launch a streaming session from a fleet instance, a new Windows user profile is created. However, if you apply an update to the same operating system on the image, users' customizations and settings from previous streaming sessions are saved. When updates to the same operating system are applied to an image, the same Windows user profile is used when users launch a streaming session from the fleet instance.

Important

WorkSpaces Applications supports applications that rely on the Microsoft Data Protection API only when the streaming instance is joined to a Microsoft Active Directory domain. In cases where a streaming instance is not joined to an Active Directory domain, the Windows user, PhotonUser, is different on each fleet instance. Due to the way in which the DPAPI security model works, users' passwords don't persist for applications that use DPAPI in this scenario. In cases where streaming instances are joined to an Active Directory domain and the user is a domain user, the Windows user name is that of the logged in user, and users' passwords persist for applications that use DPAPI.

WorkSpaces Applications automatically saves all files and folders in this path, except for the following folders:

- Contacts
- Desktop
- Documents
- Downloads
- Links
- Pictures
- Saved Games
- Searches
- Videos

Files and folders created outside of these folders are saved within the VHD and synced to Amazon S3. The default VHD maximum size is 1GB for Elastic fleets and 5GB for Always-On and On-Demand fleets. The size of the saved VHD is the total size of the files and folders that it contains. WorkSpaces Applications automatically saves the HKEY_CURRENT_USER registry hive for the user. For new users (users whose profiles don't exist in Amazon S3), WorkSpaces Applications creates the initial profile by using the default profile. This profile is created in the following location on the image builder: C:\users\default.

Note

The entire VHD must be downloaded to the streaming instance before a streaming session can begin. For this reason, a VHD that contains a large amount of data can delay the start of the streaming session. For more information, see Best Practices for Enabling Application Settings Persistence.

When you enable application settings persistence, you must specify a settings group. The settings group determines which saved application settings are used for a streaming session from this stack. WorkSpaces Applications creates a new VHD file for the settings group that is stored separately within the S3 bucket in your AWS account. If the settings group is shared between stacks, the same application settings are used in each stack. If a stack requires its own application settings, specify a unique settings group for the stack.

Enabling Application Settings Persistence

Review the following topics to learn how to enable application settings persistence for your WorkSpaces Applications users.

Contents

- Prerequisites for Enabling Application Settings Persistence
- Best Practices for Enabling Application Settings Persistence
- How to Enable Application Settings Persistence

Prerequisites for Enabling Application Settings Persistence

To enable application settings persistence, you must first do the following:

- Check that you have the correct AWS Identity and Access Management (IAM) permissions for Amazon S3 actions. For more information, see the *IAM Policies and the Amazon S3 Bucket for Home Folders* section in Identity and Access Management for Amazon WorkSpaces Applications.
- Use an image that was created from a base image published by AWS on or after December 7, 2017. For a current list of released AWS base images, see WorkSpaces Applications Base Image and Managed Image Update Release Notes.
- Associate the stack on which you plan to enable this feature with a fleet based on an image that uses a version of the WorkSpaces Applications agent released on or after August 29, 2018. For more information, see WorkSpaces Applications Agent Release Notes.
- Enable network connectivity to Amazon S3 from your virtual private cloud (VPC) by configuring internet access or a VPC endpoint for Amazon S3. For more information, see the *Home Folders and VPC Endpoints* section in Networking and Access for Amazon WorkSpaces Applications.

Best Practices for Enabling Application Settings Persistence

To enable application settings persistence without providing internet access to your instances, use a VPC endpoint. This endpoint must be in the VPC to which your WorkSpaces Applications instances are connected. You must attach a custom policy to enable WorkSpaces Applications access to the endpoint. For information about how to create the custom policy, see the *Home Folders and VPC Endpoints* section in Networking and Access for Amazon WorkSpaces Applications. For more information about private Amazon S3 endpoints, see VPC Endpoints and Endpoints for Amazon S3 in the *Amazon VPC User Guide*.

How to Enable Application Settings Persistence

You can enable or disable application settings persistence while creating a stack or after the stack is created by using the WorkSpaces Applications console, WorkSpaces Applications API, an AWS SDK, or the AWS Command Line Interface (CLI). For each AWS Region, persistent application settings are stored in an S3 bucket in your account.

The first time you enable application settings persistence for a stack in an AWS Region, WorkSpaces Applications creates an S3 bucket in your AWS account in the same Region. The same bucket stores the application settings VHD file for all users and all stacks in that AWS Region. For more information, see *Amazon S3 Bucket Storage* in Administer the VHDs for Your Users' Application Settings.

To enable application settings persistence while creating a stack

- Follow the steps in Create a Stack in Amazon WorkSpaces Applications, and make sure that **Enable Application Settings Persistence** is selected.

To enable application settings persistence for an existing stack

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.
2. In the left navigation pane, choose **Stacks**, and select the stack for which to enable application settings persistence.
3. Below the stacks list, choose **User Settings**, **Application Settings Persistence**, **Edit**.
4. In the **Application Settings Persistence** dialog box, choose **Enable Application Settings Persistence**.
5. Confirm the current settings group or type the name of a new settings group. When you're done, choose **Update**.

New streaming sessions now have application settings persistence enabled.

Administer the VHDs for Your Users' Application Settings

Review the following topics to learn how to administer the Virtual Hard Disks (VHD) files for your WorkSpaces Applications users' application settings.

Contents

- Amazon S3 Bucket Storage
- Reset a User's Application Settings
- Enable Amazon S3 Object Versioning and Revert a User's Application Settings
- Increase the Size of the Application Settings VHD

Amazon S3 Bucket Storage

When you enable application settings persistence, your users' application customizations and Windows settings are automatically saved to a Virtual Hard Disk (VHD) file that is stored in an Amazon S3 bucket created in your AWS account. For every AWS Region, WorkSpaces Applications creates a bucket in your account that is unique to your account and the Region. All application settings configured by your users are stored in the bucket for that Region.

You do not need to perform any configuration tasks to manage these S3 buckets; they are fully managed by the WorkSpaces Applications service. The VHD file that is stored in each bucket is encrypted in transit using Amazon S3's SSL endpoints and at rest using AWS Managed CMKs. The buckets are named in a specific format as follows:

```
appstream-app-settings-region-code-account-id-without-hyphens-random-identifier
```

region-code

This is the AWS Region code in which the stack is created with application settings persistence.

account-id-without-hyphens

Your AWS account ID. The random identifier ensures there is no conflict with other buckets in that Region. The first part of the bucket name, `appstream-app-settings`, does not change across accounts or Regions.

For example, if you enable application settings persistence for stacks in the US West (Oregon) Region (us-west-2) on account number 123456789012, WorkSpaces Applications creates an Amazon S3 bucket within your account in that Region with the name shown. Only an administrator with sufficient permissions can delete this bucket.

```
appstream-app-settings-us-west-2-1234567890123-abcdefg
```

Disabling application settings persistence does not delete any VHDs stored in the S3 bucket. To permanently delete settings VHDs, you or another administrator with adequate permissions must do so by using the Amazon S3 console or API. WorkSpaces Applications adds a bucket policy that prevents accidental deletion of the bucket. For more information, see *IAM Policies and the Amazon S3 Bucket for Application Settings Persistence* in Identity and Access Management for Amazon WorkSpaces Applications.

When application settings persistence is enabled, a unique folder is created for each settings group to store the settings VHD. The hierarchy of the folder in the S3 bucket depends on how the user launches a streaming session, as described in the following section.

The path for the folder where the settings VHD is stored in the S3 bucket in your account uses the following structure:

```
bucket-name/Windows/prefix/settings-group/access-mode/user-id-SHA-256-hash
```

bucket-name

The name of the S3 bucket in which users' application settings are stored. The name format is described earlier in this section.

prefix

The Windows version-specific prefix. For example, v4 for Windows Server 2012 R2.

settings-group

The settings group value. This value is applied to one or more stacks that share the same the same application settings.

access-mode

The identity method of the user: `custom` for the WorkSpaces Applications API or CLI, `federated` for SAML, and `userpool` for user pool users.

user-id-SHA-256-hash

The user-specific folder name. This name is created using a lowercase SHA-256 hash hexadecimal string generated from the user ID.

The following example folder structure applies to a streaming session that is accessed using the API or CLI with a user ID of `testuser@mydomain.com`, an AWS account ID of `123456789012`, and the settings group `test-stack` in the US West (Oregon) Region (us-west-2):

```
appstream-app-settings-us-west-2-1234567890123-abcdefg/Windows/v4/test-stack/custom/
a0bcb1da11f480d9b5b3e90f91243143eac04cfccfbdc777e740fab628a1cd13
```

You can identify the folder for a user by generating the lowercase SHA-256 hash value of the user ID using websites or open source coding libraries available online.

Reset a User's Application Settings

To reset a user's application settings, you must find and delete the VHD and associated metadata file from the S3 bucket in your AWS account. Make sure that you do not do this during a user's active streaming session. After you delete the user's VHD and the metadata file, the next time the user launches a session from a streaming instance that has application settings persistence enabled, WorkSpaces Applications creates a new settings VHD for that user.

To reset a user's application settings

1. Open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the **Bucket name** list, choose the S3 bucket that contains the application settings VHD that you want to reset.
3. Locate the folder that contains the VHD. For more information about how to navigate the S3 bucket folder structure, see *Amazon S3 Bucket Storage* earlier in this topic.
4. In the **Name** list, select the check box next to the VHD and the REG, choose **More**, and then choose **Delete**.
5. In the **Delete objects** dialog box, verify that the VHD and the REG are listed, and then choose **Delete**.

The next time the user streams from a fleet on which application settings persistence is enabled with the applicable settings group, a new application settings VHD is created. This VHD is saved to the S3 bucket at the end of the session.

Enable Amazon S3 Object Versioning and Revert a User's Application Settings

You can use Amazon S3 object versioning and lifecycle policies to manage your users' application settings when your users change them. With Amazon S3 object versioning, you can preserve, retrieve, and restore every version of the settings VHD. This enables you to recover from both unintended user actions and application failures. When versioning is enabled, after each streaming

session, a new version of the application settings VHD is synced to Amazon S3. The new version does not overwrite the previous version, so if an issue with your users' settings occurs, you can revert to a previous version of the VHD.

> **Note**
>
> Each version of the application settings VHD is saved to Amazon S3 as a separate object and is charged accordingly.

Object versioning is not enabled by default in your S3 bucket, so you must explicitly enable it.

To enable object versioning for your application settings VHD

1. Open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the **Bucket name** list, choose the S3 bucket that contains the application settings VHD on which to enable object versioning.
3. Choose **Properties**.
4. Choose **Versioning**, **Enable versioning**, and then choose **Save**.

To expire older versions of your application settings VHDs, you can use Amazon S3 lifecycle policies. For information, see How Do I Create a Lifecycle Policy for an S3 Bucket? in the *Amazon Simple Storage Service User Guide*.

To revert a user's application settings VHD

You can revert to a previous version of a user's application settings VHD by deleting newer versions of the VHD from the applicable S3 bucket. Do not do this when the user has an active streaming session.

1. Open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the **Bucket name** list, choose the S3 bucket that contains the user's application settings VHD version to revert to.
3. Locate and select the folder that contains the VHD. For information about how to navigate the S3 bucket folder structure, see *Amazon S3 Bucket Storage* earlier in this topic.

   When you select the folder, the settings VHD and associated metadata file display.
4. To display a list of the VHD and metadata file versions, choose **Show**.

5. Locate the version of the VHD to revert to.
6. In the **Name** list, select the check boxes next to the newer versions of the VHD and associated metadata files, choose **More**, and then choose **Delete**.
7. Verify that the application settings VHD that you want to revert to and the associated metadata file are the newest versions of these files.

The next time the user streams from a fleet on which application settings persistence is enabled with the applicable settings group, the reverted version of the user's settings displays.

Increase the Size of the Application Settings VHD

The default VHD maximum size is 1 GB for Elastic fleets and 5GB for Always-On and On-Demand fleets. If a user requires additional space for application settings, you can download the applicable application settings VHD to a Windows computer to expand it. Then, replace the current VHD in the S3 bucket with the larger one. Do not do this when the user has an active streaming session.

> **Note**
>
> To reduce the physical size of the virtual hard disk (VHD), clear the recycle bin before ending a session. This also reduces upload and download times, and improves the overall user experience.

To increase the size of the application settings VHD

> **Note**
>
> The full VHD must be downloaded before a user can stream applications. Increasing the size of an application settings VHD can increase the time it takes for users to start application streaming sessions.

1. Open the Amazon S3 console at https://console.aws.amazon.com/s3/.
2. In the **Bucket name** list, choose the S3 bucket that contains the application settings VHD to expand.
3. Locate and select the folder that contains the VHD. For information about how to navigate the S3 bucket folder structure, see *Amazon S3 Bucket Storage* earlier in this topic.

When you select the folder, the settings VHD and associated metadata file display.

4. Download the Profile.vhdx file to a directory on your Windows computer. Do not close your browser after the download completes, because you'll use the browser again later to upload the expanded VHD.
5. To use Diskpart to increase the size of the VHD to 7 GB, open the command prompt as an administrator, and type the following commands.

   ```
   diskpart

   select vdisk file="C:\path\to\application\settings\profile.vhdx"

   expand vdisk maximum=7000
   ```
6. Then, type the following Diskpart commands to find and attach the VHD, and display the list of volumes:

   ```
   select vdisk file="C:\path\to\application\settings\profile.vhdx"

   attach vdisk

   list volume
   ```

   In the output, make note of the volume number with the label "AppStreamUsers". In the next step, you select this volume so that you can enlarge it.
7. Type the following command:

   ```
   select volume ###
   ```

   where ### is the number in the list volume output.
8. Type the following command:

   ```
   extend
   ```
9. Type the following commands to confirm that the size of the partition on the VHD increased as expected (2 GB in this example):

   ```
   diskpart

   select vdisk file="C:\path\to\application\settings\profile.vhdx"

   list volume
   ```

10. Type the following command to detach the VHD so that it can be uploaded:

    ```
    detach vdisk
    ```

11. Return to your browser with the Amazon S3 console, choose **Upload**, **Add files**, and then select the enlarged VHD.
12. Choose **Upload**.

After the VHD is uploaded, the next time the user streams from a fleet on which application settings persistence is enabled with the applicable settings group, the larger application settings VHD is available.

Enable Regional Settings for Your WorkSpaces Applications Users

WorkSpaces Applications lets you or your users configure certain Windows settings that are specific to your users' location or language. WorkSpaces Applications also lets you configure regional settings while creating Linux images. For more information, see Tutorial: Enable Japanese Support for Your Linux Images.

Contents

- Configure Default Regional Settings for Your WorkSpaces Applications Users
- Enable Your WorkSpaces Applications Users to Configure Their Regional Settings

Configure Default Regional Settings for Your WorkSpaces Applications Users

> **Note**
>
> The instructions on this page only apply to Windows fleets. Default regional settings are not supported for Elastic fleets.

In WorkSpaces Applications, users in a Windows stack can configure their streaming sessions to use settings that are specific to their location or language. For more information, see Enable Your WorkSpaces Applications Users to Configure Their Regional Settings. You can also configure your fleets to use default settings that are specific to your users' location or language. In particular, you can apply the following Windows settings to your fleets:

- **Time Zone** — Determines the system time used by Windows and any applications that rely on the operating system time. WorkSpaces Applications makes available the same options for this setting as Windows Server 2019 or later.
- **Display Language** — Determines the display language used by the Windows operating system and certain Windows applications.
- **System Locale** — Determines the code pages (ANSI, MS-DOS, and Macintosh) and bitmap font files that Windows uses for non-Unicode applications in different languages.

- **User Locale** (also known as culture) — Determines the conventions used by Windows and any applications that query the Windows culture when formatting dates, numbers, or currencies or when sorting strings.
- **Input Method** — Determines the keystroke combinations that can be used to enter characters in another language.

Currently, WorkSpaces Applications supports English and Japanese only for these language settings.

Contents

- Specify a Default Time Zone
- Specify a Default Display Language
- Specify a Default System Locale
- Specify a Default User Locale
- Specify a Default Input Method
- Configuring Chinese and Korean input methods on the image
- Special Considerations for Application Settings Persistence
- Special Considerations for Japanese Language Settings

Specify a Default Time Zone

To specify a default time zone to be used in your users' streaming sessions, perform the steps in either of the following two procedures.

Procedures

- Specify a Default Time Zone (Windows Server 2012 R2)
- Specify a Default Time Zone (Windows Server 2016, Windows Server 2019, Windows Server 2022, and Windows Server 2025)

> **Note**
>
> Currently, WorkSpaces Applications supports only **UTC** and **(UTC+9:00) Osaka, Sapporo, Tokyo**.

Specify a Default Time Zone (Windows Server 2012 R2)

1. Connect to the image builder that you want to use and sign in with a user that has local administrator permissions. To do so, do either of the following:

   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

   > **Note**
   >
   > If the image builder that you want to connect to is joined to an Active Directory domain and your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection. For information about smart card sign in, see Smart Cards.

2. On the image builder desktop, choose the Windows **Start** button, and choose **Control Panel**.
3. Choose **Clock, Language, and Region**, then **Date and Time**, then **Change time zone**.
4. In the **Time zone** list, choose a time zone, and choose **OK**.
5. To apply any change to the time zone setting, restart your image builder. To do so, choose the Windows **Start** button, and choose **Windows PowerShell**. In PowerShell, use the **restart-computer** cmdlet.
6. While Windows restarts, the WorkSpaces Applications login prompt displays. Wait for 10 minutes before you log in to the image builder again. Otherwise, you may receive an error. After 10 minutes, you can log in as **Administrator**.
7. If required, configure additional default regional or language settings. Otherwise, on the image builder desktop, open Image Assistant and install and configure applications for streaming.
8. After you finish configuring your image builder, follow the necessary steps in Image Assistant to finish creating your image. For information about how to create an image, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.
9. Do one of the following:

   - Create a new fleet and choose your new image for the fleet. For more information, see Create an Amazon WorkSpaces Applications Fleet and Stack.
   - Update an existing fleet to use the new image.

10. Associate your fleet with the stack that is assigned to the users for whom you are configuring the default settings.

    The default time zone setting that you configured is applied to the fleet instances and user streaming sessions that are launched from those instances.

Specify a Default Time Zone (Windows Server 2016, Windows Server 2019, Windows Server 2022, and Windows Server 2025)

1. Connect to the image builder that you want to use and sign in with an account that has local administrator permissions. To do so, do either of the following:

   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

   > **Note**
   >
   > If the image builder that you want to connect to is joined to an Active Directory domain and your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection. For information about smart card sign in, see Smart Cards.

2. On the image builder desktop, choose the Windows **Start** button, and choose **Control Panel**.
3. Specify the default time zone by using PowerShell or the Windows user interface:

   - **PowerShell**
     - Open PowerShell and run the following command:

       ```
       Run Set-TimeZone -Id "Tokyo Standard Time"
       ```

       > **Note**
       >
       > To run this command, you must be logged in to the applicable computer as **Administrator**.

   - **Windows user interface**

   1. On the image builder desktop, choose the Windows **Start** button, and type **`timedate.cpl`** to open the **Date and Time** control panel item.
   2. Right-click the **Date and Time** icon, and choose **Run as administrator**.
   3. When prompted by **User Account Control** to choose whether you want to allow the app to make changes to your device, choose **Yes**.
   4. Choose **Change time zone**.
   5. In the **Time zone** list, choose a time zone, and choose **OK**.
4. If required, configure additional default regional or language settings. Otherwise, on the image builder desktop, open Image Assistant and install and configure applications for streaming.
5. After you finish configuring your image builder, follow the necessary steps in Image Assistant to finish creating your image. For information about how to create an image, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.
6. Do one of the following:
   - Create a new fleet and choose your new image for the fleet. For more information, see Create an Amazon WorkSpaces Applications Fleet and Stack.
   - Update an existing fleet to use the new image.
7. Associate your fleet with the stack that is assigned to the users for whom you are configuring the default settings.

   The default time zone setting that you configured is applied to the fleet instances and user streaming sessions that are launched from those instances.

> **Note**
>
> Your users can change their time zone from the default setting that you configured. They can configure their regional settings during an application streaming session, as described in Enable Your WorkSpaces Applications Users to Configure Their Regional Settings. Also, if a user previously selected a time zone when streaming from any fleet instance in the same AWS Region, the user-specified time zone setting automatically overrides any default time zone setting you specify through your image builder.

Specify a Default Display Language

There are two ways to specify the default display language for your users' streaming sessions. Use the WorkSpaces Applications default application and Windows settings feature, or configure your image builder while signed in with an account that has local administrator permissions. The procedure in this section describes how to specify a default display language by using the WorkSpaces Applications default application and Windows settings feature.

> **Note**
>
> Changing the display language in Windows also automatically changes the user locale and input method to match the language and region of the display language. If you want all three settings to match, you do not need to separately change the user locale or input method.

1. Connect to the image builder that you want to use and sign in with the **Template User** account. To do so, do either of the following:

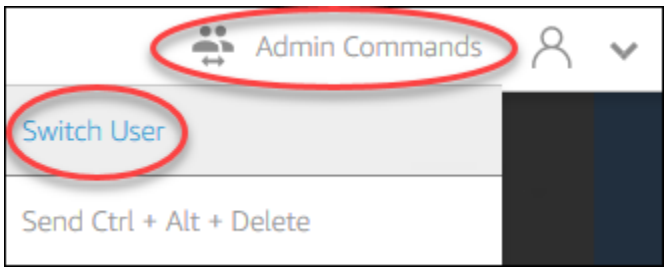
   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

   > **Note**
   >
   > If the image builder that you want to connect to is joined to an Active Directory domain and your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection. For information about smart card sign in, see Smart Cards.

   Template User lets you create default application and Windows settings for your users. For more information, see "Creating Default Application and Windows Settings for Your WorkSpaces Applications Users" in Default Application and Windows Settings and Application Launch Performance in Amazon WorkSpaces Applications.
2. On the image builder desktop, choose the Windows **Start** button, and choose **Control Panel**.
3. Choose **Clock, Language, and Region**, then **Language**, **Add a language**.
4. Choose a language, and choose **Add**.

> **Note**
>
> Currently, WorkSpaces Applications supports only **English (United States)** and **Japanese**.

5. The language that you selected appears in the list of languages you added to Windows. Choose the language that you just added. Then choose **Move up** until the language appears at the top of the language list.
6. Choose **Advanced Settings**. Under **Override for Windows display language**, choose your language from the list.
7. If you want to use the input method associated with the language that you added, under **Override for default input method**, choose the input method for the language.
8. Choose **Save**. When prompted to log off, choose **Log off now**.
9. When prompted, log in again to the image builder as **Template User**. Confirm that Windows is using the display language that you selected.
10. In the upper right area of the image builder desktop, choose **Admin Commands**, **Switch User**.



11. When prompted, log in as **Administrator**.
12. If required, configure additional default regional or language settings. Otherwise, on the image builder desktop, open Image Assistant and install and configure applications for streaming.
13. In Step 2 of the Image Assistant process, choose **Save settings**.
14. Follow the necessary steps in Image Assistant to finish creating your image. For information about how to create an image, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.
15. Do one of the following:

    - Create a new fleet and choose your new image for the fleet. For information, see Create an Amazon WorkSpaces Applications Fleet and Stack.

- Update an existing fleet to use the new image.

16. Associate your fleet with the stack that is assigned to the users for whom you are configuring the default settings.

    The default display language and associated user locale and input method settings that you configured are applied to the fleet instances and user streaming sessions that are launched from those instances.

    Alternatively, you can configure a default display language while logged in to the image builder as **Administrator**. If you chose different display languages while you were logged in under the **Template User** and **Administrator** accounts and you chose **Save settings** in Step 2 of the Image Assistant process, the **Template User** settings take precedence.

> **Note**
>
> Your users can change their user locale and input method from the default settings that you configured. They can change to any one of 11 different supported locales and nine different supported input methods. To do so, they can configure their regional settings during application streaming sessions, as described in Enable Your WorkSpaces Applications Users to Configure Their Regional Settings. Also, if a user previously selected a user locale or input method when streaming from any fleet instance in the same Region, those user-specified settings automatically override any default user locale and input method that you specify through your image builder.

Specify a Default System Locale

To specify a default system locale for your users' streaming sessions, perform the following steps.

1. Connect to the image builder that you want to use and sign in with an account that has local administrator permissions. To do so, do either of the following:

    - Use the WorkSpaces Applications console (for web connections only)
    - Create a streaming URL (for web or WorkSpaces Applications client connections)

> **Note**
>
> If the image builder that you want to connect to is joined to an Active Directory domain and your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection. For information about smart card sign in, see Smart Cards.

2. On the image builder desktop, choose the Windows **Start** button, and choose **Control Panel**.
3. Choose **Clock, Language, and Region**, then **Region**.
4. In the **Region** dialog box, choose the **Formats** tab.
5. Choose **Change system locale**.
6. In the **Region Settings** dialog box, in the **Current system locale** list, choose a language and region.

> **Note**
>
> Currently, WorkSpaces Applications supports only **English (United States)** and **Japanese (Japan)**.

7. Choose **OK** to close the **Region Settings** dialog box, and choose **OK** again to close the **Region** dialog box.
8. When prompted to restart your computer, allow Windows to restart.
9. While Windows restarts, the WorkSpaces Applications login prompt displays. Wait for 10 minutes before you log in to the image builder again. Otherwise, you may receive an error. After 10 minutes, you can log in as **Administrator**.
10. If required, configure additional default regional or language settings. Otherwise, on the image builder desktop, open Image Assistant and install and configure applications for streaming. After you finish configuring your image builder, follow the necessary steps in Image Assistant to finish creating your image. For information about how to create an image, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.
11. Do one of the following:

    - Create a new fleet and choose your new image for the fleet. For more information, see Create an Amazon WorkSpaces Applications Fleet and Stack.

- Update an existing fleet to use the new image.

12. Associate your fleet with the stack that is assigned to the users for whom you are configuring the default settings.

    The default system locale setting that you configured is applied to the fleet instances and user streaming sessions that are launched from those instances.

Specify a Default User Locale

To specify a default user locale for your users' streaming sessions, perform the following steps.

> **Note**
>
> If you plan to configure the display language and you want the user locale and display language to match, you do not need to change the user locale. Changing the display language automatically changes the user locale to match.

1. Connect to the image builder that you want to use and sign in with an account that has local administrator permissions. To do so, do either of the following:

   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

   > **Note**
   >
   > If the image builder that you want to connect to is joined to an Active Directory domain and your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection. For information about smart card sign in, see Smart Cards.

2. On the image builder desktop, choose the Windows **Start** button, and choose **Control Panel**.
3. Choose **Clock, Language, and Region**, then **Region**.
4. In the **Region** dialog box, choose the **Formats** tab.
5. In the **Format** list, choose a language and region.

> **Note**
>
> Currently, WorkSpaces Applications supports only **English (United States)** and **Japanese (Japan)**.

6. Choose **OK** to close the **Region** dialog box.
7. If required, configure additional default regional or language settings. Otherwise, on the image builder desktop, open Image Assistant and install and configure applications for streaming.
8. In Step 2 of the Image Assistant process, choose **Save settings**.
9. Follow the necessary steps in Image Assistant to finish creating your image. For information about how to create an image, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.
10. Do one of the following:
    - Create a new fleet and choose your new image for the fleet. For more information, see Create an Amazon WorkSpaces Applications Fleet and Stack.
    - Update an existing fleet to use the new image.
11. Associate your fleet with the stack that is assigned to the users for whom you are configuring the default settings.

    The default user locale setting that you configured is applied to the fleet instances and user streaming sessions that are launched from those instances.

> **Note**
>
> Your users can change their user locale from the default setting that you configured to any one of 11 different supported locales. To do so, they can configure their regional settings during application streaming sessions, as described in Enable Your WorkSpaces Applications Users to Configure Their Regional Settings. Also, if a user previously selected a user locale when streaming from any fleet instance in the same Region, that user-specified setting automatically overrides any default user locale setting that you specify through your image builder.

Specify a Default Input Method

To specify a default input method to be used in your users' streaming sessions, perform the following steps.

> **Note**
>
> If you plan to configure the display language, and you want the input method and display language to match, you do not need to change the input method. Changing the display language in Windows also automatically changes the user locale and input method to match the language and region of the display language. If you want all three settings to match, you do not need to separately change the user locale or input method.

1. Connect to the image builder that you want to use and sign in with an account that has local administrator permissions. To do so, do either of the following:

   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

   > **Note**
   >
   > If the image builder that you want to connect to is joined to an Active Directory domain and your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection. For information about smart card sign in, see Smart Cards.

2. On the image builder desktop, choose the Windows **Start** button, and choose **Control Panel**.
3. Choose **Clock, Language, and Region**, then **Language**, **Add a language**.
4. Choose a language, and choose **Add**.

   > **Note**
   >
   > Currently, WorkSpaces Applications supports only **English (United States)** and **Japanese** as admin-configurable default input methods. To configure Chinese or Korean input methods on your image, see Configuring Chinese and Korean input methods on the image.

5. The language that you chose appears in the list of languages you added to Windows.
6. Choose **Advanced Settings**. Under **Override for default input method**, choose the input method for the language you added.
7. Choose **Save**.
8. Log off and log in again. To do so, choose the Windows **Start** button on the image builder desktop. Choose **ImageBuilderAdmin**, **Sign out**. When prompted, log in as Administrator.
9. If required, configure additional default regional or language settings. Otherwise, on the image builder desktop, open Image Assistant and install and configure applications for streaming.
10. In Step 2 of the Image Assistant process, choose **Save settings**.
11. Follow the necessary steps in Image Assistant to finish creating your image. For information about how to create an image, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.
12. Do one of the following:
    - Create a new fleet and choose your new image for the fleet. For information, see Create an Amazon WorkSpaces Applications Fleet and Stack.
    - Update an existing fleet to use the new image.
13. Associate your fleet with the stack that is assigned to the users for whom you are configuring the default settings.

    The default input method that you configured is applied to the fleet instances and user streaming sessions that are launched from those instances.

Note

Your users can change their input method from the default setting that you configured to any one of nine different supported input methods. They can configure this setting by configuring their regional settings during application streaming sessions, as described in Enable Your WorkSpaces Applications Users to Configure Their Regional Settings. Also, if a user previously selected an input method when streaming from any fleet instance in the same Region, that user-specified setting automatically overrides any default input method that you specify through your image builder.

Configuring Chinese and Korean input methods on the image

Although end users can select from 9 supported input methods during their streaming sessions through **Preferences** > **Input method**, certain input methods require you to pre-install the corresponding language components on the image before they can function correctly. These input methods include Chinese (Taiwan), Chinese (Pinyin), Chinese (Traditional, DaYi), and Korean.

Without these language components, selecting these input methods from Preferences might result in keystrokes being passed through as literal English characters.

Prerequisites

Before you begin, make sure that your environment meets the following requirements:

- An Image Builder running Windows Server 2019 or later
- Administrator access to the Image Builder

Configure Chinese (Traditional, Taiwan) — Bopomofo input

1. Open the WorkSpaces Applications console, and then launch and connect to your **Image Builder**.
2. Open **Settings** > **Time & Language** > **Language**, then select **Add a language** and choose **Chinese (Traditional, Taiwan)**.

   Alternatively, open PowerShell as Administrator and run:

   ```
   $langList = Get-WinUserLanguageList
   $langList.Add("zh-Hant-TW")
   Set-WinUserLanguageList $langList -Force
   ```

3. Verify that the Basic Typing language component is installed:

   ```
   Get-WindowsCapability -Online | Where-Object { $_.Name -like "*zh-TW*" -and
    $_.State -eq "Installed" }
   ```

   You should see `Language.Basic~~~zh-TW~0.0.1.0` in the output.
4. (Optional) Verify the language is registered with the correct IME:

```
Get-WinUserLanguageList
```

Confirm that `zh-Hant-TW` appears with `InputMethodTips` containing the Bopomofo IME identifier.

5. On the Image Builder desktop, open **Image Assistant** and follow the steps to create your image.
6. Create or update your fleet using the new image.

After this configuration, end users can select **Chinese (Taiwan)** from **Preferences** > **Input method** during their streaming sessions. Bopomofo input functions correctly, and users can switch between Chinese (Taiwan) and Chinese (Pinyin) input methods through Preferences.

Configure Chinese (Simplified) — Pinyin input

1. Open the WorkSpaces Applications console, and then launch and connect to your **Image Builder**.
2. Open PowerShell as Administrator and run:

```
$langList = Get-WinUserLanguageList
$langList.Add("zh-Hans-CN")
Set-WinUserLanguageList $langList -Force
```

3. Verify installation:

```
Get-WindowsCapability -Online | Where-Object { $_.Name -like "*zh-CN*" -and
 $_.State -eq "Installed" }
```

You should see `Language.Basic~~~zh-CN~0.0.1.0` in the output.

4. Create your image and fleet.

Configure Korean input

1. Open the WorkSpaces Applications console, and then launch and connect to your **Image Builder**.
2. Open PowerShell as Administrator and run:

```
$langList = Get-WinUserLanguageList
$langList.Add("ko-KR")
Set-WinUserLanguageList $langList -Force
```

3. Verify installation:

```
Get-WindowsCapability -Online | Where-Object { $_.Name -like "*ko-KR*" -and
 $_.State -eq "Installed" }
```

   You should see `Language.Basic~~~ko-KR~0.0.1.0` in the output.

4. Create your image and fleet.

Verifying the configuration

After creating your fleet from the configured image, verify the input method works:

1. Connect to a streaming session.
2. Open **Preferences** > **Input method** and select the desired input method (for example, Chinese (Taiwan)).
3. Open a text application (for example, Notepad).
4. Type test characters:
   - For Bopomofo: Type **`su3`** — this should produce 你
   - For Pinyin: Type **`nihao`** — this should produce 你好
   - For Korean: Type using the Korean keyboard layout — Hangul characters should appear

Important notes

- **Wait for installation to complete:** After adding a language, allow time for the IME installation to finish before testing. If you attempt to type immediately, you might receive a notification that the IME is not ready.
- **Multiple languages:** You can install multiple language components on the same image. End users can switch between all configured input methods through Preferences.
- **Japanese does not require this step:** Japanese input works without additional configuration because the base image includes the required language components by default.

- **DaYi input method:** Chinese (Traditional, DaYi) uses the same `zh-Hant-TW` language pack as Bopomofo. After installing the Chinese (Traditional, Taiwan) language, both Bopomofo and DaYi IMEs are available for end users to select.
- **Admin defaults vs. user-selectable:** This procedure makes the input methods functional for end users to select during sessions. The admin-configurable default input method (configured through Template User or Image Assistant) still supports only English (United States) and Japanese.
- **Reboot consideration:** In most cases, a reboot is not required after adding a language. However, if the input method does not function correctly after installation, restart the Image Builder before creating the image.

Special Considerations for Application Settings Persistence

When you create a stack in the WorkSpaces Applications console, in **Step 3: User Settings**, if you use the same settings group under **Application settings persistence** as another stack that uses different regional settings, only one set of regional settings is used for both stacks. For each user, the default regional settings for the stack that the user logs into first automatically override the default regional settings of any other stacks in the same application settings group. To avoid this problem, do not use the same application settings group for two different stacks that have different regional settings.

Special Considerations for Japanese Language Settings

This section describes key points to keep in mind when configuring Japanese language settings for your WorkSpaces Applications users.

AWS CLI

Changing the Windows system locale to Japanese requires that your image builder have AWS Command Line Interface (AWS CLI) version 1.16.30 or later installed. To update the version of AWS CLI on your image builder, follow the steps in Installing the AWS Command Line Interface.

Japanese Keyboards

If your image builder input method is set to Japanese when you create an image, WorkSpaces Applications automatically configures your image to use a Japanese keyboard. Any fleets that use the image are also automatically configured to use Japanese keyboards. However, if you want to use a Japanese keyboard within your image builder session, update the following registry settings

for the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\i8042prt\Parameters registry key:

Name	Type	Data
LayerDriver JPN	REG_SZ	kbd106.dll
OverrideKeyboardIdentifier	REG_SZ	PCAT_106KEY
OverrideKeyboardSubtype	DWORD	2
OverrideKeyboardType	DWORD	7

After changing these settings, restart your image builder. To do so, choose the Windows **Start** button, and choose **Windows PowerShell**. In PowerShell, use the **restart-computer** cmdlet.

Enable Your WorkSpaces Applications Users to Configure Their Regional Settings

> **Note**
>
> Enabling users to configure regional settings is currently not supported in Linux-based streaming sessions.

Users can configure their Amazon WorkSpaces Applications Windows streaming sessions to use settings that are specific to their location or language. In particular, users can configure the following settings:

- **Time zone** — Determines the system time used by Windows and any applications that rely on the operating system time. WorkSpaces Applications makes available the same options for this setting as the Windows Server version used in your fleet.
  - To sync the time zone for your streaming session to match the time zone set on your device, choose **Set my time zone automatically based on my device**. This is enabled by default for both single-session and multi-session fleets, and can be disabled.
  - For single-session fleets, you can choose a specific time zone for your streaming session instead of using automatic redirection. To set a custom time zone, disable the **Set my time**

zone automatically based on my device option in **Regional settings**, and choose a preferred time zone from the available list.

- **Locale** (also known as culture) — Determines the conventions used by Windows and any applications that query the Windows culture when formatting dates, numbers, or currencies or when sorting strings. For a list of locales that WorkSpaces Applications supports, see Supported Locales.
- **Input method** — Determines the keystroke combinations that can be used to input characters in another language.

If users change regional settings during their streaming sessions, the changes are applied to any future streaming sessions in the same AWS Region.

> **Note**
>
> For guidance that you can provide your users to help them get started with configuring their regional settings, see Configure Regional Settings.

Contents

- Supported Locales
- Enable Regional Settings for Your WorkSpaces Applications Users

Supported Locales

WorkSpaces Applications supports the following locales:

Locale	Language culture name
Chinese (Simplified, China)	zh-CN
Chinese (Simplified, Singapore)	zh-SG
Chinese (Traditional)	zh-TW
Dutch (The Netherlands)	nl-NL
English (Australia)	en-AU

Locale	Language culture name
English (Canada)	en-CA
English (United Kingdom)	en-GB
English (United States)	en-US
French (France)	fr-FR
German (Germany)	de-DE
Italian (Italy)	it-IT
Japanese (Japan)	ja-JP
Korean (Korea)	ko-KR
Portuguese (Brazil)	pt-BR
Spanish (Spain, International Sort)	es-ES
Thai (Thailand)	th-TH

Enable Regional Settings for Your WorkSpaces Applications Users

To enable users to configure regional settings for a given stack during their WorkSpaces Applications streaming sessions, your stack must be associated with a fleet based on an image that uses a version of the WorkSpaces Applications agent released on or after June 6, 2018. For more information, see WorkSpaces Applications Agent Release Notes. Additionally, your image must have Windows PowerShell 5.1 or later installed. Images created from WorkSpaces Applications base images published on or after June 12, 2018 meet both criteria. Images created from WorkSpaces Applications base images published before June 12, 2018 do not have Windows PowerShell 5.1 by default.

To update an existing image to include Windows PowerShell 5.1

1. Launch a new image builder using your existing image as the base image by doing the following:

a. In the left navigation pane in the WorkSpaces Applications console, choose **Images**.

b. Choose the **Image Builder** tab, **Launch Image Builder**, and then select your existing image.

c. If you are prompted to update the WorkSpaces Applications agent when you launch the image builder, select the check box, and then choose **Start**.

2. After your image builder is running, connect to it and sign in with an account that has local administrator permissions. To do so, do either of the following:

   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

   > **Note**
   >
   > If the image builder that you want to connect to is joined to an Active Directory domain and your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection. For information about smart card sign in, see Smart Cards.

3. From the image builder desktop, open Windows PowerShell. Choose the Windows **Start** button, and then choose **Windows PowerShell**.
4. At the PowerShell command prompt, type the command `$PSVersionTable` to determine the version of Windows PowerShell that is installed on your image builder. If your image builder does not include Windows PowerShell 5.1 or later, use the following steps to install it.
5. Open a web browser and follow the steps in Install and Configure WMF 5.1 in the Microsoft documentation, making sure that you download the Windows Management Framework (WMF) 5.1 package for Windows Server 2012 R2. WMF 5.1 includes Windows PowerShell 5.1.
6. At the end of the WMF 5.1 installation process, the installer prompts you to restart your computer. Choose **Restart Now** to restart the image builder.
7. Wait about 10 minutes before logging in to your image builder, even though WorkSpaces Applications prompts you to do so immediately. Otherwise, you might encounter an error.
8. After logging in to your image builder again, open Windows PowerShell and type the command `$PSVersionTable` to confirm that Windows PowerShell 5.1 is installed on your image builder.
9. Use the image builder to create a new image. This new image now includes the latest versions of the WorkSpaces Applications agent and Windows PowerShell.

10. Update your fleet to use the new image by doing the following:

    a. In the left navigation pane in the WorkSpaces Applications console, choose **Fleets**, and then choose the fleet associated with the stack for which you want to enable regional settings.

    b. On the **Fleet Details** tab, choose **Edit**.

    c. In **Image name,** choose the new image to use for the fleet.

For more information about using image builders to create images, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

Manage Application Entitlements

Amazon WorkSpaces Applications can dynamically build the application catalog to display the WorkSpaces Applications applications that users are entitled to access. Application entitlements can be assigned based on attributes using a third-party SAML 2.0 identity provider, or by using the WorkSpaces Applications Dynamic Application Framework. The following sections describe how to manage application entitlements.

> **Note**
>
> Attribute-based application entitlements using a third-party SAML 2.0 identity provider is recommended for most scenarios. If you would like to use an existing Dynamic App Provider that manages application package delivery in addition to entitlement, so that applications do not need to be installed in an WorkSpaces Applications image, Dynamic Application Framework is recommended. For more information, see Additional Resources for Learning About Dynamic App Providers and the Dynamic Application Framework.

Contents

- Attribute-Based Application Entitlements Using a Third-Party SAML 2.0 Identity Provider
- Application Entitlements from a Dynamic App Provider Using the Dynamic Application Framework

Attribute-Based Application Entitlements Using a Third-Party SAML 2.0 Identity Provider

Application entitlements control access to specific applications within your WorkSpaces Applications stacks. This works by using SAML 2.0 attribute assertions from a third-party SAML 2.0 identity provider. The assertion is matched to a value when a user identity federates to a WorkSpaces Applications SAML application. If the entitlement is true, and the attribute name and value match, access is entitled for the user identity to one or more applications within the stack.

Attribute-based application entitlements using a third-party SAML 2.0 identity provider do not apply in the following scenarios. In other words, the entitlement is ignored in cases such as the following:

- WorkSpaces Applications user pool authentication. For more information, see Amazon WorkSpaces Applications User Pools.
- WorkSpaces Applications streaming URL authentication. For more information, see Streaming URL.
- The desktop application when WorkSpaces Applications fleets are configured for **Desktop Stream view**. For more information, see Create an Amazon WorkSpaces Applications Fleet and Stack.
- Stacks using the Dynamic Application Framework. Dynamic Application Framework provides separate application entitlement features. For more information, see Application Entitlements from a Dynamic App Provider Using the Dynamic Application Framework.
- When users federate to the WorkSpaces Applications application catalog, application entitlements will only display the applications the user is entitled to. Applications are not restricted from running within the WorkSpaces Applications session. For example, in a fleet configured for Desktop Stream view, a user can launch an application directly from the desktop.

Create Application Entitlements

Before you create application entitlements, you must do the following:

- Create an WorkSpaces Applications fleet and stack with an image containing one or more applications (Always-On or On-Demand fleet) or assigned applications (Elastic fleet) that will meet your needs. For more information, see Create an Amazon WorkSpaces Applications Fleet and Stack.
- Provide user access to the stack using a third-party SAML 2.0 identity provider. For more information, see Amazon WorkSpaces Applications Integration with SAML 2.0. If you are using an existing SAML 2.0 identity provider that you setup previously, see Step 2: Create a SAML 2.0 Federation IAM Role for the steps to add the sts:TagSession permission to your IAM role trust policy. For more information, see Passing session tags in AWS STS. This permission is required to use application entitlements.

To create an application entitlement

1. Open the WorkSpaces Applications console.
2. In the left navigation pane, choose **Stacks**, and select the stack for which to manage application entitlements.

3. In the **Application Entitlements** dialog box, choose **Create**.
4. Enter a **Name** and **Description** for your entitlement.
5. Define the attribute name and value for your entitlement.

   When mapping attributes, specify the attribute in the format https://aws.amazon.com/SAML/Attributes/PrincipalTag:{TagKey}, where {TagKey} is one of the following attributes:

   - roles
   - department
   - organization
   - groups
   - title
   - costCenter
   - userType

   The attributes that you defined are used to entitle applications in your stack to a user when they federate into an WorkSpaces Applications session. Entitlement works by matching the attribute name to a key value name in the SAML assertion created during federation. For more information, see SAML PrincipalTag Attribute.

   > **Note**
   >
   > One or more values can be included in any supported attribute, separated by a colon (:).
   > For example, groups information can be passed in a SAML attribute name https://aws.amazon.com/SAML/Attributes/PrincipalTag:groups with value “group1:group2:group3” and your entitlement can allow applications based on a single group value, i.e. “group1”. For more information, see SAML PrincipalTag Attribute.

6. Configure application settings in your stack to entitle all applications, or select applications. Choosing **All applications (*)** applies all applications available on the stack, including applications that are added in the future. Choosing **Select applications** will filter on specific application names.

7. Review your settings and create your entitlement. You can repeat the process and create additional entitlements. Entitlement to applications in a stack will be a union of all entitlements that match the user based on attribute names and values.
8. In your SAML 2.0 identity provider, configure your WorkSpaces Applications SAML application attribute mappings to send the attribute and value defined in your entitlement. When users federate to the WorkSpaces Applications application catalog, application entitlements will only display the applications the user is entitled to.

SAML 2.0 Multi-Stack Application Catalog

With attribute-based application entitlements using a third-party SAML 2.0 identity provider, you can enable access to multiple stacks from a single relay state URL. Remove the stack and app (if present) parameters from the relay state URL, as follows:

```
https://relay-state-region-endpoint?accountId=aws-account-id-without-hyphens
```

When users federate to the WorkSpaces Applications application catalog, they will be presented with all of the stacks where application entitlements have matched one or more applications to the user for the account ID and relay state endpoint associated with the Region in which your stacks are located. When a user selects a catalog, application entitlements will only display the applications the user is entitled to. For more information, see Step 6: Configure the Relay State of Your Federation.

> **Note**
>
> To use SAML 2.0 Multi-Stack Application Catalogs, you need to configure the inline policy for your SAML 2.0 Federation IAM Role. For more information, see Step 3: Embed an Inline Policy for the IAM Role.

Application Entitlements from a Dynamic App Provider Using the Dynamic Application Framework

> **Note**
>
> Managing application entitlement with the Dynamic Application Framework is currently not supported for Linux-based stacks.

Amazon WorkSpaces Applications supports dynamically building the application catalog that displays for your users when they stream from an WorkSpaces Applications stack. You can use the API operations provided by WorkSpaces Applications to develop a dynamic app provider that modifies, in real time, the applications that users can access on the streaming instance. Alternatively, you can implement a third-party dynamic app provider that uses these API operations.

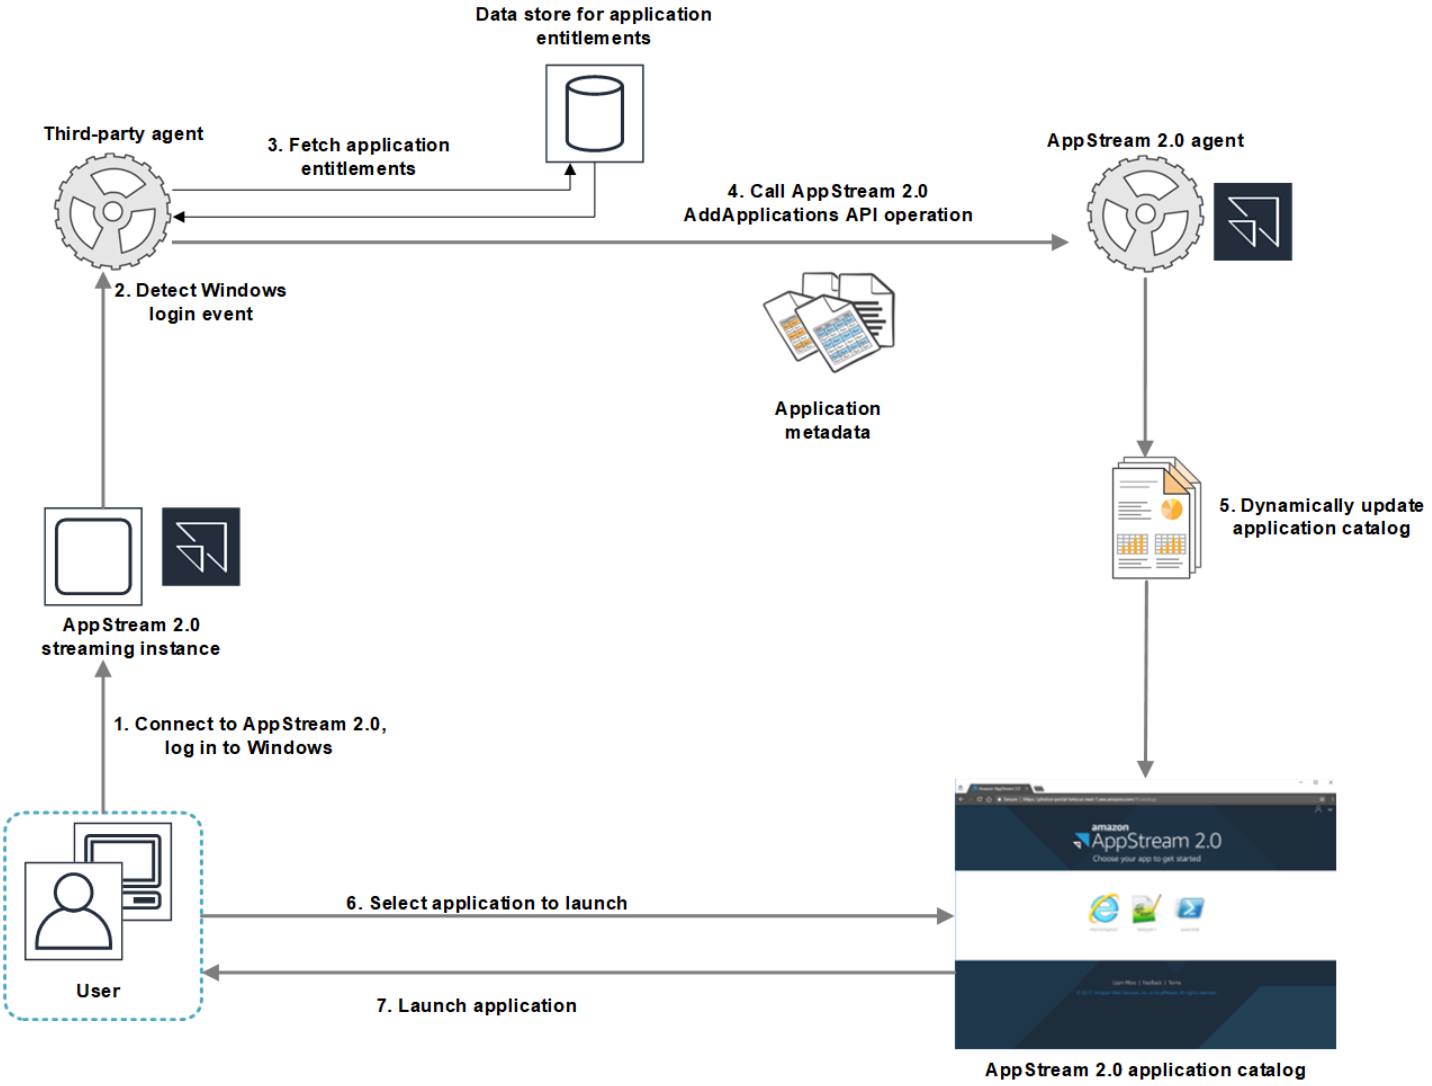
> **Note**
>
> This feature requires an WorkSpaces Applications Always-On or On-Demand fleet that is joined to a Microsoft Active Directory domain. For more information, see Using Active Directory with WorkSpaces Applications. This feature is not available on multi-session fleets.

Contents

- Example API Operations Work Flow for the Dynamic Application Framework
- Use the WorkSpaces Applications Dynamic Application Framework to Build a Dynamic App Provider
- Enable Dynamic App Providers
- Test Dynamic App Providers (Optional)
- Additional Resources for Learning About Dynamic App Providers and the Dynamic Application Framework

Example API Operations Work Flow for the Dynamic Application Framework

The following diagram is an example of the API operations flow between WorkSpaces Applications and a third-party application provider.



1. The user connects to WorkSpaces Applications. A fleet streaming instance is assigned to the user and Windows login occurs.
2. Your service or agent detects the Windows logon event and determines the user who is logging in to Windows.
3. The service or agent fetches the application entitlements for the user. In the example diagram, the application entitlements are stored in a database. This information can be stored and retrieved in different ways. For example, application entitlements may be fetched from server

software, or group names in Active Directory may be parsed to locate the application identifiers (IDs).

4. Your dynamic app provider calls the WorkSpaces Applications agent `AddApplications` API operation with the application metadata for the applications that the user should have.
5. The WorkSpaces Applications agent dynamically updates the application catalog with the modified application list.
6. The user selects an application to launch.
7. The application is launched by using the application metadata specified by your service or agent.

From the user's perspective, the process happens transparently. The user connects to WorkSpaces Applications and logs in to the fleet instance. After login, the list of applications specified in the image and provided by your dynamic app provider displays for the user.

Use the WorkSpaces Applications Dynamic Application Framework to Build a Dynamic App Provider

The WorkSpaces Applications dynamic application framework provides API operations within an WorkSpaces Applications streaming instance that you can use to build a dynamic app provider. Dynamic app providers can use the API operations provided to modify the catalog of applications that your users can access in real time. The applications managed by the dynamic app providers can be within the image, or they can be off-instance, such as from a Windows file share or an application virtualization technology.

> **Note**
>
> This feature requires an WorkSpaces Applications Always-On or On-Demand fleet that is joined to a Microsoft Active Directory domain. For more information, see Using Active Directory with WorkSpaces Applications.

Contents

- About the Dynamic Application Framework
- Dynamic Application Framework Thrift Definitions and Named Pipe Name
- API Actions for Managing App Entitlement for WorkSpaces Applications

About the Dynamic Application Framework

The dynamic application framework uses the Apache Thrift software framework for inter-process messaging. It is exposed through Named Pipes in Windows. Using the Thrift framework allows you to build your dynamic app provider in your software development language of choice. The dynamic application framework consists of three API operations: `AddApplications`, `RemoveApplications`, and `ClearApplications`.

Dynamic Application Framework Thrift Definitions and Named Pipe Name

Thrift enables you to use simple definition files provided by WorkSpaces Applications to compile RPC clients. The RPC clients let you communicate with the WorkSpaces Applications agent software running on a streaming instance. For information about how to compile the RPC client for your language, see the Apache Thrift documentation. After you compile the Thrift libraries for the language of your choice, build a Thrift client by using the Named Pipe transport. Use D56C0258-2173-48D5-B0E6-1EC85AC67893 as the pipe name.

AppStreamServer.thrift

```
namespace netstd AppStream.ApplicationCatalogService.Model

const string ServiceEndpoint = "D56C0258-2173-48D5-B0E6-1EC85AC67893";


struct AddApplicationsRequest
{
    1: required string userSid;
    2: required list<Application> applications;
}

struct AddApplicationsResponse
{
}

struct RemoveApplicationsRequest
{
    1: required string userSid;
    2: required list<string> applicationIds;
}

struct RemoveApplicationsResponse
{
```

```
}

struct ClearApplicationsRequest
{
    1: required string userSid;
}

struct ClearApplicationsResponse
{
}

struct Application
{
    1: required string id;
    2: required string displayName;
    3: required string launchPath;
    4: required string iconData;
    5: string launchParams;
    6: string workingDirectory;
}

exception AppStreamClientException
{
    1: string errorMessage,
    2: ErrorCode errorCode
}

exception AppStreamServerException
{
    1: string errorMessage,
    2: ErrorCode errorCode
}

enum ErrorCode
{
}

service ApplicationCatalogService
{
    AddApplicationsResponse AddApplications(1:AddApplicationsRequest request)
    throws (1: AppStreamClientException ce, 2: AppStreamServerException se),

    RemoveApplicationsResponse RemoveApplications(1:RemoveApplicationsRequest request)
    throws (1: AppStreamClientException ce, 2: AppStreamServerException se),
```

```
    ClearApplicationsResponse ClearApplications(1:ClearApplicationsRequest request)
    throws (1: AppStreamClientException ce, 2: AppStreamServerException se),
}
```

API Actions for Managing App Entitlement for WorkSpaces Applications

You can use the following API operations to manage application entitlement for WorkSpaces Applications.

`AddApplicationsRequest` operation

Adds applications to the application catalog for WorkSpaces Applications users. The application catalog displayed by WorkSpaces Applications includes the applications that you add by using this API operation and the applications that you add in the image. After you add applications by using one or both of these methods, your users can launch the applications.

Request syntax

```
string userSid;

list<Application> applications;
```

Request parameters

`userSid`

The SID of the user who the request applies to.

Type: String

Required: Yes

Length constraints: Minimum length of 1, maximum length of 208 characters.

`applications`

The list of applications that the request applies to.

Type: String

Required: Yes

`Application` object

Describes the application metadata required to display and launch the application. The application identifier must be unique and not in conflict with other applications specified through the API operation or the image.

`id`

The identifier of the application being specified. This value, which corresponds to the `application_name` value in an WorkSpaces Applications applications report, is provided when a user launches the application. When you enable usage reports, for each day that users launch at least one application during their streaming sessions, WorkSpaces Applications exports an applications report to your Amazon S3 bucket. For more information about applications reports, see Applications Report Fields.

Type: String

Required: Yes

Length constraints: Minimum length of 1, maximum length of 512 characters.

`displayName`

The display name of the application being specified. This name is displayed to the user in the application catalog.

Type: String

Required: Yes

Length constraints: Minimum length of 1, maximum length of 512 characters.

`launchPath`

The Windows file system path to the executable of the application to be launched.

Type: String

Required: Yes

Length constraints: Minimum length of 1, maximum length of 32,767 characters.

`iconData`

The base-64 encoded image to display in the application catalog. The image must be in one of the following formats: .png, .jpeg, or .jpg.

Type: String

Required: Yes

Length constraints: Minimum length of 1, maximum length of 1,000,000 characters.

launchParams

The parameters used to launch the application.

Type: String

Required: No

Length constraints: Maximum length of 32,000 characters.

workingDirectory

The Windows file system path to the working directory the application should be launched in.

Type: String

Required: No

Length constraints: Maximum length of 32,767 characters.

`RemoveApplicationsRequest` operation

Removes applications that were added by using the `AddApplicationsRequest` operation. The applications are removed from the application catalog for the user. After applications are removed, they can't be launched. If an application is still running, WorkSpaces Applications does not close it. Applications that are specified directly in the WorkSpaces Applications image can't be removed.

Request syntax

```
string userSid;

list<Application> applications;
```

Request parameters

userSid

The SID of the user the request applies to.

Type: String

Required: Yes

Length constraints: Minimum length of 1, maximum length of 208 characters.

applications

The list of applications that the request applies to.

Type: String

Required: Yes

`ClearApplicationsRequest` operation

Removes all applications that were added to the application catalog by using the `AddApplicationsRequest` operation. After applications are removed, they can't be launched. If the applications are running when the `ClearApplicationsRequest` operation is used, WorkSpaces Applications does not close them. Applications that are specified directly in the WorkSpaces Applications image can't be removed.

Request syntax

```
string userSid;
```

Request parameters

userSid

The SID of the user the request applies to.

Type: String

Required: Yes

Length constraints: Minimum length of 1, maximum length of 208 characters.

Enable Dynamic App Providers

Dynamic app providers must first be enabled within an WorkSpaces Applications image. After you enable these providers, they can manage applications for users on the streaming instance.

To enable this capability, you must add your dynamic app provider details to a configuration file on the image builder. The image builder must be joined to a Microsoft Active Directory domain.

Perform the following steps on an image builder, then you can test your dynamic apps to verify that they function as expected. Finally, finish creating your image.

> **Note**
>
> Third-party dynamic app providers may modify the configuration file during install. For installation instructions, see the documentation for the applicable provider.

To enable dynamic app providers

1. Connect to the image builder that you want to use and sign in with a domain account that has local administrator permissions on the image builder. To do so, do either of the following:

   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

   > **Note**
   >
   > If your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection. For information about smart card sign in, see Smart Cards.

2. Navigate to C:\ProgramData\Amazon\AppStream\AppCatalogHelper\DynamicAppCatalog\, and open the **Agents.json** configuration file.
3. In the **Agents.json** file, add the following entries:

   "DisplayName": "<Uninstall hive display name value>",

   "Path": "<C:\path\to\client\application>"

   DisplayName must match the **DisplayName** registry value for the **HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Uninstall** key created for your application.
4. Install your dynamic app provider.
5. On the image builder desktop, open Image Assistant.
6. Optionally, install any other applications that you want to include in the image.
7. In Image Assistant, on the **1. Add Apps** page, select the **Enable dynamic app providers** check box.

8. On the same page, if you installed other applications as described in step 8, choose **+Add App**, and specify the applications to add.

> **Note**
>
> When you use a dynamic app provider, you don't need to specify any applications in the image. If you specify applications in the image, they can't be removed by dynamic app providers.

9. Proceed to the steps in the next section to test your dynamic app provider.

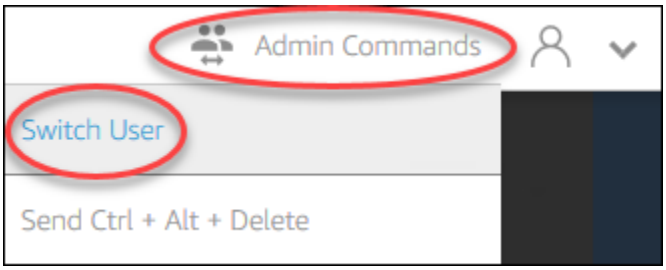
Test Dynamic App Providers (Optional)

After you enable your dynamic app provider on an image builder, you can test the provider to verify that it functions as expected. To do so, perform the following steps before you finish creating the image.

To test dynamic app providers

1. Do one of the following:

   - If you are already connected to the image builder on which you enabled dynamic app providers and you are logged on as **Administrator,** you must switch to an account that does not have local administrator permissions on the image builder. To do so, in the upper right corner of the image builder session toolbar, choose **Admin Commands**, **Switch User**.



   - If you are not already connected to the image builder, connect by either using the WorkSpaces Applications console (for web connections only) or creating a streaming URL (for web or WorkSpaces Applications client connections).

> **Note**
>
> When you are prompted to sign in, choose **Directory User**, and sign in with a domain account that does not have local administrator permissions on the image builder.

2. On the image builder desktop, open Image Assistant, if it is not already open.
3. On the **Test Apps** page, if you specified any applications in the image that are not from the dynamic app provider, they display first in the list. It may take a few moments for applications from dynamic app providers to appear in the list.
4. Choose an application from the list and open it to verify that it functions as expected.
5. After you finish testing, in the lower right corner of the **Test Apps** page, choose **Switch user**.
6. Choose **Administrator**, and log back into the image builder.
7. Follow the necessary steps in Image Assistant to finish creating your image. For information about how to create an image, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

   WorkSpaces Applications automatically optimizes the agents that are specified in the **Agents.json** configuration file.

Additional Resources for Learning About Dynamic App Providers and the Dynamic Application Framework

The following links provide information to help you learn more about dynamic app providers and the dynamic application framework.

Solution	Description
Liquidware FlexApp	FlexApp — Provides an overview of Liquidware FlexApp. FlexApp is a third-party provider that uses the WorkSpace s Applications dynamic application framework to manage application entitlements and delivery in real time. FlexApp layering delivers applications to any Windows desktop environment, independent of the Windows operating system version or delivery platform.

Solution	Description
App-V	Bring your App-V packages to WorkSpaces Applications with the dynamic application framework — Describes how to integrate App-V with WorkSpaces Applications by using the dynamic application framework.
WorkSpaces Applications	Use the WorkSpaces Applications Dynamic Application Framework to Build a Dynamic App Provider — Describes how to use the WorkSpaces Applications dynamic application framework to develop your own dynamic app provider.

Provide Your Users with Access to WorkSpaces Applications

Users can access WorkSpaces Applications streaming sessions by using either a web browser or the WorkSpaces Applications client on a supported device.

Depending on your organizational requirements, you can enable user access to WorkSpaces Applications streaming sessions by: Setting up identity federation using SAML 2.0, using an WorkSpaces Applications user pool, or creating a streaming URL. Following are recommendations for choosing a connection method.

- SAML 2.0: Use this connection method when you have an identity provider that manages your users and supports SAML 2.0 federation.

  > **Note**
  >
  > This connection method is required when your WorkSpaces Applications fleet is joined to a Microsoft Active Directory domain.

- WorkSpaces Applications user pools: Use this connection method when:
  - You want to set up a Proof-of-Concept (POC) quickly before you configure your SAML 2.0-compliant identity provider.
  - You don't have a SAML 2.0-compliant identity provider.
  - You want to manage users directly within the WorkSpaces Applications console.
- Streaming URL: Use this connection method when you want to programmatically provide access to WorkSpaces Applications by using temporary URLs. We recommend this connection method when you want to use your existing identity provider to provide programmatic access to WorkSpaces Applications.

Supported Features

The following table compares the features that are supported by the different access types.

Feature	Browser-Based Access	Client-Based Access for Windows	Client-Based Access for macOS	Notes
Enterprise Deployment Tool	✗	✓	✗	For more information, see the section called "Tutorial: Install the Client and Customize the Client Experience".
HIPAA/PCI compliance	✓	✓	✓	For more information, see Compliance.
Active Directory authentication	✓	✓	✓	For more information, see *Using Active Directory* .
MFA (multi-factor authentication)	✓	✓	✓	For WorkSpaces Applications, MFA is supported

Feature	Browser-Based Access	Client-Based Access for Windows	Client-Based Access for macOS	Notes
				via SAML 2.0.
Smart card (CAC and PIV readers)	✗	✓	✗	For more information, see the section called "Smart Card Redirection".
Certificates for access control (OS-based)	✓	✓	✓	For WorkSpaces Applications, certificate authentication is supported via SAML 2.0.

Feature	Browser-Based Access	Client-Based Access for Windows	Client-Based Access for macOS	Notes
Certificate-Based Authentication	✓	✓	✓	For more information, see the section called "Certificate-Based Authentication".
Client customization	Available with exceptions	Available with exceptions	Available with exceptions	WorkSpaces Applications supports web-based branding and custom URLs. For more information, see *Add Custom Branding*.
Desktop view connection mode	✓	✓	✓	
Classic application mode	✓	✓	✓	
Native application mode	✗	✓	✗	

Feature	Browser-Based Access	Client-Based Access for Windows	Client-Based Access for macOS	Notes
USB redirection	✗	✓	✗	Supported on WorkSpaces Applications client accessing Windows-based fleets. For more information, see the section called "USB Redirection".
Audio input (for web conferences and calls)	✓	✓	✓	Not supported on Linux. WorkSpaces Applications supports USB microphones.
Video input (conferencing applications)	✓	✓	✓	

Feature	Browser-Based Access	Client-Based Access for Windows	Client-Based Access for macOS	Notes
Storage redirection	✗	✓	✗	For more information, see the section called "Enable File System Redirection".
USB/local printer redirection	Available with exceptions	✓	✓	WorkSpaces Applications indirect printing on browser. Full redirection not supported on Linux-based stacks.
Clipboard redirection	✓	✓	✓	
Drawing tablets	✓	✓	✗	

Feature	Browser-Based Access	Client-Based Access for Windows	Client-Based Access for macOS	Notes
YubiKey support	✗	Available with exceptions	✗	Supported on WorkSpaces Applications client. For more information, see the section called "Qualify USB Devices for Use with Streaming Applications".
Monitor support	Web access dual-monitor support. For more information, see the section called "Dual-Monitor Support".	✓	✓	For more information, see the section called "Multiple Monitors".

Feature	Browser-Based Access	Client-Based Access for Windows	Client-Based Access for macOS	Notes
User Pool	✓	✓	✓	For more information, see the section called "User Pools".
Connect to an App Block Builder	✓	✓	✗	For more information, see the section called "App Block Builder".
Access Multi-Stack Application	✓	✓	✓	For more information, see the section called "Attribute-Based Application Entitlements".

The next topics provide information about how to configure user access to WorkSpaces Applications for application streaming.

For guidance that you can provide your users to help them get started with application streaming, see Guidance for WorkSpaces Applications Users.

Provide Access Through a Web Browser

Your users can start an WorkSpaces Applications streaming session by using a web browser or the WorkSpaces Applications client application for a supported device. The following topics provide information to help you provide user access through a web browser.

Contents

- System Requirements and Feature Support (Web Browser)
- Configure a Connection Method for Your WorkSpaces Applications Users (Web Browser)

For information about how to provide user access to WorkSpaces Applications through the WorkSpaces Applications client, see Provide Access Through the WorkSpaces Applications Client.

System Requirements and Feature Support (Web Browser)

This topic provides information to help you understand the requirements for providing user access to WorkSpaces Applications through a web browser. It also provides information about supported features.

Topics

- System Requirements and Considerations
- Feature and Device Support

System Requirements and Considerations

Users can access WorkSpaces Applications through an HTML5-capable web browser on a desktop computer such as a Windows, Mac, Chromebook, or Linux computer. HTML5-capable web browsers that can be used include the following:

- Google Chrome
- Mozilla Firefox
- Safari

- Microsoft Edge

No browser extensions or plugins are required to use WorkSpaces Applications in a web browser.

Users can also access WorkSpaces Applications fleet streaming sessions on the following browsers and devices:

- Chrome or Safari on an iPad (iOS 11 or later)
- Android (Android 8 or later)
- Microsoft Surface Pro (Windows 10) tablet

WorkSpaces Applications is not supported on devices that have screen resolutions smaller than 1024x768 pixels.

Feature and Device Support

WorkSpaces Applications provides the following feature and peripheral device support for users who access WorkSpaces Applications through a web browser.

Topics

- Dual-Monitor Support
- Touchscreen Device Support
- Drawing Tablet Support
- Relative Mouse Offset

Dual-Monitor Support

WorkSpaces Applications supports the use of multiple monitors during streaming sessions, including monitors that have different resolutions. To help ensure an optimal streaming experience, we recommend that users who have monitors with different resolutions set the display scale for their monitors to 100 percent.

Dual monitors are supported for streaming sessions that are started on the following web browsers:

- Google Chrome

- Mozilla Firefox
- Safari

For browser-based streaming sessions on dual monitors, a maximum display resolution of 2560x1600 pixels is supported per monitor. If your users require more than two monitors, or a display resolution that is greater than 2560x1600 pixels per monitor, the WorkSpaces Applications client is available.

> **Note**
>
> Dual monitors are not supported on mobile devices or for embedded WorkSpaces Applications streaming sessions.

In addition to user connections for streaming sessions, WorkSpaces Applications also supports the use of dual monitors for administrative connections to image builders.

Touchscreen Device Support

WorkSpaces Applications supports gestures on touch-enabled iPads, Android tablets, and Windows devices. All touch events are passed through to the streaming session and handled according to Windows conventions. Examples of supported touch gestures include long-tap to right-click, swipe to scroll, pinch to zoom, and two-finger rotation for supporting applications.

> **Note**
>
> To enable support for gestures on touch-enabled devices, your WorkSpaces Applications image must use a version of the WorkSpaces Applications agent released on or after March 7, 2019. For more information, see WorkSpaces Applications Agent Release Notes.

For guidance that you can provide your users to help them get started with touch-enabled devices during their WorkSpaces Applications streaming sessions, see Touchscreen Devices.

Drawing Tablet Support

Drawing tablets, also known as pen tablets, are computer input devices that let users draw with a stylus (pen). With WorkSpaces Applications, your users can connect a drawing tablet, such

as a Wacom drawing tablet, to their local computer and use the tablet with their streaming applications.

Following are requirements and considerations for enabling your users to use drawing tablets with their streaming applications.

- To enable your users to use this feature, you must configure your WorkSpaces Applications fleet to use an image that runs Windows Server 2019.
- To use this feature, users must access WorkSpaces Applications through the Google Chrome or Mozilla Firefox browsers only, or the WorkSpaces Applications client.
- Streaming applications must support Windows Ink technology. For more information, see Pen interactions and Windows Ink in Windows apps.
- Some applications, such GIMP, must detect drawing tablets on the streaming instance to support pressure sensitivity. If this is the case, your users must use the WorkSpaces Applications client to access WorkSpaces Applications and stream these applications. In addition, you must qualify your users' drawing tablets, and users must share their drawing tablets with WorkSpaces Applications every time they start a new streaming session. For step-by-step guidance, see Qualify USB Devices for Use with Streaming Applications.
- This feature is not supported on Chromebooks.

To get started with using drawing tablets during application streaming sessions, users connect their drawing tablet to their local computer with USB and use a supported web browser or the WorkSpaces Applications client, if it is installed, to start a streaming session. No USB redirection is required to use this feature.

Relative Mouse Offset

By default, during users' streaming sessions, WorkSpaces Applications transmits information about mouse movements to the streaming instance by using absolute coordinates and rendering the mouse movements locally. For graphics-intensive applications, such as computer-aided design (CAD)/computer-aided manufacturing (CAM) software or video games, mouse performance improves when relative mouse mode is enabled. Relative mouse mode uses relative coordinates, which represent how far the mouse moved since the last frame, rather than the absolute x-y coordinate values within a window or screen. When relative mouse mode is enabled, WorkSpaces Applications renders the mouse movements remotely.

Users can enable this feature during their WorkSpaces Applications streaming sessions by doing either of the following:

- Windows: Pressing Ctrl+Shift+F8
- Mac: Pressing Control+Fn+Shift+F8

Configure a Connection Method for Your WorkSpaces Applications Users (Web Browser)

Depending on your organizational requirements, you can provide users with access to WorkSpaces Applications through a web browser by doing one of the following: Setting up identity federation using SAML 2.0, using an WorkSpaces Applications user pool, or creating a streaming URL.

Contents

- SAML 2.0
- WorkSpaces Applications User Pool
- Streaming URL
- Next Steps

SAML 2.0

Users enter the URL that you provide for them to access your internal organizational portal. After they enter their organizational credentials, they're redirected to WorkSpaces Applications.

For more information, see Setting Up SAML.

> **Note**
>
> If your organization requires a smart card for Windows sign in to Active Directory-joined streaming instances and in-session authentication for streaming applications, your users must install and use the WorkSpaces Applications client. For more information, see Smart Cards.

WorkSpaces Applications User Pool

When you create a new user in the WorkSpaces Applications user pool, or assign a user pool user to an WorkSpaces Applications stack, WorkSpaces Applications sends email to users on your behalf.

Users enter the URL that was provided to them in the welcome email, enter their credentials, and then choose **Connect**.

For more information, see Amazon WorkSpaces Applications User Pools.

Streaming URL

To create a streaming URL, use one of the following methods:

- WorkSpaces Applications console
- The CreateStreamingURL API action
- The create-streaming-url AWS CLI command

To create a streaming URL by using the WorkSpaces Applications console, complete the steps in the following procedure.

To create a streaming URL by using the WorkSpaces Applications console

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the navigation pane, choose **Fleets**.
3. In the list of fleets, choose the fleet that is associated with the stack for which you want to create a streaming URL. Verify that the status of the fleet is **Running**.
4. In the navigation pane, choose **Stacks**. Choose the stack, and then choose **Actions**, **Create streaming URL**.
5. In **User id**, enter the user ID.
6. For **URL Expiration**, choose an expiration time, which determines how long the generated URL is valid. This URL is valid for a maximum of seven days.
7. Choose **Get URL**.
8. Copy the URL, save it to an accessible location, and then provide it to your users.

Next Steps

After you configure a web browser connection method, you can provide your users with the following step-by-step guidance to help them connect to WorkSpaces Applications and start a streaming session: Connect to WorkSpaces Applications.

Provide Access Through the WorkSpaces Applications Client

Your users can start WorkSpaces Applications streaming sessions by using the WorkSpaces Applications client application for a supported device or by using a web browser.

The WorkSpaces Applications client is a native application that is designed for users who require the following functionality during their WorkSpaces Applications streaming sessions:

- Require support for more than two monitors or 4K resolution.
- Use their USB devices with applications streamed through WorkSpaces Applications.
- Use their local webcam for video conferencing within their streaming sessions and the browser in use doesn't support video or audio input.
- Use keyboard shortcuts during their streaming sessions.
- Require seamless access to local drives and folders during their streaming sessions.
- Require the ability to redirect print jobs from their streaming application to a printer that is connected to their local computer.
- Prefer to interact with remote streaming applications in much the same way as they interact with locally installed applications.

The following topics provide information to help you provide user access through the WorkSpaces Applications client. For information about how to provide user access to WorkSpaces Applications through a web browser, see Provide Access Through a Web Browser.

Contents

- System Requirements and Feature Support (WorkSpaces Applications Client)
- Install and Configure the WorkSpaces Applications Client

System Requirements and Feature Support (WorkSpaces Applications Client)

This topic provides information to help you understand the requirements for the WorkSpaces Applications client and supported features.

Topics

- System Requirements and Considerations

- Feature and Device Support

System Requirements and Considerations

The WorkSpaces Applications client requires the following:

- Follow the principle of least privilege when launching the WorkSpaces Applications client. The client should only run with the level of privilege required to complete a task.
- Client requirements
  - Windows client
    - Operating system — Windows 10 (32-bit or 64-bit), Windows 11 (64-bit)
    - Microsoft Visual C++ 2019 version 14.20.xx Redistributable or later. For more information, see the Latest Microsoft Visual C++ Redistributable Version in the Microsoft Support documentation.
    - RAM — 2 GB minimum
    - Hard drive space — 200 MB minimum
  - macOS client
    - Operating system — macOS 13 (Ventura), macOS 14 (Sonoma), macOS 15 (Sequoia)
    - Hard drive space — 200 MB minimum
- Local administrator rights — Used if you want to install the WorkSpaces Applications USB driver for USB driver support.

  > **Note**
  >
  > Local administrator rights are not supported for the macOS client.

- An WorkSpaces Applications image that uses the latest WorkSpaces Applications agent or agent versions published on or after November 14, 2018. For information about WorkSpaces Applications agent versions, see WorkSpaces Applications Agent Release Notes.
- The client supports UDP as well as the default TCP-based streaming over NICE DCV. For more information about NICE DCV and UDP, see Enabling the QUIC UDP transport protocol. If you want to enable UDP streaming for the client, make sure you meet the following requirements. If you don't meet the following requirements, the client will default back to TCP-based streaming.
  - Your Stack has been configured to prefer UDP in the **Streaming Setting Experience** section. For more information, see Create an Amazon WorkSpaces Applications Fleet and Stack.

- Your network allows UDP traffic on port 8433 for the AWS IP Ranges. For more information, see AWS IP address ranges.
- You are using the latest base image when creating your fleet. For more information, see WorkSpaces Applications Base Image and Managed Image Update Release Notes.
- Your end users are using the latest client. For more information, see Supported clients.

> **Note**
>
> We recommend an internet connection for WorkSpaces Applications client installation. In some cases, the client can't be installed on a computer that is not connected to the internet, or USB devices might not work with applications streamed from WorkSpaces Applications. For more information, see Troubleshooting WorkSpaces Applications User Issues.

Feature and Device Support

The WorkSpaces Applications client supports the following features and devices.

Topics

- Native Application Mode
- Automatic and On-Demand Diagnostic Log Uploads
- Peripheral Devices
- Host-to-client URL redirection

Native Application Mode

> **Note**
>
> Native application mode is not available when streaming from Linux instances, streaming in Desktop mode, or when using the WorkSpaces Applications macOS client application.

Native application mode provides a familiar experience for your users during their WorkSpaces Applications streaming sessions. When your users connect to WorkSpaces Applications in this mode, they can work with their remote streaming applications in much the same way that they

work with applications that are installed on their local computer. Each streaming application in native application mode opens in its own window, and application icons appear on the taskbar on your users' local PC.

If you want your users to connect to WorkSpaces Applications in classic mode only, you can configure the `NativeAppModeDisabled` registry value to disable native application mode. For more information, see Choose Whether to Disable Native Application Mode.

For more information about native application mode and classic mode, and for guidance that you can provide to your users, see WorkSpaces Applications Client Connection Modes.

Feature requirements

To enable this feature for your users, you must use an image that uses a version of the WorkSpaces Applications agent released on or after February 19, 2020. In addition, version 1.1.129 or later of the WorkSpaces Applications client must be installed on your users' PCs. For more information about client versions, see WorkSpaces Applications Windows Client Release Notes.

If WorkSpaces Applications client version 1.1.129 or later is installed on your users' computer, but you are not using an image that uses an agent version released on or after February 19, 2020, the client falls back to classic mode even if native application mode is selected.

Testing requirements

Applications must be tested thoroughly in native application mode before production deployment. Testing in classic mode is not sufficient, because applications might have compatibility issues with native application mode. Key testing areas include the following:

- Core application functionality
- Network performance
- Local device interactions
- File handling operations
- Print capabilities
- Multi-monitor support
- Audio/video features

We recommend starting with a pilot user group and documenting any application-specific limitations before full deployment. Application behavior and performance might vary between

streaming modes. Comprehensive testing helps ensure optimal user experience and identifies any potential limitations before production rollout.

Known Issues

When users try to dock or undock tabs in one browser window into separate windows during a streaming session in native application mode, their remote streaming browser doesn't work the same way as a local browser. To perform this task during a streaming session in native application mode, users must press the Alt key until their browser tabs are docked into separate browser windows.

Automatic and On-Demand Diagnostic Log Uploads

To help with troubleshooting issues that might occur when your users are using the WorkSpaces Applications client, you can enable automatic or on-demand diagnostic log uploads, or let your users do so themselves.

> **Note**
>
> Diagnostic logs do not contain sensitive information. You can disable automatic and on-demand diagnostic log uploads on user PCs that you manage, or allow your users to disable these features themselves.

Automatic diagnostic log uploads

When you install the client on PCs that you manage, you can configure the WorkSpaces Applications client to upload diagnostic logs automatically. That way, when a client issue occurs, the logs are sent to WorkSpaces Applications (AWS) without user interaction. For more information, see Configure Additional WorkSpaces Applications Client Settings for Your Users.

Or, you can let your users choose whether to enable automatic diagnostic log uploads when they install the WorkSpaces Applications client, or after client installation. For guidance that you can provide your users to help them perform this task, see Setup for Windows.

On-demand diagnostic log uploads

If you require more control over logging, you can disable automatic logging and enable on-demand diagnostic log uploads. If you let your users upload diagnostic logs on demand, they can also choose whether to send minidumps (error reports) to WorkSpaces Applications (AWS) if an exception occurs or the client stops responding.

For guidance that you can provide your users to help them perform these tasks, see Logging.

Peripheral Devices

The WorkSpaces Applications client provides the following support for peripheral devices such as monitors, webcams, mice, keyboards, and drawing tablets.

> **Note**
>
> With certain exceptions, USB redirection is required for the WorkSpaces Applications client to support USB devices. And in most cases, when USB redirection is required for a device, you must qualify the device before it can be used with WorkSpaces Applications streaming sessions. For more information, see USB Redirection.

Topics

- Multiple Monitors
- Real-Time Audio-Video
- USB Devices
- Drawing Tablets
- Keyboard Shortcuts
- Relative Mouse Offset

Multiple Monitors

WorkSpaces Applications supports the use of multiple monitors during streaming sessions, including monitors that have different resolutions. To help ensure an optimal streaming experience, we recommend that users who have monitors with different resolutions set the display scale for their monitors to 100 percent.

> **Note**
>
> For WorkSpaces Applications streaming sessions that use native application mode, monitors with up to 4K resolution are supported. If higher-resolution monitors are used for streaming sessions, the WorkSpaces Applications client falls back to classic mode. In this scenario, the WorkSpaces Applications classic mode streaming view occupies 4K of the screen, and the remaining portion of the screen is black.

Multiple Monitors (up to 4K Resolution)

Default behavior: By default, all instance types support up to 4 monitors with a maximum display resolution of 4096x2160 pixels per monitor, except for instances with small or medium sizes.

For all instance types with small and medium sizes: While we do not recommend using 4K resolution on small and medium instances due to performance limitations, you can enable it by making the following changes in your images:

- For Windows: Configure the following registry setting on image builder.
  - Using PowerShell:

    ```
    Set-ItemProperty "HKLM:\SOFTWARE\Amazon\AppStream" -Name EnableNonGraphics4K -Value
     1
    ```

  - Using registry edit (regedit):
    - Registry path: `HKLM:Software\Amazon\AppStream`
    - Registry value name: `EnableNonGraphics4K`
    - Registry value data: `1`
    - Registry value type: `DWord`
- For Linux: Edit `/etc/euc/workspaces-applications.conf` and update the following `[display]` configuration:

  ```
  EnableNonGraphics4K=1
  ```

- Create a new image using this setting.
- Use this new image with your fleets.
- Your end users will see a max of 4K resolution on their display devices, depending upon the display resolution of their devices.

> **Important**
>
> High-resolution monitors (such as 4K displays) require significantly more GPU and encoding resources, with actual performance depending on display configuration (resolution and number of monitors) and the compute instance being used. If your users experience suboptimal performance on 4K monitors, either switch your fleets to more

powerful instances or recommend your users to reduce their display resolution to improve responsiveness.
Ultrawide monitors with resolutions exceeding 4096 pixels in either dimension (e.g., 5120x2160) will display black bars on the sides, as the maximum supported resolution is limited to 4096 pixels per dimension.

Disable 4K resolution for non-graphics instances

If your users are facing performance issues with 4K monitor resolutions on instances without graphics hardware, you can follow the steps below to limit the resolution to 2K (2560x1440).

- For Windows: Configure the following registry setting on image builder.
  - Using PowerShell:

```
Set-ItemProperty "HKLM:\SOFTWARE\Amazon\AppStream" -Name EnableNonGraphics4K -Value
 0
```

  - Using registry edit (regedit):
    - Registry path: `HKLM:Software\Amazon\AppStream`
    - Registry value name: `EnableNonGraphics4K`
    - Registry value data: `0`
    - Registry value type: `DWord`
- For Linux: Edit `/etc/euc/workspaces-applications.conf` and update the following `[display]` configuration:

```
EnableNonGraphics4K=0
```

- Create a new image using this setting.
- Use this new image with your fleets.
- Your end users will see a max of 2K resolution on their display devices.

Real-Time Audio-Video

WorkSpaces Applications supports real-time audio-video (AV) by redirecting local webcam video input to WorkSpaces Applications streaming sessions. This capability enables your users to use

their local webcam for video and audio conferencing within an WorkSpaces Applications streaming session. With real-time AV and support for real-time audio, your users can collaborate by using familiar video and audio conferencing applications without having to leave their WorkSpaces Applications streaming session.

When a user starts a video conference from within an WorkSpaces Applications streaming session, WorkSpaces Applications compresses the webcam video and microphone audio input locally before transmitting this data over a secure channel to a streaming instance. During their streaming sessions, users can enable audio and video input by using the WorkSpaces Applications toolbar. If users have more than one webcam (for example, if they have a USB webcam that is connected to their local computer and a built-in webcam), they can also choose which webcam to use during their streaming session.

> **Note**
>
> For multi-session fleets, only in/out functionalities are accessible. Video in (webcam support) is not yet available for multi-session fleets.

To configure and test support for real-time AV, complete the following steps.

Configure and test support for real-time AV

1. Create a new image builder or connect to an existing image builder that meets the following requirements:

   - The image builder must run Windows Server 2016 or Windows Server 2019.
   - The image builder must use a version of the WorkSpaces Applications agent released on or after June 1, 2021.
   - For WorkSpaces Applications agents released on or after May 17, 2021, real-time AV is enabled by default. To create a streaming URL for testing, you can skip steps 3 through 6 and disconnect from the image builder. If you need to disable real-time AV, complete all of the steps, and disable webcam permissions in step 4.
   - The image builder must use a version of the WorkSpaces Applications agent released on or after June 24, 2021 to support video when connecting using web browser access. For more information about supported web browsers, see the section called “Web Browser Access”.

For information about how to create an image builder, see Launch an Image Builder to Install and Configure Streaming Applications.

2. Connect to the image builder that you want to use and sign in as Administrator. To connect to the image builder, do either of the following:

   - Use the WorkSpaces Applications console (for web connections only)
   - Create a streaming URL (for web or WorkSpaces Applications client connections)

   > **Note**
   >
   > If the image builder that you want to connect to is joined to an Active Directory domain and your organization requires smart card sign in, you must create a streaming URL and use the WorkSpaces Applications client for the connection. For information about smart card sign in, see Smart Cards.

3. On the image builder, open Registry Editor. To do so, on the image builder desktop, in the search box on the taskbar, type **`regedit`**. Then, select the top result for **Registry Editor**.
4. Under **HKEY_LOCAL_MACHINE\SOFTWARE\Amazon\AppStream**, create a new registry value that has the following type, name, and value data:

   - Registry value type: DWORD
   - Registry value name: WebcamPermission
   - Registry value data (Hexademical): 1 to enable or 0 to disable webcam permissions
5. After you create the registry value, switch to **Template User** or to a domain account that does not have administrator permissions on the image builder. To switch to **Template User**, in the toolbar on the top right of the session window, choose **Admin Commands**, **Switch User**, **Template User**.
6. Switch back to **Administrator**.
7. Disconnect from the image builder and create a streaming URL for the image builder. To do so:

   a. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2.

   b. In the navigation pane, choose **Images**, then choose **Image Builder**.

   c. Select the image builder from which you just disconnected, and choose **Actions**, **Create streaming URL**.

   d. Choose **Copy Link** and save the link to a secure and accessible location. You will use the link in the next step to connect to the image builder.

8. Using the streaming URL that you just created, connect to the image builder by using the WorkSpaces Applications client or web browser access.
9. Test the real-time AV experience on the image builder by following the steps in Video and Audio Conferencing.
10. After you verify that real-time AV is working as expected, disconnect from your streaming session, reconnect to the image builder and follow the necessary steps in Image Assistant to finish creating your image. For information about how to create an image, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

After you finish configuring your image builder and creating an image that supports real-time AV, you can make this feature available to your users on WorkSpaces Applications fleets. Ensure that version 1.1.257 or later of the WorkSpaces Applications client is installed on your users' computers.

> **Note**
>
> To use real-time AV with the WorkSpaces Applications client, your WorkSpaces Applications base image and agent version should be June 1, 2021 or later. We recommend using the latest WorkSpaces Applications client. For guidance that you can provide to your users to help them use real-time AV, see Video and Audio Conferencing.
> To use real-time AV with web browser access, your WorkSpaces Applications image must use a version of the WorkSpaces Applications agent released on or after June 24, 2021. For more information on supported web browsers, see the section called “Web Browser Access”.

USB Devices

The following sections provide information about WorkSpaces Applications support for USB devices.

Contents

- USB Redirection
- Smart Cards

USB Redirection

USB redirection is required for most local USB devices to be used during WorkSpaces Applications streaming sessions. When USB redirection is required, you must qualify the device before your users can use it during their WorkSpaces Applications streaming sessions. After you qualify the device, users must share the device with WorkSpaces Applications. With USB redirection, during WorkSpaces Applications streaming sessions, users' devices are not accessible for use with local applications.

In other cases, USB devices are already enabled for use with WorkSpaces Applications and no further configuration is required. For example, smart card redirection is already enabled by default when the WorkSpaces Applications client is installed. Because USB redirection isn't used when this feature is enabled, you don't need to qualify smart card readers, and users don't need to share these devices with WorkSpaces Applications to use them during streaming sessions.

> **Note**
>
> USB redirection is currently not supported for Linux-based fleet instances, or when using the WorkSpaces Applications macOS client application.

Smart Cards

WorkSpaces Applications supports using a smart card for Windows sign in to Active Directory-joined streaming instances and in-session authentication for streaming applications. Because smart card redirection is enabled by default, users can use smart card readers that are connected to their local computer and their smart cards without USB redirection.

Contents

- Windows Sign In and In-Session Authentication
- Smart Card Redirection

Windows Sign In and In-Session Authentication

WorkSpaces Applications supports the use of Active Directory domain passwords or smart cards such as Common Access Card (CAC) and Personal Identity Verification (PIV) smart cards for Windows sign in to WorkSpaces Applications streaming instances (fleets and image builders). Your users can use smart card readers connected to their local computer and their smart cards to sign in to an WorkSpaces Applications streaming instance that is joined to a Microsoft Active Directory

domain. They can also use their local smart card readers and smart cards to sign in to applications within their streaming session.

To ensure that your users can use their smart cards for Windows sign in to Active Directory-joined streaming instances and for in-session authentication for streaming applications, you must:

- Use an image that meets the following requirements:
  - The image must be created from a base image published by AWS on or after December 28, 2020. For more information, see WorkSpaces Applications Base Image and Managed Image Update Release Notes.
  - The image must use a version of the WorkSpaces Applications agent released on or after January 4, 2021. For more information, see WorkSpaces Applications Agent Release Notes.
- Enable **Smart card sign in for Active Directory** on the WorkSpaces Applications stack that your users access for streaming sessions, as described in this section.

> **Note**
>
> This setting controls only the authentication method that can be used for Windows sign in to an WorkSpaces Applications streaming instance (fleet or image builder). It doesn't control the authentication method that can be used for in-session authentication, after a user signs in to a streaming instance.

- Ensure that your users have WorkSpaces Applications client version 1.1.257 or later installed. For more information, see WorkSpaces Applications Windows Client Release Notes.

By default, password sign in for Active Directory is enabled on WorkSpaces Applications stacks. You can enable smart card sign in for Active Directory by performing the following steps in the WorkSpaces Applications console.

To enable smart card sign in for Active Directory by using the WorkSpaces Applications console

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the left navigation pane, choose **Stacks**.
3. Choose the stack for which you want to enable smart card authentication for Active Directory.
4. Choose the **User Settings** tab, and then expand the **Clipboard, file transfer, print to local device, and authentication permissions** section.

5. For **Smart card sign in for Active Directory**, choose **Enabled**.

   You can also enable **Password sign in for Active Directory**, if it's not already enabled. At least one authentication method must be enabled.
6. Choose **Update**.

Alternatively, you can enable smart card sign in for Active Directory by using the WorkSpaces Applications API, an AWS SDK, or the AWS Command Line Interface (AWS CLI).

Smart Card Redirection

When the WorkSpaces Applications client is installed, smart card redirection is enabled by default. When this feature is enabled, users can use smart card readers that are connected to their local computer and their smart cards during WorkSpaces Applications streaming sessions without USB redirection. During WorkSpaces Applications streaming sessions, users' smart card readers and smart cards remain accessible for use with local applications. The WorkSpaces Applications client redirects the smart card API calls from users' streaming applications to their local smart card.

> **Note**
>
> Smart card redirection is currently not supported for Linux-based fleet instances or multi-session fleet instances, or when using the WorkSpaces Applications macOS client application.

> **Note**
>
> If your smart card requires middleware software to operate, the middleware software must be installed on both the user's device, and the WorkSpaces Applications streaming instance.

You can disable smart card redirection during client installation on managed devices. For more information, see Choose Whether to Disable Smart Card Redirection. If you disable smart card redirection, your users can't use their smart card reader and smart card during a WorkSpaces Applications streaming session without USB redirection. In this case, you must qualify the device. After you qualify the device, users must share the device with WorkSpaces Applications. When smart card redirection is disabled, during users' WorkSpaces Applications streaming sessions, their smart card reader and smart card are not accessible for use with local applications.

Drawing Tablets

Drawing tablets, also known as pen tablets, are computer input devices that let users draw with a stylus (pen). With WorkSpaces Applications, your users can connect a drawing tablet, such as a Wacom drawing tablet, to their local computer and use the tablet with their streaming applications.

> **Note**
>
> Drawing tablets are not supported when using the WorkSpaces Applications macOS client application.

Following are requirements and considerations for enabling your users to use drawing tablets with their streaming applications.

- To enable your users to use this feature, you must configure your WorkSpaces Applications fleet to use an image that runs Windows Server 2019.
- To use this feature, users must access WorkSpaces Applications by using the WorkSpaces Applications client, or through the Google Chrome or Mozilla Firefox browsers only.
- Streaming applications must support Windows Ink technology. For more information, see Pen interactions and Windows Ink in Windows apps.
- Some applications, such GIMP, must detect drawing tablets on the streaming instance to support pressure sensitivity. If this is the case, your users must use the WorkSpaces Applications client to access WorkSpaces Applications and stream these applications. In addition, you must qualify your users' drawing tablets, and users must share their drawing tablets with WorkSpaces Applications every time they start a new streaming session. For more information, see Qualify USB Devices for Use with Streaming Applications.
- This feature is not supported on Chromebooks.

To get started with using drawing tablets during application streaming sessions, users connect their drawing tablet to their local computer with USB, share the device with WorkSpaces Applications if required for pressure sensitivity detection, and then use the WorkSpaces Applications client or a supported web browser to start a WorkSpaces Applications streaming session.

Keyboard Shortcuts

For the Windows client, most operating system keyboard shortcuts are supported. Supported keyboard shortcuts include Alt + Tab, Clipboard shortcuts (Ctrl + X, Ctrl + C, Ctrl+ V), Esc, and Alt + F4.

For the macOS client, supported keyboard shortcuts include Clipboard shortcuts (Command + X, Command + C, Command + V, Command + A, Command-Z).

Relative Mouse Offset

By default, during users' streaming sessions, WorkSpaces Applications transmits information about mouse movements to the streaming instance by using absolute coordinates and rendering the mouse movements locally. For graphics-intensive applications, such as computer-aided design (CAD)/computer-aided manufacturing (CAM) software or video games, mouse performance improves when relative mouse mode is enabled. Relative mouse mode uses relative coordinates, which represent how far the mouse moved since the last frame, rather than the absolute x-y coordinate values within a window or screen. When relative mouse mode is enabled, WorkSpaces Applications renders the mouse movements remotely.

Users can enable this feature during their WorkSpaces Applications streaming sessions by doing either of the following:

- Pressing Ctrl+Shift+F8 on the Windows client application or Ctrl+Shift+Fn+F8 on the macOS client application
- Choosing **Relative Mouse Position [Ctrl+Shift+F8]** from the **Settings** menu on the WorkSpaces Applications toolbar in the top left area of their streaming session window. This method works when they use classic mode or **Desktop View**.

Host-to-client URL redirection

Host-to-client URL redirection opens URLs accessed on a WorkSpaces Applications streaming session in the web browser on the user's local device instead of on the streaming instance. This offloads resource-intensive web content such as video to the local device, which reduces latency and server-side resource consumption.

You control which URLs are redirected by configuring URL patterns. Only approved links open locally while all others remain in the remote session. You can also configure an exception list to exclude specific URLs from redirection even if they match a configured pattern.

When the feature is enabled, WorkSpaces Applications automatically installs the Amazon DCV URL redirection browser extension on Google Chrome and Microsoft Edge in the streaming session when a new session starts. The extension redirects URLs that users open in web browsers. Links opened from apps within the session do not require the browser extension for redirection.

Topics

- Prerequisites
- Enable host-to-client URL redirection
- Configure URL patterns (allowlist)
- Configure the exception list (optional)
- Troubleshooting
- Security considerations and best practices

Prerequisites

To use host-to-client URL redirection, you must meet the following requirements:

- One of the following supported WorkSpaces Applications clients:
  - Windows client version 1.2.1778 or later
  - macOS client version 1.4.0.1012 or later
  - Web client
- A streaming instance that runs Windows Server 2019 or later
- An WorkSpaces Applications agent released on or after March 30, 2026
- A fleet with internet access enabled

Enable host-to-client URL redirection

Complete the following steps to enable host-to-client URL redirection on a new or existing stack.

To enable host-to-client URL redirection

1. Open the WorkSpaces Applications AWS Management Console.
2. In the left navigation pane, choose **Stacks**.
3. Do one of the following:
   - To configure a new stack, choose **Create Stack**.

- To modify an existing stack, select the stack and choose **Edit**.

4. In the **Content Redirection** section, select **Enable host to client URL redirection**.
5. Configure the URL patterns and optional exception list as described in the following sections.
6. Choose **Update** to save your stack configuration.

> **Note**
>
> After you save the stack settings, the changes take effect when a new streaming session is created. Existing sessions are not affected.

After you enable the feature, two configuration fields become available: **Configure host to client URL patterns** (required) and **Configure exception list** (optional).

Configure URL patterns (allowlist)

The URL patterns field defines which links redirect to the local browser. Only URLs that match these patterns are redirected. All other URLs open in the remote session.

Enter one URL pattern per line in the text field. The following example shows three URL patterns.

```
https://www.youtube.com/*
https://www.linkedin.com
https://docs.aws.amazon.com
```

Wildcard support

URL patterns support two wildcard characters for flexible matching.

Wildcard characters for URL pattern matching

Wildcard	Matches	Example	Matches URLs
*	Multiple characters (zero or more)	`https://www.youtube.com/*`	`https://www.youtube.com/watch?v=abc` , `https://www.youtube.com/playlist/xyz`

Wildcard	Matches	Example	Matches URLs
?	Single character (exactly one)	`https://www.example.com/page?`	`https://www.example.com/page1` , `https://www.example.com/pageA`

Common wildcard patterns

- Redirect all pages within a domain: `https://www.youtube.com/*`
- Redirect all subdomains: `https://*.example.com/*`

Pattern matching rules

- Case sensitivity: Patterns are case-insensitive for domains and paths.
- Protocol matching: The protocol (`http://` or `https://`) must match exactly.

Configure the exception list (optional)

The exception list lets you exclude specific URL patterns from redirection, even if they match the allowlist.

How URL evaluation works

URLs are evaluated in the following order:

1. **Check exception list first.** If a URL matches an exception pattern, it opens in the remote session (no redirection).
2. **Check allowlist second.** If a URL matches an allowlist pattern and is not in the exception list, it redirects to the local browser.
3. **Default behavior.** If a URL matches neither list, it opens in the remote session.

Example scenarios

Scenario 1: Allow YouTube except for specific channels

- Allowlist: `https://www.youtube.com/*`

- Exception list: `https://www.youtube.com/channel/UCrestrictedChannel`
- Result: `https://www.youtube.com/watch?v=abc123` redirects to the local browser. `https://www.youtube.com/channel/UCrestrictedChannel` opens in the remote session.

Scenario 2: Exclude specific subdomains

- Allowlist: `https://*.example.com/*`
- Exception list: `https://admin.example.com/*`, `https://secure.example.com/*`
- Result: `https://www.example.com/page` redirects to the local browser. `https://admin.example.com/dashboard` opens in the remote session.

Troubleshooting

URLs not redirecting when expected

- Verify that the URL pattern includes the protocol (`https://`).
- Check for typos in the pattern.
- Verify that wildcards are used correctly.
- Confirm that the feature is enabled on the stack.

URLs redirecting when they are not expected to

- Add the URL pattern to the exception list.
- Verify that exception list patterns are more specific than allowlist patterns.

Local browser not launching

- Verify that the local device has a default browser configured.
- Verify that the WorkSpaces Applications client is up to date.
- Verify network connectivity on the local device.
- On the web client, verify that popups are allowed on the streaming URL.

Security considerations and best practices

Keep the following security considerations in mind when you configure host-to-client URL redirection.

- Add URLs that contain sensitive data to the exception list to keep them in the remote session.
- Redirected URLs might require separate authentication in the local browser.
- Cookies and session data are not shared between the remote session and the local browser.
- Use HTTPS patterns to help ensure encrypted communication.

Best practices

- Verify that local devices can access redirected URLs. Check corporate firewall policies if needed.
- If users connect through a proxy, verify that redirected URLs are accessible.
- Start small: Begin with a limited set of trusted domains and expand based on user feedback.
- Use HTTPS: Always prefer `https://` patterns over `http://` for security.
- Be specific: Use specific paths rather than broad wildcards when possible.
- Review regularly: Review and update URL patterns regularly to remove unused entries.
- Test thoroughly: Validate configuration with pilot users before organization-wide deployment.

Install and Configure the WorkSpaces Applications Client

You can have your users install the WorkSpaces Applications client themselves, or you can install the WorkSpaces Applications client for them by running PowerShell scripts remotely.

You must qualify the USB devices that you want to enable your users to use with their streaming session. If their USB device is not qualified, it won't be detected by WorkSpaces Applications and can't be shared with the session.

The following topics describe how to install and configure the WorkSpaces Applications client.

Contents

- Have Your Users Install the WorkSpaces Applications Client Themselves
- Tutorial: Install the Amazon WorkSpaces Applications Client and Customize the Client Experience for Your Users

- Update the WorkSpaces Applications Enterprise Deployment Tool, Client, and USB Driver Manually
- Qualify USB Devices for Use with Streaming Applications
- Configure a Connection Method for Your WorkSpaces Applications Client Users
- Enable Users to Share a USB Device with an WorkSpaces Applications Streaming Session
- Redirect a Streaming Session from the Web Browser to the WorkSpaces Applications Client
- Enable File System Redirection for Your WorkSpaces Applications Users
- Enable Local Printer Redirection for Your WorkSpaces Applications Users

Have Your Users Install the WorkSpaces Applications Client Themselves

For step-by-step guidance that you can provide your users to help them install the WorkSpaces Applications client, see Setup for Windows or the section called "Setup and installation for macOS".

> **Important**
>
> For the Windows client, if your organization has deployed antivirus software that prevents users from running .exe files, you must add an exception to allow your users to run the WorkSpaces Applications client installation .exe program. Otherwise, when users try to install the client, either nothing happens, or they receive an error after they start the installation program.

After users install the client, if you plan to let your users use USB devices during their WorkSpaces Applications streaming sessions, the following requirements must be met:

- You must qualify the USB devices that can be used with WorkSpaces Applications. For more information, see Qualify USB Devices for Use with Streaming Applications.
- After their devices are qualified, your users must share the devices with WorkSpaces Applications every time they start a new streaming session. For guidance that you can provide your users to them complete this task, see USB Devices.

Tutorial: Install the Amazon WorkSpaces Applications Client and Customize the Client Experience for Your Users

The following sections describe how to install the WorkSpaces Applications client and customize the client experience for your users. If you plan to download and install the client for your users, first download the Enterprise Deployment Tool. You can then run PowerShell scripts to install the WorkSpaces Applications client and configure client settings remotely.

> **Note**
>
> Using the Enterprise Deployment Tool with the WorkSpaces Applications macOS client is not supported.

Contents

- Download the Enterprise Deployment Tool
- Install the WorkSpaces Applications Client and USB Driver
- Accessing WorkSpaces Applications with the WorkSpaces Applications Client
- Set the StartURL Registry Value for WorkSpaces Applications Client Users
- Set the TrustedDomains Registry Value to Enable Other Domains for the WorkSpaces Applications Client
- Create the AS2TrustedDomains DNS TXT Record to Enable Your Domain for the WorkSpaces Applications Client Without Registry Changes
- Disable DNS TXT Record Lookup for Trusted Domains
- Choose Whether to Disable Automatic Client Updates
- Choose Whether to Disable On-Demand Diagnostic Log Uploads
- Choose Whether to Disable Native Application Mode
- Choose Whether to Disable Local Printer Redirection
- Choose Whether to Disable Smart Card Redirection
- Configure Additional WorkSpaces Applications Client Settings for Your Users
- Using Group Policy to Customize WorkSpaces Applications Client Experience

Download the Enterprise Deployment Tool

The Enterprise Deployment Tool includes the WorkSpaces Applications client installation files and a Group Policy administrative template.

1. To download the Enterprise Deployment Tool, on the bottom right of the WorkSpaces Applications supported clients page, select the **Enterprise Deployment Tool** link. This link opens a .zip file that contains the required files for the latest version of the tool.
2. To extract the required files, navigate to the location where you downloaded the tool, right-click the **AmazonWorkSpacesApplicationsClient_EnterpriseSetup_<version>.zip** file, and choose **Extract All**. The .zip file contains two installation programs and a Group Policy administrative template:

   - WorkSpaces Applications client installer (AmazonWorkSpacesApplicationsClientSetup_<version>.msi) — Installs the WorkSpaces Applications client.
   - WorkSpaces Applications USB driver installer (AmazonAppStreamUsbDriverSetup_<version>.exe) — Installs the WorkSpaces Applications USB driver that is required to use USB devices with applications streamed through WorkSpaces Applications.
   - WorkSpaces Applications client Group Policy administrative template (as2_client_config.adm) — Lets you configure the WorkSpaces Applications client through Group Policy.

Install the WorkSpaces Applications Client and USB Driver

After you download the WorkSpaces Applications client installation files, run the following PowerShell script on your users' computers to install the WorkSpaces Applications client and the USB driver silently.

> **Note**
>
> To run this script, you must be logged in to the applicable computer with Administrator permissions. You can also run the script remotely under the System account on startup.

Note

The WorkSpaces Applications client requires a 64-bit (x64) version of Windows.

```
Start-Process msiexec.exe -Wait -ArgumentList '/i
 AmazonWorkSpacesApplicationsClientSetup_<version>.msi ALLUSERS=1 /quiet'

Start-Process AmazonAppStreamUsbDriverSetup_<version>.exe -Wait -ArgumentList '/quiet'
```

After installation completes, the installer places the WorkSpaces Applications client in the following location on your users' computers: `C:\Program Files\Amazon Web Services, Inc\Amazon WorkSpaces Applications\`. The installer automatically registers and starts the auto-update service.

Accessing WorkSpaces Applications with the WorkSpaces Applications Client

By default, when users launch the WorkSpaces Applications client, they can connect only to URLs that include the WorkSpaces Applications domain or domains that include a DNS TXT record that enables the connection. You can let client users access domains other than the WorkSpaces Applications domain by doing any of the following:

- Set the `StartURL` registry value to specify a custom URL that users can access, such as the URL for your organization's login portal.
- Set the `TrustedDomains` registry value to specify trusted domains that users can access.
- Create the `AS2TrustedDomains` DNS TXT record to specify trusted domains that users can access. This method lets you avoid registry changes.

Note

The WorkSpaces Applications client and DNS TXT record configuration do not prevent users from using other connection methods to access the domains or URLs that you specify. For example, users can access specified domains or URLs by using a web browser, if they have network access to the domains or URLs.

Set the StartURL Registry Value for WorkSpaces Applications Client Users

You can use the `StartUrl` registry value to set a custom URL that is populated in the WorkSpaces Applications client when a user launches the client. You can create this HKLM registry key while installing the client so that your users don't need to specify a URL when they launch the client.

After the WorkSpaces Applications client is installed, you can run the following PowerShell script to create this registry key, or you can use the administrative template that is included in the WorkSpaces Applications client Enterprise Deployment Tool to configure the client through Group Policy.

Replace the `StartUrl` value with a URL for your identity provider (IdP). The URL must use a certificate that is trusted by the device. This means that the certificate that is used by the `StartUrl` webpage must contain a Subject Alternative Name (SAN) that includes the URL's domain name. For example, if your `StartUrl` is set to https://appstream.example.com, the SSL certificate must have a SAN that includes appstream.example.com.

> **Note**
>
> To run this script, you must be logged in to the applicable computer with Administrator permissions. You can also run the script remotely under the System account on startup.

```
$registryPath="HKLM:\Software\Amazon\AppStream Client"
New-Item -Path "HKLM:\Software\Amazon" -Name "AppStream Client" -Force

New-ItemProperty -Path $registryPath -Name "StartUrl" -Value "https://www.example.com"
 -PropertyType String -Force | Out-Null
```

Set the TrustedDomains Registry Value to Enable Other Domains for the WorkSpaces Applications Client

You can configure the WorkSpaces Applications client to connect to URLs in trusted domains that you specify. For example, you might want to let users connect to any URL in your organizational domain or to any URL in one or more of your IdP domains. When you specify the URL, use the following format: *.*example-idp.com*.

You can specify a list of trusted domains in a comma-separated format. Add this list as a registry value to the WorkSpaces Applications `TrustedDomains` HKLM registry key. We recommend that you create this registry key and specify the list of trusted domains when you install the WorkSpaces

Applications client or, if you are using Microsoft Active Directory, through Group Policy. That way, your users can connect to a URL in any trusted domain immediately after the client is installed.

After the WorkSpaces Applications client is installed, you can run the following PowerShell script to create this registry key. Or, you can use the administrative template that is included in the WorkSpaces Applications client Enterprise Deployment Tool to configure the client through Group Policy.

Replace the `TrustedDomains` value with a comma-separated list for one or more of your organizational or IdP domains. The certificate used by the trusted domain webpage must contain a SAN that includes the URL's domain. For example, if your trusted domain includes *.example.com, and users specify https://appstream.example.com, the SSL certificate must have a SAN that includes appstream.example.com.

> **Note**
>
> To run this script, you must be logged in to the applicable computer with Administrator permissions. You can also run the script remotely under the System account on startup.

```
$registryPath="HKLM:\Software\Amazon\AppStream Client"
New-Item -Path "HKLM:\Software\Amazon" -Name "AppStream Client" -Force

New-ItemProperty -Path $registryPath -Name "TrustedDomains" -Value "*.example1.com,
 *.example2.com, aws.amazon.com" -PropertyType String -Force | Out-Null
```

The following are requirements and considerations for formatting trusted domain names.

- The following characters are supported: a-z, 0-9, -, *
- DNS treats the * character either as a wildcard or as an asterisk character (ASCII 42), depending on where it appears in the domain name. Following are restrictions when using * as a wildcard in the name of a DNS record:
  - The * must replace the leftmost label in a domain name. For example, *.example.com or *.prod.example.com. If you include * in any other position, such as prod.*.example.com, DNS treats it as an asterisk character (ASCII 42), not as a wildcard.
  - The * must replace the entire label. For example, you can't specify *prod.example.com or prod*.example.com.

- The * applies to the subdomain level that includes the *, and to all the subdomains of that subdomain. For example, if an entry is named *.example.com, the WorkSpaces Applications client allows zenith.example.com, acme.zenith.example.com, and pinnacle.acme.zenith.example.com.

Create the AS2TrustedDomains DNS TXT Record to Enable Your Domain for the WorkSpaces Applications Client Without Registry Changes

You can enable users to connect to any URL in your organizational domain (for example, *.example.com) or to any URL in your IdP domains (for example, *.example-idp.com) by creating a DNS TXT record in that domain. When you create the DNS TXT record, the `StartURL` or `TrustedDomains` registry values are not required to allow a user to connect to a URL.

You can specify a list of trusted subdomains in a comma-separated format, prefixed with `AS2TrustedDomains=`. Then, create a DNS TXT record for the appropriate domain. The `AS2TrustedDomains` DNS TXT record can only enable the same domain, or subdomains, of the domain in which the DNS TXT record is created. You cannot use the DNS TXT record to enable other domains.

For more information about setting up the DNS record, see Enable your organizational domain for the AppStream 2.0 client with a Route 53 DNS TXT record and Creating an AS2TrustedDomains DNS TXT record to redirect the AppStream 2.0 native client to a third-party identity provider.

> **Note**
>
> When you create DNS TXT records, any users can stream from enabled domains that are not included in the `StartURL` or `TrustedDomains` registry values. The WorkSpaces Applications client and DNS TXT record configuration do not prevent users from using other connection methods to access the domains or URLs that you specify. For example, users can access specified domains or URLs by using a web browser, if they have network access to the domains or URLs.

DNS TXT Record Configuration Example

The following is an example of a DNS TXT record configuration. With the configuration for this example, users can launch the WorkSpaces Applications client and connect to

appstream.example.com or appstream-dev.example.com. However, they cannot connect to example.com.

- `Domains to enable` — appstream.example.com, appstream-dev.example.com
- `DNS TXT record location` — example.com
- `DNS TXT record value` — AS2TrustedDomains=appstream.example.com,appstream-dev.example.com

Requirements and Considerations

Keep in mind the following requirements and considerations for creating a DNS TXT record:

- You must create the TXT record at the second-level domain. For example, if your domain is prod.appstream.example.com, you must create the DNS TXT record at example.com.
- The TXT record value must start with `AS2TrustedDomains=`
- The following characters are supported: a-z, 0-9, -, *
- DNS treats the * character either as a wildcard or as an asterisk character (ASCII 42), depending on where it appears in the domain name. Following are restrictions when using * as a wildcard in the name of a DNS record:
  - The * must replace the leftmost label in a domain name. For example, *.example.com or *.prod.example.com. If you include * in any other position, such as prod.*.example.com, DNS treats it as an asterisk character (ASCII 42), not as a wildcard.
  - The * must replace the entire label. For example, you can't specify *prod.example.com or prod*.example.com.
  - The * applies to the subdomain level that includes the *, and to all the subdomains of that subdomain. For example, if an entry is named *.example.com, the WorkSpaces Applications client allows connections to the following domains: zenith.example.com, acme.zenith.example.com, and pinnacle.acme.zenith.example.com.

Disable DNS TXT Record Lookup for Trusted Domains

By default, when users launch the WorkSpaces Applications client and specify a URL that is not an WorkSpaces Applications domain, the client performs a DNS TXT record lookup. The lookup is performed on the second-level domain of the URL so that the client can determine whether the domain is included in the `AS2TrustedDomains` list. This behavior lets users connect to

domains that are not specified in the `StartURL` or `TrustedDomains` registry keys, or WorkSpaces Applications domains.

You can disable this behavior by setting the value for the `DnsTxtRecordQueryDisabled` registry key to `true`. You can create this registry key when you install the WorkSpaces Applications client. That way, the client connects only to URLs that are specified by the `StartURL` or `TrustedDomains` registry keys.

After the WorkSpaces Applications client is installed, you can run the following PowerShell script to create this registry key. Or, you can use the administrative template that is included in the WorkSpaces Applications client Enterprise Deployment Tool to configure the client through Group Policy.

> **Note**
>
> To run this script, you must be logged in to the applicable computer with Administrator permissions. You can also run the script remotely under the System account on startup.

```
$registryPath="HKLM:\Software\Amazon\AppStream Client"
New-Item -Path "HKLM:\Software\Amazon" -Name "AppStream Client" -Force

New-ItemProperty -Path $registryPath -Name "DnsTxtRecordQueryDisabled" -Value "true" -
PropertyType String -Force | Out-Null
```

Choose Whether to Disable Automatic Client Updates

By default, when a new version of the WorkSpaces Applications client is available, the client updates automatically to the latest version. You can disable automatic updates by setting the value for the `AutoUpdateDisabled` registry key to `true`. You can create this registry key when you install the WorkSpaces Applications client. That way, the client is not updated automatically whenever a new version is available.

After the WorkSpaces Applications client is installed, you can run the following PowerShell script to create this registry key. Or, you can use the administrative template that is included in the WorkSpaces Applications client Enterprise Deployment Tool to configure the client through Group Policy.

> **Note**
>
> To run this script, you must be logged in to the applicable computer with Administrator permissions. You can also run the script remotely under the System account on startup.

```
$registryPath="HKLM:\Software\Amazon\AppStream Client"
New-Item -Path "HKLM:\Software\Amazon" -Name "AppStream Client" -Force

New-ItemProperty -Path $registryPath -Name "AutoUpdateDisabled" -Value "True" -
PropertyType String -Force | Out-Null
```

Choose Whether to Disable On-Demand Diagnostic Log Uploads

By default, the WorkSpaces Applications client allows users to upload diagnostic logs and minidumps on demand to WorkSpaces Applications (AWS). In addition, if an exception occurs or the WorkSpaces Applications client stops responding, users are prompted to choose whether they want to upload the minidump and associated logs. For more information about on-demand diagnostic logging, see Automatic and On-Demand Diagnostic Log Uploads.

You can disable these behaviors by setting the value for the `UserUploadOfClientLogsAllowed` registry key to `false`. You can create this HKLM registry key when you install the WorkSpaces Applications client.

After the WorkSpaces Applications client is installed, you can run the following PowerShell script to create this registry key. Or, you can use the administrative template that is included in the WorkSpaces Applications client Enterprise Deployment Tool to configure the client through Group Policy.

> **Note**
>
> To run this script, you must be logged in to the applicable computer with Administrator permissions. You can also run the script remotely under the System account on startup.

```
$registryPath="HKLM:\Software\Amazon\AppStream Client"
New-Item -Path "HKLM:\Software\Amazon" -Name "AppStream Client" -Force
```

```
New-ItemProperty -Path $registryPath -Name "UserUploadOfClientLogsAllowed" -Value
 "false" -PropertyType String -Force | Out-Null
```

Choose Whether to Disable Native Application Mode

By default, the WorkSpaces Applications client can run in either classic mode or native application mode. You can disable native application mode by setting the value for the `NativeAppModeDisabled` registry key to `true`. You can create this HKLM registry key when you install the WorkSpaces Applications client. If the value is set to `true`, the client runs in classic mode only. For information about native application mode, see Native Application Mode.

After the WorkSpaces Applications client is installed, you can run the following PowerShell script to create this registry key. Or, you can use the administrative template that is included in the WorkSpaces Applications client Enterprise Deployment Tool to configure the client through Group Policy.

> **Note**
>
> To run this script, you must be logged in to the applicable computer with Administrator permissions. You can also run the script remotely under the System account on startup.

```
$registryPath="HKLM:\Software\Amazon\AppStream Client"
New-Item -Path "HKLM:\Software\Amazon" -Name "AppStream Client" -Force

New-ItemProperty -Path $registryPath -Name "NativeAppModeDisabled" -Value "True" -
PropertyType String -Force | Out-Null
```

Choose Whether to Disable Local Printer Redirection

By default, the WorkSpaces Applications client enables users to redirect print jobs from their streaming applications to a printer that is connected to their local computer. You can disable local printer redirection by setting the value for the `PrinterRedirectionDisabled` registry key to `true`. You can create this HKLM registry key when you install the WorkSpaces Applications client. If the value is set to `true`, the client does not redirect print jobs from users' streaming applications to a printer that is connected to their local computer.

After you install the WorkSpaces Applications client, you can run the following PowerShell script to create this registry key. Or, you can use the administrative template that is included in the

WorkSpaces Applications client Enterprise Deployment Tool to configure the client through Group Policy.

> **Note**
>
> To run this script, you must be logged in to the applicable computer with Administrator permissions. You can also run the script remotely under the System account on startup.

```
$registryPath="HKLM:\Software\Amazon\AppStream Client"
New-Item -Path "HKLM:\Software\Amazon" -Name "AppStream Client" -Force

New-ItemProperty -Path $registryPath -Name "PrinterRedirectionDisabled" -Value "True" -
PropertyType String -Force | Out-Null
```

Choose Whether to Disable Smart Card Redirection

By default, smart card redirection is enabled for the WorkSpaces Applications client. When this feature is enabled, users can use smart card readers that are connected to their local computers and their smart cards during WorkSpaces Applications streaming sessions without USB redirection. During WorkSpaces Applications streaming sessions, users' smart card readers and smart cards remain accessible for use with local applications. The client redirects the smart card API calls from users' streaming applications to their local smart card. You can disable smart card redirection by setting the value for the `SmartCardRedirectionDisabled` registry key to `true`. You can create this HKLM registry key when you install the WorkSpaces Applications client.

If the value is set to `true`, your users can't use their smart card readers and smart cards during a WorkSpaces Applications streaming session without USB redirection. In this case, users can't sign in to their streaming applications by using a smart card that is connected to their local computer unless you qualify the device. After you qualify the device, users must share the device with WorkSpaces Applications. When smart card redirection is disabled, during users' WorkSpaces Applications streaming sessions, their smart card readers and smart cards are not accessible for use with local applications.

After you install the WorkSpaces Applications client, you can run the following PowerShell script to create this registry key. Or, you can use the administrative template that is included in the WorkSpaces Applications client Enterprise Deployment Tool to configure the client through Group Policy.

> **Note**
>
> To run this script, you must be logged in to the applicable computer with Administrator permissions. You can also run the script remotely under the System account on startup.

```
$registryPath="HKLM:\Software\Amazon\AppStream Client"
New-Item -Path "HKLM:\Software\Amazon" -Name "AppStream Client" -Force

New-ItemProperty -Path $registryPath -Name "SmartCardRedirectionDisabled" -Value "True"
 -PropertyType String -Force | Out-Null
```

Configure Additional WorkSpaces Applications Client Settings for Your Users

The WorkSpaces Applications client uses registry keys to configure the following additional client settings:

- WorkSpaces Applications client End-User License Agreement (EULA) acceptance
- WorkSpaces Applications client EULA version accepted
- Automatic diagnostic log uploads for the WorkSpaces Applications client
- Automatic updates for the USB driver that is used to pass USB devices to WorkSpaces Applications
- Enabling hardware rendering in the WorkSpaces Applications client
- Setting custom folder paths for file system redirection in the WorkSpaces Applications client
- Opening URL for your identity provider (IdP) in system default browser
- Adding an indicator to provide real-time visual feedback on streaming data loss

The following table summarizes the registry values for additional client settings that you can use to customize the WorkSpaces Applications client experience for your users.

> **Note**
>
> These values are case sensitive.

Value	Registry path	Type	Description	Data
`EULAAccepted`	`HKCU\Software\Amazon\AppStream Client`	String	Set this value to `true` to accept the WorkSpaces Applications client EULA on behalf of your users.	**`true`/`false`**
`AcceptedEULAVersion`	`HKCU\Software\Amazon\AppStream Client`	String	The version of EULA that is accepted. If the current version of the WorkSpaces Applications client EULA is different from the version of the EULA that is accepted, users are prompted to accept the current version of the EULA.	**`1.0`**
`DiagnosticInfoCollectionAllowed`	`HKCU\Software\Amazon\AppStream Client`	String	Set this value to `true` to enable WorkSpaces Applications to automatically send diagnostic logs from	**`true`/`false`**

Value	Registry path	Type	Description	Data
			the WorkSpaces Applications client to WorkSpaces Applications (AWS).	
USBDriverOptIn	HKCU\Software\Amazon\AppStream Client	String	Set this value to true to enable WorkSpaces Applications to automatically update the USB driver that is used to pass USB devices to WorkSpaces Applications.	**true/false**
HardwareRenderingEnabled	HKCU\Software\Amazon\AppStream Client	String	Set this value to true to enable hardware rendering in the WorkSpaces Applications client.	**true/false**

Value	Registry path	Type	Description	Data
`FileRedirectionCustomDefaultFolders`	`HKCU\Software\Amazon\AppStream Client`	String	Set this value to include at least one folder path for file system redirection. Separate multiple folder paths by using '|'. By default, the following folder paths are specified: %USERPROFILE%\Desktop|%USERPROFILE%\Documents|%USERPROFILE%\Downloads	*Valid folder path*

Value	Registry path	Type	Description	Data
`OpenIdpUrlInSystemBrowser`	`HKLM\Software\Amazon\AppStream Client`	String	Set this value to `true` to enable the WorkSpaces Applications client to open the IdP URL in a system default browser. This setting applies only when the URL that the client connects to is not an WorkSpaces Applications URL. In that case, you must also set the `StartUrl` or `TrustedDomains` registry value for your IdP domain. When the client opens the IdP URL in the system default browser, the client closes. After authentication	**`true`/`false`**

Value	Registry path	Type	Description	Data
			completes in the browser, the streaming session opens the client again. This feature is supported on client version 1.1.1360 and later.	
`DataLossIndicator`	`HKLM\Software\Amazon\AppStream Client`	String	Set this value to SHOW_ON_LOSSY to include a red warning indicator when there is streaming data loss. Set this value to SHOW_ON_LOSSLESS to include a green healthy indicator when there is no streaming data loss.	**DISABLED/SHOW_ON_LOSSY/SHOW_ON_LOSSLESS**

After the WorkSpaces Applications client is installed, you can run the following PowerShell script to create these registry keys. If you don't want to create all of the registry keys, modify the script as

needed to create only the registry keys that you want. Or, you can use the administrative template that is provided in the WorkSpaces Applications client Enterprise Deployment Tool to configure the client through Group Policy.

Note

You must set the following entries for each user.

```
$registryPath="HKCU:\Software\Amazon\AppStream Client"
New-Item -Path "HKCU:\Software\Amazon" -Name "AppStream Client" -Force
New-ItemProperty -Path $registryPath -Name "EULAAccepted" -Value "true" -PropertyType
 String -Force | Out-Null
New-ItemProperty -Path $registryPath -Name "AcceptedEULAVersion" -Value "1.0" -
PropertyType String -Force | Out-Null
New-ItemProperty -Path $registryPath -Name "DiagnosticInfoCollectionAllowed" -Value
 "true" -PropertyType String -Force | Out-Null
New-ItemProperty -Path $registryPath -Name "USBDriverOptIn" -Value "true" -PropertyType
 String -Force | Out-Null
New-ItemProperty -Path $registryPath -Name "HardwareRenderingEnabled" -Value "true" -
PropertyType String -Force | Out-Null
New-ItemProperty -Path $registryPath -Name "FileRedirectionCustomDefaultFolders" -Value
 "%USERPROFILE%\Desktop|%USERPROFILE%\Documents|%USERPROFILE%\Downloads" -PropertyType
 String -Force | Out-Null
```

To set the `OpenIdpUrlInSystemBrowser` and `DataLossIndicator` registry values, run the following PowerShell script. These values are stored in the HKLM registry path and require Administrator permissions.

Note

To run this script, you must be logged in to the applicable computer with Administrator permissions. You can also run the script remotely under the System account on startup.

```
$registryPath="HKLM:\Software\Amazon\AppStream Client"
New-Item -Path "HKLM:\Software\Amazon" -Name "AppStream Client" -Force

New-ItemProperty -Path $registryPath -Name "OpenIdpUrlInSystemBrowser" -Value "true" -
PropertyType String -Force | Out-Null
```

```
New-ItemProperty -Path $registryPath -Name "DataLossIndicator" -Value "SHOW_ON_LOSSY" -
PropertyType String -Force | Out-Null
```

Using Group Policy to Customize WorkSpaces Applications Client Experience

You can use the administrative template that is provided in the WorkSpaces Applications client Enterprise Deployment Tool to configure the client through Group Policy. To learn how to load administrative templates into the Group Policy Management Console, see Recommendations for managing Group Policy administrative template (.adm) files in the Microsoft Support documentation.

Update the WorkSpaces Applications Enterprise Deployment Tool, Client, and USB Driver Manually

By default, the WorkSpaces Applications client and USB driver update automatically when we release a new client version. If you used the Enterprise Deployment Tool to install the WorkSpaces Applications client and disabled automatic updates, you must update the client and USB driver manually. To do so, run the following PowerShell commands on your users' computers.

> **Note**
>
> To run these commands, you must either be logged in to the applicable computer as Administrator, or you can run the script remotely under the SYSTEM account on startup. Using the Enterprise Deployment Tool to manage the WorkSpaces Applications macOS client is not supported.

1. Install the new version of the WorkSpaces Applications client over the existing version:

   ```
   Start-Process msiexec.exe -Wait -ArgumentList '/i
    AmazonWorkSpacesApplicationsClientSetup_<new_version>.msi ALLUSERS=1 /quiet'
   ```

   > **Note**
   >
   > You don't need to uninstall the previous version or reboot. The new version automatically replaces the previous version.

2. (Optional) Update the WorkSpaces Applications USB driver:

```
Start-Process AmazonAppStreamUsbDriverSetup_<new_version>.exe -Wait -ArgumentList
 '/quiet'
```

Upgrading from versions 1.2.1830 and earlier

If you previously installed the client using the Enterprise Deployment Tool (versions 1.2.1830 and earlier), the new MSI automatically removes the legacy installation during upgrade. No manual uninstall or reboot is required. The install location changes from `C:\Program Files (x86)\Amazon WorkSpaces Applications Client Installer\` to `C:\Program Files\Amazon Web Services, Inc\Amazon WorkSpaces Applications\`. The MSI requires a 64-bit version of Windows.

Qualify USB Devices for Use with Streaming Applications

There are two methods for specifying which USB devices your users can redirect into their WorkSpaces Applications streaming instances:

Note

USB redirection is currently only supported on Windows WorkSpaces Applications streaming instances. It is not supported on the macOS client.

- You can create the USB device filter strings within the configuration file saved on an image. This method can only be used with Always-On and On-Demand fleets.
- You can specify USB device filter strings when you create the fleet, either with the AWS Management Console or with the `CreateFleet` API. For detailed information about these strings, see the section below. This method can only be used with Elastic fleets.

You can create a file on your WorkSpaces Applications image that specifies which USB devices a user can make available for their streaming applications. To qualify your users' USB devices so that the devices can be used with streaming applications, perform these steps.

> **Note**
>
> The device qualification process creates a device compatibility filter. This filter determines which USB devices are recognized for use during streaming sessions. This filter is not a security access control. Do not rely on it as a security boundary. Limit qualification to the devices your users need to avoid unexpected device behavior during streaming sessions.

1. If you haven't already done so, install the WorkSpaces Applications client. For information, see Install and Configure the WorkSpaces Applications Client.
2. Connect the USB device that you want to qualify to your computer.
3. Navigate to **C:\Users\<logged-in-user>\AppData\Local\AppStreamClient**, and double-click **dcvusblist.exe**.
4. In the **DCV - USB devices** dialog box, the list of USB devices that are connected to your local computer displays. The **Filter** column displays the filter string for every USB device. Right-click the list entry for a USB device that you want to enable, and choose **Copy filter string**.
5. On your desktop, choose the Windows **Start** button, and search for Notepad. Double-click **Notepad** to open a new file, copy the filter string to the file, and save it. Later, you'll use the filter string to qualify the USB device.
6. Launch a new image builder. For more information, see Launch an Image Builder to Install and Configure Streaming Applications.
7. After your image builder is in the **Running** state, perform the following steps to create a streaming URL and connect to the image builder by using the WorkSpaces Applications client.

    a. With your image builder selected in the list, choose **Actions**, **Create streaming URL**.

    b. In the **Create streaming URL** dialog box, choose **Copy link**, and copy and paste the web address into a separate file for later use. You'll use this URL to reconnect to the image builder in step 12.

    c. Choose **Launch in Client**.

    d. If the **Launch Application** dialog box opens and prompts you to choose the application to use when opening the link, choose **Amazon AppStream**, **Open link**. To prevent this dialog box from displaying the next time you perform this step to connect to an image builder, select the **Remember my choice for amazonappstream links** check box.

e. If the WorkSpaces Applications client displays links to the AWS Customer Agreement, AWS Service Terms, and the AWS Privacy Notice, and third-party notices, review this information, and then choose **Finish**.

f. If the client sign-in page is displayed, the web address field is prepopulated with the streaming URL. Choose **Connect**.

g. If prompted, log in to the image builder as Administrator.

8. After you are connected to the image builder, if your USB device requires you to install drivers before you use it, download and install the drivers on the image builder. For example, if you use the Connexion 3D mouse, you must download and install the required Connexion drivers on the image builder.

9. On your image builder desktop, choose the Windows **Start** button, and search for Notepad. Right-click **Notepad**, and choose **Run as Administrator**.

10. Choose **File**, **Open**, and open the following file: `C:\ProgramData\Amazon\Photon\DCV\usb_device_allowlist.txt`. You can also allow an entire category of devices or all devices from a specific manufacturer by using wildcard expressions in the `usb_device_allowlist.txt` file.

11. Copy the filter string from your local computer to the image builder. The filter string for a specific USB device is a comma-separated string of the following fields: **Name**, **Base Class**, **SubClass**, **Protocol**, **ID Vendor**, **ID Product**, **Support Autoshare**, and **Skip Reset**. For detailed information about these strings, see Working with USB Device Filter Strings.

12. Disconnect from your image builder, restart it, and reconnect to it by using the WorkSpaces Applications client. To do so, open the WorkSpaces Applications client and paste the streaming URL that you created in step 7 into the client sign-in web address field, and choose **Connect**.

13. On the image builder, test your USB device to confirm that it works as expected.

14. Before your users can use the USB device in an WorkSpaces Applications session, they must first share the device with their session. For guidance that you can provide your users to help them perform this task, see USB Devices.

15. If the USB device works with the image builder as expected, create an image. For more information, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

16. After you finish creating the image, update your WorkSpaces Applications fleet to use the new image.

Working with USB Device Filter Strings

This section describes the filter strings that are available for qualifying USB devices for WorkSpaces Applications streaming sessions. It also provides guidance for working with these strings. The following filter strings are available:

- `Name` — By default, the value for this filter string is the name of the device, but you can specify your own value.
- `Base Class,SubClass,Protocol` — The USB class code for the device. For more information, see Defined Class Codes.
- `ID Vendor (VID)` — A unique identifier that is assigned by the USB organization to the manufacturer of the USB device.
- `ID Product (PID)` — A unique identifier that assigned by the manufacturer to the USB device.
- `Support Autoshare` — Lets the WorkSpaces Applications client automatically share the device when a streaming session starts. Set this value to `1` to allow automatic device sharing. Set this value to `0` to not allow automatic device sharing.
- `Skip Reset` — By default, when a USB device is shared by WorkSpaces Applications with a streaming session, the device is reset to ensure that it functions correctly. However, some USB devices don't function correctly during the streaming session if they are reset. To prevent this problem from occurring, set the value for this filter string to `1` to instruct the WorkSpaces Applications client not to reset the device while it is shared with a streaming session. To ensure that the device is reset while it is shared with a streaming session, set this value to `0`. When you set a value for `Skip Reset`, make sure that you set the value for `Support Autoshare` to `0` or `1`.

The filter string that is copied from the local computer is specific to a USB device. In some cases, you might want to allow an entire class of devices instead of allowing every possible USB device. For example, you might want to allow your users to use any kind of Wacom design tablets or use any USB mass storage device. In such scenarios, you can provide wildcard characters for specific filter string fields. If you don't know the VID and PID for your USB devices, you can search for this information in the USB ID database.

The following examples show how to configure filter strings for USB device sharing during streaming sessions:

- Allow all mass storage devices automatically on starting a streaming session — "Mass storage, 8, *, *, *, *,1,0"

- Allow all Wacom devices automatically on starting a streaming session — "Wacom tablets, 3, *, *, 1386, *,1,0"
- Allow all devices that provide an audio interface — "Audio, 1, *, *, *, *,1,0"
- Allow device X, but don't reset it while the device is shared. Don’t share the device automatically on starting a streaming session — "X, Y, *, *, 1386, *,0,1"

Configure a Connection Method for Your WorkSpaces Applications Client Users

After you install the WorkSpaces Applications client on your users' local computers, they can use the WorkSpaces Applications client to connect to a streaming session. Depending on your organizational requirements, you can provide client users with access to WorkSpaces Applications by doing one of the following: Setting up identity federation using SAML 2.0, using an WorkSpaces Applications user pool, or creating a streaming URL.

Contents

- SAML 2.0
- WorkSpaces Applications User Pool
- Streaming URL
- Next Steps

SAML 2.0

If you use external identity providers to federate your users to an WorkSpaces Applications stack, you must create a registry value to configure the WorkSpaces Applications client with a prepopulated URL whenever the client is launched. The URL must use a certificate that is trusted by the device. The certificate must contain a Subject Alternative Name (SAN) that includes the URL's domain name.

For more information, see:

- Setting Up SAML
- Set the StartURL Registry Value for WorkSpaces Applications Client Users

WorkSpaces Applications User Pool

When you create a new user in the WorkSpaces Applications user pool, or assign a user pool user to an WorkSpaces Applications stack, WorkSpaces Applications sends email to users on your behalf. Users enter the URL that was provided to them in the welcome email, enter their credentials, and then choose **Connect**.

For more information, see Amazon WorkSpaces Applications User Pools.

Streaming URL

To create a streaming URL, use one of the following methods:

- WorkSpaces Applications console
- The CreateStreamingURL API action
- The create-streaming-url AWS CLI command

To create a streaming URL by using the WorkSpaces Applications console, complete the steps in the following procedure.

To create a streaming URL by using the WorkSpaces Applications console

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the navigation pane, choose **Fleets**.
3. In the list of fleets, choose the fleet that is associated with the stack for which you want to create a streaming URL. Verify that the status of the fleet is **Running**.
4. In the navigation pane, choose **Stacks**. Choose the stack, and then choose **Actions**, **Create streaming URL**.
5. In **User id**, enter the user ID.
6. For **URL Expiration**, choose an expiration time, which determines how long the generated URL is valid. This URL is valid for a maximum of seven days.
7. Choose **Get URL**.
8. Copy the URL, save it to an accessible location, and then provide it to your users.

   In the WorkSpaces Applications client sign-in page, users enter the streaming URL that you created as the web address, and then choose **Connect**.

Next Steps

After you configure a client connection method, you can provide your users with the following step-by-step guidance to help them connect to WorkSpaces Applications and start a streaming session: Connect to WorkSpaces Applications on Windows Client or the section called "Connect to WorkSpaces Applications on macOS client".

Enable Users to Share a USB Device with an WorkSpaces Applications Streaming Session

Before users share their USB devices with an WorkSpaces Applications session, the USB devices must be qualified. Otherwise, when users start a streaming session, their USB device is not detected by WorkSpaces Applications and cannot be shared with the session. For more information, see Qualify USB Devices for Use with Streaming Applications.

> **Note**
>
> Sharing a USB device with a WorkSpaces Applications streaming session is not supported on the macOS client.

Redirect a Streaming Session from the Web Browser to the WorkSpaces Applications Client

You can configure WorkSpaces Applications to redirect a streaming session from a web browser to the WorkSpaces Applications client. That way, when your users sign in to WorkSpaces Applications and start a streaming session in their web browser, their session is redirected to the WorkSpaces Applications client. To do so, perform these steps.

1. Use the WorkSpaces Applications `CreateStreamingURL` API action to generate a streaming URL.
2. Add the following prefix for the custom WorkSpaces Applications client handler to the streaming URL: **`amazonappstream:`**

   Together, the prefix and streaming URL are formatted as follows:

   `amazonappstream:`*`base64encoded(streamingURL)`*

> **Note**
>
> When encoding the URL, make sure that the encoding is in UTF-8.
> Powershell sample to encode:
> `[Convert]::ToBase64String([System.Text.Encoding]::UTF8.GetBytes("Streamin`
> `IdpURL"))`

3. When users are redirected to the streaming URL, their browser detects that the link must be opened by the WorkSpaces Applications client.
4. Users are prompted to choose whether they want to start the streaming session by using the WorkSpaces Applications client.
5. After the prompt, either of the following occurs:

   - If the WorkSpaces Applications client is installed, the user can choose to continue the streaming session by using the WorkSpaces Applications client.
   - If the WorkSpaces Applications client is not installed, the browser behavior varies as follows:
     - Chrome — No message is displayed.
     - Firefox — A message states that the user needs a new app to open Amazon AppStream.
     - Microsoft Edge — No message is displayed.
     - Internet Explorer — A message notifies the user that the WorkSpaces Applications client is not installed.

       In this case, users can select the **Download AppStream Client** link to download the client. After they download the client, they can install it, and refresh their browser to start the streaming session by using the client.

Create a Windows desktop shortcut using the default browser

To create a Windows desktop shortcut using the default browser to launch the client, use the following sample Powershell script.

```
$StringToEncode = 'your URL string'

$encodedUrl =
 [Convert]::ToBase64String([System.Text.Encoding]::UTF8.GetBytes($StringToEncode))
```

```
$shortcutContent = "
[{000214A0-0000-0000-C000-000000000046}]
Prop3=19,0
[InternetShortcut]
IDList=
URL=amazonappstream:$encodedUrl
IconIndex=0
HotKey=0
IconFile=$env:USERPROFILE\AppData\Local\AppStreamClient\appstreamclient.exe
"

Set-Content -Path "$env:USERPROFILE\Desktop\AppStream 2.0 Client Launcher.url" -Value
 $shortcutContent
```

Enable File System Redirection for Your WorkSpaces Applications Users

WorkSpaces Applications file system redirection lets users who have the WorkSpaces Applications client installed access files on their local computer from within their streaming session. When you enable file system redirection, you can specify the list of local drives and folders that your users can choose to access. When users sign in to WorkSpaces Applications and start a streaming session, they can choose the drive or folder that they want to access from the list. Then they can share the drive or folder with WorkSpaces Applications. The drive or folder remains available for them to access during their streaming sessions. Users can stop sharing their local drives or folders at any time.

Note

File system redirection is currently not supported for Linux-based fleet instances, or when using the macOS client.

Note

If your image uses an WorkSpaces Applications agent released on or after September 25, 2025, the drive for file system redirection will be available as a removable disk instead of a network drive.

Topics

- Prerequisites for File System Redirection
- How to Enable File System Redirection
- Make Default Drives and Folders Available for Your Users to Share
- Provide Your WorkSpaces Applications Users with Guidance for Working with File System Redirection

Prerequisites for File System Redirection

To enable WorkSpaces Applications file redirection:

- You must use an image that uses a version of the WorkSpaces Applications agent released on or after August 8, 2019. For more information, see WorkSpaces Applications Agent Release Notes.
- Your users must have WorkSpaces Applications client version 1.0.480 or later installed. For more information, see WorkSpaces Applications Windows Client Release Notes.
- File upload and download must be enabled on the stack that your users access for streaming sessions. See the following procedure.

How to Enable File System Redirection

Perform the following steps to enable both file upload and download on the stack that your users access for streaming sessions.

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the left navigation pane, choose **Stacks**.
3. Choose the stack for which you want to enable file system redirection.
4. Choose the **User Settings** tab, and then expand the **Clipboard, file transfer, and local print permissions** section.
5. For **File transfer**, verify that **Upload and download** is selected. If not, choose **Edit**, and then choose **Upload and download**.
6. Choose **Update**.

Make Default Drives and Folders Available for Your Users to Share

By default, when you enable file direction for users of a stack, the following drives and folders are made available for those users to share in their streaming session:

- Drives:
  - All local hard disks (physical drives, such as the C Drive and D Drive)
  - All virtual drives (network and virtual drives such as mapped drive letters, Google Drive, and OneDrive)
  - All local USB drives
- Folders:
  - %USERPROFILE%\Desktop
  - %USERPROFILE%\Documents
  - %USERPROFILE%\Downloads

These drive and folder paths prepopulate the **Share your local drives and folders** dialog box. This dialog box is displayed when users sign in to WorkSpaces Applications, start a streaming session, and choose **Settings**, **Local Resources**, and **Local Drives and Folders**.

You can change or define your own default drive and folder paths by editing the registry. You can also use the administrative template file that is provided in the WorkSpaces Applications client Enterprise Deployment Tool. This template lets you configure the client by using Group Policy. For more information, see Install and Configure the WorkSpaces Applications Client.

When users access their shared local drives and folders during a streaming session, the corresponding paths appear with backslashes replaced by underscores. They are also suffixed with the name of the local computer and a drive letter. For example, for a user with the user name janedoe and a computer name of ExampleCorp-123456, the default Desktop, Documents, and Downloads folder paths appear as follows:

C_Users_janedoe_Desktop (\\ExampleCorp-123456) (F:)

C_Users_janedoe_Documents (\\ExampleCorp-123456) (G:)

C_Users_janedoe_Downloads (\\ExampleCorp-123456) (H:)

Provide Your WorkSpaces Applications Users with Guidance for Working with File System Redirection

To help your users understand how to work with file redirection during their streaming sessions, you can provide them with the information in Local File Access.

Enable Local Printer Redirection for Your WorkSpaces Applications Users

With local printer redirection, your WorkSpaces Applications users can redirect print jobs from their streaming application to a printer that is connected to their local computer, including any network printers that the users have mapped. You don't need a printer driver installed on the WorkSpaces Applications streaming instance to enable users to print documents during their streaming sessions.

> **Note**
>
> Enabling local printer redirection is currently not supported for Linux-based stacks.

Topics

- Prerequisites for Local Printer Redirection
- How to Enable Local Printer Redirection
- How to Disable Local Printer Redirection

Prerequisites for Local Printer Redirection

To ensure that your users can use local printer redirection, you must:

- Use an image that uses a version of the WorkSpaces Applications agent released on or after July 30, 2020. For more information, see WorkSpaces Applications Agent Release Notes.
- Ensure that your users have WorkSpaces Applications client version 1.1.179 or later installed. For more information, see WorkSpaces Applications Windows Client Release Notes.
- Ensure printer redirection is enabled on the stack that your users access for streaming sessions.

How to Enable Local Printer Redirection

By default, local printer redirection is enabled when the WorkSpaces Applications client is installed. However, if local printer redirection is not enabled on the stack that your users access for streaming

sessions, you can enable it in the WorkSpaces Applications console by performing the following steps.

To enable local printer redirection by using the WorkSpaces Applications console

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the left navigation pane, choose **Stacks**.
3. Choose the stack for which you want to enable local printer redirection.
4. Choose the **User Settings** tab, and then expand the **Clipboard, file transfer, print to local device, and authentication permissions** section.
5. For **Print to local device**, verify that **Enabled** is selected. If not, choose **Edit**, and then choose **Enabled**.
6. Choose **Update**.

Alternatively, you can enable local printer redirection by using the WorkSpaces Applications API, an AWS SDK, or the AWS Command Line Interface (AWS CLI).

How to Disable Local Printer Redirection

To disable local printer redirection follow these steps.

To disable local printer redirection

You can disable local printer redirection in any of the following ways:

- During client installation on managed devices. For more information, see Choose Whether to Disable Local Printer Redirection.
- By using the WorkSpaces Applications console to disable this option on an WorkSpaces Applications stack.
- By using the WorkSpaces Applications API, an AWS SDK, or the AWS Command Line Interface (AWS CLI) to disable this option on an WorkSpaces Applications stack.

Tagging Your Amazon WorkSpaces Applications Resources

AWS enables you to assign metadata to your AWS resources in the form of tags. You can use these tags to help manage your WorkSpaces Applications image builders, images, fleets, and stacks, and also organize data, including billing data.

You can:

- Logically group resources in different ways (for example, by purpose, owner, or environment).

  This is useful when you have many resources of the same type.
- Quickly identify a specific resource based on the tags that you've assigned to it
- Identify and control AWS costs

For example, you can identify and group WorkSpaces Applications fleets that are in different environments (such as Development or Production) or that are assigned to different business units (such as HR or Marketing). You can then track the associated AWS costs for these fleets on a detailed level. To do this, sign up to get your Amazon Web Services account bill with tag key values included. For more information about setting up a cost allocation report with tags, see Monthly Cost Allocation Report in the *AWS Billing User Guide*.

Contents

- Tagging Basics for Amazon WorkSpaces Applications
- Tag Restrictions for Amazon WorkSpaces Applications
- Adding Tags during Resource Creation in the Amazon WorkSpaces Applications Console
- Adding, Editing, and Deleting Tags for Existing Resources in the Amazon WorkSpaces Applications Console
- Working with Tags by Using the Amazon WorkSpaces Applications API, an AWS SDK, or AWS CLI

Tagging Basics for Amazon WorkSpaces Applications

Tags consist of a key-value pair, similar to other AWS services tags. To tag a resource, you specify a *key* and a *value* for each tag. A key can be a general category, such as "project", "owner", or

"environment", with specific associated values, and you can share the same key and value across multiple resources. You can tag an WorkSpaces Applications resource immediately after you create it or later on. If you delete a resource, the tags are removed from that resource on deletion. However, other WorkSpaces Applications and AWS resources that have the same tag key are not impacted.

You can edit tag keys and values, and you can remove tags from a resource at any time. You can set the value of a tag to an empty string, but you can't set the name of a tag to null. If you add a tag that has the same key as an existing tag on that resource, the new value overwrites the old value. If you delete a resource, any tags for the resource are also deleted.

> **Note**
>
> If you plan to set up a monthly cost allocation report to track AWS costs for WorkSpaces Applications resources, keep in mind that tags added to existing WorkSpaces Applications resources appear in your cost allocation report on the first of the following month for resources that are renewed in that month.

Tag Restrictions for Amazon WorkSpaces Applications

- The maximum number of tags per WorkSpaces Applications resource is 50.
- The maximum key length is 128 Unicode characters in UTF-8.
- The maximum value length is 256 Unicode characters in UTF-8.
- Tag keys and values are case-sensitive.
- Do not use the "aws:" prefix in your tag names or values because it is a system tag that is reserved for AWS use. You cannot edit or delete tag names or values with this prefix. Tags with this prefix do not count against your tags per resource limit.
- Generally allowed characters are: letters, numbers, and spaces representable in UTF-8, and the following special characters: + - = . _ : / @.
- Although you can share the same key and value across multiple resources, you cannot have duplicate keys on the same resource.
- You can add tags for resources during resource creation. You can also add, edit, and delete tags for resources that are already created.

Adding Tags during Resource Creation in the Amazon WorkSpaces Applications Console

When you create a resource in the WorkSpaces Applications console, you can add one or more tags to manage the resource. For more information, see the following topics:

- Image builders — Launch an Image Builder to Install and Configure Streaming Applications, step 4
- Images — Step 6: Finish Creating Your Image, step 1
- Fleets — Create a Fleet in Amazon WorkSpaces Applications, step 3
- Stacks — Create a Stack in Amazon WorkSpaces Applications, step 2

Adding, Editing, and Deleting Tags for Existing Resources in the Amazon WorkSpaces Applications Console

You can add, edit, and delete tags for existing resources by using the WorkSpaces Applications console.

To add, edit, or delete tags for an existing WorkSpaces Applications resource

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. From the navigation bar, select the Region that contains the resource for which you want to add, edit, or delete tags.
3. In the navigation pane, select the resource type. The resource type can be an image builder, image, fleet, or stack.
4. Select the resource from the resource list.
5. Choose **Tags**, **Add/Edit Tags**, and then do one or more of the following:
   - To add a tag, choose **Add Tag**, and type the key and value for each tag.
   - To edit a tag, modify the key and value for the tag as needed.
   - To delete a tag, choose the **Delete** icon (X) for the tag.

6. Choose **Save**.

Working with Tags by Using the Amazon WorkSpaces Applications API, an AWS SDK, or AWS CLI

If you're using the WorkSpaces Applications API, an AWS SDK, or the AWS Command Line Interface (AWS CLI), you can use the following WorkSpaces Applications operations with the `tags` parameter to add tags when you create new resources.

> **Note**
>
> You can use spaces in tag keys and values. To indicate a space when you use the AWS CLI, use "\s" (without the quotation marks).

Task	AWS CLI	API Operation
Add one or more tags for a new fleet	create-fleet	CreateFleet
Add one or more tags for a new image builder	create-imagebuilder	CreateImageBuilder
Add one or more tags for a new stack	create-stack	CreateStack

You can use the following WorkSpaces Applications operations to add, edit, remove, or list tags for existing resources:

Task	AWS CLI	API Operation
Add or overwrite one or more tags for a resource	tag-resource	TagResource
Remove one or more tags for a resource	untag-resource	UntagResource

Task	AWS CLI	API Operation
List one or more tags for a resource	list-tags-for-resource	ListTagsForResource

When you use the WorkSpaces Applications API, an AWS SDK, or AWS CLI actions to add, edit, remove, or list tags for an existing WorkSpaces Applications resource, specify the resource by using its Amazon Resource Name (ARN). An ARN uniquely identifies an AWS resource and uses the following general syntax.

```
arn:aws:appstream:region:account:resourceType/resourceName
```

region

The AWS Region in which the resource was created (for example, `us-east-1`).

account

The AWS account ID, with no hyphens (for example, `123456789012`).

resourceType

The type of resource. You can tag the following WorkSpaces Applications resource types: `image-builder`, `image`, `fleet`, and `stack`.

resourceName

The name of the resource.

For example, you can obtain the ARN for an WorkSpaces Applications fleet by using the AWS CLI describe-fleets command. Copy the following command.

```
aws appstream describe-fleets
```

For an environment that contains a single fleet named `TestFleet`, the ARN for this resource would appear in the JSON output similar to the following.

```
"Arn": "arn:aws:appstream:us-east-1:123456789012:fleet/TestFleet"
```

After you obtain the ARN for this resource, you can add two tags by using the tag-resource command:

```
aws appstream tag-resource --resource arn:awsappstream:us-east-1:123456789012:fleet/
TestFleet --tags Environment=Test,Department=IT
```

The first tag, `Environment=Test`, indicates that the fleet is in a test environment. The second tag, `Department=IT`, indicates that the fleet is in the IT department.

You can use the following command to list the two tags that you added to the fleet.

```
aws appstream list-tags-for-resource --resource arn:aws:appstream:us-
east-1:123456789012:fleet/TestFleet
```

For this example, the JSON output appears as follows:

```
{
    "Tags": {
       "Environment" : "Test",
       "Department" : "IT"
    }
}
```

Monitoring and Reporting

Monitoring and reporting are an important part of maintaining the reliability, availability, and performance of your Amazon WorkSpaces Applications streaming instances and providing your users with a responsive streaming experience.

Contents

- Monitoring Amazon WorkSpaces Applications Resources
- WorkSpaces Applications Usage Reports
- Logging Amazon WorkSpaces Applications API Calls with AWS CloudTrail

For more information, see the following:

- Creating custom logging and CloudWatch alerting in WorkSpaces Applications
- Getting started with your AWS Health Dashboard – Your account health
- Monitoring AWS Health events with EventBridge

Monitoring Amazon WorkSpaces Applications Resources

WorkSpaces Applications publishes metrics to Amazon CloudWatch to enabled detailed tracking and deep dive analysis. These statistics are recorded for an extended period so you can access historical information and gain a better perspective on how your fleets are performing. For more information, see the Amazon CloudWatch User Guide.

Contents

- Viewing Fleet Usage Using the Console
- Viewing instance and session metrics using the console
- WorkSpaces Applications Metrics and Dimensions
- Manage Multi-Session Fleet Instances
- Create custom CloudWatch dashboards using CloudFormation templates
- Monitor your WorkSpaces Applications health using the CloudWatch automatic dashboard

Viewing Fleet Usage Using the Console

You can monitor your Amazon WorkSpaces Applications fleet usage using the WorkSpaces Applications or CloudWatch console.

To view fleet usage in the WorkSpaces Applications console

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the left pane, choose **Fleets**.
3. Select a fleet and choose its **Fleet Usage** tab.
4. By default, the graph displays the following metrics:

   - `ActualCapacity`, `InUseCapacity`, `DesiredCapacity`, `AvailableCapacity`, `PendingCapacity`, and `CapacityUtilization` for single-session fleets.
   - `ActualUserSessionCapacity`, `ActiveUserSessionCapacity`, `AvailableUserSessionCapacity`, `DesiredUserSessionCapacity`, `PendingUserSessionCapacity`, `CapacityUtilization`, `DrainingCapacity`, `DrainModeActiveUserSessionCapacity`, and `DrainModeUnusedUserSessionCapacity` for multi-session fleets.

To view fleet usage in the CloudWatch console

1. Open the CloudWatch console at https://console.aws.amazon.com/cloudwatch/.
2. In the left pane, choose **Metrics**.
3. Choose the **AppStream** namespace and then choose **Fleet Metrics**.
4. Select the metrics to graph.

Viewing instance and session metrics using the console

You can monitor Amazon WorkSpaces Applications fleet instance and session metrics using the WorkSpaces Applications console or the CloudWatch console.

These metrics are collected at a 5-minute interval. After a new session is provisioned, the first metric data point appears within 5 minutes. Subsequent metric data points are available at every 5-minute interval.

Required permissions for in-console monitoring

To view instance and session performance metrics on the WorkSpaces Applications console, your IAM user or role must have the following CloudWatch permissions:

- `cloudwatch:GetMetricData` – Retrieves metric data points. This permission is required to display metric graphs (CPU, memory, GPU, frame rate, latency, and so on).
- `cloudwatch:ListMetrics` – Lists available metrics. This permission is required to discover and populate available metrics for the selected fleet or session.

Example IAM Policy Example

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "cloudwatch:GetMetricData",
                "cloudwatch:ListMetrics"
            ],
            "Resource": "*"
        }
    ]
}
```

> **Note**
>
> Without these permissions, the monitoring panels on the WorkSpaces Applications console display an "Insufficient permissions" error or show no data. The fleet and session details are still visible (governed by `appstream:*` permissions), but the metric graphs do not render.

To view instance and session in the WorkSpaces Applications console

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the left pane, choose **Fleets**.
3. Select a fleet and choose **View Details**.

4. View fleet utilization information under **Sessions on fleet**.
5. View the list of all active sessions under **Instances with sessions**.
6. Select a session to view the metrics.
7. You can sort and filter the table to find specific user sessions.

To view instance and session metrics in the CloudWatch console

1. Open the CloudWatch console at https://console.aws.amazon.com/cloudwatch/.
2. In the left pane, choose **Metrics**.
3. Choose the **AppStream** namespace and then choose **Fleet Instance Metrics** or **Fleet Session Metrics**.
4. Select the metrics to graph.

WorkSpaces Applications Metrics and Dimensions

Amazon WorkSpaces Applications sends the following metrics and dimension information to Amazon CloudWatch.

All of the following metrics except `InsufficientConcurrencyLimitError` apply to Always-On and On-Demand fleets. The only metrics that apply to Elastic fleets are `InUseCapacity` and `InsufficientCapacityError`.

WorkSpaces Applications sends fleet usage metrics to CloudWatch one time every minute, and instance and session performance metrics one time every 5 minutes. The `AWS/AppStream` namespace includes the following metrics.

Topics

- Fleet Usage Metrics for Single-session Fleets
- Fleet Usage Metrics for Multi-session Fleets
- Instance and Session Performance Metrics for Single-session and Multi-session Fleets
- Dimensions for Amazon WorkSpaces Applications Metrics

Fleet Usage Metrics for Single-session Fleets

The following are fleet usage metrics for single-session fleets.

Metric	Description	Dimensions	Statistics	Units
ActualCapacity	The total number of instances that are available for streaming or are currently streaming. `ActualCapacity = AvailableCapacity + InUseCapacity`	[Fleet]	Average, Minimum, Maximum	Count
AvailableCapacity	The number of idle instances currently available for user sessions. `AvailableCapacity = ActualCapacity - InUseCapacity`	[Fleet]	Average, Minimum, Maximum	Count
CapacityUtilization	The percentage of instances in a fleet that are being used, using the following formula. `CapacityUtilization = (InUseCapacity/ActualCapacity) * 100` Monitoring this metric helps with decisions about increasing or decreasing the value of a fleet's desired capacity.	[Fleet]	Average, Minimum, Maximum	Percent
DesiredCapacity	The total number of instances that are either running or pending. This represents the total number of concurrent streaming sessions your fleet can support in a steady state. `DesiredCapacity = ActualCapacity + PendingCapacity`	[Fleet]	Average, Minimum, Maximum	Count

Metric	Description	Dimensions	Statistics	Units
`InUseCapacity`	The number of instances currently being used for streaming sessions. One `InUseCapacity` count represents one streaming session.	[Fleet]	Average, Minimum, Maximum	Count
`PendingCapacity`	The number of instances being provisioned by WorkSpaces Applications. Represents the additional number of streaming sessions the fleet can support after provisioning is complete. When provisioning starts, it usually takes 10-20 minutes for an instance to become available for streaming.	[Fleet]	Average, Minimum, Maximum	Count
`RunningCapacity`	The total number of instances currently running. Represents the number of concurrent streaming sessions that can be supported by the fleet in its current state. This metric is provided for Always-On fleets only, and has the same value as the `ActualCapacity` metric.	[Fleet]	Average, Minimum, Maximum	Count
`InsufficientCapacityError`	The number of session requests rejected due to lack of capacity. You can set alarms to use this metric to be notified of users waiting for streaming sessions.	[Fleet]	Average, Minimum, Maximum, Sum	Count

Metric	Description	Dimensions	Statistics	Units
InsufficientConcurrencyLimitError	The number of Elastic fleet session requests rejected due to reaching max concurrent streaming capacity. You can set alarms to use this metric to be notified of users waiting for streaming sessions.	[Fleet]	Average, Minimum, Maximum, Sum	Count

Fleet Usage Metrics for Multi-session Fleets

The following are fleet usage metrics for multi-session fleets.

Metric	Description	Dimensions	Statistics	Units
CapacityUtilization	The percentage of sessions in a fleet that are being used, using the following formula. `CapacityUtilization = (ActiveUserSessionCapacity / ActualUserSessionCapacity) * 100` Monitoring this metric helps with decisions about increasing or decreasing the value of a fleet's desired capacity.	[Fleet]	Average, Minimum, Maximum	Percent
ActualUserSessionCapacity	The total number of session slots that are available for streaming or are currently streaming. `ActualUserSessionCapacity = AvailableUserSessionCapacity + ActiveUserSessionCapacity`	[Fleet]	Average, Minimum, Maximum	Count

Metric	Description	Dimensions	Statistics	Units
AvailableUserSessionCapacity	The number of idle session slots currently available for user sessions. `AvailableUserSessionCapacity = ActualUserSessionCapacity - ActiveUserSessions`	[Fleet]	Average, Minimum, Maximum	Count
DesiredUserSessionCapacity	The total number of session slots that are either running or pending. This represents the total number of concurrent streaming sessions your fleet can support in a steady state. `DesiredUserSessionCapacity = ActualUserSessionCapacity + PendingUserSessionCapacity`	[Fleet]	Average, Minimum, Maximum	Count
ActiveUserSessionCapacity	The number of user sessions currently being used for streaming sessions.	[Fleet]	Average, Minimum, Maximum	Count
PendingUserSessionCapacity	The number of session slots being provisioned by WorkSpaces Applications. Represents the additional number of streaming sessions the fleet can support after provisioning is complete. When provisioning starts, it usually takes 10-20 minutes for an instance to become available for streaming.	[Fleet]	Average, Minimum, Maximum	Count

Metric	Description	Dimensions	Statistics	Units
RunningUserSessionCapacity	The total number of session slots currently that are available for streaming or are currently streaming . Represents the number of concurrent streaming sessions that can be supported by the fleet in its current state. This metric is provided for Always-On fleets only, and has the same value as the ActualUserSessionCapacity metric.	[Fleet]	Average, Minimum, Maximum	Count
DrainingCapacity	The number of instances in Drain Mode.	[Fleet]	Average, Minimum, Maximum	Count
DrainModeActiveUserSessionCapacity	The number of active user sessions on instances in Drain Mode.	[Fleet]	Average, Minimum, Maximum	Count
DrainModeUnusedUserSessionCapacity	The number of unused session slots on instances in Drain Mode which cannot be used for user session provisioning.	[Fleet]	Average, Minimum, Maximum	Count

Metric	Description	Dimensions	Statistics	Units

`ActualCapacity`	The total number of instances that are available for streaming or are currently streaming. For multi-session fleets, this metric represents instance-level capacity. The relationship between instance-level and session-level capacity is as follows:
	`ActualUserSessionCapacity = ActualCapacity × (Maximum number of user sessions per instance)`
	`ActualCapacity = ActualUserSessionCapacity / (Maximum number of user sessions per instance)`
		[Fleet]	Average, Minimum, Maximum	Count

Metric	Description	Dimensions	Statistics	Units

DesiredCapacity	The total number of instances that are either running or pending. This represents the total number of instances your fleet targets in a steady state. For multi-session fleets, the relationship between instance-level and session-level capacity is as follows:
	`DesiredUserSessionCapacity = DesiredCapacity × (Maximum number of user sessions per instance)`
	`DesiredCapacity = DesiredUserSessionCapacity / (Maximum number of user sessions per instance)`
		[Fleet]	Average, Minimum, Maximum	Count

InUseCapacity	The number of instances currently being used for streaming sessions. For multi-session fleets, an instance is considered in use if it has at least one active user session. The relationship between instance-level and session-level capacity is as follows:
	`ActiveUserSessionCapacity = InUseCapacity × (Maximum number of user sessions per instance)`
	`InUseCapacity = ActiveUserSessionCapacity / (Maximum number of user sessions per instance)`
		[Fleet]	Average, Minimum, Maximum	Count

Metric	Description	Dimensions	Statistics	Units

PendingCapacity	The number of instances being provisioned by WorkSpaces Applications. Represents the additional number of instances the fleet will have after provisioning is complete. When provisioning starts, it usually takes 10-20 minutes for an instance to become available for streaming. For multi-session fleets, the relationship between instance-level and session-level capacity is as follows:
	`PendingUserSessionCapacity = PendingCapacity × (Maximum number of user sessions per instance)`
	`PendingCapacity = PendingUserSessionCapacity / (Maximum number of user sessions per instance)`
		[Fleet]	Average, Minimum, Maximum	Count

Metric	Description	Dimensions	Statistics	Units

RunningCapacity	The total number of instances currently running. Represents the number of instances that can support streaming sessions in the fleet's current state. For multi-session fleets, the relationship between instance-level and session-level capacity is as follows:
	`RunningUserSessionCapacity = RunningCapacity × (Maximum number of user sessions per instance)`
	`RunningCapacity = RunningUserSessionCapacity / (Maximum number of user sessions per instance)`
	This metric is provided for Always-On fleets only, and has the same value as the ActualCapacity metric.
		[Fleet]	Average, Minimum, Maximum	Count

InsufficientCapacityError	The number of session requests rejected due to lack of capacity. This metric applies the same way for both single-session and multi-session fleets. You can set alarms to use this metric to be notified of users waiting for streaming sessions.	[Fleet]	Average, Minimum, Maximum, Sum	Count

Instance and Session Performance Metrics for Single-session and Multi-session Fleets

The following are instance and session performance metrics for single-session and multi-session fleets.

Important

All metrics are available for the Windows operating system. For the remaining set of supported operating systems, only the following metrics are available: `UDPPacketLossRate`, `TCPRetransmissionRate`, `BandwidthInbound`, `CongestionWindow`, `ConnectionDuration`, `MetadataNoToken`, and `MetadataNoTokenRejected`.

Additionally, some metrics may not be available depending on the software versions used for your images and client. We recommend always using the latest available image and client versions to ensure full metric availability.

For images, if you are using bundle-based images (Windows only), select the LATEST bundle or a version dated 2025-02-07 or later. If you are using managed image updates, Windows requires version 2025-02-07 or later, and RHEL and Rocky Linux require version 2025-09-05 or later. For a complete list of available managed image versions, refer to the base image version history. If you are using custom images with the WorkSpaces Applications agent, ensure your agent version is current. For details on agent versions and update options, refer to the agent software versions documentation.

For clients, the Mac native client requires version 1.2.0 or later, and the Windows native client requires version 1.2.1581 or later. The web client always delivers the latest version automatically, so no action is required.

If your image or client version does not meet these minimum requirements, some metrics will not be reported. Updating to the latest versions will enable full metric availability.

Metric	Description	Dimensions	Statistics	Units
`CpuUtilizationInstance`	The percentage of allocated compute units that are currently in use on the instance.	[Fleet] [UserId] [FleetName, InstanceId]	Average, Minimum, Maximum	Percent

Metric	Description	Dimensions	Statistics	Units
		[FleetName, InstanceId, SessionId, UserId]		
MemoryUtilizationInstance	The percentage of allocated physical memory units that are currently in use on the instance.	[Fleet] [UserId] [FleetName, InstanceId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Percent
PagingFileUtilizationInstance	The percentage of the paging file that is currently in use to extend the memory (RAM) capacity.	[Fleet] [UserId] [FleetName, InstanceId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Percent
DiskUtilizationInstance	The percentage of the instance's available disk storage capacity that is currently in use.	[Fleet] [UserId] [FleetName, InstanceId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Percent

Metric	Description	Dimensions	Statistics	Units
GpuUtilizationInstance	The percentage of the GPU resource used on the instance.	[Fleet] [UserId] [FleetName, InstanceId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Percent
CpuQueueLength	The number of threads waiting for CPU time on the instance.	[Fleet] [UserId] [FleetName, InstanceId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Count
DiskIoQueueLength	The number of pending I/O requests waiting to be processed by the disk on the instance.	[Fleet] [UserId] [FleetName, InstanceId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Count
MemoryPageHardFaults	The rate of hard page faults on the instance.	[Fleet] [UserId] [FleetName, InstanceId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Count/Second

Metric	Description	Dimensions	Statistics	Units
CpuUtilizationSession	The percentage of allocated compute units that are currently in use by the session.	[Fleet] [UserId] [FleetName, InstanceId, SessionId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Percent
MemoryUtilizationSession	The percentage of allocated physical memory units that are currently in use by the session.	[Fleet] [UserId] [FleetName, InstanceId, SessionId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Percent
DiskReadOperations	The number of disk read operations per instance.	[Fleet] [UserId] [FleetName, InstanceId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Count
DiskWriteOperations	The number of disk write operations per instance.	[Fleet] [UserId] [FleetName, InstanceId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Count

Metric	Description	Dimensions	Statistics	Units
InSessionLatency	The round-trip time between the client and the instance.	[Fleet] [UserId] [FleetName, InstanceId, SessionId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Milliseconds
UDPPacketLossRate	The percentage of UDP packets lost in traffic from the gateway to the client.	[Fleet] [UserId] [FleetName, InstanceId, SessionId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Percent
TCPRetransmissionRate	The percentage of TCP segments that were retransmitted from the gateway to the client.	[Fleet] [UserId] [FleetName, InstanceId, SessionId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Percent
Bandwidth	The rate of data transferred from the gateway to the client (outbound).	[Fleet] [UserId] [FleetName, InstanceId, SessionId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Kilobits/Second

Metric	Description	Dimensions	Statistics	Units
BandwidthInbound	The rate of data transferred from the client to the gateway (inbound).	[Fleet] [UserId] [FleetName, InstanceId, SessionId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Kilobits/Second
CongestionWindow	The size of the congestion window at the gateway for traffic flowing to the client.	[Fleet] [UserId] [FleetName, InstanceId, SessionId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Bytes
ConnectionDuration	The duration of the streaming connection.	[Fleet] [UserId] [FleetName, InstanceId, SessionId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Seconds
FramesPerSecond	The number of frames sent per second from the instance to the client.	[Fleet] [UserId] [FleetName, InstanceId, SessionId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Count

Metric	Description	Dimensions	Statistics	Units
MetadataNoToken	The number of times the instance metadata service was accessed without a token (IMDSv1). This metric helps identify workloads that have not yet migrated to IMDSv2.	[Fleet] [ImageBuilder] [AppBlockBuilder]	Sum	Count
MetadataNoTokenRejected	The number of times an IMDSv1 request to the instance metadata service was rejected. This metric is available when the instance is configured to require IMDSv2.	[Fleet] [ImageBuilder] [AppBlockBuilder]	Sum	Count
SessionDisconnects	The number of times a streaming session was disconnected.	[Fleet] [UserId] [FleetName, InstanceId, SessionId] [FleetName, InstanceId, SessionId, UserId]	Sum	Count
SessionLaunchLatency	The time taken for the session to launch, measured from the initial connection request to the first frame rendered on the client.	[Fleet] [UserId] [FleetName, InstanceId, SessionId] [FleetName, InstanceId, SessionId, UserId]	Average, Minimum, Maximum	Milliseconds

Metric	Description	Dimensions	Statistics	Units
ConnectionAttempts	The number of connection attempts to the instance.	[Fleet] [UserId] [FleetName, InstanceId, SessionId] [FleetName, InstanceId, SessionId, UserId]	Sum	Count
ConnectionSuccesses	The number of successful connections to the instance.	[Fleet] [UserId] [FleetName, InstanceId, SessionId] [FleetName, InstanceId, SessionId, UserId]	Sum	Count
ConnectionFailures	The number of failed connections to the instance.	[Fleet] [UserId] [FleetName, InstanceId, SessionId] [FleetName, InstanceId, SessionId, UserId]	Sum	Count

Dimensions for Amazon WorkSpaces Applications Metrics

The AWS/AppStream namespace includes the following dimensions and dimension groups.

Dimension	Description
Fleet	Filters the metric data by name of the Fleet.

Dimension	Description
FleetName	Filters the metric data by name of the Fleet.
SessionId	Filters the metric data by session identifier.
InstanceId	Filters the metric data by instance identifier.
UserId	Filters the metric data by user identifier.
ImageBuilder	Filters the metric data by name of the image builder.
AppBlockBuilder	Filters the metric data by name of the app block builder.

Dimension	Where Available in Amazon CloudWatch Metrics
[Fleet]	Fleet Metrics
[FleetName, InstanceId]	Fleet Instance Metrics
[FleetName, InstanceId, SessionId]	Fleet Session Metrics
[UserId]	UserId
[FleetName, InstanceId, SessionId, UserId]	FleetName, InstanceId, SessionId, UserId
[ImageBuilder]	Image Builder Metrics
[AppBlockBuilder]	App Block Builder Metrics

Manage Multi-Session Fleet Instances

When running multi-session fleets on Amazon WorkSpaces Applications, multiple user sessions share a fleet instance. Unlike single-session fleets — where each instance serves one user and is

automatically recycled after the session is terminated — multi-session instances can remain active for extended periods as long as users continue to occupy sessions on them.

This creates an important operational consideration: a multi-session instance cannot be terminated while active user sessions are running. If an instance is never fully vacated, it may never be reclaimed or rebooted, which means:

- **Image updates may not propagate** — new image versions containing patches, software updates, or security fixes will not reach an instance until all sessions on it have ended and the instance is recycled.
- **Fleet hygiene is harder to maintain** — without a mechanism to gracefully cycle instances, administrators must manually restart fleets or terminate all active sessions to apply updates, which can be disruptive to end users.
- **Long-running instances may experience performance degradation** — instances that remain active for extended periods without a restart can accumulate resource contention, memory pressure, and other system-level issues over time, leading to a degraded experience for end users sharing that instance.

To address this, WorkSpaces Applications provides a method to manage the lifecycle of individual fleet instances through a process called draining (also referred to as Drain Mode).

Putting an instance in Drain Mode impacts the capacity of the fleet. The fleet capacity metrics are available at the section called “Viewing Fleet Usage Using the Console”.

Drain a Fleet Instance

Draining an instance means instructing the fleet to stop accepting new user sessions while allowing existing sessions to continue until users naturally disconnect. Once all sessions have ended, the instance is automatically reclaimed and replaced with a fresh instance running the latest image.

This approach lets you:

- **Apply image updates without disrupting users** — existing sessions run to completion, while new sessions are routed to other instances.
- **Manage maximum instance lifetime** — by periodically draining older instances, you ensure that no instance runs indefinitely without being refreshed.
- **Maintain fleet hygiene proactively** — rather than waiting for a full fleet restart, you can cycle individual instances on a rolling basis.

Note

Draining an instance does not disconnect or terminate active user sessions. Users already connected will not be affected.

Important

Once an instance is placed in Drain Mode, this action cannot be reversed. The instance will continue to drain until all active sessions have ended and it is reclaimed.

When to Drain an Instance

Consider draining a fleet instance when:

- A new image version has been published and you want to roll it out without a full fleet restart.
- An instance has been running for an extended period and you want to refresh it.
- You observe performance degradation on sessions running on a specific instance and want to recycle it without terminating active sessions.

How to Drain an Instance

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the left pane, choose **Fleets**.
3. Select a multi-session fleet and choose **View Details** and **View Sessions**.
4. Select a session running on the instance you want to drain.
5. Choose **Drain Session Instance**.

The instance will immediately stop accepting new sessions. Existing sessions will continue uninterrupted. Once the last session ends, the instance is automatically terminated and replaced.

Create custom CloudWatch dashboards using CloudFormation templates

AWS provides CloudFormation templates that you can use to create custom CloudWatch dashboards for WorkSpaces Applications. Choose from the following CloudFormation template options to create custom dashboards for your WorkSpaces Applications workloads in the CloudFormation console.

Considerations before getting started

Consider the following before you get started with custom CloudWatch dashboards:

- Create your dashboards in the same AWS Region as the deployed WorkSpaces you want to monitor.
- You can also create custom dashboards using the CloudWatch console.
- A cost might be associated with custom CloudWatch dashboards. For information about pricing, see Amazon CloudWatch Pricing.

Prerequisites

Before you begin, ensure you have access to the AWS Management Console and appropriate permissions to create CloudFormation stacks and CloudWatch dashboards.

Procedure

This procedure covers three CloudFormation templates for monitoring WorkSpaces Applications fleets: the Fleet Dashboard for fleet-level health metrics with drill-down capabilities, the User Dashboard for individual user session experience troubleshooting, and the Fleet Performance Alerts template for automated email notifications when performance thresholds are breached.

Fleet Dashboard Template

The Fleet Dashboard gives administrators a drill-down view of a specific fleet's health. A Fleet dropdown at the top auto-populates with all fleets in the account, shows metrics for each fleet and instances associated to that fleet.

Steps to create:

1. Open the Create Stack page in the CloudFormation console. This link opens the page with the Amazon S3 bucket location of the custom CloudWatch dashboard template pre-populated.

2. Review the default selections on the **Create stack** page, and choose **Next**.
3. On the **Specify stack details** page, do the following:

    a. In the **Stack name** field, use the default name or enter a custom name for your stack.
    b. In the **DashboardName** field, use the default name or enter a custom name for your dashboard.
    c. Choose **Next**.
4. On the **Configure stack options** page, review the default selections, and choose **Next**.
5. On the **Review** page, scroll down to the **Transforms** section and select the acknowledgement checkboxes to confirm that CloudFormation might create IAM resources.
6. Choose **Submit** to create the stack and the custom CloudWatch dashboard.
7. After the stack creation is complete, open the CloudWatch console to view your newly created custom dashboard.

    a. Open the CloudWatch console at https://console.aws.amazon.com/cloudwatch/.
    b. In the left navigation bar, choose **Dashboards**.
    c. Under **Custom Dashboards**, choose the dashboard with the dashboard name you entered earlier in this procedure.

User Dashboard Template

The User Dashboard allows administrators to troubleshoot individual user experience. A UserId dropdown at the top auto-populates with all user IDs from CloudWatch metrics, shows metrics for each User.

Steps to create:

1. Open the Create Stack page in the CloudFormation console. This link opens the page with the Amazon S3 bucket location of the custom CloudWatch dashboard template pre-populated.
2. Review the default selections on the **Create stack** page, and choose **Next**.
3. On the **Specify stack details** page, do the following:

    a. In the **Stack name** field, use the default name or enter a custom name for your stack.
    b. In the **DashboardName** field, use the default name or enter a custom name for your dashboard.
    c. Choose **Next**.

4. On the **Configure stack options** page, review the default selections, and choose **Next**.
5. On the **Review** page, scroll down to the **Transforms** section and select the acknowledgement checkboxes to confirm that CloudFormation might create IAM resources.
6. Choose **Submit** to create the stack and the custom CloudWatch dashboard.
7. After the stack creation is complete, open the CloudWatch console to view your newly created custom dashboard.

    a. Open the CloudWatch console at https://console.aws.amazon.com/cloudwatch/.

    b. In the left navigation bar, choose **Dashboards**.

    c. Under **Custom Dashboards**, choose the dashboard with the dashboard name you entered earlier in this procedure.

Fleet Performance Alerts Template

This template creates 6 CloudWatch alarms for a specific fleet and an Amazon SNS topic that sends email notifications when alarms trigger or recover. After deployment, the administrator will receive a confirmation email from Amazon SNS — they must confirm the subscription to start receiving alerts.

Steps to create:

1. Open the Create Stack page in the CloudFormation console. This link opens the page with the Amazon S3 bucket location of the custom CloudWatch dashboard template pre-populated.
2. Review the default selections on the **Create stack** page, and choose **Next**.
3. On the **Specify stack details** page, do the following:

    a. In the **Stack name** field, use the default name or enter a custom name for your stack.

    b. In the **FleetName** field, enter the Fleet name that you want alerting for.

    c. In the **NotificationEmail** field, enter an email address through which you want to get notified.

    d. Choose **Next**.
4. On the **Configure stack options** page, review the default selections, and choose **Next**.
5. On the **Review** page, scroll down to the **Transforms** section and select the acknowledgement checkboxes to confirm that CloudFormation might create IAM resources.
6. Choose **Submit** to create the stack and the custom CloudWatch dashboard.

7. After the stack creation is complete, open the CloudWatch console to view your newly created custom dashboard.

   a. Open the CloudWatch console at https://console.aws.amazon.com/cloudwatch/.

   b. In the left navigation bar, choose **Dashboards**.

   c. Under **Custom Dashboards**, choose the dashboard with the dashboard name you entered earlier in this procedure.

Next steps

After creating your dashboard, you can customize it further by adding or removing widgets, adjusting time ranges, and configuring alarms as needed.

Monitor your WorkSpaces Applications health using the CloudWatch automatic dashboard

You can monitor WorkSpaces Applications using the CloudWatch automatic dashboard, which collects raw data and processes it into readable, near real-time metrics. The metrics are kept for 15 months to access historical information and to monitor the performance of your fleets, instances, and sessions. You can also set alarms that watch for certain thresholds, and send notifications or take actions when those thresholds are met. For more information, see the Amazon CloudWatch User Guide.

The CloudWatch automatic dashboard is created when you use your AWS account to configure your WorkSpaces Applications. The dashboard allows you to monitor your WorkSpaces Applications metrics across Regions, including:

- **Capacity metrics** — Monitor overall fleet capacity status, including capacity utilization, actual capacity, available capacity, desired capacity, in-use capacity, pending capacity, and running capacity. Track insufficient capacity errors to identify scaling issues.
- **Performance metrics** — View fleet, instance, and session performance data including CPU utilization, memory utilization, disk utilization, paging file utilization, and disk I/O operations.
- **Session quality metrics** — Monitor in-session latency, frames per second (FPS), and bandwidth to assess streaming performance and user experience.

You can use the dashboard for the following purposes:

- Identify fleets experiencing capacity constraints or insufficient capacity errors.
- Monitor resource utilization (CPU, memory, disk) over time to optimize instance sizing.
- Track session quality metrics to identify performance degradation or network issues.
- Identify anomalies to help with troubleshooting streaming or connectivity problems.
- View metrics by fleet to compare performance across different configurations.

The automatic dashboard provides time range filters (1 hour, 3 hours, 12 hours, 1 day, 1 week, or custom) and displays metrics in UTC timezone by default.

To use the WorkSpaces Applications CloudWatch automatic dashboard

1. Open the CloudWatch console at https://console.aws.amazon.com/cloudwatch/.
2. In the navigation pane, choose **Dashboards**.
3. Choose the **Automatic dashboards** tab.
4. Choose **AppStream**.

WorkSpaces Applications Usage Reports

You can subscribe to Amazon WorkSpaces Applications usage reports to receive detailed reports about how your users are using the service. Two .csv files are exported to an Amazon Simple Storage Service (Amazon S3) bucket in your account every day.

> **Note**
>
> To enable WorkSpaces Applications usage reports, you must use an image that uses a version of the WorkSpaces Applications agent released on or after May 7, 2019.

Contents

- Enable WorkSpaces Applications Usage Reports
- WorkSpaces Applications Usage Reports Fields
- Create Custom Reports and Analyze WorkSpaces Applications Usage Data

Enable WorkSpaces Applications Usage Reports

To receive usage reports, you subscribe to them by using the WorkSpaces Applications console, the AWS Command Line Interface (AWS CLI), or the `CreateUsageReportSubscription` API operation. You must enable usage reports separately for each AWS Region for which you want to receive usage data.

> **Note**
>
> You can start or stop your subscription to usage reports at any time. There is no charge for subscribing to usage reports, but standard Amazon S3 charges may apply to reports that are stored in your S3 bucket. For more information, see Amazon S3 Pricing.

To subscribe to usage reports for WorkSpaces Applications by using the WorkSpaces Applications console, perform the following steps.

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. Choose the AWS Region for which you want to enable usage reports.
3. In the navigation pane, choose **Usage Reports**.
4. Choose **Enabled**, and then choose **Apply**.

If you enabled on-instance session scripts and Amazon S3 logging for your session script configuration, WorkSpaces Applications created an S3 bucket to store the script output. The bucket is unique to your account and Region. When you enable usage reporting in this case, WorkSpaces Applications uses the same bucket to store your usage reports. If you haven't already enabled on-instance session scripts, when you enable usage reports, WorkSpaces Applications creates a new S3 bucket in the following location:

```
appstream-logs-region-code-account-id-without-hyphens-random-identifier
```

region-code

The AWS Region code for the Region in which usage reporting is enabled.

account-id-without-hyphens

Your Amazon Web Services account identifier. The random ID ensures that there is no conflict with other buckets in the same Region. The first part of the bucket name, `appstream-logs`, does not change across accounts or Regions.

For example, if you enable usage reporting in the US West (Oregon) Region (us-west-2) on account number 123456789012, WorkSpaces Applications creates an Amazon S3 bucket within your account in that Region similar to the name shown in the following example:

```
appstream-logs-us-west-2-1234567890123-abcdefg
```

Only an administrator with sufficient permissions can delete this bucket.

Topics

- WorkSpaces Applications Sessions Reports
- WorkSpaces Applications Applications Reports

WorkSpaces Applications Sessions Reports

For each day that users launch at least one streaming session in your Amazon Web Services account, WorkSpaces Applications exports a sessions report to your Amazon S3 bucket. The report, named **daily-session-report-[YYYY]-[MM]-[DD].csv**, is stored in a nested folder structure in your Amazon S3 account, using the following folder path:

[bucket_name]/sessions/schedule=DAILY/year=[YYYY]/month=[MM]/day=[DD]/

This nesting structure facilitates partitioning if you choose to query your reports by using Amazon Athena. Athena is a serverless, interactive query service that you can use to analyze data stored in your S3 buckets using standard SQL. For more information, see Create Custom Reports and Analyze WorkSpaces Applications Usage Data.

Each user session is described in a single record in a sessions report. Sessions reports are generated daily according to UTC time within 24 hours of the close of the day that is the subject of the report. If a session spans more than one day, the session record appears in the sessions report corresponding to the day in which the session ends. For information about the data included in sessions reports, see Sessions Report Fields.

WorkSpaces Applications Applications Reports

For each day that users launch at least one application during their streaming sessions, WorkSpaces Applications exports an applications report to your Amazon S3 bucket. The report, named **daily-app-report-[YYYY]-[MM]-[DD].csv**, is stored in a nested folder structure in your Amazon S3 account, using the following folder path:

[bucket_name]/applications/schedule=DAILY/year=[YYYY]/month=[MM]/day=[DD]/

This nesting structure facilitates partitioning if you choose to query your reports by using Amazon Athena. Athena is a serverless, interactive query service that you can use to analyze data stored in your S3 buckets using standard SQL. For more information, see Create Custom Reports and Analyze WorkSpaces Applications Usage Data.

Each application launch is described in a single record in an applications report. For example, if a user launches five separate applications during a session, five separate records appear in the relevant applications report. An application is recorded as launched if any of the following events occurs:

- The application is launched directly when the session begins, because the application ID is embedded in either the streaming URL or the relay state.
- A user chooses the application from the application catalog when launching a new streaming session.
- A user chooses the application from the application catalog list during a streaming session.

The applications report doesn't include applications that are launched in other ways. For example, if you provide users with access to Windows Explorer, PowerShell, or the Windows desktop **Start** menu, and users use those tools to launch applications directly, or if another program or script launches an application, those application launches are not included in the applications report.

Applications reports are generated daily according to UTC time within 24 hours of the close of the day that is the subject of the report. If a session spans more than one day, applications launched during the session are reflected in the applications report corresponding to the day in which the session ends. For information about the data included in applications reports, see Applications Report Fields.

WorkSpaces Applications Usage Reports Fields

This topic provides information about the fields included in WorkSpaces Applications usage reports.

Contents

- Sessions Report Fields
- Applications Report Fields

Sessions Report Fields

The following table describes the fields included in WorkSpaces Applications sessions reports.

Field name	Description
`user_session_id`	The unique identifier (ID) for the session.
`aws_account_id`	The Amazon Web Services account ID.
`region`	The AWS Region.
`session_start_time`	The date and time that the session started. Must be specified in ISO 8601 format and as UTC.
`session_end_time`	The date and time that the session ended. Must be specified in ISO 8601 format and as UTC.
`session_duration_in_seconds`	The duration of the session in seconds.

Field name	Description
user_id	The unique ID for the user within the authentication type.
user_arn	The Amazon Resource Name (ARN) for the user.
authentication_type	The method used to authenticate the user. Possible values: CUSTOM | SAML | USERPOOL
authentication_type_user_id	The concatenation of the user ID and authentication type, which uniquely identifies the user for the purpose of assessing user fees. For more information, see WorkSpaces Applications Pricing.
fleet_name	The name of the fleet associated with the session.
stack_name	The name of the stack associated with the session.
instance_type	The WorkSpaces Applications instance type used for the session. For a list of instance types, see WorkSpaces Applications Pricing.

Field name	Description
`eni_private_ip_address`	The IP address of the elastic network interface used by the WorkSpaces Applications instance for network communications.
`connected_at_least_once`	Indicates whether the user connected to the session at least once. Possible values: `true` | `false`
`client_ip_addresses`	The IP addresses associated with the user device or devices used to connect to the session. If the user connected and then disconnected from the session more than once, up to the last 10 distinct IP addresses are stored, separated by semicolons.
`google_drive_enabled`	Indicates whether Google Drive was enabled as a persistent storage option for the session. For more information, see Enable and Administer Google Drive for Your WorkSpaces Applications Users. Possible values: `true` | `false`

Field name	Description
one_drive_enabled	Indicates whether OneDrive was enabled as a persisten t storage option for the session. For more informati on, see Enable and Administe r Google Drive for Your WorkSpaces Applications Users. Possible values: true | false
home_folders_storage_location	The Amazon S3 bucket used for files that are stored using home folders.
user_settings_clipboard_copy_from_lo cal_device	Indicates whether the user was able to copy data from the local device to the streaming session using the clipboard during the session. Possible values: ENABLED | DISABLED
user_settings_clipboard_copy_to_loca l_device	Indicates whether the user was able to copy data from the streaming session to the local device using the clipboard during the session. Possible values: ENABLED | DISABLED

Field name	Description
user_settings_file_upload	Indicates whether the user was able to upload files from the local device to the streaming session during the session. Possible values: ENABLED | DISABLED
user_settings_file_download	Indicates whether the user was able to download files from the streaming session to the local device during the session. Possible values: ENABLED | DISABLED
user_settings_printing_to_local_device	Indicates whether the user was able to print files from the streaming session to the local device during the session. Possible values: ENABLED | DISABLED
application_settings_enabled	Indicates whether applicati on settings persistence was enabled for the session. Possible values: true | false

Field name	Description
domain_joined	Indicates whether the WorkSpaces Applications streaming instance was joined to an Active Directory domain at session launch. For more information, see Using Active Directory with WorkSpaces Applications. Possible values: Y | N
max_session_duration	The maximum allowed duration of the session, in seconds.
session_type	The session type. Possible values: ALWAYS_ON | ON_DEMAND
stream_view	The stream view. Possible values: APPLICATION | DESKTOP
streaming_experience_settings_protocol	The protocol that the session ended streaming with. Possible values: UDP | TCP
instance_id	The instance ID associated with the user session.

Field name	Description
`is_multisession`	Indicates whether the session belongs to a multi-session fleet. Possible values: `true` | `false`

Applications Report Fields

The following table describes the fields included in WorkSpaces Applications applications reports.

Field name	Description
user_session_id	The unique identifier (ID) for the session.
application_name	The name of the applicati on, as specified in Image Assistant. This value is provided when a user launches an application through the WorkSpaces Applications interface.
schedule	The frequency with which reports are generated. Possible value: DAILY
year	The year of the report.
month	The month of the report.
day	The day of the report.

Create Custom Reports and Analyze WorkSpaces Applications Usage Data

Amazon Athena is a serverless, interactive query service that you can use to analyze data stored in your S3 buckets using standard SQL queries. You can use Athena to aggregate your usage reports or generate other types of custom reports.

Contents

- Create an AWS Glue Crawler
- Create a Data Catalog by Using the AWS Glue Crawler
- Create and Run Athena Queries
- Working with Athena Queries

Create an AWS Glue Crawler

AWS Glue is a fully managed extract, transform, and load (ETL) service that lets you create a database from your Amazon S3 data and query that database by using Athena. This database is also referred to as an AWS Glue Data Catalog. An AWS Glue crawler can automatically detect the schema of your Amazon S3 data and create the corresponding database and tables. WorkSpaces Applications provides an CloudFormation template that you can use to create the necessary AWS Glue resources.

> **Important**
>
> Completing the steps in the following procedure creates an AWS Glue crawler. However, these steps don't start the crawler. To start the crawler, you must perform the steps in the next procedure. For more information about AWS Glue crawlers, see Defining Crawlers.

To create an AWS Glue crawler

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. Choose the AWS Region for which you have subscribed to usage reports.
3. In the navigation pane, choose **Usage Reports**, and verify that usage reports logging is enabled.

4. On the **Report Details** tab, in the paragraph next to **Analytics**, choose the **CloudFormation template** link.

   Choosing the link opens the CloudFormation console, where you can review the parameters of the CloudFormation stack specified by the template before you run it. The template, when run, creates an AWS Glue crawler and several sample Athena queries.

5. On the **Specify Details** page, next to **ScheduleExpression**, either keep the default value or specify a different cron expression value for the frequency that you want to run the crawler. Do not change any other default value. When you're done, choose **Next**.

   By default, the crawler is scheduled to run on a daily basis, but you can configure the crawler to run weekly, monthly, or on another frequency. For information about cron syntax, see Cron Expressions.

6. On the **Options** page, keep the default values, and choose **Next**.

7. On the **Review** page, select the check box next to "I acknowledge that CloudFormation might create IAM resources with custom names," and then choose **Create**.

   You must have sufficient AWS Glue and AWS Identity and Access Management (IAM) permissions to create and run the CloudFormation stack. If you don't have the required permissions, ask your Amazon Web Services account administrator either to perform these steps in your account or to grant you the following permissions.

   JSON

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "athena:CreateNamedQuery",
                "athena:BatchGetNamedQuery",
                "athena:GetNamedQuery",
                "athena:StartQueryExecution",
                "athena:GetQueryResults",
                "athena:GetQueryExecution",
                "athena:ListNamedQueries",
                "cloudformation:DescribeStacks",
                "cloudformation:GetStackPolicy",
```

```
                "cloudformation:DescribeStackEvents",
                "cloudformation:CreateStack",
                "cloudformation:GetTemplate",
                "cloudformation:ListChangeSets",
                "cloudformation:ListStackResources",
                "iam:GetRole",
                "iam:CreateRole",
                "iam:GetRolePolicy",
                "s3:GetBucketLocation",
                "s3:ListBucketMultipartUploads",
                "s3:ListBucket",
                "s3:ListMultipartUploadParts",
                "s3:PutObject",
                "s3:GetObject",
                "s3:AbortMultipartUpload"
            ],
            "Resource": [
                "arn:aws:iam::*:role/AppStreamUsageReports-
AppStreamUsageReportGlueRole*",
                "arn:aws:cloudformation:*:*:stack/AppStreamUsageReports/*",
                "arn:aws:athena:*:*:workgroup/primary",
                "arn:aws:s3:::aws-athena-query-results-*"
            ]
        },
        {
            "Effect": "Allow",
            "Action": [
                "iam:AttachRolePolicy",
                "iam:PutRolePolicy",
                "s3:GetObject",
                "s3:ListBucket"
            ],
            "Resource": [
                "arn:aws:s3:::appstream-logs-*",
                "arn:aws:iam::*:role/AppStreamUsageReports-
AppStreamUsageReportGlueRole*"
            ]
        },
        {
            "Effect": "Allow",
            "Action": [
                "iam:PassRole"
            ],
            "Resource": [
```

```
                "arn:aws:iam::*:role/AppStreamUsageReports-
AppStreamUsageReportGlueRole*"
            ],
            "Condition": {
                "StringEquals": {
                    "iam:PassedToService": "glue.amazonaws.com"
                }
            }
        },
        {
            "Effect": "Allow",
            "Action": [
                "cloudformation:GetTemplateSummary",
                "glue:GetResourcePolicy",
                "glue:GetCrawlers",
                "glue:BatchGetCrawlers",
                "glue:GetClassifiers",
                "glue:CreateClassifier",
                "glue:ListCrawlers",
                "glue:GetTags",
                "glue:GetCrawlerMetrics",
                "glue:GetClassifier",
                "tag:GetResources"
            ],
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": "athena:RunQuery",
            "Resource": "arn:aws:athena:*:*:workgroup/primary"
        },
        {
            "Effect": "Allow",
            "Action": [
                "glue:GetTables",
                "glue:GetPartitions",
                "glue:GetTable"
            ],
            "Resource": [
                "arn:aws:glue:*:*:table/appstream-usage/*",
                "arn:aws:glue:*:*:database/appstream-usage",
                "arn:aws:glue:*:*:catalog"
            ]
        },
```

```
        {
            "Effect": "Allow",
            "Action": [
                "glue:GetDatabase",
                "glue:CreateDatabase",
                "glue:GetDatabases"
            ],
            "Resource": [
                "arn:aws:glue:*:*:database/appstream-usage",
                "arn:aws:glue:*:*:catalog"
            ]
        },
        {
            "Effect": "Allow",
            "Action": [
                "glue:GetCrawler",
                "glue:StartCrawler",
                "glue:CreateCrawler"
            ],
            "Resource": "arn:aws:glue:*:*:crawler/appstream-usage*"
        },
        {
            "Effect": "Allow",
            "Action": "glue:GetCatalogImportStatus",
            "Resource": "arn:aws:glue:*:*:catalog"
        }
    ]
}
```

Create a Data Catalog by Using the AWS Glue Crawler

The AWS Glue crawler, when run, creates a Data Catalog and schema that are mapped to the structure of your sessions and applications reports. Each time a new report is stored in your Amazon S3 bucket, you must run the crawler to update your AWS Glue Data Catalog with the data from the new report.

> **Note**
>
> Charges may apply to the running of your AWS Glue crawler. For more information, see AWS Glue Pricing.

1. Open the AWS Glue console at https://console.aws.amazon.com/glue/.
2. Choose the AWS Region for which you have subscribed to usage reports.
3. Select the check box next to the crawler named **appstream-usage-sessions-crawler**, and then choose **Run crawler**. Repeat this step for the crawler named **appstream-usage-apps-crawler**.

   Performing these steps runs the crawlers and schedules them to run automatically according to the schedule specified in the CloudFormation stack.
4. After both crawlers finish running, in the navigation pane, choose **Databases**. A database named **appstream-usage**, which represents your usage reports, displays. This database is an AWS Glue Data Catalog that was created when **appstream-usage-sessions-crawler** and **appstream-usage-apps-crawler** were run.
5. To view the tables in the database, choose **appstream-usage**, **Tables**. Two tables, **applications** and **sessions**, which represent your applications and sessions usage reports respectively, display. Choose either table to view its schema.

   You can now query these tables in Athena by using SQL.

Create and Run Athena Queries

To query your usage reports by using Athena, perform the following steps.

> **Note**
>
> Charges may apply to Athena queries that you run. For more information, see Amazon Athena Pricing.

1. Open the Athena console at https://console.aws.amazon.com/athena/.
2. Under **Database**, choose **appstream-usage**.
3. In the query pane, enter a SQL query, and choose **Run query**.

Working with Athena Queries

This section provides SQL queries that you can run in Athena to analyze the usage reports data in your Amazon S3 bucket.

To create a consolidated report of all sessions in a given month, run the following query:

```
SELECT *
FROM "appstream-usage"."sessions"
WHERE year='four-digit-year'
AND month='two-digit-month'
```

You can also perform join operations between the **applications** and **sessions** tables in your query. For example, to view the distinct users who launched each application in a given month, run the following query:

```
SELECT DISTINCT apps.application_name, sessions.user_id
FROM "appstream-usage"."applications" apps
   INNER JOIN "appstream-usage"."sessions" sessions ON (apps.user_session_id =
 sessions.user_session_id AND sessions.year='four-digit-year' AND sessions.month='two-
digit-month')
WHERE apps.year='four-digit-year'
  AND apps.month='two-digit-month'
ORDER BY 1, 2
```

Athena query results are stored as .csv files in an Amazon S3 bucket in your account that is named `aws-athena-query-results-`*`account-id-without-hyphens`*`-`*`region-code`*. For ease in locating query results, choose **Save as** and provide a name for your query before you run it. You can also choose the download icon in the **Athena Results** pane to download the results of the query as a .csv file.

To enhance performance and reduce costs, Athena uses partitioning to reduce the amount of data scanned in queries. For more information, see Partitioning Data. Usage reports are partitioned in your Amazon S3 buckets by year, month, and day. You can restrict your queries to certain date range partitions using the **year**, **month**, and **day** fields as conditions in your queries. For example, the following query ingests only the sessions reports for the week of May 19, 2019.

```
SELECT SUBSTRING(session_start_time, 1, 10) AS report_date,
    COUNT(DISTINCT user_session_id) AS num_sessions
FROM "appstream-usage"."sessions"
WHERE year='2019'
   AND month='05'
   AND day BETWEEN '19' and '25'
GROUP BY 1
ORDER BY 1
```

In contrast, the following query produces identical results, but because it isn't restricted to any partitions, it ingests all sessions reports stored in your Amazon S3 bucket.

```
SELECT SUBSTRING(session_start_time, 1, 10) AS report_date,
    COUNT(DISTINCT user_session_id) AS num_sessions
FROM "appstream-usage"."sessions"
WHERE session_end_time BETWEEN '2019-05-19' AND '2019-05-26'
GROUP BY 1
ORDER BY 1
```

If a session spans more than one day, the session and application records appear in the sessions and applications reports, respectively, corresponding to the day in which the session ended. For this reason, if you need to find records that relate to all sessions that were active during a given date range, consider expanding the partition set of your query by the maximum session length you have configured for your fleets.

For example, to view all sessions that were active for a given fleet during a calendar month, where the fleet had a maximum session duration of 100 hours, run the following query to expand your partition set by five days.

```
SELECT *
FROM "appstream-usage"."sessions"
WHERE fleet_name = 'fleet_name'
   AND session_start_time BETWEEN '2019-05-01' AND '2019-06-01'
   AND year='2019'
   AND (month='05' OR (month='06' AND day<='05'))
ORDER BY session_start_time
```

The CloudFormation template that created the AWS Glue crawlers also created and saved several sample queries in your Athena account that you can use to analyze your usage data. These sample queries include the following:

- Aggregated monthly session report
- Average session length per stack
- Number of sessions per day
- Total streaming hours per user

> **Note**
>
> On-demand usage charges are rounded up to the next hour for each session.

- Distinct users per app

To use any of these queries, perform the following steps.

1. Open the Athena console at https://console.aws.amazon.com/athena/.
2. Choose **Saved Queries**. The five queries noted before this procedure should display. The name of each query begins with "AS2." For example, "AS2_users_per_app_curr_mo."
3. To run a query, choose the query name rather than the option next to the name.
4. The text of the query appears in the query pane. Choose **Run query**.

To view these queries in a separate CloudFormation template, see athena-sample-queries-appstream-usage-data_template.yml in the *AWS Code Sample Catalog.*

Logging Amazon WorkSpaces Applications API Calls with AWS CloudTrail

Amazon WorkSpaces Applications is integrated with AWS CloudTrail. CloudTrail is a service that provides a record of actions taken by a user, role, or an AWS service in WorkSpaces Applications. CloudTrail captures API calls for WorkSpaces Applications as events. The calls captured include calls from the WorkSpaces Applications console and code calls to the WorkSpaces Applications API operations. If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for WorkSpaces Applications. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. You can use the information collected by CloudTrail to determine details such as request information. For example, CloudTrail collects the following information: What request was made to WorkSpaces Applications, the IP address from which the request was made, who made the request, and when it was made.

To learn more about CloudTrail, including how to configure and enable it, see the AWS CloudTrail User Guide.

Topics

- WorkSpaces Applications Information in CloudTrail
- Example: WorkSpaces Applications Log File Entries

WorkSpaces Applications Information in CloudTrail

CloudTrail is enabled on your AWS account when you create the account. When supported event activity occurs in WorkSpaces Applications, that activity is recorded in a CloudTrail event along with other AWS service events in **Event history**. You can view, search, and download recent events in your AWS account. For more information, see Viewing Events with CloudTrail Event History.

For an ongoing record of events in your AWS account, including events for WorkSpaces Applications, create a trail. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all AWS Regions. The trail logs events from all Regions in the AWS partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other AWS services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following:

- Overview for Creating a Trail
- CloudTrail Supported Services and Integrations
- Configuring Amazon SNS Notifications for CloudTrail
- Receiving CloudTrail Log Files from Multiple Regions and Receiving CloudTrail Log Files from Multiple Accounts

WorkSpaces Applications supports logging the following actions as events in CloudTrail log files:

- AssociateFleet
- BatchAssociateUserStack
- BatchDisassociateUserStack
- CopyImage
- CreateDirectoryConfig
- CreateFleet
- CreateImageBuilder
- CreateImageBuilderStreamingURL
- CreateStack

- CreateStreamingURL
- DeleteDirectoryConfig
- DeleteFleet
- DeleteImage
- DeleteImageBuilder
- DeleteImagePermissions
- DeleteStack
- DescribeDirectoryConfigs
- DescribeFleets
- DescribeImageBuilders
- DescribeImagePermissions
- DescribeImages
- DescribeSessions
- DescribeStacks
- DescribeUserStackAssociations
- ExpireSession
- ListAssociatedFleets
- ListAssociatedStacks
- ListTagsForResource
- StartFleet
- StartImageBuilder
- StopFleet
- StopImageBuilder
- TagResource
- UntagResource
- UpdateDirectoryConfig
- UpdateFleet
- UpdateImagePermissions
- UpdateStack

Every event or log entry contains information about who generated the request. The identity information helps you determine the following:

- Whether the request was made with root or IAM user credentials.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

For more information, see the CloudTrail userIdentity Element.

Example: WorkSpaces Applications Log File Entries

A trail is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the `AssociateFleet` event.


```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "AssumedRole",
    "principalId": "AIDACKCEVSQ6C2EXAMPLE:janeroe",
    "arn": "arn:aws:sts:: 123456789012:assumed-role/Admin/janeroe",
    "accountId": "123456789012",
    "accessKeyId": "AKIAIOSFODNN7EXAMPLE",
    "sessionContext": {
      "attributes": {
        "mfaAuthenticated": "false",
        "creationDate": "2019-03-12T06:41:50Z"
      },
      "sessionIssuer": {
        "type": "Role",
        "principalId": "AIDACKCEVSQ6C2EXAMPLE",
        "arn": "arn:aws:iam:: 123456789012:role/Admin",
        "accountId": "123456789012",
        "userName": "Admin"
      }
    }
```

```
  },
  "eventTime": "2019-03-12T06:58:09Z",
  "eventSource": "appstream.amazonaws.com",
  "eventName": "AssociateFleet",
  "awsRegion": "us-east-1",
  "sourceIPAddress": "198.51.100.15",
  "userAgent": "Mozilla/5.0 (Macintosh; Intel Mac OS X 10_11_6) AppleWebKit/537.36
 (KHTML, like Gecko) Chrome/72.0.3626.121 Safari/537.36",
  "requestParameters": {
    "fleetName": "ExampleFleet1",
    "stackName": "ExampleStack1"
  },
  "responseElements": null,
  "requestID": "3210a159-4494-11e9-8017-873084baf125",
  "eventID": "a6fbde60-a55a-46fe-87d4-89ead558dffd",
  "eventType": "AwsApiCall",
  "recipientAccountId": "123456789012"
}
```

The following example shows a CloudTrail log entry that demonstrates the `CreateImage` event when an image is created using the WorkSpaces Applications image builder.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "arn": "arn:aws:appstream:us-east-1: 123456789012:image-builder/
ExampleImageBuilder",
    "accountId": "123456789012"
  },
  "eventTime": "2019-03-21T22:32:05Z",
  "eventSource": "appstream.amazonaws.com",
  "eventName": "CreateImage",
  "awsRegion": "us-east-1",
  "requestParameters": null,
  "responseElements": null,
  "eventID": "12b2d6e2-c9a9-402e-8886-2c388d3df610",
  "readOnly": false,
  "eventType": "AwsServiceEvent",
  "recipientAccountId": "123456789012",
  "serviceEventDetails": {
    "ImageName": "ExampleImage1",
    "ImagePlatform": "WINDOWS",
    "PublicBaseImageReleasedDate": "Tue Jan 15 22:19:56 UTC 2019",
```

```
    "ImageDisPlayName": "Example Image 1",
    "ImageBuilderSupported": "True",
    "ImageCreatedTime": "Thu Mar 21 22:32:05 UTC 2019",
    "ImageDescription": "Example image for testing",
    "ImageState": "PENDING"
  }
}
```

Security in Amazon WorkSpaces Applications

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The shared responsibility model describes this as security of the cloud and security in the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the AWS Compliance Programs. For information about the compliance programs that apply to WorkSpaces Applications, see AWS Services in Scope by Compliance Program.
- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations

This documentation helps you understand how to apply the shared responsibility model when using WorkSpaces Applications. It shows you how to configure WorkSpaces Applications to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your WorkSpaces Applications resources.

Contents

- Data Protection in Amazon WorkSpaces Applications
- Identity and Access Management for Amazon WorkSpaces Applications
- Logging and Monitoring in Amazon WorkSpaces Applications
- Compliance Validation for Amazon WorkSpaces Applications
- Resilience in Amazon WorkSpaces Applications
- Infrastructure Security in Amazon WorkSpaces Applications
- Security Groups in Amazon WorkSpaces Applications
- Update Management in Amazon WorkSpaces Applications
- Amazon WorkSpaces Applications Cross-Service Confused Deputy Prevention
- Security Best Practices in Amazon WorkSpaces Applications

Data Protection in Amazon WorkSpaces Applications

The AWS shared responsibility model applies to data protection in Amazon WorkSpaces Applications. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. This content includes the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the Data Privacy FAQ. For information about data protection in Europe, see the AWS Shared Responsibility Model and GDPR blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual users with AWS Identity and Access Management (IAM). That way each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources. We recommend TLS 1.2.
- Set up API and user activity logging with AWS CloudTrail.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing personal data that is stored in Amazon S3.
- If you require FIPS 140-3 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see Federal Information Processing Standard (FIPS) 140-3.

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form fields such as a **Name** field. This includes when you work with WorkSpaces Applications or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into tags or free-form fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Topics

- Encryption at Rest
- Encryption in Transit
- Administrator Controls

- Application Access

Encryption at Rest

WorkSpaces Applications fleet instances are ephemeral in nature. After a user's streaming session is finished, the underlying instance and its associated Amazon Elastic Block Store (Amazon EBS) volume are terminated. In addition, WorkSpaces Applications periodically recycles unused instances for freshness.

When you enable application settings persistence, home folders, session scripts, or usage reports your users, the data that is generated by your users and stored in Amazon Simple Storage Service buckets is encrypted at rest. AWS Key Management Service is a service that combines secure, highly available hardware and software to provide a key management system scaled for the cloud. Amazon S3 uses AWS Managed CMKs to encrypt your Amazon S3 object data.

When you turn on Replay Session for agent sessions, recordings are encrypted at rest while they are processed, and you can optionally encrypt them with a customer managed AWS KMS key. For more information, see Encryption at rest for Replay Session.

Encryption in Transit

The following table provides information about how data is encrypted in transit. Where applicable, other data protection methods for WorkSpaces Applications are also listed.

Data	Network path	How protected
Web assets This traffic includes assets such as images and JavaScript files.	Between WorkSpaces Applications users and WorkSpaces Applications	Encrypted using TLS 1.2
Pixel and related streaming traffic	Between WorkSpaces Applications users and WorkSpaces Applications	Encrypted using 256-bit Advanced Encryption Standard (AES-256) Transported using TLS 1.2

Data	Network path	How protected
API traffic	Between WorkSpaces Applications users and WorkSpaces Applications	Encrypted using TLS 1.2 Requests to create a connection are signed using SigV4
Application settings and home folder data generated by users Applicable when application settings persistence and home folders are enabled.	Between WorkSpaces Applications users and Amazon S3	Encrypted using Amazon S3 SSL endpoints
WorkSpaces Applications-managed traffic	Between WorkSpaces Applications streaming instances and: • WorkSpaces Applications management services • AWS services and resources in your Amazon Web Services account • Non-AWS services and resources (such as Google Drive and Microsoft OneDrive)	Encrypted using TLS 1.2 Requests to create a connection are signed using SigV4 where applicable

Administrator Controls

WorkSpaces Applications provides administrative controls that you can use to limit the ways in which users can transfer data between their local computer and an WorkSpaces Applications fleet instance. You can limit or disable the following when you create or update an WorkSpaces Applications stack:

- Clipboard/copy and paste actions
- File upload and download, including folder and drive redirection

- Printing

When you create an WorkSpaces Applications image, you can specify which USB devices are available to redirect to WorkSpaces Applications fleet instances from the WorkSpaces Applications client for Windows. The USB devices that you specify will be available for use during users' WorkSpaces Applications streaming sessions. For more information, see Qualify USB Devices for Use with Streaming Applications.

Application Access

By default, WorkSpaces Applications enables the applications that you specify in your image to launch other applications and executable files on the image builder and fleet instance. This ensures that applications with dependencies on other applications (for example, an application that launches the browser to navigate to a product website) function as expected. Make sure that you configure your administrative controls, security groups, and other security software to grant users the minimum permissions required to access resources and transfer data between their local computers and fleet instances.

You can use application control software, such as Microsoft AppLocker, and policies to control which applications and files your users can run. Application control software and policies help you control the executable files, scripts, Windows installer files, dynamic-link libraries, and application packages that your users can run on WorkSpaces Applications image builders and fleet instances.

> **Note**
>
> The WorkSpaces Applications agent software relies on the Windows command prompt and Windows Powershell to provision streaming instances. If you choose to prevent users from launching the Windows command prompt or Windows Powershell, the policies must not apply to the Windows NT AUTHORITY\SYSTEM or users in the Administrators group.

Rule type	Action	Windows user or group	Name/Path	Condition	Description
Executable	Allow	NT AUTHORITY \System	*	Path	Required for the WorkSpaces Applications agent software
Executable	Allow	BUILTIN\Administrators	*	Path	Required for the WorkSpaces Applications agent software
Executable	Allow	Everyone	%PROGRAMFILES%\nodejs*	Path	Required for the WorkSpaces Applications agent software
Executable	Allow	Everyone	%PROGRAMFILES%\NICE*	Path	Required for the WorkSpaces Applications agent software
Executable	Allow	Everyone	%PROGRAMFILES%\Amazon*	Path	Required for the WorkSpaces Applicati

Rule type	Action	Windows user or group	Name/Path	Condition	Description
					ons agent software
Executable	Allow	Everyone	%PROGRAMFILES%\<*default-browser* >*	Path	Required for the WorkSpaces Applications agent software when persistent storage solutions, such as Google Drive or Microsoft OneDrive for Business, are used. This exception is not required when WorkSpaces Applications home folders are used.

Identity and Access Management for Amazon WorkSpaces Applications

Your security credentials identify you to services in AWS and grant you unlimited use of your AWS resources, such as your WorkSpaces Applications resources. You can use features of WorkSpaces

Applications and AWS Identity and Access Management (IAM) to allow other users, services, and applications to use your WorkSpaces Applications resources without sharing your security credentials.

You can use IAM to control how other users use resources in your Amazon Web Services account, and you can use security groups to control access to your WorkSpaces Applications streaming instances. You can allow full use or limited use of your WorkSpaces Applications resources.

Contents

- Network Access to Your Streaming Instance
- Using AWS Managed Policies and Linked Roles to Manage Administrator Access to WorkSpaces Applications Resources
- Using IAM Policies to Manage Administrator Access to Application Auto Scaling
- Using IAM Policies to Manage Administrator Access to the Amazon S3 Bucket for Home Folders and Application Settings Persistence
- Using an IAM Role to Grant Permissions to Applications and Scripts Running on WorkSpaces Applications Streaming Instances
- SELinux on Red Hat Enterprise Linux and Rocky Linux
- Cookie-Based Authentication in Amazon WorkSpaces Applications

Network Access to Your Streaming Instance

A security group acts as a stateful firewall that controls what traffic is allowed to reach your streaming instances. When you launch an WorkSpaces Applications streaming instance, assign it to one or more security groups. Then, add rules to each security group that control traffic for the instance. You can modify the rules for a security group at any time. The new rules are automatically applied to all instances to which the security group is assigned.

For more information, see Security Groups in Amazon WorkSpaces Applications.

Using AWS Managed Policies and Linked Roles to Manage Administrator Access to WorkSpaces Applications Resources

By default, IAM users don't have the permissions required to create or modify WorkSpaces Applications resources, or perform tasks by using the WorkSpaces Applications API. This means that these users can't perform these actions in the WorkSpaces Applications console or by using

WorkSpaces Applications AWS CLI commands. To allow IAM users to create or modify resources and perform tasks, attach an IAM policy to the IAM users or groups that require those permissions.

When you attach a policy to a user, group of users, or IAM role, it allows or denies the users permission to perform the specified tasks on the specified resources.

Contents

- AWS Managed Policies Required to Access WorkSpaces Applications Resources
- Roles Required for WorkSpaces Applications, Application Auto Scaling, and AWS Certificate Manager Private CA
- Checking for the AmazonAppStreamServiceAccess Service Role and Policies
- Checking for the ApplicationAutoScalingForAmazonAppStreamAccess Service Role and Policies
- Checking for the AWSServiceRoleForApplicationAutoScaling_AppStreamFleet Service-Linked Role and Policies
- Checking for the AmazonAppStreamPCAAccess Service Role and Policies

AWS Managed Policies Required to Access WorkSpaces Applications Resources

To provide full administrative or read-only access to WorkSpaces Applications, you must attach one of the following AWS managed policies to the IAM users or groups that require those permissions. An *AWS managed policy* is a standalone policy that is created and administered by AWS. For more information, see AWS Managed Policies in the *IAM User Guide*.

> **Note**
>
> In AWS, IAM roles are used to grant permissions to an AWS service so it can access AWS resources. The policies that are attached to the role determine which AWS resources the service can access and what it can do with those resources. For WorkSpaces Applications, in addition to having the permissions defined in the **AmazonAppStreamFullAccess** policy, you must also have the required roles in your AWS account. For more information, see the section called "Roles Required for WorkSpaces Applications, Application Auto Scaling, and AWS Certificate Manager Private CA".

AmazonAppStreamFullAccess

This managed policy provides full administrative access to WorkSpaces Applications resources. To manage WorkSpaces Applications resources and perform API actions through the AWS Command Line Interface (AWS CLI), AWS SDK, or AWS Management Console, you must have the permissions defined in this policy.

If you sign into the WorkSpaces Applications console as an IAM user, you must attach this policy to your AWS account. If you sign in through console federation, you must attach this policy to the IAM role that was used for federation.

To view the permissions for this policy, see AmazonAppStreamFullAccess.

AmazonAppStreamReadOnlyAccess

This identity-based policy grants users read-only permissions to view and monitor WorkSpaces Applications resources and related service configurations. Users can access the WorkSpaces Applications console to view streaming applications, fleet status, usage reports, and associated resources, but cannot make any changes. The policy also includes necessary read permissions for supporting services like IAM, Application Auto Scaling, and CloudWatch to enable comprehensive monitoring and reporting capabilities.

To view the permissions for this policy, see AmazonAppStreamReadOnlyAccess.

The WorkSpaces Applications console uses an additional action that provides functionality that is not available through the AWS CLI or AWS SDK. The **AmazonAppStreamFullAccess** and **AmazonAppStreamReadOnlyAccess** policies both provide permissions for the following action.

Action	Description	Access Level
`DescribeImageBuilders`	Grants permission to retrieve a list that describes one or more specified image builders, if the image builder names are provided. Otherwise, all image builders in the account are described.	Read

AmazonAppStreamPCAAccess

This managed policy provides full administrative access to AWS Certificate Manager Private CA resources in your AWS account for certificate-based authentication.

To view the permissions for this policy, see AmazonAppStreamPCAAccess.

AmazonAppStreamServiceAccess

This managed policy is the default policy for the WorkSpaces Applications service role.

This role permissions policy allows WorkSpaces Applications to complete the following actions:

- When using subnets in your account for your WorkSpaces Applications fleets, WorkSpaces Applications is able to describe subnets, VPCs, and availability zones, as well as create and manage the lifecycle of all elastic network interfaces associated with the fleet instances in those subnets. This also includes being able to attach Security Groups and IP addresses from those subnets to those elastic network interfaces.
- When using features such as UPP and HomeFolders, WorkSpaces Applications is able to create and manage Amazon S3 buckets, objects and their lifecyles, policies, and encryption configuration in the account. These buckets include the following naming prefixes:
  - `"arn:aws:s3:::appstream2-36fb080bb8-",`
  - `"arn:aws:s3:::appstream-app-settings-",`
  - `"arn:aws:s3:::appstream-logs-"`

To view the permissions for this policy, see AmazonAppStreamServiceAccess.

ApplicationAutoScalingForAmazonAppStreamAccess

This managed policy enables application autoscaling for WorkSpaces Applications.

To view the permissions for this policy, see ApplicationAutoScalingForAmazonAppStreamAccess .

AWSApplicationAutoscalingAppStreamFleetPolicy

This managed policy grants permissions for Application Auto Scaling to access WorkSpaces Applications and CloudWatch .

To view the permissions for this policy, see AWSApplicationAutoscalingAppStreamFleetPolicy .

WorkSpaces Applications updates to AWS managed policies

View details about updates to AWS managed policies for WorkSpaces Applications since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the *Document History* page.

Change	Description	Date
AmazonAppStreamServiceAccess – Change	Added allow permissions for `"ec2:DescribeImages"` to the policy JSON policy document	November 17, 2025
AmazonAppStreamReadOnlyAccess – Change	Removed `"appstream:Get*",` from the JSON policy document	October 22, 2025
WorkSpaces Applications started tracking changes	WorkSpaces Applications started tracking changes for its AWS managed policies	October 31, 2022

Roles Required for WorkSpaces Applications, Application Auto Scaling, and AWS Certificate Manager Private CA

In AWS, IAM roles are used to grant permissions to an AWS service so it can access AWS resources. The policies that are attached to the role determine which AWS resources the service can access and what it can do with those resources. For WorkSpaces Applications, in addition to having the permissions defined in the **AmazonAppStreamFullAccess** policy, you must also have the following roles in your AWS account.

Roles

- AmazonAppStreamServiceAccess
- ApplicationAutoScalingForAmazonAppStreamAccess
- AWSServiceRoleForApplicationAutoScaling_AppStreamFleet
- AmazonAppStreamPCAAccess

AmazonAppStreamServiceAccess

This role is a service role that is created for you automatically when you get started with WorkSpaces Applications in an AWS Region. For more information about services roles, see Creating a role to delegate permissions to an AWS service in the *IAM User Guide*.

While WorkSpaces Applications resources are being created, the WorkSpaces Applications service makes API calls to other AWS services on your behalf by assuming this role. To create fleets, you must have this role in your account. If this role is not in your AWS account and the required IAM permissions and trust relationship policies are not attached, you cannot create WorkSpaces Applications fleets.

For more information, see Checking for the AmazonAppStreamServiceAccess Service Role and Policies to check whether the **AmazonAppStreamServiceAccess** service role is present and has the correct policies attached.

> **Note**
>
> This service role can have permissions that are different from the first user that is getting started with WorkSpaces Applications. For details on the permissions of this role see "AmazonAppStreamServiceAccess" in the section called "AWS Managed Policies Required to Access WorkSpaces Applications Resources".

ApplicationAutoScalingForAmazonAppStreamAccess

This role is a service role that is created for you automatically when you get started with WorkSpaces Applications in an AWS Region. For more information about services roles, see Creating a role to delegate permissions to an AWS service in the *IAM User Guide*.

Automatic scaling is a feature of WorkSpaces Applications fleets. To configure scaling policies, you must have this service role in your AWS account. If this service role is not in your AWS account and the required IAM permissions and trust relationship policies are not attached, you cannot scale WorkSpaces Applications fleets.

For more information, see Checking for the ApplicationAutoScalingForAmazonAppStreamAccess Service Role and Policies.

AWSServiceRoleForApplicationAutoScaling_AppStreamFleet

This role is a service-linked role that is created for you automatically. For more information, see Service-linked roles in the *Application Auto Scaling User Guide*.

Application Auto Scaling uses a service-linked role to perform automatic scaling on your behalf. A *service-linked role* is an IAM role that is linked directly to an AWS service. This role includes all the permissions that the service requires to call other AWS services on your behalf.

For more information, see Checking for the `AWSServiceRoleForApplicationAutoScaling_AppStreamFleet` Service-Linked Role and Policies.

AmazonAppStreamPCAAccess

This role is a service role that is created for you automatically when you get started with WorkSpaces Applications in an AWS Region. For more information about services roles, see Creating a role to delegate permissions to an AWS service in the *IAM User Guide*.

Certificate-based authentication is a feature of WorkSpaces Applications fleets joined to Microsoft Active Directory domains. To enable and use certificate-based authentication, you must have this service role in your AWS account. If this service role is not in your AWS account and the required IAM permissions and trust relationship policies are not attached, you cannot enable or use certificate-based authentication.

For more information, see the section called “Checking for the AmazonAppStreamPCAAccess Service Role and Policies”.

Checking for the AmazonAppStreamServiceAccess Service Role and Policies

Complete the steps in this section to check whether the **AmazonAppStreamServiceAccess** service role is present and has the correct policies attached. If this role is not in your account and must be created, you or an administrator with the required permissions must perform the steps to get started with WorkSpaces Applications in your Amazon Web Services account.

To check whether the AmazonAppStreamServiceAccess IAM service role is present

1. Open the IAM console at https://console.aws.amazon.com/iam/.
2. In the navigation pane, choose **Roles**.

3. In the search box, type **amazonappstreamservice** to narrow the list of roles to select, and then choose **AmazonAppStreamServiceAccess**. If this role is listed, select it to view the role **Summary** page.
4. On the **Permissions** tab, confirm whether the **AmazonAppStreamServiceAccess** permissions policy is attached.
5. Return to the role **Summary** page.
6. On the **Trust relationships** tab, choose **Show policy document**, and then confirm whether the **AmazonAppStreamServiceAccess** trust relationship policy is attached and follows the correct format. If so, the trust relationship is correctly configured. Choose **Cancel** and close the IAM console.

AmazonAppStreamServiceAccess trust relationship policy

The **AmazonAppStreamServiceAccess** trust relationship policy must include the WorkSpaces Applications service as the principal. A *principal* is an entity in AWS that can perform actions and access resources. This policy must also include the `sts:AssumeRole` action. The following policy configuration defines WorkSpaces Applications as a trusted entity.

JSON

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
      "Service": "appstream.amazonaws.com"
    },
    "Action": "sts:AssumeRole"
    }
  ]
}
```

Checking for and creating the service role (AWS CLI)

To make sure that the required service role exists before you provision resources, perform the steps in this section when you get started with Amazon WorkSpaces Applications (WorkSpaces

Applications) through the WorkSpaces Applications API or AWS Command Line Interface (AWS CLI) instead of the console. You check whether the `AmazonAppStreamServiceAccess` service role exists and verify that it has the correct policies attached. If this role is not in your account, create it before you provision resources with WorkSpaces Applications.

Creating this role requires administrator-level AWS Identity and Access Management (IAM) permissions. If you are not an administrator, ask an administrator to perform these steps or to grant you the required permissions.

> **Note**
>
> WorkSpaces Applications automatically creates the `AmazonAppStreamServiceAccess` service role the first time you get started with WorkSpaces Applications in the console. To provision resources with the WorkSpaces Applications API or AWS CLI without using the console, first create this role manually with the API or AWS CLI.

Checking whether the service role is present

Before you begin, make sure that the following are true:

- You installed and configured the AWS CLI with credentials for your AWS account.
- Your AWS Identity and Access Management (IAM) principal has the `iam:GetRole` and `iam:ListAttachedRolePolicies` permissions.

Checking whether the role exists (AWS CLI)

1. Run the following AWS CLI command to check whether the `AmazonAppStreamServiceAccess` role is present in your account.

   ```
   PROMPT> aws iam get-role --role-name AmazonAppStreamServiceAccess
   ```

2. If the role exists, continue to the next procedure. If the role does not exist, you receive the following error.

   ```
   An error occurred (NoSuchEntity) when calling the GetRole operation:
   The role with name AmazonAppStreamServiceAccess cannot be found.
   ```

   If you receive this error, create the role as described in Creating the service role (AWS CLI).

Verifying that the permissions policy is attached (AWS CLI)

1. Run the following AWS CLI command to list the policies that are attached to the role.

```
PROMPT> aws iam list-attached-role-policies --role-name
 AmazonAppStreamServiceAccess
```

2. Verify that the following policy is listed in the output.

```
{
    "AttachedPolicies": [
        {
            "PolicyName": "AmazonAppStreamServiceAccess",
            "PolicyArn": "arn:aws:iam::aws:policy/service-role/
AmazonAppStreamServiceAccess"
        }
    ]
}
```

3. If this policy is not in the output, run the following AWS CLI command to attach it.

```
PROMPT> aws iam attach-role-policy \
    --role-name AmazonAppStreamServiceAccess \
    --policy-arn arn:aws:iam::aws:policy/service-role/AmazonAppStreamServiceAccess
```

> **Note**
>
> If you receive an `AccessDenied` error, verify that your IAM principal has the `iam:AttachRolePolicy` permission. If you do not have this permission, ask your AWS account administrator to attach the policy or to grant you the permission.

Verifying the trust relationship (AWS CLI)

1. Run the following AWS CLI command to view the trust relationship policy for the role.

```
PROMPT> aws iam get-role \
    --role-name AmazonAppStreamServiceAccess \
    --query "Role.AssumeRolePolicyDocument" \
    --output json
```

2. Confirm that the output matches the following format.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "Service": "appstream.amazonaws.com"
            },
            "Action": "sts:AssumeRole"
        }
    ]
}
```

If the trust relationship specifies `appstream.amazonaws.com` as the trusted principal, the service role is ready and no further action is required.

> **Note**
>
> If the trust relationship principal does not match `appstream.amazonaws.com`, someone might have modified the role. Contact your AWS account administrator to review and correct the trust policy before you use WorkSpaces Applications.

Creating the service role (AWS CLI)

If the `AmazonAppStreamServiceAccess` role is not in your account, complete the following steps to create it.

Before you begin, make sure that the following are true:

- You installed and configured the AWS CLI with credentials for your AWS account.
- Your IAM principal has the `iam:CreateRole`, `iam:AttachRolePolicy`, `iam:GetRole`, and `iam:ListAttachedRolePolicies` permissions.

Creating the trust policy document (AWS CLI)

- To grant the WorkSpaces Applications service permission to assume the role, create a file named `trust-policy.json` with the following content.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "Service": "appstream.amazonaws.com"
            },
            "Action": "sts:AssumeRole"
        }
    ]
}
```

Creating the role (AWS CLI)

- Run the following AWS CLI command to create the `AmazonAppStreamServiceAccess` role with the trust policy document that you created. The command returns the Amazon Resource Name (ARN) and details of the new role.

```
PROMPT> aws iam create-role \
    --role-name AmazonAppStreamServiceAccess \
    --assume-role-policy-document file://trust-policy.json \
    --description "Service role that lets AppStream manage resources on your
 behalf"
```

Note

If you receive an `EntityAlreadyExists` error, the role already exists in your account. If you receive an `AccessDenied` error, verify that your IAM principal has the `iam:CreateRole` permission, or ask your AWS account administrator to create the role.

Attaching the permissions policy (AWS CLI)

- Run the following AWS CLI command to attach the `AmazonAppStreamServiceAccess` AWS managed policy to the role.

```
PROMPT> aws iam attach-role-policy \
    --role-name AmazonAppStreamServiceAccess \
    --policy-arn arn:aws:iam::aws:policy/service-role/AmazonAppStreamServiceAccess
```

> **Note**
>
> If you receive an `AccessDenied` error, verify that your IAM principal has the `iam:AttachRolePolicy` permission, or ask your AWS account administrator to attach the policy.

Verifying that the role exists (AWS CLI)

- Run the following AWS CLI command to confirm that you created the role.

```
PROMPT> aws iam get-role --role-name AmazonAppStreamServiceAccess
```

Verifying that the permissions policy is attached (AWS CLI)

- Run the following AWS CLI command to verify that the permissions policy is attached.

```
PROMPT> aws iam list-attached-role-policies --role-name
 AmazonAppStreamServiceAccess
```

Verifying the trust relationship (AWS CLI)

- Run the following AWS CLI command to confirm that the trust relationship specifies `appstream.amazonaws.com` as the trusted principal.

```
PROMPT> aws iam get-role \
    --role-name AmazonAppStreamServiceAccess \
    --query "Role.AssumeRolePolicyDocument" \
```

```
    --output json
```

The `AmazonAppStreamServiceAccess` service role is ready to use with WorkSpaces Applications when all of the following are true:

- The role exists.
- You see the permissions policy listed.
- The trust relationship shows `appstream.amazonaws.com` as the trusted principal.

> **Important**
>
> WorkSpaces Applications must have this role to manage resources on your behalf. Keep the role name, permissions policy, and trust relationship unchanged so that the service continues to function correctly. To make changes, contact your AWS account administrator.

Finding related service role resources

For more information about controlling access to WorkSpaces Applications resources, see the following resources:

- For information about the AWS managed policies for WorkSpaces Applications, see AWS Managed Policies Required to Access WorkSpaces Applications Resources.
- For information about getting started with WorkSpaces Applications, see Create an Amazon WorkSpaces Applications Fleet and Stack.
- For information about IAM roles, see IAM roles in the *IAM User Guide*.

Checking for the ApplicationAutoScalingForAmazonAppStreamAccess Service Role and Policies

Complete the steps in this section to check whether the **ApplicationAutoScalingForAmazonAppStreamAccess** service role is present and has the correct policies attached. If this role is not in your account and must be created, you or an administrator with the required permissions must perform the steps to get started with WorkSpaces Applications in your Amazon Web Services account.

To check whether the ApplicationAutoScalingForAmazonAppStreamAccessIAM service role is present

1. Open the IAM console at https://console.aws.amazon.com/iam/.
2. In the navigation pane, choose **Roles**.
3. In the search box, type **applicationautoscaling** to narrow the list of roles to select, and then choose **ApplicationAutoScalingForAmazonAppStreamAccess**. If this role is listed, select it to view the role **Summary** page.
4. On the **Permissions** tab, confirm whether the **ApplicationAutoScalingForAmazonAppStreamAccess** permissions policy is attached.
5. Return to the role **Summary** page.
6. On the **Trust relationships** tab, choose **Show policy document**, and then confirm whether the **ApplicationAutoScalingForAmazonAppStreamAccess** trust relationship policy is attached and follows the correct format. If so, the trust relationship is correctly configured. Choose **Cancel** and close the IAM console.

ApplicationAutoScalingForAmazonAppStreamAccess trust relationship policy

The **ApplicationAutoScalingForAmazonAppStreamAccess** trust relationship policy must include the Application Auto Scaling service as the principal. This policy must also include the `sts:AssumeRole` action. The following policy configuration defines Application Auto Scaling as a trusted entity.

JSON

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "application-autoscaling.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

Checking for the AWSServiceRoleForApplicationAutoScaling_AppStreamFleet Service-Linked Role and Policies

Complete the steps in this section to check whether the AWSServiceRoleForApplicationAutoScaling_AppStreamFleet service-linked role is present and has the correct policies attached. If this role is not in your account and must be created, you or an administrator with the required permissions must perform the steps to get started with WorkSpaces Applications in your Amazon Web Services account.

To check whether the AWSServiceRoleForApplicationAutoScaling_AppStreamFleet IAM service-linked role is present

1. Open the IAM console at https://console.aws.amazon.com/iam/.
2. In the navigation pane, choose **Roles**.
3. In the search box, type **applicationautoscaling** to narrow the list of roles to select, and then choose AWSServiceRoleForApplicationAutoScaling_AppStreamFleet. If this role is listed, select it to view the role **Summary**page.
4. On the **Permissions** tab, confirm whether the AWSApplicationAutoscalingAppStreamFleetPolicy permissions policy is attached.
5. Return to the **Role** summary page.
6. On the **Trust relationships** tab, choose **Show policy document**, and then confirm whether the AWSServiceRoleForApplicationAutoScaling_AppStreamFleet trust relationship policy is attached and follows the correct format. If so, the trust relationship is correctly configured. Choose **Cancel** and close the IAM console.

AWSServiceRoleForApplicationAutoScaling_AppStreamFleet trust relationship policy

The AWSServiceRoleForApplicationAutoScaling_AppStreamFleet trust relationship policy must include **appstream.application-autoscaling.amazonaws.com** as the principal. This policy must also include the sts:AssumeRole action. The following policy configuration defines **appstream.application-autoscaling.amazonaws.com** as a trusted entity.

JSON

```
{
```

```
  "Version":"2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "appstream.application-autoscaling.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

Checking for the AmazonAppStreamPCAAccess Service Role and Policies

Complete the steps in this section to check whether the **AmazonAppStreamPCAAccess** service role is present and has the correct policies attached. If this role is not in your account and must be created, you or an administrator with the required permissions must perform the steps to get started with WorkSpaces Applications in your Amazon Web Services account.

To check whether the AmazonAppStreamPCAAccess IAM service role is present

1. Open the IAM console at https://console.aws.amazon.com/iam/.
2. In the navigation pane, choose **Roles**.
3. In the search box, type **appstreampca** to narrow the list of roles to select, and then choose **AmazonAppStreamPCAAccess**. If this role is listed, select it to view the role **Summary**page.
4. On the **Permissions** tab, confirm whether the **AmazonAppStreamPCAAccess** permissions policy is attached.
5. Return to the **Role** summary page.
6. On the **Trust relationships** tab, choose **Show policy document**, and then confirm whether the **AmazonAppStreamPCAAccess** trust relationship policy is attached and follows the correct format. If so, the trust relationship is correctly configured. Choose **Cancel** and close the IAM console.

AmazonAppStreamPCAAccess trust relationship policy

The **AmazonAppStreamPCAAccess** trust relationship policy must include prod.euc.ecm.amazonaws.com as the principal. This policy must also include the `sts:AssumeRole` action. The following policy configuration defines ECM as a trusted entity.

To create the AmazonAppStreamPCAAccess trust relationship policy using the AWS CLI

1. Create a JSON file named `AmazonAppStreamPCAAccess.json` with the following text.

   JSON

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "Service": [
                    "prod.euc.ecm.amazonaws.com"
                ]
            },
            "Action": "sts:AssumeRole",
            "Condition": {}
        }
    ]
}
```

2. Adjust the `AmazonAppStreamPCAAccess.json` path as needed and run the following AWS CLI commands to create the trust relationship policy and attach the AmazonAppStreamPCAAccess managed policy. For more information about the managed policy, see the section called “AWS Managed Policies Required to Access WorkSpaces Applications Resources”.

```
aws iam create-role --path /service-role/ --role-name AmazonAppStreamPCAAccess --
assume-role-policy-document file://AmazonAppStreamPCAAccess.json
```

```
aws iam attach-role-policy –role-name AmazonAppStreamPCAAccess –policy-arn
 arn:aws:iam::aws:policy/AmazonAppStreamPCAAccess
```

Using IAM Policies to Manage Administrator Access to Application Auto Scaling

Automatic scaling for fleets is made possible by a combination of the WorkSpaces Applications, Amazon CloudWatch, and Application Auto Scaling APIs. WorkSpaces Applications fleets are created with WorkSpaces Applications, alarms are created with CloudWatch, and scaling policies are created with Application Auto Scaling.

In addition to having the permissions defined in the AmazonAppStreamFullAccess policy, the IAM user that accesses fleet scaling settings must have the required permissions for the services that support dynamic scaling. IAM users must have permissions to use the actions shown in the following example policy.

JSON

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
          "appstream:*",
          "application-autoscaling:*",
          "cloudwatch:DeleteAlarms",
          "cloudwatch:DescribeAlarmsForMetric",
          "cloudwatch:DisableAlarmActions",
          "cloudwatch:DescribeAlarms",
          "cloudwatch:EnableAlarmActions",
          "cloudwatch:ListMetrics",
          "cloudwatch:PutMetricAlarm",
          "iam:ListRoles"
      ],
      "Resource": "*"
    },
    {
      "Sid": "iamPassRole",
      "Effect": "Allow",
      "Action": [
          "iam:PassRole"
      ],
      "Resource": "*",
```

```
      "Condition": {
         "StringEquals": {
             "iam:PassedToService": "application-autoscaling.amazonaws.com"
          }
      }
    }
  ]
}
```

You can also create your own IAM policies to set more specific permissions for calls to the Application Auto Scaling API. For more information, see Authentication and Access Control in the *Application Auto Scaling User Guide*.

Using IAM Policies to Manage Administrator Access to the Amazon S3 Bucket for Home Folders and Application Settings Persistence

The following examples show how you can use IAM policies to manage access to the Amazon S3 bucket for home folders and application settings persistence.

Examples

- Deleting the Amazon S3 Bucket for Home Folders and Application Settings Persistence
- Restricting Administrator Access to the Amazon S3 Bucket for Home Folders and Application Settings Persistence

Deleting the Amazon S3 Bucket for Home Folders and Application Settings Persistence

WorkSpaces Applications adds an Amazon S3 bucket policy to the buckets that it creates to prevent them from being accidentally deleted. To delete an S3 bucket, you must first delete the S3 bucket policy. Following are the bucket policies that you must delete for home folders and application settings persistence.

Home folders policy

JSON

```
{
```

```
  "Version":"2012-10-17",
  "Statement": [
    {
      "Sid": "PreventAccidentalDeletionOfBucket",
      "Effect": "Deny",
      "Principal": "*",
      "Action": "s3:DeleteBucket",
      "Resource": "arn:aws:s3:::appstream2-36fb080bb8-region-code-123456789012-
without-hyphens"
    }
  ]
}
```

Application settings persistence policy

JSON

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Sid": "PreventAccidentalDeletionOfBucket",
      "Effect": "Deny",
      "Principal": "*",
      "Action": "s3:DeleteBucket",
      "Resource": "arn:aws:s3:::appstream-app-settings-region-code-123456789012-
without-hyphens-unique-identifier"
     }
   ]
}
```

For more information, see Deleting or Emptying a Bucket in the *Amazon Simple Storage Service User Guide*.

Restricting Administrator Access to the Amazon S3 Bucket for Home Folders and Application Settings Persistence

By default, administrators who can access the Amazon S3 buckets created by WorkSpaces Applications can view and modify content that is part of users' home folders and persistent

application settings. To restrict administrator access to the S3 buckets that contain user files, we recommend applying the S3 bucket access policy based on the following template:

```
{
  "Sid": "RestrictedAccess",
  "Effect": "Deny",
  "NotPrincipal":
  {
    "AWS": [
      "arn:aws:iam::account:role/service-role/AmazonAppStreamServiceAccess",
      "arn:aws:sts::account:assumed-role/AmazonAppStreamServiceAccess/PhotonSession",
      "arn:aws:iam::account:user/IAM-user-name"
    ]
  },
    "Action": "s3:*",
    "Resource": "arn:aws:s3:::home-folder-or-application-settings-persistence-s3-
bucket-region-account"
  }
 ]
}
```

This policy allows S3 bucket access only to the users specified and to the WorkSpaces Applications service. For every IAM user who should have access, replicate the following line:

```
"arn:aws:iam::account:user/IAM-user-name"
```

In the following example, the policy restricts access to the home folder S3 bucket for anyone other than IAM users marymajor and johnstiles. It also allows access to the WorkSpaces Applications service, in AWS Region US West (Oregon) for account ID 123456789012.

```
{
  "Sid": "RestrictedAccess",
  "Effect": "Deny",
  "NotPrincipal":
  {
    "AWS": [
      "arn:aws:iam::123456789012:role/service-role/AmazonAppStreamServiceAccess",
      "arn:aws:sts::123456789012:assumed-role/AmazonAppStreamServiceAccess/
PhotonSession",
      "arn:aws:iam::123456789012:user/marymajor",
      "arn:aws:iam::123456789012:user/johnstiles"
    ]
```

```
  },
    "Action": "s3:*",
    "Resource": "arn:aws:s3:::appstream2-36fb080bb8-us-west-2-123456789012"
  }
 ]
}
```

Using an IAM Role to Grant Permissions to Applications and Scripts Running on WorkSpaces Applications Streaming Instances

Applications and scripts that run on WorkSpaces Applications streaming instances must include AWS credentials in their AWS API requests. You can create an IAM role to manage these credentials. An IAM role specifies a set of permissions that you can use to access AWS resources. This role is not uniquely associated with one person, however. Instead, it can be assumed by anyone that needs it.

You can apply an IAM role to an WorkSpaces Applications streaming instance. When the streaming instance switches to (assumes) the role, the role provides temporary security credentials. Your application or scripts use these credentials to perform API actions and management tasks on the streaming instance. WorkSpaces Applications manages the temporary credential switch for you.

Contents

- Best Practices for Using IAM Roles With WorkSpaces Applications Streaming Instances
- Configuring an Existing IAM Role to Use With WorkSpaces Applications Streaming Instances
- How to Create an IAM Role to Use With WorkSpaces Applications Streaming Instances
- How to Use the IAM Role With WorkSpaces Applications Streaming Instances

Best Practices for Using IAM Roles With WorkSpaces Applications Streaming Instances

When you use IAM roles with WorkSpaces Applications streaming instances, we recommend that you follow these practices:

- Limit the permissions that you grant to AWS API actions and resources.

  Follow least privilege principles when you create and attach IAM policies to the IAM roles associated with WorkSpaces Applications streaming instances. When you use an application or script that requires access to AWS API actions or resources, determine the specific actions and

resources that are required. Then, create policies that allow the application or script to perform only those actions. For more information, see Grant Least Privilege in the *IAM User Guide*.

- Create an IAM role for each WorkSpaces Applications resource.

  Creating a unique IAM role for each WorkSpaces Applications resource is a practice that follows least privilege principles. Doing so also lets you modify permissions for a resource without affecting other resources.

- Limit where the credentials can be used.

  IAM policies let you define the conditions under which your IAM role can be used to access a resource. For example, you can include conditions to specify a range of IP addresses that requests can come from. Doing so prevents the credentials from being used outside of your environment. For more information, see Use Policy Conditions for Extra Security in the *IAM User Guide*.

Configuring an Existing IAM Role to Use With WorkSpaces Applications Streaming Instances

This topic describes how to configure an existing IAM role so that you can use it with image builders and fleet streaming instances.

Prerequisites

The IAM role that you want to use with an WorkSpaces Applications image builder or fleet streaming instance must meet the following prerequisites:

- The IAM role must be in the same Amazon Web Services account as the WorkSpaces Applications streaming instance.
- The IAM role cannot be a service role.
- The trust relationship policy that is attached to the IAM role must include the WorkSpaces Applications service as the principal. A *principal* is an entity in AWS that can perform actions and access resources. The policy must also include the `sts:AssumeRole` action. This policy configuration defines WorkSpaces Applications as a trusted entity.

- If you are applying the IAM role to an image builder, the image builder must run a version of the WorkSpaces Applications agent released on or after September 3, 2019. If you are applying the IAM role to a fleet, the fleet must use an image that uses a version of the agent released on or after the same date. For more information, see WorkSpaces Applications Agent Release Notes.

To enable the WorkSpaces Applications service principal to assume an existing IAM role

To perform the following steps, you must sign into the account as an IAM user who has the permissions required to list and update IAM roles. If you don't have the required permissions, ask your Amazon Web Services account administrator either to perform these steps in your account or to grant you the required permissions.

1. Open the IAM console at https://console.aws.amazon.com/iam/.
2. In the navigation pane, choose **Roles**.
3. In the list of roles in your account, choose the name of the role that you want to modify.
4. Choose the **Trust relationships** tab, and then choose **Edit trust relationship**.
5. Under **Policy Document**, verify that the trust relationship policy includes the `sts:AssumeRole` action for the `appstream.amazonaws.com` service principal:

   JSON

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "appstream.amazonaws.com"
        ]
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

6. When you are finished editing your trust policy, choose **Update Trust Policy** to save your changes.
7. The IAM role that you selected will display in the WorkSpaces Applications console. This role grants permissions to applications and scripts to perform API actions and management tasks on streaming instances.

How to Create an IAM Role to Use With WorkSpaces Applications Streaming Instances

This topic describes how to create a new IAM role so that you can use it with image builders and fleet streaming instances.

1. Open the IAM console at https://console.aws.amazon.com/iam/.
2. In the navigation pane, choose **Roles**, and then choose **Create role**.
3. For **Select type of trusted entity**, choose **AWS service**.
4. From the list of AWS services, choose **WorkSpaces Applications**.
5. Under **Select your use case**, **WorkSpaces Applications — Allows WorkSpaces Applications instances to call AWS services on your behalf** is already selected. Choose **Next: Permissions**.
6. If possible, select the policy to use for the permissions policy or choose **Create policy** to open a new browser tab and create a new policy from scratch. For more information, see step 4 in the procedure Creating IAM Policies (Console) in the *IAM User Guide*.

   After you create the policy, close that tab and return to your original tab. Select the check box next to the permissions policies that you want WorkSpaces Applications to have.
7. (Optional) Set a permissions boundary. This is an advanced feature that is available for service roles, but not service-linked roles. For more information, see Permissions Boundaries for IAM Entities in the *IAM User Guide*.
8. Choose **Next: Tags**. You can optionally attach tags as key-value pairs. For more information, see Tagging IAM Users and Roles in the *IAM User Guide*.
9. Choose **Next: Review**.
10. For **Role name**, type a role name that is unique within your Amazon Web Services account. Because other AWS resources might reference the role, you can't edit the name of the role after it has been created.
11. For **Role description**, keep the default role description or type a new one.
12. Review the role, and then choose **Create role**.

How to Use the IAM Role With WorkSpaces Applications Streaming Instances

After you create an IAM role, you can apply it to an image builder or fleet streaming instance when you launch the image builder or create a fleet. You can also apply an IAM role to existing fleets. For

information about how to apply IAM role when you launch an image builder, see Launch an Image Builder to Install and Configure Streaming Applications. For information about how to apply IAM role when you create a fleet, see Create a Fleet in Amazon WorkSpaces Applications.

When you apply an IAM role to your image builder or fleet streaming instance, WorkSpaces Applications retrieves temporary credentials and creates the **appstream_machine_role** credential profile on the instance. The temporary credentials are valid for 1 hour, and new credentials retrieved every hour. The previous credentials do not expire, so you can use them for as long as they are valid. You can use the credential profile to call AWS services programmatically by using the AWS Command Line Interface (AWS CLI), AWS Tools for PowerShell, or the AWS SDK with the language of your choice.

When you make the API calls, specify **appstream_machine_role** as the credential profile. Otherwise, the operation fails due to insufficient permissions.

WorkSpaces Applications assumes the specified role while the streaming instance is provisioned. Because WorkSpaces Applications uses the elastic network interface that is attached to your VPC for AWS API calls, your application or script must wait for the elastic network interface to become available before making AWS API calls. If API calls are made before the elastic network interface is available, the calls fail.

The following examples show how you can use the **appstream_machine_role** credential profile to describe streaming instances (EC2 instances) and to create the Boto client. Boto is the Amazon Web Services (AWS) SDK for Python.

Describe Streaming Instances (EC2 instances) by Using the AWS CLI

```
aws ec2 describe-instances --region us-east-1 --profile appstream_machine_role
```

Describe Streaming Instances (EC2 instances) by Using AWS Tools for PowerShell

You must use AWS Tools for PowerShell version 3.3.563.1 or later, with the Amazon Web Services SDK for .NET version 3.3.103.22 or later. You can download the AWS Tools for Windows installer, which includes AWS Tools for PowerShell and the Amazon Web Services SDK for .NET, from the AWS Tools for PowerShell website.

```
Get-EC2Instance -Region us-east-1 -ProfileName appstream_machine_role
```

Creating the Boto Client by Using the AWS SDK for Python

```
session = boto3.Session(profile_name='appstream_machine_role')
```

SELinux on Red Hat Enterprise Linux and Rocky Linux

By default, Security Enhanced Linux (SELinux) is `enabled` and set to `enforcing` mode for WorkSpaces Applications image builders and streaming instances powered by Red Hat Enterprise Linux and Rocky Linux. In `enforcing` mode, permission denials are enforced. SELinux is a collection of kernel features and utilities to provide a strong, flexible, mandatory access control (MAC) architecture to the major subsystems of the kernel.

SELinux provides an enhanced mechanism to enforce the separation of information based on confidentiality and integrity requirements. This separation of information reduces threats of tampering and bypassing of application security mechanisms. It also confines damage that can be caused by malicious or flawed applications.

SELinux includes a set of sample security policy configuration files that's designed to meet everyday security goals. For more information about SELinux features and functionality, see What is SELinux?

Cookie-Based Authentication in Amazon WorkSpaces Applications

WorkSpaces Applications uses browser cookies to authenticate streaming sessions and allow users to reconnect to an active session without re-entering their sign-in credentials every time. Authentication tokens are stored in browser cookies for every authentication scenario. While cookies are necessary for many online services, they can potentially be vulnerable to cookie theft attacks. We strongly recommend that you take proactive measures to prevent cookie theft, such as implementing robust endpoint protection solutions for your users' devices. Furthermore, to mitigate the potential impact in the event of cookie theft, we advise you to consider the following actions:

- **Enforce single-session limit**: For your WorkSpaces Applications Windows images, create a registry key under `HKEY_USERS\S-1-5-18\Software\GSettings\com\nicesoftware\dcv\session-management` with the name **max-concurrent-clients** set to 1 to only allow one connection at a time. This limits the number of concurrent session to one, and blocks mirroring of active sessions. For more information, see session-management Parameters.
- **Enforce session expiry and re-authentication**
  - Reduce the SessionDuration value so that the authentication token expires after the user successfully starts the streaming session. Reusing authentication cookies after the

sessionDuration expires requires users to re-authenticate themselves. SessionDuration specifies the maximum amount of time that a federated streaming session for a user can remain active before re-authentication is required. The default value is 60 minutes. For more information, see the section called “Step 5: Create Assertions for the SAML Authentication Response”.

- To help maximize security, users should end sessions properly with the toolbar (terminate session), instead of closing the streaming window. Ending the session through the toolbar terminates both the user session and the streaming instance. This requires re-authentication for future access, preventing cookie misuse. If a user closes the streaming window without ending the session, the session and instance remains active for a configurable disconnect timeout period (in minutes). The disconnect timeout must be a number between 1 and 5760, with a default value of 15 minutes. To prevent misuse of inactive sessions, we recommend setting a short disconnect timeout. For more information, see the section called “Create a Fleet”.

- **Limit access to stream WorkSpaces Applications applications to your IP ranges**: We recommend that you implement IP-based IAM policies. This ensures that WorkSpaces Applications sessions can only be accessed from clients whose IP address belongs to an authorized IP range. All connection attempts initiated by a user whose client's IP address is outside an authorized range will be denied, even if they are presenting an otherwise valid authentication cookie (potentially stolen from a user). For more information, see Limit access to stream Amazon AppStream 2.0 applications to your IP ranges.
- **Add additional authentication**: To launch domain-joined streaming instances, you can join your WorkSpaces Applications Always-On and On-Demand Windows fleets and image builders to domains in Microsoft Active Directory, and use your existing Active Directory domains, either cloud-based or on-premises. After the initial SAML-based authentication, your users will be prompted to provide their domain credentials for additional authentication against the organizational domain. For more information, see *Using Active Directory* .

If you have any concerns or need help, contact AWS Support Center.

Logging and Monitoring in Amazon WorkSpaces Applications

Monitoring is an important part of maintaining the reliability, availability, and performance of Amazon WorkSpaces Applications. This topic describes the services and tools that AWS provides for monitoring your WorkSpaces Applications resources and responding to potential incidents.

Amazon CloudWatch Alarms

Amazon CloudWatch alarms let you watch a single metric over a time period that you specify. If the metric exceeds a given threshold, a notification is sent to an Amazon Simple Notification Service topic or AWS Auto Scaling policy. CloudWatch alarms do not invoke actions that are in a particular state. Instead, the state must have changed and been maintained for a specified number of periods. For more information, see Monitoring Amazon WorkSpaces Applications Resources.

> **Note**
>
> WorkSpaces Applications currently can't be configured as a target for CloudWatch Events. For a list of services that you can configure as targets for CloudWatch events, see What Is Amazon CloudWatch Events.

AWS CloudTrail

AWS CloudTrail provides a record of actions taken by a user, role, or an AWS service in WorkSpaces Applications. This record lets you determine the request that was made to WorkSpaces Applications, the IP address from which the request was made, who made the request, when it was made, and additional details. For more information, see Logging Amazon WorkSpaces Applications API Calls with AWS CloudTrail.

AWS Trusted Advisor

AWS Trusted Advisor inspects your AWS environment and then recommends ways to save money, improve system availability and performance, or help close security gaps. Trusted Advisor uses best practices collected from a wide variety of AWS customers. All AWS customers have access to five Trusted Advisor checks. If you have a Business or Enterprise support plan, you can view all Trusted Advisor checks.

When you enable application settings persistence or home folders for your users, the data that is generated by your users is stored in Amazon S3 buckets. Trusted Advisor contains the following checks related to Amazon S3:

- Logging configuration of Amazon S3 buckets.
- Security checks for Amazon S3 buckets that have open access permissions.
- Fault tolerance checks for Amazon S3 buckets that don't have versioning enabled, or have versioning suspended.

For more information, see AWS Trusted Advisor in the *AWS Support User Guide*.

Amazon S3 Access Logs

If your users have application settings data or home folders data stored in Amazon S3 buckets, consider viewing Amazon S3 server access logs to monitor access. These logs provide detailed records about requests that are made to a bucket. Server access logs are useful for many applications. For example, access log information can be useful in security and access audits. For more information, see Amazon S3 Server Access Logging in the *Amazon Simple Storage Service User Guide*.

WorkSpaces Applications Usage Reports

You can subscribe to WorkSpaces Applications usage reports to receive detailed reports about how your users are using the service. The reports include how long users stream and which applications they launch. For more information, see WorkSpaces Applications Usage Reports.

Compliance Validation for Amazon WorkSpaces Applications

Third-party auditors assess the security and compliance of Amazon WorkSpaces Applications as part of multiple AWS compliance programs. These include the following: SOC, PCI, ISO, FedRAMP, HIPAA, MTCS, ENS High, HITRUST CSF, VPAT, and others.

> **Note**
>
> WorkSpaces Applications supports FIPS 140-3. For information about how to use WorkSpaces Applications FIPS endpoints for administrative use or streaming, see the section called "FIPS Endpoints".
> WorkSpaces Applications is also undergoing assessment for the Department of Defense (DoD) Cloud Computing Security Requirements Guide (SRG).

For a list of AWS services in scope of specific compliance programs, see AWS Services in Scope by Compliance Program. For general information, see AWS Compliance Programs.

You can download third-party audit reports using AWS Artifact. For more information, see Downloading Reports in AWS Artifact.

Your compliance responsibility when using WorkSpaces Applications is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. AWS provides the following resources to help with compliance:

- Security and Compliance Quick Start Guides – These deployment guides discuss architectural considerations and provide steps for deploying security- and compliance-focused baseline environments on AWS.
- Architecting for HIPAA Security and Compliance Whitepaper – This whitepaper describes how companies can use AWS to create HIPAA-compliant applications.
- AWS Compliance Resources – This collection of workbooks and guides might apply to your industry and location.
- Evaluating Resources with Rules in the *AWS Config Developer Guide* – The AWS Config service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- AWS Security Hub CSPM – This AWS service provides a comprehensive view of your security state within AWS that helps you check your compliance with security industry standards and best practices.

Resilience in Amazon WorkSpaces Applications

The AWS global infrastructure is built around AWS Regions and Availability Zones. Regions provide multiple physically separated and isolated Availability Zones, which are connected through low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure.

Infrastructure Security in Amazon WorkSpaces Applications

As a managed service, Amazon WorkSpaces Applications is protected by AWS global network security. For information about AWS security services and how AWS protects infrastructure, see AWS Cloud Security. To design your AWS environment using the best practices for infrastructure security, see Infrastructure Protection in *Security Pillar AWS Well-Architected Framework*.

You use AWS published API calls to access WorkSpaces Applications through the network. Clients must support the following:

- Transport Layer Security (TLS). We require TLS 1.2 and recommend TLS 1.3.
- Cipher suites with perfect forward secrecy (PFS) such as DHE (Ephemeral Diffie-Hellman) or ECDHE (Elliptic Curve Ephemeral Diffie-Hellman). Most modern systems such as Java 7 and later support these modes.

The following topics provide additional information about WorkSpaces Applications infrastructure security.

Contents

- Network Isolation
- Isolation on Physical Hosts
- Controlling Network Traffic
- WorkSpaces Applications Interface VPC Endpoints
- Protecting Data in Transit with FIPS Endpoints

Network Isolation

A virtual private cloud (VPC) is a virtual network in your own logically isolated area in the Amazon Web Services Cloud. Use separate VPCs to isolate infrastructure by workload or organizational entity.

A subnet is a range of IP addresses in a VPC. When you launch an instance, you launch it into a subnet in your VPC. Use subnets to isolate the tiers of your application (for example, web, application, and database) within a single VPC. Use private subnets for your instances if they should not be accessed directly from the internet.

You can stream from WorkSpaces Applications streaming instances in your VPC without going through the public internet. To do so, use an interface VPC endpoint (interface endpoint). For more information, see Tutorial: Creating and Streaming from Interface VPC Endpoints.

You can also call WorkSpaces Applications API operations from your VPC without sending traffic over the public internet by using an interface endpoint. For information, see Access WorkSpaces Applications API Operations and CLI Commands Through an Interface VPC Endpoint.

Isolation on Physical Hosts

Different streaming instances on the same physical host are isolated from each other as though they are on separate physical hosts. The hypervisor isolates CPU and memory, and the instances are provided virtualized disks instead of access to the raw disk devices.

When you stop or terminate a streaming instance, the memory allocated to it is scrubbed (meaning, it's set to zero) by the hypervisor before it is allocated to a new instance, and every block of storage is reset. This ensures that your data is not exposed to another instance.

Controlling Network Traffic

To help control network traffic to your WorkSpaces Applications streaming instances, consider these options:

- When you launch an Amazon AppStream streaming instance, you launch it into a subnet in your VPC. You can deploy streaming instances in a private subnet if they should not be accessible from the internet.
- To provide internet access to your streaming instances in a private subnet, use a NAT gateway. For more information, see Configure a VPC with Private Subnets and a NAT Gateway.
- Security groups that belong to your VPC let you control the network traffic between WorkSpaces Applications streaming instances and VPC resources such as license servers, file servers, and database servers. Security groups also isolate traffic between your streaming instances and WorkSpaces Applications management services.

  Use security groups to restrict access to your streaming instances. For example, you can allow traffic only from the address ranges for your corporate network. For more information, see Security Groups in Amazon WorkSpaces Applications.
- You can stream from WorkSpaces Applications streaming instances in your VPC without going through the public internet. To do so, use an interface VPC endpoint (interface endpoint). For more information, see Tutorial: Creating and Streaming from Interface VPC Endpoints.

  You can also call WorkSpaces Applications API operations from your VPC without sending traffic over the public internet by using an interface endpoint. For more information, see Access WorkSpaces Applications API Operations and CLI Commands Through an Interface VPC Endpoint.
- Use IAM roles and policies to manage administrator access to WorkSpaces Applications, Application Auto Scaling, and Amazon S3 buckets. For more information, see the following topics:

- Using AWS Managed Policies and Linked Roles to Manage Administrator Access to WorkSpaces Applications Resources
- Using IAM Policies to Manage Administrator Access to Application Auto Scaling
- Restricting Administrator Access to the Amazon S3 Bucket for Home Folders and Application Settings Persistence

- You can use SAML 2.0 to federate authentication to WorkSpaces Applications. For more information, see Amazon WorkSpaces Applications Service Quotas.

> **Note**
>
> For smaller WorkSpaces Applications deployments, you can use WorkSpaces Applications user pools. By default, user pools support a maximum of 50 users. For more information about WorkSpaces Applications quotas (also referred to as limits), see Amazon WorkSpaces Applications Service Quotas. For deployments that must support 100 or more WorkSpaces Applications users, we recommend using SAML 2.0.

WorkSpaces Applications Interface VPC Endpoints

A virtual private cloud (VPC) is a virtual network in your own logically isolated area in the Amazon Web Services Cloud. If you use Amazon Virtual Private Cloud to host your AWS resources, you can establish a private connection between your VPC and WorkSpaces Applications. You can use this connection to enable WorkSpaces Applications to communicate with your resources on your VPC without going through the public internet.

Interface endpoints are powered by AWS PrivateLink, a technology that lets you keep streaming traffic within a VPC that you specify by using private IP addresses. When you use the VPC with an Direct Connect or AWS Virtual Private Network tunnel, you can keep the streaming traffic within your network.

The following topics provide information about WorkSpaces Applications interface endpoints.

Contents

- Tutorial: Creating and Streaming from Interface VPC Endpoints
- Access WorkSpaces Applications API Operations and CLI Commands Through an Interface VPC Endpoint

Tutorial: Creating and Streaming from Interface VPC Endpoints

You can use an interface VPC endpoint in your Amazon Web Services account to restrict all network traffic between your Amazon VPC and WorkSpaces Applications to the Amazon network. After you create this endpoint, you configure your WorkSpaces Applications stack or image builder to use it.

> **Note**
>
> Streaming from Interface VPC Endpoints is currently not supported for fleets powered by Windows Desktop operating system (such as Windows 11).

Prerequisites

Before you set up interface VPC endpoints for WorkSpaces Applications, be aware of the following prerequisites:

- Internet connectivity is required to authenticate users and deliver the web assets that WorkSpaces Applications requires to function. The streaming interface endpoint maintains the streaming traffic within your VPC. Streaming traffic includes pixel, USB, user input, audio, clipboard, file upload and download, and printer traffic. To allow this traffic, you must allow the domains listed in Allowed Domains. After creating the VPC endpoint, you must allow the WorkSpaces Applications user authentication domains. However, for the streaming gateways, you can restrict access to just <vpc-endpoint-id>.streaming.appstream.<aws-region>.vpce.amazonaws.com. Allow listing to *.amazonappstream.com is not required. The VPC endpoint fully qualified domain name replaces that dependency.
- The network to which your users' devices are connected must be able to route traffic to the interface endpoint.
- The security groups that are associated with the interface endpoint must allow inbound access to port 443 (TCP) and ports 1400-1499 (TCP) from the IP address range from which your users connect.
- The network access control list for the subnets must allow outbound traffic from ephemeral network ports 1024-65535 (TCP) to the IP address range from which your users connect.
- You must have an IAM permissions policy in your AWS account that provides permissions to perform the `ec2:DescribeVpcEndpoints` API action. By default, this permission is defined in the IAM policy that is attached to the AmazonAppStreamServiceAccess role. If you have the required permissions, this service role is automatically created by WorkSpaces Applications, with

the required IAM policies attached, when you get started with the WorkSpaces Applications service in an AWS Region. For more information, see Identity and Access Management for Amazon WorkSpaces Applications.

To create an interface endpoint

1. Open the Amazon VPC console at https://console.aws.amazon.com/vpc/.
2. In the navigation pane, choose **Endpoints**, **Create Endpoint**.
3. Choose **Create Endpoint**.
4. For **Service category**, ensure that **AWS services** is selected.
5. For **Service Name**, choose **`com.amazonaws.`***`<AWS Region>`*`.appstream.streaming`**.
6. Specify the following information. When you're done, choose **Create endpoint**.

   - For **VPC**, choose a VPC in which to create the interface endpoint. You can choose a different VPC than the VPC with WorkSpaces Applications resources.
   - For **Subnets**, choose the subnets (Availability Zones) in which to create the endpoint network interfaces. We recommend that you choose subnets in at least two Availability Zones.
   - For **IP address type**, choose either IPV6 or IPV4.
   - Ensure that the **Enable Private DNS Name** check box is selected.

     > **Note**
     >
     > If your users use a network proxy to access streaming instances, disable any proxy caching on the domain and DNS names that are associated with the private endpoint. The VPC endpoint DNS name should be allowed through the proxy.

   - For **Security group**, choose the security groups to associate with the endpoint network interfaces.

     > **Note**
     >
     > The security groups must provide inbound access to the ports from the IP address range from which your users connect.

While your interface endpoint is being created, the status of the endpoint in the console appears as **Pending**. After your endpoint is created, the status changes to **Available.**

To update a stack to use the interface endpoint that you created for streaming sessions, perform the following steps.

To update a stack to use a new interface endpoint

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.

   Ensure that you open the console in the same AWS Region as the interface endpoint that you want to use.
2. In the navigation pane, choose **Stacks**, and then choose the stack that you want.
3. Choose the **VPC Endpoints** tab, and then choose **Edit**.
4. In the **Edit VPC Endpoint** dialog box, for **Streaming Endpoint**, choose the endpoint through which to stream traffic.
5. Choose **Update**.

Traffic for new streaming sessions will be routed through this endpoint. However, traffic for current streaming sessions continues to be routed through the previously specified endpoint.

> **Note**
>
> Users cannot stream using the internet endpoint when an interface endpoint is specified.

Access WorkSpaces Applications API Operations and CLI Commands Through an Interface VPC Endpoint

If you use Amazon Virtual Private Cloud to host your AWS resources, you can connect directly to WorkSpaces Applications API operations or command line interface (CLI) commands through an interface VPC endpoint (interface endpoint) in your virtual private cloud (VPC) instead of connecting over the internet. Interface endpoints are powered by AWS PrivateLink, a technology that lets you keep streaming traffic within a VPC that you specify by using private IP addresses. When you use an interface endpoint, communication between your VPC and WorkSpaces Applications is conducted entirely and securely within the AWS network.

> **Note**
>
> This topic describes how to access the WorkSpaces Applications API operations and CLI commands through an interface endpoint. For information about how to create and stream from WorkSpaces Applications interface endpoints, see Tutorial: Creating and Streaming from Interface VPC Endpoints.

Prerequisites

To use interface endpoints, you must meet the following prerequisites:

- The security groups that are associated with the interface endpoint must allow inbound access to port 443 (TCP) from the IP address range from which your users connect.
- The network access control list for the subnets must allow outbound traffic from ephemeral network ports 1024-65535 (TCP) to the IP address range from which your users connect.

Topics

- Create an Interface Endpoint to Access WorkSpaces Applications API Operations and CLI Commands
- Use an Interface Endpoint to Access WorkSpaces Applications API Operations and CLI Commands

Create an Interface Endpoint to Access WorkSpaces Applications API Operations and CLI Commands

Perform the following steps to create an interface endpoint.

1. Open the Amazon VPC console at https://console.aws.amazon.com/vpc/.
2. In the navigation pane, choose **Endpoints**, **Create Endpoint**.
3. Choose **Create Endpoint**.
4. For **Service category**, ensure that **AWS services** is selected.
5. For **Service Name**, choose **`com.amazonaws.`***`<AWS Region>`*`.appstream.api`**.
6. Specify the following information. When you're done, choose **Create endpoint**.
   - For **VPC**, select a VPC in which to create the interface endpoint.

- For **Subnets**, select the subnets (Availability Zones) in which to create the endpoint network interfaces. We recommend that you choose subnets in at least two Availability Zones.
- Optionally, you can select the **Enable Private DNS Name** check box.

  > **Note**
  >
  > If you select this option, ensure that you configure VPC and DNS as needed to support private DNS. For more information, see Private DNS in the *Amazon VPC User Guide*.

- For **Security group**, select the security groups to associate with the endpoint network interfaces.

  > **Note**
  >
  > The security groups must provide inbound access to the ports from the IP address range from which your users connect.

While your interface endpoint is being created, the status of the endpoint in the console appears as **Pending**. After your endpoint is created, the status changes to **Available.**

Use an Interface Endpoint to Access WorkSpaces Applications API Operations and CLI Commands

After the status of the interface VPC endpoint that you create changes to **Available**, you can use the endpoint to access WorkSpaces Applications API operations and CLI commands. To do so, specify the `endpoint-url` parameter with the DNS name of the interface endpoint when you use these operations and commands. The DNS name is publicly resolvable, but it only successfully routes traffic in your VPC.

The following example shows how to specify the DNS name of the interface endpoint when you use the **describe-fleets** CLI command:

```
aws appstream describe-fleets --endpoint-url <vpc-endpoint-id>.api.appstream.<aws-
region>.vpce.amazonaws.com
```

The following example shows how to specify the DNS name of the interface endpoint when you instantiate the WorkSpaces Applications Boto3 Python client:

```
appstream2client = boto3.client('appstream',region_name='<aws-
region>',endpoint_url='<vpc-endpoint-id>.api.appstream.<aws-region>.vpce.amazonaws.com'
```

Subsequent commands using the `appstream2client` object automatically use the interface endpoint that you specified.

If you enabled the private DNS host names on the interface endpoint, you don't need to specify the endpoint URL. The WorkSpaces Applications API DNS host name that the API and CLI use by default resolves within your VPC. For more information about private DNS host names, see Private DNS in the *Amazon VPC User Guide*.

Protecting Data in Transit with FIPS Endpoints

By default, when you communicate with the WorkSpaces Applications service, whether as an administrator using the WorkSpaces Applications console, the AWS Command Line Interface (AWS CLI), or an AWS SDK, or as a user streaming from an image builder or a fleet instance, all data in transit is encrypted using TLS 1.2.

If you require FIPS 140-3 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. WorkSpaces Applications offers FIPS endpoints in all United States AWS Regions where WorkSpaces Applications is available. When you use a FIPS endpoint, all data in transit is encrypted using cryptographic standards that comply with Federal Information Processing Standard (FIPS) 140-3. For information about FIPS endpoints, including a list of WorkSpaces Applications endpoints, see Federal Information Processing Standard (FIPS) 140-3.

Topics

- FIPS Endpoints for Administrative Use
- FIPS Endpoints for User Streaming Sessions
- Exceptions

FIPS Endpoints for Administrative Use

To specify a FIPS endpoint when you run an AWS CLI command for WorkSpaces Applications, use the `endpoint-url` parameter. The following example uses the WorkSpaces Applications FIPS endpoint in the US West (Oregon) Region to retrieve a list of all stacks in the Region:

```
aws appstream describe-stacks --endpoint-url https://appstream2-fips.us-
west-2.amazonaws.com
```

To specify a FIPS endpoint for WorkSpaces Applications API operations, use the procedure in your AWS SDK for specifying a custom endpoint.

FIPS Endpoints for User Streaming Sessions

If you use SAML 2.0 or a streaming URL to authenticate users, you can configure FIPS-compliant connections for your users' streaming sessions.

To use a FIPS-compliant connection for users who authenticate using SAML 2.0, specify an WorkSpaces Applications FIPS endpoint when you configure the relay state of your federation. For more information about constructing a relay state URL for identity federation using SAML 2.0, see Setting Up SAML.

To configure a FIPS-compliant connection for users who authenticate through a streaming URL, specify an WorkSpaces Applications FIPS endpoint when you call the CreateStreamingURL or CreateImageBuilderStreamingURL operation from the AWS CLI or an AWS SDK. A user who connects to a streaming instance using the resulting URL is connected through a FIPS-compliant connection. The following example uses the WorkSpaces Applications FIPS endpoint in the US East (Virginia) Region to generate a FIPS-compliant streaming URL:

```
aws appstream create-streaming-url --stack-name stack-name --fleet-name fleet-name --
user-id user-id --endpoint-url https://appstream2-fips.us-east-1.amazonaws.com
```

Exceptions

FIPS-compliant connections are not supported in the following scenarios:

- Administration of WorkSpaces Applications through the WorkSpaces Applications console
- Streaming sessions for users who authenticate using the WorkSpaces Applications user pool feature
- Streaming using an interface VPC endpoint
- Generating FIPS-compliant streaming URLs through the WorkSpaces Applications console
- Connections to your Google Drive or OneDrive storage accounts where your storage provider does not provide a FIPS endpoint

Security Groups in Amazon WorkSpaces Applications

You can provide additional access control to your VPC from streaming instances in a fleet or an image builder in Amazon WorkSpaces Applications by associating them with VPC security groups. Security groups that belong to your VPC allow you to control the network traffic between WorkSpaces Applications streaming instances and VPC resources such as license servers, file servers, and database servers. For more information, see Security Groups for your VPC in the *Amazon VPC User Guide*.

The rules that you define for your VPC security group are applied when the security group is associated with a fleet or image builder. The security group rules determine what network traffic is allowed from your streaming instances. For more information, see Security Group Rules in the *Amazon VPC User Guide*.

You can associate up to five security groups while launching a new image builder or while creating a new fleet. You can also associate security groups with an existing fleet or change the security groups for a fleet (to change security groups for a fleet, you must first stop the fleet). For more information, see Working with Security Groups in the *Amazon VPC User Guide*.

If you don't select a security group, your image builder or fleet is associated with the default security group for your VPC. For more information, see Default Security Group for Your VPC in the *Amazon VPC User Guide*.

Use these additional considerations when using security groups with WorkSpaces Applications.

- All end user data, such as internet traffic, home folder data, or application communication with VPC resources, are affected by the security groups associated with the streaming instance.
- Streaming pixel data is not affected by security groups.
- If you have enabled default internet access for your fleet or image builder, the rules of the associated security groups must allow internet access.

You can create or edit rules for your security groups or create new security groups using the Amazon VPC console.

- **To associate security groups with an image builder** — Follow the instructions at Launch an Image Builder to Install and Configure Streaming Applications.
- **To associate security groups with a fleet**

- *While creating the fleet* — Follow the instructions at Create a Fleet in Amazon WorkSpaces Applications.
- *For an existing fleet* — Edit the fleet settings using the AWS Management Console.

You can also associate security groups to your fleets using the AWS CLI and SDKs.

- **AWS CLI** — Use the create-fleet and update-fleet commands.
- **AWS SDKs** — Use the CreateFleet and UpdateFleet API operations.

For more information, see the AWS Command Line Interface User Guide and Tools for Amazon Web Services.

Update Management in Amazon WorkSpaces Applications

WorkSpaces Applications provides an automated way to update your image builder with newer WorkSpaces Applications software. When your images are configured to always use the latest WorkSpaces Applications agent version, your streaming instances are automatically updated with the latest features, performance improvements, and security updates that are available from AWS. For information about how to manage WorkSpaces Applications agent versions, see Manage WorkSpaces Applications Agent Versions.

You are responsible for installing and maintaining the updates for the Windows operating system, your applications, and their dependencies. For more information, see Keep Your Amazon WorkSpaces Applications Image Up-to-Date.

You can keep your WorkSpaces Applications image up-to-date by using managed WorkSpaces Applications image updates. This update method provides the latest Windows operating system updates and driver updates, and the latest WorkSpaces Applications agent software. For more information, see Update an Image by Using Managed WorkSpaces Applications Image Updates.

To manage updates for applications on your streaming instances, you can use any automatic update services provided. You can also follow the recommendations for installing updates provided by the application vendor.

Amazon WorkSpaces Applications Cross-Service Confused Deputy Prevention

The confused deputy problem is a security issue where an entity that doesn't have permission to perform an action coerces a more-privileged entity to perform the action. In AWS, cross-service impersonation can leave account resources vulnerable to the confused deputy problem. Cross-service impersonation occurs when one service (the *calling service*) calls another service (the *called service*). The calling service can manipulate the called service to use its permissions to act on a customer's resources in ways that the calling service doesn't have permission to perform for itself. To prevent this, AWS provides tools that helps you protect your data for all services with service principals that have access to resources in your account.

We recommend using the `aws:SourceArn` and `aws:SourceAccount` global condition context keys in resource policies to limit permissions when accessing these resources. The following guidelines detail recommendations and requirements when you use these keys to protect your resources:

- Use `aws:SourceArn` if you want only one resource associated with cross-service access.
- Use `aws:SourceAccount` if you want to allow any resource in the specified account associated with cross-service use.
- If the `aws:SourceArn` key doesn't contain an account ID, you must use both global condition context keys (`aws:SourceArn` and `aws:SourceAccount`) to limit permissions.
- If you use both global condition context keys and the `aws:SourceArn` value contains an account ID, the `aws:SourceAccount` key must use the same account ID when used in the same policy statement.

The most effective way to protect against the confused deputy problem is to use the exact Amazon Resource Name (ARN) of the resource you want to allow. If you don't know the full ARN of the resource, use the `aws:SourceArn` global context condition key with wildcards (such as *) for the unknown portions of the ARN. You can also use a wildcard in the ARN if you want to specify multiple resources. For example, you can format the ARN as `arn:aws:`*`servicename`*`::`*`region-name`*`::`*`your AWS account ID`*`:*`.

Topics

- Example: WorkSpaces Applications service role cross-service confused deputy prevention
- Example: WorkSpaces Applications fleet machine role cross-service confused deputy prevention

- Example: WorkSpaces Applications Elastic fleets session script Amazon S3 bucket policy cross-service confused deputy prevention
- Example: WorkSpaces Applications Application Amazon S3 bucket policy cross-service confused deputy prevention

Example: WorkSpaces Applications service role cross-service confused deputy prevention

WorkSpaces Applications assumes a service role using a variety of resource ARNs, which leads to a complicated conditional statement. We recommend using a wildcard resource type to prevent any unexpected WorkSpaces Applications resources failures.

Example `aws:SourceAccount` Conditional:

JSON

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "Service": [
                    "appstream.amazonaws.com"
                ]
            },
            "Action": "sts:AssumeRole",
            "Condition": {
                "StringEquals": {
                    "aws:SourceAccount": "your AWS account ID"
                }
            }
        }
    ]
}
```

Example aws:SourceArn Conditional:

JSON

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "Service": [
                    "appstream.amazonaws.com"
                ]
            },
            "Action": "sts:AssumeRole",
            "Condition": {
                "ArnLike": {
                "aws:SourceArn": "arn:aws:appstream:us-east-1:111122223333:*"
                }
            }
        }
    ]
}
```

Example: WorkSpaces Applications fleet machine role cross-service confused deputy prevention

Example aws:SourceAccount Conditional:

JSON

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "Service": [
                    "appstream.amazonaws.com"
                ]
```

```
            },
            "Action": "sts:AssumeRole",
            "Condition": {
                "StringEquals": {
                    "aws:SourceAccount": "your AWS account ID"
                }
            }
        }
    ]
}
```

Example aws:SourceArn Conditional:

> **Note**
>
> If you want to use one IAM role for multiple fleets, we recommend using the aws:SourceArn global context condition key with wildcards (*) to match multiple WorkSpaces Applications fleet resources.

JSON

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "Service": [
                    "appstream.amazonaws.com"
                ]
            },
            "Action": "sts:AssumeRole",
            "Condition": {
                "ArnLike": {
                "aws:SourceArn": "arn:aws:appstream:us-
east-1:111122223333:fleet/your-fleet-name"
                }
            }
        }
    ]
```

```
}
```

Example: WorkSpaces Applications Elastic fleets session script Amazon S3 bucket policy cross-service confused deputy prevention

Example `aws:SourceAccount` Conditional:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "Service": [
                    "appstream.amazonaws.com"
                ]
            },
            "Action": "s3:GetObject",
            "Resource": "arn:aws:s3:::your-bucket-name/your-session-script-path",
            "Condition": {
                "StringEquals": {
                    "aws:SourceAccount": "111122223333"
                }
            }
        }
    ]
}
```

Example `aws:SourceArn` Conditional:

JSON

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "Service": [
                    "appstream.amazonaws.com"
```

```
                ]
            },
            "Action": "s3:GetObject",
            "Resource": "arn:aws:s3:::bucket/AppStream2/*",
            "Condition": {
                "ArnLike": {
                "aws:SourceArn": "arn:aws:appstream:us-
east-1:111122223333:fleet/yourFleetName"
                }
            }
        }
    ]
}
```

Example: WorkSpaces Applications Application Amazon S3 bucket policy cross-service confused deputy prevention

When you store data in an Amazon S3 bucket, the bucket might be exposed to confused deputy issues. This can leave data such as Elastic fleets, app blocks, setup scripts, application icons, and session scripts vulnerable to malicious actors.

To prevent confused deputy issues, you can specify the `aws:SourceAccount` condition or the `aws:SourceArn` condition in the Amazon S3 bucket policy for `ELASTIC-FLEET-EXAMPLE-BUCKET`.

The resource policies below show how to prevent the confused deputy problem with either of the following:

- The `aws:SourceAccount` with your AWS account ID
- The global condition context key `aws:SourceArn`

WorkSpaces Applications currently doesn't support confused deputy prevention for application icons. WorkSpaces Applications only supports VHD files and setup scripts. If you try to add additional conditions for application icons, the icons won't be displayed to end users.

In the following example, the bucket policy only allows WorkSpaces Applications Elastic fleet resources in the owner's account to access `ELASTIC-FLEET-EXAMPLE-BUCKET`.

```
{
```

```
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ConfusedDeputyPreventionExamplePolicy",
            "Effect": "Allow",
            "Principal": {
                "Service": "appstream.amazonaws.com"
            },
            "Action": "s3:GetObject",
            "Resource": [
                "arn:aws:s3:::ELASTIC-FLEET-EXAMPLE-BUCKET/vhd-folder/*",
                "arn:aws:s3:::ELASTIC-FLEET-EXAMPLE-BUCKET/scripts/*"
            ],
            "Condition": {
                "StringEquals": {
                    "aws:SourceAccount": "111122223333"
                }
            }
        },
        {
            "Sid": "AllowRetrievalPermissionsToS3AppIconsForAppStream",
            "Effect": "Allow",
            "Principal": {
                "Service": "appstream.amazonaws.com"
            },
            "Action": "s3:GetObject",
            "Resource": "arn:aws:s3:::ELASTIC-FLEET-EXAMPLE-BUCKET/app-icons/*"
        }
    ]
}
```

You can also use the `aws:SourceArn` condition to limit resource access for specific resources.

> **Note**
>
> If you don't know the full ARN of a resource, or you want to specify multiple resources, use the `aws:SourceArn` global context condition key with wildcards (*) for the unknown portions of the ARN.

JSON

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Sid": "ConfusedDeputyPreventionExamplePolicy",
            "Effect": "Allow",
            "Principal": {
                "Service": "appstream.amazonaws.com"
            },
            "Action": "s3:GetObject",
            "Resource": [
                "arn:aws:s3:::ELASTIC-FLEET-EXAMPLE-BUCKET/vhd-folder/*",
                "arn:aws:s3:::ELASTIC-FLEET-EXAMPLE-BUCKET/scripts/*"
            ],
            "Condition": {
                "ArnLike": {
                    "aws:SourceArn": "arn:aws:appstream:us-
east-1:111122223333:app-block/*"
                }
            }
        },
        {
            "Sid": "AllowRetrievalPermissionsToS3AppIconsForAppStream",
            "Effect": "Allow",
            "Principal": {
                "Service": "appstream.amazonaws.com"
            },
            "Action": "s3:GetObject",
            "Resource": "arn:aws:s3:::ELASTIC-FLEET-EXAMPLE-BUCKET/app-icons/*"
        }
    ]
}
```

You can use the `aws:SourceArn` and `aws:SourceAccount` conditions to limit the resource access for specific resources and accounts.

Note

If you don't know the full ARN of a resources, or if you want to specify multiple resources, use the `aws:SourceArn` global context condition key with wildcards (*) for the unknown portions of the ARN.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "ConfusedDeputyPreventionExamplePolicy",
            "Effect": "Allow",
            "Principal": {
                "Service": "appstream.amazonaws.com"
            },
            "Action": "s3:GetObject",
            "Resource": [
                "arn:aws:s3:::ELASTIC-FLEET-EXAMPLE-BUCKET/vhd-folder/*",
                "arn:aws:s3:::ELASTIC-FLEET-EXAMPLE-BUCKET/scripts/*"
            ],
            "Condition": {
                "ArnLike": {
                    "aws:SourceArn": "arn:aws:appstream:us-east-1:111122223333:app-
block/*"
                },
                "StringEquals": {
                    "aws:SourceAccount": "111122223333"
                }
            }
        },
        {
            "Sid": "AllowRetrievalPermissionsToS3AppIconsForAppStream",
            "Effect": "Allow",
            "Principal": {
                "Service": "appstream.amazonaws.com"
            },
            "Action": "s3:GetObject",
            "Resource": "arn:aws:s3:::ELASTIC-FLEET-EXAMPLE-BUCKET/app-icons/*"
        }
    ]
}
```

Security Best Practices in Amazon WorkSpaces Applications

Cloud security at Amazon Web Services (AWS) is the highest priority. Security and compliance is a shared responsibility between AWS and the customer. For more information, refer to the Shared Responsibility Model. As an AWS and WorkSpaces Applications customer, it is important to implement security measures on different layers such as stack, fleet, image, and networking.

Due to its ephemeral nature, WorkSpaces Applications is often preferred as a secure solution to application and desktop delivery. Consider whether antivirus solutions that are commonplace in Windows deployments are relevant in your use cases for an environment that is predefined and purged at the end of a user session. Antivirus adds overhead to virtualized instances, making it is a best practice to mitigate unnecessary activities. For example, scanning the system volume (which is ephemeral) at boot, for instance, does not add to the overall security of WorkSpaces Applications.

The two key questions for security WorkSpaces Applications are centered on:

- Is persisting user state beyond the session a requirement?
- How much access should a user have within a session?

Topics

- Securing Persistent Data
- Endpoint Security and Antivirus
- Network Exclusions
- Securing an WorkSpaces Applications Session
- Firewalls and Routing
- Data Loss Prevention
- Controlling egress traffic
- Using AWS services

Securing Persistent Data

Deployments of WorkSpaces Applications can require the user state to persist in some form. It might be to persist data for individual users, or to persist data for collaboration using a shared folder. AppStream 2.0 instance storage is ephemeral and has no encryption option.

WorkSpaces Applications provides user state persistence through home folders and application settings in Amazon S3. Some use cases require greater control over user state persistence. For these use cases, AWS recommends using a Server Message Block (SMB) file share.

User state and data

Because most Windows applications perform best and most securely when co-located with application data created by the user, it is a best practice to keep this data in the same AWS Region as WorkSpaces Applications fleets. Encrypting this data is a best practice. The default behavior of the user home folder is to encrypt files and folders at rest using Amazon S3-managed encryption keys from the AWS key management services (AWS KMS). It is important to note that AWS Administrative Users with access to the AWS Console or Amazon S3 bucket will be able to access those files directly.

In designs that require a Server Message Block (SMB) target from a Windows File Share to store user files and folders, the process is either automatic or requires configuration.

Table 5 — Options for securing user data

SMB target	Encryption-at-rest	Encryption-in-transit	Antivirus (AV)
FSx for Windows File Server	Automatic through AWS KMS	Automatic through SMB encryption	AV installed on a remote instance performs scan on mapped drive
File Gateway, AWS Storage Gateway	By default, all data stored by AWS Storage Gateway in S3 is encrypted server-side with Amazon S3-Manage d Encryption Keys (SSE-S3). You can optionally configure different gateway types to encrypt stored data with AWS	All data transferred between any type of gateway appliance and AWS storage is encrypted using SSL.	AV installed on a remote instance performs scan on mapped drive

SMB target	Encryption-at-rest	Encryption-in-transit	Antivirus (AV)
	Key Management Service (KMS)		
EC2-based Windows File Servers	Enable EBS encryption	PowerShell; `Set-SmbServer Configuration – EncryptData $True`	AV installed on server performs scan on local drives

Endpoint Security and Antivirus

The brief ephemeral nature of WorkSpaces Applications instances and the lack of persistency of data means a different approach is required to ensure user experience and performance is not compromised by activities that would be required on a persistent desktop. Endpoint Security agents are installed in WorkSpaces Applications images when there is an organizational policy or when used with external data ingress e.g. e-mail, files ingress, external web browsing.

Removing unique identifiers

Endpoint Security agents may have a globally unique identifier (GUID) which must be reset during the fleet instance creation process. Vendors have instructions on installing their products in images which will ensure a new GUID is generated for each instance generated from an image.

To ensure the GUID is not generated, install the Endpoint Security agent as the last action before running the WorkSpaces Applications Assistant to generate the image.

Performance optimization

Endpoint Security Vendors provide switches and setting that optimize the performance of WorkSpaces Applications. The settings vary between vendors and can be found in their documentation, typically in a section on VDI. Some common settings include but are not limited to are:

- Turn off boot up scans to ensure instance creation, startup and login times are minimized
- Turn off scheduled scans to prevent unnecessary scans
- Turn off signature caches to prevent file enumeration

- Enable VDI optimized IO settings
- Exclusions required by applications to ensure performance

Endpoint security vendors provide instructions for use with virtual desktop environments which optimize performance.

- Trend Micro Office Scan Support for Virtual Desktop Infrastructure - Apex One/OfficeScan (trendmicro.com)
- CrowdStrike and How to Deploy the CrowdStrike Falcon Sensor on AWS on the CrowdStrike website
- Sophos and Sophos Central Endpoint: How to install on a gold image to avoid duplicate identities on the Sophos website and Sophos Central: Best practices when installing Windows Endpoints in Virtual Desktop Environments
- McAfee and Trellix (formerly McAfee) Agent provisioning and deployment on Virtual Desktop Infrastructure systems on the Trellix website
- Microsoft Endpoint Security and Configure Microsoft Defender Antivirus on VDI on Microsoft Learn

Scanning exclusions

If security software is installed in WorkSpaces Applications instances, the security software must not interfere with the following processes.

Table 6 — Security software must not interfere with the following processes. This may impact the service availability and performance.

Service	Processes
AmazonCloudWatchAgent	"C:\Program Files\Amazon\AmazonCloudWatchAgent\amazon-cloudwatch-agent.exe"
AmazonSSMAgent	"C:\Program Files\Amazon\SSM\amazon-ssm-agent.exe"
Amazon DCV	"C:\Program Files\NICE\DCV\Server\bin\dcvserver.exe"

Service	Processes
	"C:\Program Files\NICE\DCV\Server\bin\dcvagent.exe" "C:\Program Files\NICE\DCV\Server\bin\dcvdrivehelper.exe" "C:\Program Files\NICE\DCV\Server\bin\dcvprinterhelper.exe" "C:\Program Files\NICE\DCV\Server\bin\dcvwebauthnnativemsghost.exe" "C:\Program Files\NICE\DCV\Server\bin\dcvwebrtcnativemsghost.exe" "C:\Program Files\NICE\DCV\Server\bin\dcvlogonhelper.exe" "C:\Program Files\NICE\DCV\Server\bin\xpstopdf.exe"

Service	Processes
WorkSpaces Applications	"C:\Program Files\Amazon\Workspaces\AmazonWorkSpaces.exe" "C:\Program Files\Amazon\AppStream2\StorageConnector\StorageConnector.exe" In the folder "C:\Program Files\Amazon\Photon\" ".\Agent\PhotonAgent.exe" ".\Agent\s5cmd.exe" ".\WebServer\PhotonAgentWebServer.exe" ".\CustomShell\PhotonWindowsAppSwitcher.exe" ".\CustomShell\PhotonWindowsCustomShell.exe" ".\CustomShell\PhotonWindowsCustomShellBackground.exe"
Windows Driver Foundation	"C:\Windows\System32\WUDFHost.exe"

Folders

If security software is installed in WorkSpaces Applications instances, the software must not interfere with the following folders:

Example

```
C:\Program Files\Amazon\*
C:\ProgramData\Amazon\*
C:\Program Files (x86)\AWS Tools\*
C:\Program Files (x86)\AWS SDK for .NET\*
```

```
C:\Program Files\NICE\*
C:\ProgramData\NICE\*
C:\AppStream\*
C:\Program Files\WindowsPowerShell\Modules\AWSPowerShell\*
```

Endpoint security console hygiene

WorkSpaces Applications will create new unique instances each time a user connects beyond the idle and disconnect timeouts. The instances will have a unique name and will build up in endpoint security management consoles. Setting unused aged machines over 4 or more days old (or lower depending on WorkSpaces Applications session timeouts) to be deleted will minimize the number of expired instances in the console.

Network Exclusions

The WorkSpaces Applications management network range (`198.19.0.0/16`) and following ports and addresses should not be blocked by any security / firewall or antivirus solutions within WorkSpaces Applications instances.

Table 7 — Ports in WorkSpaces Applications streaming instances security software must not interfere with

Port	Usage
8300	This is used for establishing the streaming connection
3128	This is used for managing the streaming instance by WorkSpaces Applications
8000	This is used for managing the streaming instance by WorkSpaces Applications
8443	This is used for managing the streaming instance by WorkSpaces Applications
53	DNS

Table 8 — WorkSpaces Applications managed service addresses security software must not interfere with

Port	Usage
169.254.169.123	NTP
169.254.169.249	NVIDIA GRID License Service
169.254.169.250	KMS
169.254.169.251	KMS
169.254.169.253	DNS
169.254.169.254	Metadata

Securing an WorkSpaces Applications Session

Limiting application and operating system controls

WorkSpaces Applications gives the administrator the ability to specify exactly which applications can be launched from the web page in application streaming mode. This does not, however, guarantee that only those applications specified can be run.

Windows utilities and applications can be launched through the operating system through additional means. AWS recommends using Microsoft AppLocker to ensure that only the applications that your organization requires can be run. The default rules must be modified, as they grant everyone path access to critical system directories.

> **Note**
>
> Windows Server 2016 and 2019 require the Windows Application Identity service to be running to enforce AppLocker rules. Application access from WorkSpaces Applications using Microsoft AppLocker is detailed in the AppStream Admin Guide.

For fleet instances joined to an Active Directory domain, use Group Policy Objects (GPOs) to deliver user and system settings to secure the users application and resource access.

Firewalls and Routing

When creating an WorkSpaces Applications fleet, subnets and a Security Group must be assigned. Subnets have existing assignments of Network Access Control Lists (NACLs) and route table(s). You can associate up to five security groups while launching a new image builder or while creating a new fleet Security Groups can have up to five assignments from the existing Security Groups. For each security group, you add rules that control the outbound and inbound network traffic from and to your instances

A NACL is an optional layer of security for your VPC that acts as a stateless firewall for controlling traffic in and out of one or more subnets. You might set up network ACLs with rules similar to your security groups in order to add an additional layer of security to your VPC. For more information about the differences between security groups and network ACLs, see the compare security groups and NACLs page.

When designing and applying Security Group and NACL rules, consider the AWS Well-Architected best practices for least privilege. *Least privilege* is a principle of granting only the permissions required to complete a task.

For customers who have a high-speed private network connecting their on premise environment to AWS (via an AWS Direct Connect), you may consider using the VPC Endpoints for AppStream, which will mean the streaming traffic will be routed via your private network connectivity rather than going across the public internet. For more information on this topic, see the WorkSpaces Applications streaming interface VPC endpoint section of this document.

Data Loss Prevention

We'll look at two kinds of data loss prevention.

Client to AppStream 2.0 Instance Data Transfer Controls

Table 9 — Guidance for controlling data ingress and egress

Setting	Options	Guidance
Clipboard	• Copy and paste to remote session only • Copy to local device only • Disabled	Disabling this setting does not disable copy and paste within the session. If copying data into the session is required, choose Paste to

Setting	Options	Guidance
		remote session only to minimize the potential for data leakage.
File transfer	• Upload and download • Upload only • Download only • Disabled	Avoid enabling this setting to prevent data leakage.
Print to local device	• Enabled • Disabled	If printing is required, use network mapped printers that are controlled and monitored by your organization.

Consider the advantages of the existing organizational data transfer solution over the stack settings. These configurations are not designed to replace a comprehensive secure data transfer solution.

Controlling egress traffic

Where data loss is a concern, it's important to cover off what a User can access once they are inside of their WorkSpaces Applications instance. What does the network exit (or egress) path look like? It is a common requirement to have public internet access available to the end user inside their WorkSpaces Applications instance, so placing a WebProxy or Content Filtering Solution in the network path needs to be considered. Other considerations include a local Antivirus application and other endpoint security measures inside the AppStream instance (see the section "Endpoint Security and Antivirus" for more information).

Using AWS services

AWS Identity and Access Management

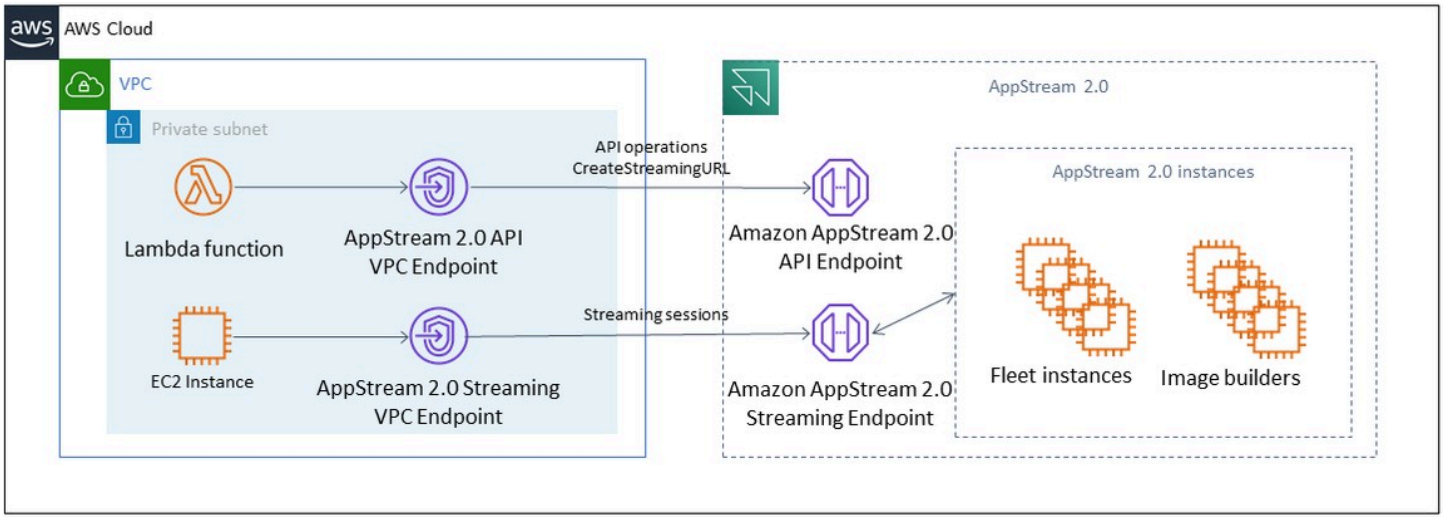
Using an IAM role to access AWS services, and being specific in the IAM policy attached to it, is a best practice that provides only the users in WorkSpaces Applications sessions have access without managing additional credentials. Follow the best practices for using IAM Roles with WorkSpaces Applications.

Create IAM policies to protect Amazon S3 buckets that are created to persist user data in both home folders and application settings persistence. This prevents non-WorkSpaces Applications administrators from access.

VPC endpoints

A VPC endpoint enables private connections between your VPC and supported AWS services and VPC endpoint services powered by AWS PrivateLink. AWS PrivateLink is a technology that enables you to privately access services by using private IP addresses. Traffic between your VPC and the other service does not leave the Amazon network. If public internet access is required only for AWS services, VPC endpoints remove the requirement for NAT gateways and internet gateways altogether.

In environments where automation routines or developers require making API calls for WorkSpaces Applications, create an interface VPC endpoint for WorkSpaces Applications API operations. For example, if there are EC2 instances in private subnets without public internet access, a VPC endpoint for WorkSpaces Applications API can be used to call AppStream 2.0 API operations such as CreateStreamingURL. The following diagram shows an example setup where WorkSpaces Applications API and streaming VPC endpoints are consumed by Lambda functions and EC2 instances.



VPC endpoint

The streaming VPC endpoint allows you to stream sessions through a VPC endpoint. The streaming interface endpoint maintains the streaming traffic within your VPC. Streaming traffic includes pixel, USB, user input, audio, clipboard, file upload and download, and printer traffic. To use the VPC endpoint, the VPC endpoint setting must be enabled at the WorkSpaces Applications stack. This

serves as an alternative to streaming user sessions over the public internet from locations that have limited internet access and would benefit from accessing through a Direct Connect instance. Streaming user sessions through a VPC endpoint require the following:

- The Security Groups that are associated with the interface endpoint must allow inbound access to port 443 (TCP) and ports 1400–1499 (TCP) from the IP address range from which your users connect.
- The Network Access Control List for the subnets must allow outbound traffic from ephemeral network ports 1024-65535 (TCP) to the IP address range from which your users connect.
- Internet connectivity is required to authenticate users and deliver the web assets that WorkSpaces Applications requires to function.

To learn more about restricting traffic to AWS services with WorkSpaces Applications, see the administration guide for creating and streaming from VPC endpoints.

When full public internet access is required, it's a best practice to disable Internet Explorer Enhanced Security Configuration (ESC) on the Image Builder. For more information, see the WorkSpaces Applications administration guide to disable Internet Explorer enhanced security configuration.

Configuring the Instance Metadata Service (IMDS) on your instances

This topic describes the Instance Metadata Service (IMDS).

Instance metadata is data that's related to an Amazon Elastic Compute Cloud (Amazon EC2) instance that applications can use to configure or manage the running instance. The instance metadata service (IMDS) is an on-instance component that code on the instance uses to securely access instance metadata. For more information, see Instance metadata and user data in the *Amazon EC2 User Guide*.

Code can access instance metadata from a running instance using one of two methods: Instance Metadata Service Version 1 (IMDSv1) or Instance Metadata Service Version 2 (IMDSv2). IMDSv2 uses session-oriented requests and mitigates several types of vulnerabilities that could be used to try to access the IMDS. For information about these two methods, see Configuring the instance metadata service in the *Amazon EC2 User Guide*.

Resource support for IMDS

Always-On, On-Demand, Single-Session, and Multi-Session Fleets, and all Image Builders support both IMDSv1 and IMDSv2 when running WorkSpaces Applications images with the agent version or managed image update released on or after January 16, 2024.

Elastic Fleets and AppBlock Builders instances also support both IMDSv1 and IMDSv2.

Example of IMDS attribute settings

Below are two examples of choosing the IMDS method:

Java v2 SDK example

Below example request disable IMDSv1 using `disableIMDSV1` attributes

```
CreateFleetRequest request = CreateFleetRequest.builder()
 .name("TestFleet")
 .imageArn("arn:aws:appstream:us-east-1::image/TestImage")
 .instanceType("stream.standard.large")
 .fleetType(FleetType.ALWAYS_ON)
 .computeCapacity(ComputeCapacity.builder()
 .desiredInstances(5)
 .build())
 .description("Test fleet description")
 .displayName("Test Fleet Display Name")
 .enableDefaultInternetAccess(true)
 .maxUserDurationInSeconds(3600)
 .disconnectTimeoutInSeconds(900)
 .idleDisconnectTimeoutInSeconds(600)
 .iamRoleArn("arn:aws:iam::123456789012:role/TestRole")
 .streamView(StreamView.APP)
 .platform(PlatformType.WINDOWS)
 .maxConcurrentSessions(10)
 .maxSessionsPerInstance(2)
 .tags(tags)
 .disableIMDSV1(true)
 .build();
```

Set **disableIMDSV1** to true to disable IMDSv1 and enforce IMDSv2.

Set **disableIMDSV1** to false to enable both IMDSv1 and IMDSv2.

CLI Example

Below example request disable IMDSv1 using `--disable-imdsv1` attributes

```
aws appstream create-fleet --name test-fleet --image-arn "arn:aws:appstream:us-
east-1::image/test-image" --disable-imdsv1 --instance-type stream.standard.small --
compute-capacity DesiredInstances=2 --max-user-duration-in-seconds 57600 --disconnect-
timeout-in-seconds 57600 --region us-east-1
```

Set `--disable-imdsv1` to true to disable IMDSv1 and enforce IMDSv2.

Set `--no-disable-imdsv1` to false to enable both IMDSv1 and IMDSv2.

Troubleshooting

If you encounter difficulties when working with Amazon WorkSpaces Applications, consult the following troubleshooting resources.

Contents

- General Troubleshooting
- Troubleshooting Image Builders
- Troubleshooting Fleets
- Troubleshooting Active Directory
- Troubleshooting WorkSpaces Applications User Issues
- Troubleshooting Persistent Storage Issues
- Troubleshooting Notification Codes

General Troubleshooting

The following are general issues that might occur when you use Amazon WorkSpaces Applications.

Issues

- SAML federation is not working. The user is not authorized to view WorkSpaces Applications applications.
- After federating from an ADFS portal, my streaming session doesn't start. I am getting the error "Sorry connection went down".
- I get an invalid redirect URI error.
- My image builders and fleets never reach the running state. My DNS servers are in a Simple AD directory.
- I've enabled application settings persistence for my users, but their persistent application settings aren't being saved or loaded.
- I've enabled application settings persistence for my users, but for certain streaming applications, my users' passwords aren't persisting across sessions.
- Google Chrome data is filling the VHD file that contains my users' persistent application settings. This is preventing their settings from persisting. How can I manage the Chrome profile?

- I set up a custom domain for my embedded WorkSpaces Applications streaming sessions, but my WorkSpaces Applications streaming URLs aren't redirecting to my custom domain.
- I launched an app on a smartcard-enabled WorkSpaces Applications fleet, and there are a limited number of certificates (or none) available to the app for authentication.
- The Certification Propagation service isn't starting on my smartcard-enabled WorkSpaces Applications fleet.
- I can't log in with my Active Directory username or password after SAML authentication.

SAML federation is not working. The user is not authorized to view WorkSpaces Applications applications.

This might happen because the inline policy that is embedded for the SAML 2.0 federation IAM role does not include permissions to the stack ARN. The IAM role is assumed by the federated user who is accessing an WorkSpaces Applications stack. Edit the role permissions to include the stack ARN. For more information, see Amazon WorkSpaces Applications Integration with SAML 2.0 and Troubleshooting SAML 2.0 Federation with AWS in the *IAM User Guide*.

After federating from an ADFS portal, my streaming session doesn't start. I am getting the error "Sorry connection went down".

Set the claim rule's **Incoming Claim Type** for the **NameID** SAML attribute to **UPN** and try the connection again.

I get an invalid redirect URI error.

This error occurs due to a malformed or invalid WorkSpaces Applications stack relay state URL. Make sure that the relay state configured in your federation setup is the same as the stack relay state that is displayed in the stack details through the WorkSpaces Applications console. If they are the same and the problem still persists, contact AWS Support. For more information, see Amazon WorkSpaces Applications Integration with SAML 2.0.

My image builders and fleets never reach the running state. My DNS servers are in a Simple AD directory.

WorkSpaces Applications relies on the DNS servers within your VPC to return a non-existent domain (NXDOMAIN) response for local domain names that don’t exist. This enables the

WorkSpaces Applications-managed network interface to communicate with the management servers.

When you create a directory with Simple AD, AWS Directory Service creates two domain controllers that also function as DNS servers on your behalf. Because the domain controllers don't provide the NXDOMAIN response, they can't be used with WorkSpaces Applications.

I've enabled application settings persistence for my users, but their persistent application settings aren't being saved or loaded.

WorkSpaces Applications automatically saves application settings that are created in certain locations on the Windows instance. The settings are saved only if your application saves them to one of these locations. For a list of supported locations, see How Application Settings Persistence Works. If your application is configured to save to C:\Users\%username% and your users' settings for the application aren’t persisting between sessions, the mount point might not be created. This prevents the settings from being saved to the VHD file that contains your users’ persistent application settings.

To resolve this issue, follow these steps:

1. On the fleet instance, open File Explorer and browse to the user profile directory at C:\Users\%username%.
2. Confirm whether this directory contains a symlink, and then do either of the following:
   - If there is a symlink, confirm that it points to D:\%username%.
   - If there isn't a symlink, try to delete the C:\Users\%username% directory.

     If you can’t delete this directory, identify the file in the directory that is preventing it from being deleted and the application that created the file. Then contact the application vendor for information about how to change the file permissions or move the file.

     If you can delete this directory, contact AWS Support for further guidance to resolve this issue. For more information, see AWS Support Center.

I've enabled application settings persistence for my users, but for certain streaming applications, my users' passwords aren't persisting across sessions.

This issue occurs when:

- Users are streaming applications such as Microsoft Outlook, which use the Microsoft Data Protection API.
- App settings persistence is enabled for streaming instances that are not joined to Active Directory domains.

In cases where a streaming instance is not joined to an Active Directory domain, the Windows user, PhotonUser, is different on each fleet instance. Due to the way in which the DPAPI security model works, users' passwords don't persist for applications that use DPAPI in this scenario. In cases where streaming instances are joined to an Active Directory domain and the user is a domain user, the Windows user name is that of the logged in user, and users' passwords persist for applications that use DPAPI.

Google Chrome data is filling the VHD file that contains my users' persistent application settings. This is preventing their settings from persisting. How can I manage the Chrome profile?

By default, Google Chrome stores both user data and the local disk cache in the Windows user profile. To prevent the local disk cache data from filling the VHD file that contains users' persistent application settings, configure Chrome to save only the user data. To do so, on the fleet instance, open the command line as an administrator and start Chrome with the following parameters to change the location of the disk cache:

```
chrome.exe --disk-cache-dir C:\path-to-unsaved-location\
```

Running Chrome with these parameters prevents the disk cache from being persisted between WorkSpaces Applications sessions.

I set up a custom domain for my embedded WorkSpaces Applications streaming sessions, but my WorkSpaces Applications streaming URLs aren't redirecting to my custom domain.

To resolve this issue, verify that when you created your WorkSpaces Applications streaming URL, you replaced the WorkSpaces Applications endpoint with your custom domain. By default, WorkSpaces Applications streaming URLs are formatted as follows:

```
https://appstream2.region.aws.amazon.com/authenticate?parameters=authenticationcode
```

To replace the default WorkSpaces Applications endpoint in your streaming URL, replace **`https://appstream2.`***`region`* in the URL with your custom domain. For example, if your custom domain is **`training.example.com`**, your new streaming URL must follow this format:

```
https://training.example.com/authenticate?parameters=authenticationcode
```

For more information about configuring custom domains for embedded WorkSpaces Applications streaming sessions, see Configuration Requirements for Using Custom Domains.

I launched an app on a smartcard-enabled WorkSpaces Applications fleet, and there are a limited number of certificates (or none) available to the app for authentication.

This happens when the application is launched before the Certificate Propagation service is in a running state.

To resolve this issue, use the PowerShell module Get-Service to query the Certificate Propagation service's status, and make sure that it's in a running state before launching your application.

For example, the following script won't launch the application until the Certificate Propagation service is running:

```
$logFile = "$Env:TEMP\AS2\Logging\$(Get-Date -Format "yyyy-MM-dd-HH-mm-
ss")_applaunch.log"
New-Item -path $logfile -ItemType File -Force | Out-Null

Function Write-Log {
    Param ([string]$message)
```

```
    $stamp = Get-Date -Format "yyyy/MM/dd HH:mm:ss"
    $logoutput = "$stamp $message"
    Add-content $logfile -value $logoutput
}

if (Get-Service -Name "CertPropSvc" | Where-Object -Property Status -eq Running) {

    Write-Log "The Certificate Propagation Service is running. Launching
 Application..."
    try {
        Start-Process -FilePath "Path to Application" -WindowStyle Maximized -
ErrorAction Stop
    }
    catch {
        Write-Log "There was an error launching the application: $_"

    }

}
else {

    do {

        $status = Get-Service "CertPropSvc" | select-object -ExpandProperty Status
        Write-Log "The Certificate Propagation service status is currently $status"
        Start-Sleep -Seconds 2


    } until (Get-Service -Name "CertPropSvc" | Where-Object -Property Status -eq
 Running)

    write-log "The Certificate Propagation Service is running. Launching
 Application..."
    try {
        Start-Process -FilePath "Path to Application" -WindowStyle Maximized -
ErrorAction Stop
    }
    catch {
        Write-Log "There was an error launching the application: $_"

    }
}
```

The Certification Propagation service isn't starting on my smartcard-enabled WorkSpaces Applications fleet.

If the Certificate Propagation service isn't starting, the service's startup type might be set to **Disabled**. To resolve this, on the WorkSpaces Applications image builder used to create your fleet's image, launch the Windows Services Microsoft Management Console, and make sure that the startup type of the Certificate Propagation service isn't set to **Disabled**.

If the startup type isn't set to **Disabled**, and the service is still not starting on your WorkSpaces Applications fleet, use the PowerShell module Start-Service to start the Certificate Propagation service when your fleet instance starts.

For example, the following PowerShell script will start the service if it detects that it's in a stopped state:

```
$logFile = "C:\AppStream\Logging\$(Get-Date -Format "yyyy-MM-dd-HH-mm-
ss")_certpropcheck.log"
New-Item -path $logfile -ItemType File -Force | Out-Null

Function Write-Log {
    Param ([string]$message)
    $stamp = Get-Date -Format "yyyy/MM/dd HH:mm:ss"
    $logoutput = "$stamp $message"
    Add-content $logfile -value $logoutput
}

if (Get-Service -Name "CertPropSvc" | Where-Object -Property Status -eq Running) {

    Write-Log "The Certificate Propagation Service is running. Exiting..."
    Exit
}
else {
    do {

        if (Get-Service -Name "CertPropSvc" | Where-Object -Property Status -eq
 Stopped) {

            Write-Log "The Certificate Propagation Service is stopped, attepmting to
 start..."
            try {
                Start-Service -Name "CertPropSvc" -ErrorAction Stop
```

```
            }
            catch {
                Write-Log "There was a problem starting the service: $_"
                break
            }

            $status = Get-Service "CertPropSvc" | select-object -ExpandProperty Status
            Write-Log "The Certificate Propagation service status is currently $status"


        }
        else {

            $status = Get-Service "CertPropSvc" | select-object -ExpandProperty Status
            Write-Log "The Certificate Propagation service status is currently $status"
            break
        }

    } until (Get-Service -Name "CertPropSvc" | Where-Object -Property Status -eq
 Running)
}
```

I can't log in with my Active Directory username or password after SAML authentication.

The nameID in the SAML claim needs to match the username in Active Directory. Some IdPs require an update, refresh, or redeploy after adjusting certain attributes. If you make an adjustment and it is not reflected in your SAML capture, refer to your IdP's documentation or support program regarding the specific steps required to make the change take effect.

Troubleshooting Image Builders

The following are issues that might occur when you use Amazon WorkSpaces Applications image builders.

Issues

- I cannot connect to the internet from my image builder.
- When I tried installing my application, I see an error that the operating system version is not supported.

- I want to use a Windows PowerShell script to open my applications.
- I want to make ClickOnce applications available to users.
- When I connect to my image builder, I see a login screen asking me to enter Ctrl+Alt+Delete to log in. However, my local machine intercepts the key strokes.
- When I switched between admin and test modes, I saw a request for a password. I don't know how to get a password.
- I get an error when I add my installed application.
- I accidentally quit a background service on the image builder and got disconnected. I am now unable to connect to that image builder.
- The application fails to launch in test mode.
- The application could not connect to a network resource in my VPC.
- I customized my image builder desktop, but my changes are not available when connecting to a session after launching a fleet from the image I created.
- My application is missing a command line parameter when launching.
- I am unable to use my image with a fleet after installing an antivirus application.
- My image creation failed.
- The Image Assistant create-image operation failed with an error message that access to the PrewarmManifest.txt is denied
- The image assistant reports that the SSM Agent version is below the recommended minimum

I cannot connect to the internet from my image builder.

Image builders cannot communicate to the internet by default. To resolve this issue, launch your image builder in a VPC subnet that has internet access. You can enable internet access from your VPC subnet using a NAT gateway. Or, you can configure an internet gateway in your VPC, and attach an Elastic IP address to your image builder. For more information, see Networking and Access for Amazon WorkSpaces Applications.

When I tried installing my application, I see an error that the operating system version is not supported.

Only applications that can be installed on Windows Server 2016, Windows Server 2019, Windows Server 2022, and Windows Server 2025 can be added to an WorkSpaces Applications image. Check

if your application is supported on one of these operating systems, as applicable for your image builder.

I want to use a Windows PowerShell script to open my applications.

You can use Windows PowerShell scripts to open your applications in the fleet instance. You might want to do this to configure the application or environment before the application opens. To launch a Windows PowerShell script for your application, specify the PowerShell .exe file in Image Assistant. Navigate to `C:\Windows\System32\WindowsPowerShell\v1.0\powershell.exe`, and specify the following launch parameters:

-file "C:\Path\To\PowerShell\Script.ps1"

> **Note**
>
> To allow the specified script to open the application, you must override the PowerShell script execution policy. To do so, add **`-ExecutionPolicy Bypass`** to the launch parameter.

I want to make ClickOnce applications available to users.

To make a ClickOnce application available to your WorkSpaces Applications users, you must install the application on your image builder first as an Administrator, and then as a Template User. Because ClickOnce applications require a user-specific installation, you must install your application as a Template User to enable users to launch the application from fleet instances. To install the ClickOnce application as an Administrator and then as a Template User, perform these steps.

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the left navigation pane, choose **Images**, **Image Builder**.
3. In the list, select the image builder that you want to use, and log into it as an Administrator.
4. Create a batch file that calls the `appref-ms` file within the user profile. Use the %APPDATA% environment variable to replace C:\Users\username\AppData\Roaming. Here is an example batch file call:

   ```
   explorer "%APPDATA%\Microsoft\Windows\Start Menu\Programs\Company\ClickOnce.appref-ms"
   ```

5. On the image builder desktop, open Image Assistant.
6. On the **Configure Apps** page, choose **Switch user**.
7. On the **Local User** tab, choose **Template User**.
8. After you log in as a Template User, install the application again.
9. On the image builder desktop, open Image Assistant.
10. On the **Configure Apps** page, open the ClickOnce application to verify that it functions correctly. After you finish testing, choose **Switch user**.
11. Log back in as an Administrator and perform the required steps in Image Assistant to finish creating your image.

When I connect to my image builder, I see a login screen asking me to enter Ctrl+Alt+Delete to log in. However, my local machine intercepts the key strokes.

Your client might intercept certain key combinations locally instead of sending them to the image builder session. To reliably send the Ctrl+Alt+Delete key combination to the image builder, choose **Admin Commands**, **Send Ctrl+Alt+Delete**. The **Admin Commands** menu is available on the top right corner of the image builder session toolbar.

When I switched between admin and test modes, I saw a request for a password. I don't know how to get a password.

WorkSpaces Applications usually logs you into the user mode that you choose automatically. On some occasions, the switch might not happen automatically. If a password is requested, choose **Admin Commands**, **Log me in**. This sends a one-time password, securely, to your image builder and pastes it into the **Password** field.

I get an error when I add my installed application.

Check if your application type is supported. You can add applications of the types `.exe`, `.lnk`, and `.bat`.

Check if your application is installed under the `C:\Users` folder hierarchy. Any application installed under `C:\Users` is not supported. Select a different installation folder under `C:\` when installing the application.

I accidentally quit a background service on the image builder and got disconnected. I am now unable to connect to that image builder.

Try stopping the image builder, restarting it and connecting to it again. If the problem persists, you must launch (create) a new image builder. Do not stop any background services running on the image builder instance. Doing so might interrupt your image builder session or interfere with the image creation.

The application fails to launch in test mode.

Check if your application requires elevated user privileges or any special permissions that are usually available only to an administrator. The Image Builder Test mode has the same limited permissions on the image builder instance as your end users have on an WorkSpaces Applications test fleet. If your applications require elevated permissions, they do not launch in the Image Builder Test mode.

The application could not connect to a network resource in my VPC.

Check if the image builder was launched in the correct VPC subnet. You might also need to verify that the route tables in your VPC are configured to allow a connection.

I customized my image builder desktop, but my changes are not available when connecting to a session after launching a fleet from the image I created.

Changes that are saved as part of a local user session, such as time settings, are not persisted when creating an image. To persist any local user session changes, add them to the local group policy on the image builder instance.

My application is missing a command line parameter when launching.

You can provide a command line parameter when using image builder to add an application to an image. If the launch parameters for the application do not change for each user, you can enter them while adding an application to the image in the image builder instance.

If the launch parameters are different for every launch, you can pass them programmatically when using the `CreateStreamingURL` API. Set the `sessionContext` and `applicationID` parameters

in the API fields. The sessionContext is included as a command line option when launching the application.

If the launch parameters must be computed on the fly, you can launch your application using a script. You can parse the `sessionContext` parameter within your script before launching your application with a computed parameter.

I am unable to use my image with a fleet after installing an antivirus application.

You can install any tools, including antivirus programs, on your WorkSpaces Applications stack by using the image builder before creating an image. However, these programs should not block any network ports or stop any processes that are used by the WorkSpaces Applications service. We recommend testing your application in Image Builder Test mode before creating an image and attempting to use it with a fleet.

My image creation failed.

Verify that you did not make any changes to WorkSpaces Applications services before starting the image creation. Try creating your image again; if it fails, contact AWS Support. For more information, see AWS Support Center.

The Image Assistant `create-image` operation failed with an error message that access to the PrewarmManifest.txt is denied

The application optimization manifest was created with elevated privileges. To create the image, do either of the following, and then try again:

- Run the Image Assistant command line interface (CLI) executable file (Image-Assistant.exe) with administrator privileges.
- Delete the application optimization manifest file.

The image assistant reports that the SSM Agent version is below the recommended minimum

The image assistant in the image builder displays the following error message: "The installed SSM Agent version x is below the recommended minimum version 3.3.0.0. Consider updating the SSM Agent to ensure full functionality."

AWS Systems Manager is deprecating the ec2messages (Amazon Message Delivery Service) endpoints that older SSM Agent versions use for Run Command execution. SSM Agent 3.3.0.0 and above uses the newer ssmmessages (Amazon Message Gateway Service) endpoints, which provide improved security and reliability.

To resolve this issue, complete the following steps:

1. Manually update the SSM Agent by following the instructions in the Systems Manager User Guide.
2. Confirm that the SSM Agent version is 3.3.0.0 or above by running the following PowerShell command:

```
$currentVersion = (Get-ChildItem 'HKLM:\SOFTWARE\Microsoft\Windows\CurrentVersion
\Uninstall' | Get-ItemProperty | Where-Object { $_.DisplayName -like
 '*SSM*' }).DisplayVersion
```

3. Confirm that the SSM Agent proxy is correctly configured by running the following PowerShell commands. The output must include `https_proxy=http://squid-proxy.appstream.local:3128`.

```
$registryPath = 'HKLM:\SYSTEM\CurrentControlSet\Services\AmazonSSMAgent'
$existingEnv = (Get-ItemProperty -Path $registryPath -Name 'Environment' -ErrorAction
 SilentlyContinue).Environment
```

4. If the SSM Agent proxy is not correctly configured, run the following PowerShell commands to configure it:

```
$registryPath = 'HKLM:\SYSTEM\CurrentControlSet\Services\AmazonSSMAgent'
$proxyVars = @(
  'http_proxy=http://squid-proxy.appstream.local:3128',
  'https_proxy=http://squid-proxy.appstream.local:3128',
  'no_proxy=169.254.169.254'
)
Set-ItemProperty -Path $registryPath -Name 'Environment' -Value $proxyVars -Type
 MultiString
Restart-Service AmazonSSMAgent
```

5. Test image creation on the image builder to confirm that the workflow succeeds.

Troubleshooting Fleets

The following are issues that might occur when users connect to Amazon WorkSpaces Applications streaming sessions launched from fleet instances.

Issues

- I tried to increase my fleet capacity, but the update isn't taking effect.
- My applications won't work correctly unless I use the Internet Explorer defaults. How do I restore the Internet Explorer default settings?
- I need to persist environment variables across my fleet instances.
- I want to change the default Internet Explorer home page for my users.
- When my users end a streaming session and then start a new one, they see a message that says no streaming resources are available.

I tried to increase my fleet capacity, but the update isn't taking effect.

You can increase your fleet capacity in either of the two following ways:

- Manually, by increasing the **Minimum capacity** value on the **Scaling Policies** tab for the fleet in the WorkSpaces Applications console.
- Automatically, by configuring a fleet scaling policy that manages your capacity for the fleet.

If your manual modification or scaling policy exceeds your current WorkSpaces Applications quota for your fleet instance type and size, the new values will not take effect. If you experience this issue, you can use the AWS Command Line Interface (CLI) describe-scaling-activities command to verify whether your capacity request exceeds your quota for the applicable fleet instance type and size. This command uses the following format:

```
aws application-autoscaling describe-scaling-activities
  --service-namespace appstream \
  --resource-id fleet/fleetname \
```

For example, the following command provides information for the **`TestFleet`** fleet in the **`us-west-2`** AWS Region.

```
aws application-autoscaling describe-scaling-activities --service-namespace appstream
 --resource-id fleet/TestFleet --region us-west-2
```

The following JSON output shows that a scaling policy for **`TestFleet`** with a **Minimum capacity** value of 150 was set. This value exceeds the limit (quota) for **`TestFleet`**, which is 100, so the new scaling policy doesn't take effect. In the output, the **StatusMessage** parameter provides detailed information about the cause of the error, including the fleet instance type (in this case, stream.standard.medium), and the current quota, which is 100.

> **Note**
>
> WorkSpaces Applications instance type and size quotas are per Amazon Web Services account, per AWS Region. If you have multiple fleets in the same Region that use the same instance type and size, the total number of instances in all fleets in that Region must be less than or equal to the applicable quota.

```
{
    "ScalingActivities": [
        {
            "ActivityId": "id",
            "ServiceNamespace": "appstream",
            "ResourceId": "fleet/TestFleet",
            "ScalableDimension": "appstream:fleet:DesiredCapacity",
            "Description": "Setting desired capacity to 150.",
            "Cause": "minimum capacity was set to 150",
            "StartTime": 1596828816.136,
            "EndTime": 1596828816.646,
            "StatusCode": "Failed",
            "StatusMessage": "Failed to set desired capacity to 150. Reason: The
 Instance type 'stream.standard.medium' capacity limit for fleet TestFleet' was
 exceeded. Requested: 150, Limit: 100 (Service: AmazonAppStream; Status Code: 400;
 Error Code: LimitExceededException; Request ID: id; Proxy: null)."
```

If you run the `describe-scaling-activities` command and the output indicates that your capacity request exceeds your current quota, you can resolve the issue by:

- Changing your capacity request to a value that doesn't exceed your quota.

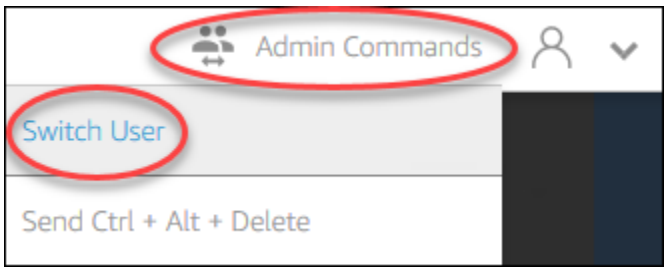
- Requesting a quota increase. To request a quota increase, use the WorkSpaces Applications Limits form.

My applications won't work correctly unless I use the Internet Explorer defaults. How do I restore the Internet Explorer default settings?

If your WorkSpaces Applications environment includes applications that render elements, you might need to restore the Internet Explorer default settings to enable full enable access to the internet.

To automatically restore the Internet Explorer default settings

1. Open the WorkSpaces Applications console at https://console.aws.amazon.com/appstream2/home.
2. In the left navigation pane, choose **Images**, **Image Builder**.
3. Choose the image builder on which to restore the Internet Explorer default settings, verify that it is in the **Running** state, and choose **Connect**.
4. Log in to the image builder by doing either of the following:
   - If your image builder is not joined to an Active Directory domain, on the **Local User** tab, choose **Template User**.
   - If your image builder is joined to an Active Directory domain, choose the **Directory User** tab, enter the credentials for a domain user who does not have local administrator permissions on the image builder, and then choose **Log in**.
5. Open Internet Explorer and reset your settings by doing the following:
   a. In the upper right area of the Internet Explorer browser window, choose the **Tools** icon, then choose **Internet options**.
   b. Choose the **Advanced** tab, and then choose **Reset**.
   c. When prompted to confirm your choice, choose **Reset** again.
   d. When the **Reset Internet Explorer Settings** message is displayed, choose **Close**.
6. In the upper right area of the image builder desktop, choose **Admin Commands**, **Switch User**.



7. This disconnects your current session and opens the login menu. Do either of the following:
   - If your image builder is not joined to an Active Directory domain, on the **Local User** tab, choose **Administrator**.
   - If your image builder is joined to an Active Directory domain, choose the **Directory User** tab, and log in as a domain user who has local administrator permissions on the image builder.
8. On the image builder desktop, open Image Assistant.
9. Follow the required steps in Image Assistant to finish creating your image. For more information, see Tutorial: Create a Custom WorkSpaces Applications Image by Using the WorkSpaces Applications Console.

I need to persist environment variables across my fleet instances.

Environment variables enable you to dynamically pass settings across applications. You can make user environment variables and system environment variables available across your fleet instances. You can also create environment variables with limited scope, which is useful when you need to use the same environment variable with different values across different applications. For more information, see Persist Environment Variables in Amazon WorkSpaces Applications.

I want to change the default Internet Explorer home page for my users.

You can use Group Policy to set the default home page in Internet Explorer for your users. You can also enable users to change the default page that you set. For more information, see Change the Default Internet Explorer Home Page for Users' Streaming Sessions in Amazon WorkSpaces Applications.

When my users end a streaming session and then start a new one, they see a message that says no streaming resources are available.

When a user ends a session, WorkSpaces Applications terminates the underlying instance and creates a new instance if needed to meet the desired capacity of the fleet. If a user tries to start a

new session before WorkSpaces Applications creates the new instance and all other instances are in use, the user will receive an error stating that no streaming resources are available. If your users start and stop sessions frequently, consider increasing your fleet capacity. For more information, see Fleet Auto Scaling for Amazon WorkSpaces Applications. Or, consider increasing the maximum session duration for your fleet and instructing your users to close their browser during periods of inactivity rather than ending their session.

Troubleshooting Active Directory

The following are issues that might occur when you set up and use Active Directory with Amazon WorkSpaces Applications. For help troubleshooting notification codes, see Troubleshooting Notification Codes.

Issues

- My image builders and fleet instances are stuck in the PENDING state.
- My users aren't able to log in with the SAML application.
- My fleet instances work for one user but don't cycle correctly.
- My user Group Policy objects aren't being successfully applied.
- My WorkSpaces Applications streaming instances aren't joining the Active Directory domain.
- User login is taking a long time to complete on a domain-joined streaming session.
- My users can't access a domain resource in a domain-joined streaming session, but they can access the resource from a domain-joined image builder.
- My users receive the error “Certificate-Based Authentication not available” and are prompted to enter their domain password. Or users receive the error “Disconnected from session” when they are starting a session enabled with certificate-based authentication.
- I'm experiencing domain join failures after changing the Active Directory (AD) service account.

My image builders and fleet instances are stuck in the PENDING state.

Image builders and fleet instances can take up to 25 minutes to move into a ready state and become available. If your instances are taking longer than 25 minutes to become available, in Active Directory, verify whether new computer objects were created in the correct organizational units (OUs). If there are new objects, the streaming instances will be available soon. If the objects aren't there, check the directory configuration details in your WorkSpaces Applications Directory

Config: Directory name (the fully qualified domain name of the directory, service account sign-in credentials, and the OU distinguished name.

Image builder and fleet errors are displayed in the WorkSpaces Applications console on the **Notifications** tab for the fleet or image builder. Fleet errors are also available using the WorkSpaces Applications API through the DescribeFleets operation or the CLI command describe-fleets.

My users aren't able to log in with the SAML application.

WorkSpaces Applications relies on the SAML_Subject "NameID" attribute from your identity provider to populate the username field to log in your user. The username can either be formatted as "*domain*\username", or "user@domain.com". If you are using "*domain*\username" format, *domain* can either be the NetBIOS name or the fully qualified domain name. If using "user@domain.com" format, the UserPrincipalName attribute can be used. If you've verified your SAML_Subject attribute is configured correctly and the problem persists, contact AWS Support. For more information, see AWS Support Center.

My fleet instances work for one user but don't cycle correctly.

Fleet instances are cycled after a user completes a session, ensuring that each user has a new instance. When the cycled fleet instance is brought online, it joins the domain using the computer name of the previous instance. To ensure that this operation happens successfully, the service account requires **Change Password** and **Reset Password** permissions on the organizational unit (OU) to which the computer object is joining. Check the service account permissions and try again. If the problem persists, contact AWS Support. For more information, see AWS Support Center.

My user Group Policy objects aren't being successfully applied.

By default, computer objects apply computer-level policies based on the OU in which the computer object resides, while applying user-level policies based on the OU in which the user resides. If your user-level policies aren't being applied, you can do one of the following:

- Move the user-level policies to the OU in which the user Active Directory object resides
- Enable computer-level loopback processing, which applies the user-level policies in the computer object OU.

For more information, see Loopback processing of Group Policy at Microsoft Support.

My WorkSpaces Applications streaming instances aren't joining the Active Directory domain.

The Active Directory domain to use with WorkSpaces Applications must be accessible through its fully qualified domain name (FQDN) through the VPC in which your streaming instances are launched.

To test that your domain is accessible

1. Launch an Amazon EC2 instance in the same VPC, subnet, and security groups that you use with WorkSpaces Applications.
2. Manually join the EC2 instance to your Active Directory domain by using the FQDN (for example, `yourdomain.example.com`) with the service account that you intend to use with WorkSpaces Applications. Use the following command in a Windows PowerShell console:

   ```
   netdom join computer /domain:FQDN /OU:path /ud:user /pd:password
   ```

   If this manual join fails, go to the next step.
3. If you cannot manually join to your domain, open a command prompt and verify that you can resolve the FQDN using the `nslookup` command. For example:

   ```
   nslookup yourdomain.exampleco.com
   ```

   Successful name resolution returns a valid IP address. If you are unable to resolve your FQDN, you might need to update your VPC DNS servers by using a DHCP option set for your domain. Then, come back to this step. For more information, see DHCP Options Sets in the *Amazon VPC User Guide*.
4. If the FQDN resolves, use the `telnet` command to validate connectivity.

   ```
   telnet yourdomain.exampleco.com 389
   ```

   A successful connection shows a blank command prompt window without any connection errors. You might need to install the Telnet Client feature on your EC2 instance. For more information, see Install Telnet Client in the Microsoft documentation.

If you were not able to manually join the EC2 instance to your domain, but were successful in resolving the FQDN and testing connectivity with the Telnet Client, your VPC security groups might be preventing access. Active Directory requires certain network port settings. For more information, see Active Directory and Active Directory Domain Services Port Requirements in the Microsoft documentation.

User login is taking a long time to complete on a domain-joined streaming session.

WorkSpaces Applications performs a Windows login action after users provide their domain password. After successful authentication, WorkSpaces Applications launches the application. The login and launch times are impacted by many variables, such as network contention for the domain controllers or the time it takes to apply Group Policy settings to the streaming instance. If domain authentication takes too long to complete, try performing the following actions.

- Minimize the network latency from your WorkSpaces Applications Region to your domain controllers by choosing the correct domain controllers. For example, if your fleet is in `us-east-1`, use domain controllers with high bandwidth and low latency to `us-east-1` through Active Directory Sites and Services zone mappings. For more information, see Active Directory Sites and Services in the Microsoft documentation.
- Ensure that your Group Policy settings and user login scripts don't take prohibitively long to apply or run.

If your domain users' login to WorkSpaces Applications fails with the message "An unknown error occurred," you might need to update the Group Policy settings described in Before You Begin Using Active Directory with Amazon WorkSpaces Applications. Otherwise, these settings might prevent WorkSpaces Applications from authenticating and logging in your domain users.

My users can't access a domain resource in a domain-joined streaming session, but they can access the resource from a domain-joined image builder.

Confirm that your fleet is created in the same VPC, subnets, and security groups as your image builder, and that your user has the permissions required to access and use the domain resource.

My users receive the error “Certificate-Based Authentication not available” and are prompted to enter their domain password. Or users receive the error “Disconnected from session” when they are starting a session enabled with certificate-based authentication.

These errors occur if certificate-based authentication was unsuccessful for the session. The “Certificate-Based Authentication not available” error is displayed when certificate-based authentication is enabled to allow fallback to password logon. The “Disconnected from session” error is displayed when certificate-based authentication is enabled without fallback.

The user can refresh the page on the web client or reconnect from the client for Windows, as this may be an intermittent issue with certificate-based authentication. If the problem continues, certificate-based authentication failure can result from one of the following issues:

- WorkSpaces Applications could not communicate with AWS Private CA, or AWS Private CA did not issue the certificate. Check CloudTrail to determine if a certificate was issued. For more information, see What Is AWS CloudTrail? and the section called “Manage Certificate-based Authentication”.
- The domain controller has no domain controller certificate for smart card logon, or it is expired. For more information, see step 7.a in the section called “Prerequisites”.
- The certificate is not trusted. For more information, see step 7.c in the section called “Prerequisites”.
- The userPrincipalName format for the SAML_Subject NameID is not formatted properly, or does not resolve to the actual domain for the user. For more information, see step 1 in the section called “Prerequisites”.
- The (optional) ObjectSid attribute in your SAML assertion does not match the Active Directory security identifier (SID) for the user specified in the SAML_Subject NameID. Confirm that the attribute mapping is correct in your SAML federation, and that your SAML identity provider is synchronizing the SID attribute for the Active Directory user.
- The WorkSpaces Applications agent does not support certificate-based authentication. Use WorkSpaces Applications agent version 10-13-2022 or later.
- There are Group Policy settings that are modifying the default Active Directory settings for smart card logon, or taking action if a smart card is removed from a smart card reader. These settings may cause additional unexpected behavior other than the errors listed above. Certificate-based authentication presents a virtual smart card to the instance operating system, and removes it

after logon is complete. For more information, see Primary Group Policy settings for smart cards and Additional smart card Group Policy settings and registry keys. Do not enable **Smart card sign in for Active Directory** in your stack if you want to use certificate-based authentication. For more information, see the section called “Smart Cards”.

- The CRL distribution point for the private CA is not online or accessible from either the WorkSpaces Applications fleet instance or the domain controller. For more information, see step 5 in the section called “Prerequisites”.

Additional troubleshooting steps involve reviewing the WorkSpaces Applications instance Windows event logs. A common event to review for logon failure is 4625(F): An account failed to log on. For more information about capturing log information, see Persisting application and Windows event logs. Alternatively, to troubleshoot an active WorkSpaces Applications session as an administrator, you can connect to the logs using an Event Viewer on another computer. For more information, see How to Select Computers in Event Viewer. Or, you can connect by using Remote Desktop to connect to the instance private IP address from another computer that can connect to Remote Desktop Services in your WorkSpaces Applications virtual private cloud (VPC). Use the AWS CLI to determine the IP address for the session based on the AWS Region, WorkSpaces Applications stack name, fleet name, user ID, and authentication type. For more information, see the AWS Command Line Interface.

If the problem persists, contact AWS Support. For more information, see AWS Support Center.

I'm experiencing domain join failures after changing the Active Directory (AD) service account.

If you have an existing fleet with an image based on the August 2024 Microsoft Windows Server operating system update, and if you change your Active Directory (AD) service account for that fleet, your fleet instances might encounter domain join failures during provisioning.

Microsoft has released a patch KB5020276, which modifies the behavior of domain join operations. WorkSpaces Applications reuses existing computer objects when joining your streaming instances to your AD domains. This computer object is generated using the AD service account that you provide when you create a fleet or Directory Config with WorkSpaces Applications. Prior to this Microsoft patch, new AD service accounts could reuse existing computer objects created by WorkSpaces Applications, as long as they had "Create Computer Object" permissions configured in the organizational unit (OU).

When the Microsoft patch is enforced, starting on August 13, 2024, and if you change your AD service account for an existing WorkSpaces Applications fleet, the new service account will no longer be able to reuse the existing computer objects in the AD. This results in domain join failures on WorkSpaces Applications fleets, with one of the following error messages under fleet notifications:

- DOMAIN_JOIN_INTERNAL_SERVICE_ERROR "The group name could not be found."
- An account with the same name exists in Active Directory. Re-using the account was blocked by security policy

To control which account can reuse the existing computer objects, Microsoft has implemented a new Group Policy setting called **Domain controller: Allow computer account re-use during domain join**. This setting allows you to specify a list of trusted service accounts that bypass the check during the domain join operation. For your self-managed AD configuration, we recommend following the Microsoft documented steps to add your AD service account to the new allow list policy, using Group Policies on a domain controller.

For Managed Active Directory (MAD), you must restart your WorkSpaces Applications fleet after you make changes to your WorkSpaces Applications domain join service account.

If the problem persists, contact AWS Support. For more information, see AWS Support Center.

Troubleshooting WorkSpaces Applications User Issues

Enable advanced logging

To help troubleshoot issues that your users might experience, you can enable advanced logging on any WorkSpaces Applications client. Advanced logging generates log files that contain diagnostic information and debugging-level details, including verbose performance data.

> **Note**
>
> To get an AWS review of advanced logging files, and receive technical support for issues with your WorkSpaces Applications clients, contact Support. For more information, see the AWS Support Center Console.

Enable advanced logging for web access

If users are using SAML, User Pool, or have access to the Application Catalogue page, follow these steps:

1. Load the catalog page.
2. Open **Developer tools** and choose the **Console** tab.
3. In the browser console, enter **`window.siteConfig.logLevel = "INFO"`** and choose **Enter**.
4. Launch the application, and you should see **Logging** on the **Console** tab.
5. Reproduce the issue.
6. Right-click on the **Console** tab and choose **Save all Messages to File**.

Enable advanced logging for Windows clients

To enable advanced logging for Windows clients, follow these steps:

1. On the client machine, go to `%localappdata%\AppStreamClient\app-<versionID>`.
2. Open `Log4Net.config` in Notepad.
3. Change the root level of logging from **INFO** to **DEBUG**.
4. Save the file.
5. Restart the WorkSpaces Applications Client and try connecting again.
6. Gather the logs from `C:\Users\%USERNAME%\AppData\Local\Amazon\AppStreamClient\'` by zipping the complete folder.

The following are specific issues that might occur for your users when they use WorkSpaces Applications.

Issues

- My users' WorkSpaces Applications client installations fail, and they're getting a message stating that .NET Framework 4.6 is required.
- My users' USB driver installations fail when they install the WorkSpaces Applications client, and now they can't use their USB devices with WorkSpaces Applications.
- My WorkSpaces Applications client users are getting disconnected from their WorkSpaces Applications session after every 60 minutes.
- My users can’t copy and paste between their local device and their streaming session.

- Some keyboard shortcuts aren't working for users during their streaming sessions.
- My users' drawing tablets are not working with the streaming applications I deployed.
- The Japanese language input method doesn't work for my users during their streaming sessions
- My user sees an error about reaching the max number of streaming sessions when they try to launch an application from the application catalog.
- My user sees a black screen or the desktop, and their application doesn't launch on an Elastic fleet. No error appears.

My users' WorkSpaces Applications client installations fail, and they're getting a message stating that .NET Framework 4.6 is required.

When users install the WorkSpaces Applications client, WorkSpaces Applications also installs .NET Framework version 4.6.2, if that version or a later version is not already installed. If the PC on which the client is being installed is not connected to the internet, .NET Framework can't be installed. In this case, a message prompts users to install .NET Framework version 4.6 manually. However, when users choose **Install**, an error message is displayed stating that the installation failed. Users are then prompted to try installing the latest version of the .NET Framework manually. When they choose **Close**, they exit the installation.

To resolve this issue, users must establish an internet connection from the PC on which they plan to install the client, and then download and install .NET Framework version 4.6.2 or later on the same PC. For a list of the .NET Framework versions available for download, see Download .NET Framework.

> **Note**
>
> Users who have version 1.1.156 of the WorkSpaces Applications client installed must have .NET Framework version 4.7.2 or later installed on the same PC.

My users' USB driver installations fail when they install the WorkSpaces Applications client, and now they can't use their USB devices with WorkSpaces Applications.

When users install the WorkSpaces Applications client, they choose whether to install the WorkSpaces Applications USB driver. The driver is required to use USB devices with applications

streamed through WorkSpaces Applications. However, the USB driver installation fails if both of the following occur:

- The root certificate used to sign the `AppStreamUsbDriver.exe` file is not present in the Windows certificate store.
- The PC on which the client is being installed is not connected to the internet.

In this case, the certificate for the Amazon AppStream USB driver can't be validated, and an error message notifies users that the USB driver installation failed. When users choose **OK**, the WorkSpaces Applications client installation is completed without the USB driver. Although users can still use the WorkSpaces Applications client for application streaming, their USB devices won't work with applications streamed through WorkSpaces Applications.

To resolve this issue, users must establish an internet connection from the PC on which they plan to install the WorkSpaces Applications client, and reinstall the client.

My WorkSpaces Applications client users are getting disconnected from their WorkSpaces Applications session after every 60 minutes.

If you have configured identity federation using SAML 2.0 for access to WorkSpaces Applications, depending on your identity provider (IdP), you may need to configure the information that the IdP passes as SAML attributes to AWS as part of the authentication response. This includes configuring the **Attribute** element with the `SessionDuration` attribute set to `https://aws.amazon.com/SAML/Attributes/SessionDuration`.

`SessionDuration` specifies the maximum amount of time that a federated streaming session for a user can remain active before reauthentication is required. Although `SessionDuration` is an optional attribute, we recommend that you include it in the SAML authentication response. If you do not specify this attribute, the session duration is set to a default value of 60 minutes.

To resolve this issue, configure your SAML-compatible IdP to include the `SessionDuration` value in the SAML authentication response, and set the value as required. For more information, see Step 5: Create Assertions for the SAML Authentication Response.

> **Note**
>
> If your users access their streaming applications in WorkSpaces Applications by using the WorkSpaces Applications native client or by using the web browser on the new experience,

their sessions are disconnected after their session duration expires. If your users access their streaming applications in WorkSpaces Applications by using a web browser on the old/classic experience, after the users' session duration expires and they refresh their browser page, their sessions are disconnected.

If your users sign in to the new portal experience with a SAML-compatible IdP, and they continue to have random disconnections, it might be due to the session cookies used by the WorkSpaces Applications session being invalidated by other web applications using `aws.amazon.com` as a subdomain. The following are common user scenarios:

- If a user initiates a new WorkSpaces Applications session in the same browser, the existing WorkSpaces Applications session will be disconnected.
- If a user initiates any other web applications in the same browser, resulting in a new user authentication under the `aws.amazon.com` domain, the existing WorkSpaces Applications session will be disconnected.
- If a user signs into an AWS Management Console with new IAM credentials in the same browser, the existing WorkSpaces Applications session will be disconnected.

You can resolve this issue by using the new relay state endpoints to configure your SAML 2.0 federation, and by using the WorkSpaces Applications client version 1.1.1300 and later. For more information, see Table 1 on the section called “Step 6: Configure the Relay State of Your Federation”.

My users can’t copy and paste between their local device and their streaming session.

WorkSpaces Applications takes advantage of the W3C specification for enabling asynchronous clipboard operations in web applications. This enables users to copy and paste content between their local device and their streaming session in the same ways that they copy and paste between applications on their local device, including using keyboard shortcuts.

The only browser that currently supports the W3C asynchronous clipboard specification is Google Chrome version 66 or later, which supports copying and pasting only for text. For all other browsers, users can use the clipboard feature in the WorkSpaces Applications web portal, which provides a dialog box for copying or pasting text.

If your users run into issues using the clipboard during their streaming sessions, you can provide them with the following information:

- **I'm using Chrome version 66 or later, and keyboard shortcuts aren't working.**

  Chrome displays a prompt for you to choose whether to allow WorkSpaces Applications to access content copied to the clipboard. Choose **Allow** to enable pasting to your remote session. If you're copying text from your remote session to your local device, both the Chrome application and the tab containing your streaming session must stay in focus on your local device long enough for the text to be copied from your streaming session. Small amounts of text should be copied almost immediately, but for large amounts of text, you might need to wait 1 to 2 seconds before switching away from Chrome or from the tab containing your streaming session. The time required to copy the text varies based on network conditions.

- **Copying and pasting doesn't work when I try to copy and paste a large amount of text.**

  WorkSpaces Applications has a default limit of 20 MB for the amount of text that you can copy and paste between your local device and your streaming session. If you try to copy more than 20 MB, no text is copied. However, the text will be truncated if your admin set a limit and you go beyond that limit. This limit doesn't apply if you try to copy and paste text between applications on your local device or between applications in your streaming session. Administrators can also limit the number of characters that you copy/paste in/out of your streaming sessions. If you need to copy or paste text more than 20 MB or the specified limit between your local device and your streaming session, you can divide it into smaller chunks or upload it as a file instead.

- **I'm using the WorkSpaces Applications web portal clipboard feature to paste text to my streaming session and it's not working.**

  In some cases, after you paste text into the clipboard dialog box and the dialog box closes, nothing happens when you try to use keyboard shortcuts to paste the text in your streaming session. This issue occurs because when the clipboard dialog box appears, it takes the focus away from your streaming application. After the dialog box closes, the focus might not automatically return to your streaming application. Clicking your streaming application should return the focus to it and enable you to use keyboard shortcuts to paste your text into your streaming session.

Some keyboard shortcuts aren’t working for users during their streaming sessions.

The following keyboard shortcuts work on users' local computers, but are not passed to WorkSpaces Applications streaming sessions:

Windows:

- Win+L
- Ctrl+Alt+Del

Mac:

- Ctrl+F3
- All shortcuts that use Alt or Option key combinations

This issue is due to the following limitations on users’ local computers:

- The keyboard shortcuts are filtered by the operating system that is running on users’ local computers and not propagated to the browsers on which users are accessing WorkSpaces Applications. This behavior applies to the Windows Win+L and Ctrl+Alt+Del keyboard shortcuts and Mac Ctrl+F3 keyboard shortcut.
- When used with web applications, some keyboard shortcuts are filtered by the browser and don’t generate an event for the web applications. As a result, the web applications can’t respond to the keyboard shortcuts typed by users.
- The keyboard shortcuts are translated by the browser before a keyboard event is generated and so are not translated correctly. For example, Alt key combinations and Option key combinations on Mac computers are translated as if they are Alt Graph key combinations on Windows. When this occurs, the results are not as the users intend when they use these key combinations.

My users' drawing tablets are not working with the streaming applications I deployed.

If your users' drawing tablets are not working with streaming applications, make sure that you meet the requirements and understand additional considerations for enabling this feature.

Following are the requirements and considerations for enabling your users to use drawing tablets during WorkSpaces Applications streaming sessions.

> **Note**
>
> Drawing tablets are supported for users who access WorkSpaces Applications by using the WorkSpaces Applications client, or through a supported web browser.

- To enable your users to use this feature, you must configure your WorkSpaces Applications fleet to use an image that runs Windows Server 2019.
- To use this feature, users must access WorkSpaces Applications by using the WorkSpaces Applications client, or through the Google Chrome or Mozilla Firefox browsers only.
- Streaming applications must support Windows Ink technology. For more information, see Pen interactions and Windows Ink in Windows apps.
- Some applications, such GIMP, must detect drawing tablets on the streaming instance to support pressure sensitivity. If this is the case, your users must use the WorkSpaces Applications client to access WorkSpaces Applications and stream these applications. In addition, you must qualify your users' drawing tablets, and users must share their drawing tablets with WorkSpaces Applications every time they start a new streaming session.
- This feature is not supported on Chromebooks.

The Japanese language input method doesn't work for my users during their streaming sessions

To enable your users to use the Japanese language input method during their WorkSpaces Applications streaming sessions, do the following:

- Configure your fleet to use the Japanese input method. To do so, enable the Japanese input method on your image builder when you create an image, and then configure your fleet to use the image. For more information, see Specify a Default Input Method. Doing so enables WorkSpaces Applications to automatically configure your image to use a Japanese keyboard. For more information, see Japanese Keyboards.
- Ensure that the Japanese input method is also enabled on the user's local computer.

If the fleet instance and the user's local computer don't use the same language input method, the mismatch might result in unexpected keyboard inputs on the fleet instance during the user's streaming sessions. For example, if the fleet instance uses the Japanese input method and the user's local computer uses the English input method, during a streaming session, the local computer will send keys to the fleet instance that have different key mappings than the fleet instance.

To verify whether the Japanese input method is enabled for a fleet instance, enable the **Desktop** stream view for the fleet. For more information, see Step 6 in Create a Fleet in Amazon WorkSpaces Applications.

Windows Keyboard Shortcuts

Following are Windows keyboard shortcuts for switching Japanese input modes and for Japanese conversions. For these keyboard shortcuts to work, the WorkSpaces Applications streaming session must be active.

Windows keyboard shortcuts for switching Japanese input modes

Keyboard shortcut	Description
半角/全角/漢字 (Hankaku/Zenkaku/Kanji) Or Alt+`	Switches the input mode between alphanumeric and Japanese mode
無変換 (Muhenkan)	Converts characters to Hiragana, full-width Katakana, and half-width Katakana in sequence
カタカナ/ひらがな/ローマ字 (Katakana/Hiragana/Romaji)	Changes the input mode to Hiragana
Shift+カタカナ/ひらがな/ローマ字 (Katakana/Hiragana/Romaji)	Changes the input mode to Katakana
Alt+カタカナ/ひらがな/ローマ字	Switches the input mode between Japanese Romaji and Japanese Kana

Keyboard shortcut	Description
(Katakana/Hiragana/Romaji)	

Windows keyboard shortcuts for Japanese conversions

Keyboard shortcut	Description
変換 (Henkan) + Space	Lists conversion options
F6	Converts to Hiragana
F7	Converts to full-width Katakana
F8	Converts to half-width Katakana
F9	Converts to full-width Romaji
F10	Converts to half-width Romaji

Mac Keyboard Shortcuts

For information about Mac keyboard shortcuts for switching Japanese input methods and for Japanese conversions, see the following articles in the Mac Support documentation.

> **Note**
>
> Because WorkSpaces Applications streaming sessions run on Windows instances, Mac users might experience different key mappings.

- Keyboard shortcuts for switching Japanese input methods — Set up and switch to a Japanese input source on Mac
- Keyboard short link cuts for Japanese conversions — Keyboard shortcuts for Japanese conversions on Mac

My user sees an error about reaching the max number of streaming sessions when they try to launch an application from the application catalog.

With WorkSpaces Applications Elastic fleets, you specify a maximum number of users that can stream concurrently using the max concurrency parameter. Any user that tries to stream beyond that value receives this error. To resolve this issue, you can increase the maximum number of concurrent streams, or advise your user to wait for another user to complete their streaming session.

> **Note**
>
> You might need to request a limit increase to increase the instance type and size limit.

My user sees a black screen or the desktop, and their application doesn't launch on an Elastic fleet. No error appears.

This can happen if the application launch path is incorrect, and AppStream 2.0 can't launch the application. You can validate the application launch path by using Desktop View on the fleet to navigate the root volume. Validate that the application executable exists at the path specified.

If you're not able to find the app block's VHD or setup script on the streaming instance, AppStream 2.0 might not have been able to download them from the S3 bucket. Validate that the VPC you specified has access to S3. For more information, see Using Amazon S3 VPC Endpoints for WorkSpaces Applications Features.

Troubleshooting Persistent Storage Issues

Amazon WorkSpaces Applications supports the following options for persistent storage: Home folders, Google Drive for G Suite, and OneDrive for Business. Because content synchronization behaviors are consistent across these persistent storage solutions, we recommend that you review Home Folder Content Synchronization for information about expected behavior.

The following are issues that might occur when you or your users use WorkSpaces Applications persistent storage.

Issues

- My stack's home folders aren't working correctly.
- My users can't access their home folder directory from one of our applications.
- My users receive a “Device is not ready” error message when they access their home folder from one of our applications.
- I removed or replaced a file in a user’s home folder in Amazon S3, but my users don’t see the changes in their home folder on the fleet instance during their streaming sessions.
- Persistent storage isn't performing as expected. My users' files are taking longer than expected to save to persistent storage.
- My users are getting errors that files are already in use when their files are not in use.
- When a folder contains thousands of files, WorkSpaces Applications might take a long time to display the list of files.

My stack's home folders aren't working correctly.

Problems with home folder backup to an S3 bucket can occur in the following scenarios:

- There is no internet connectivity from the streaming instance, or there is no access to the private Amazon S3 VPC endpoint, if applicable.
- Network bandwidth consumption is too high. For example, multiple large files are being downloaded or streamed by the user while the service is trying to back up a home folder that contains large files to Amazon S3.
- A file is larger than 5 GB.
- An administrator deleted the bucket created by the service.
- An administrator incorrectly edited the Amazon S3 permissions for the `AmazonAppStreamServiceAccess` service role.

For more information, see the Amazon Simple Storage Service User Guide.

My users can't access their home folder directory from one of our applications.

Some applications do not recognize the redirect that displays the home folder as a top-level folder in File Explorer. If this is the case, your users can access their home folder from within an

application during a streaming session by choosing **File Open** from the application interface and browsing to either of the following directories:

- Non-domain-joined Windows instances: C:\Users\PhotonUser\My Files\Home Folder
- Domain-joined Windows instances: C:\Users\%username%\My Files\Home Folder
- Linux instances: ~/MyFiles/HomeFolder

My users receive a “Device is not ready” error message when they access their home folder from one of our applications.

Persistent storage mounting happens after a user logs in, and it can take several seconds. A “Device is not ready” error can happen if your application is trying to access the files from the home folder before persistent storage mounting is complete. We recommend that you try again after waiting for a few minutes.

To avoid this issue, you can use session scripts and monitor the storage mounting status. Then, start the streaming session after mounting is complete. This also improves your end-users' experience. For more information, see the section called “Session Scripts to Manage Your Users' Streaming Experience”.

I removed or replaced a file in a user’s home folder in Amazon S3, but my users don’t see the changes in their home folder on the fleet instance during their streaming sessions.

Differences between content that is stored in a user’s home folder in an S3 bucket and content that is available to a user on a fleet instance during their streaming sessions may be due to the way in which home folder content stored in Amazon S3 buckets is synchronized with home folder content stored on AppStream 2.0 fleet instances.

At the beginning of a user’s WorkSpaces Applications streaming session, WorkSpaces Applications catalogs the user’s home folder files stored in the Amazon S3 bucket for your Amazon Web Services account and Region. When a user uses a streaming application to open a file in their home folder on their fleet instance, AppStream 2.0 downloads the file to the fleet instance.

Changes that a user makes to files on a fleet instance during their active streaming session are uploaded to their home folder in the S3 bucket every few seconds, or at the end of the user’s streaming session.

If a user opens a file in their home folder on a fleet instance during a streaming session and then closes the file without making any changes or saving the file, and you remove the file from that user's home folder in an S3 bucket during the streaming session, the file is removed from the fleet instance if the user refreshes the folder. If the user modifies the file and saves the file locally, the file remains available to the user on the fleet instance during their current streaming session. The file is also uploaded to the S3 bucket again. However, the file may or may not be available to the user on the fleet instance during their next streaming session.

The availability of the file on the fleet instance during a user's next streaming session depends on whether the user changed the file on the fleet instance before or after you changed the file in the S3 bucket.

For more information, see Home Folder Content Synchronization.

Persistent storage isn't performing as expected. My users' files are taking longer than expected to save to persistent storage.

During WorkSpaces Applications streaming sessions, saving large files and directories associated with compute-intensive applications to persistent storage can take longer than saving files and directories required for basic productivity applications. For example, it might take longer for applications to save a large amount of data or frequently modify the same files than it would to save files created by applications that perform a single write action. It might also take longer to save many small files.

If your users save files and directories associated with compute-intensive applications and WorkSpaces Applications persistent storage options aren't performing as expected, we recommend that you use a Server Message Block (SMB) solution such as Amazon FSx for Windows File Server or an AWS Storage Gateway file gateway. Following are examples of files and directories associated with compute-intensive applications that are more suitable for use with these SMB solutions:

- Workspace folders for integrated development environments (IDEs)
- Local database files
- Scratch space folders created by graphics simulation applications

For more information, see:

- *Amazon FSx for Windows File Server Windows User Guide*
- Using Amazon FSx with Amazon WorkSpaces Applications

- File gateways in the *AWS Storage Gateway User Guide*

> **Note**
>
> Before proceeding with further troubleshooting, first ensure that the issue your users are experiencing with saving files and directories is associated with WorkSpaces Applications persistent storage only, and not another cause. To rule out other causes, have your users try saving the files or directories to the Temporary Files directory that is available on their streaming instance.

My users are getting errors that files are already in use when their files are not in use.

This behavior is typically occurs in the following cases:

- When users' files are still being uploaded after the files were last saved
- Files that are modified frequently (for example, database files)

Large file uploads might take significant time. Also, each attempt to upload the file might result in another file update, which can lead to repeated file upload attempts.

To resolve this issue, we recommend that you use a Server Message Block (SMB) solution such as Amazon FSx for Windows File Server or an AWS Storage Gateway file gateway. For more information, see:

- *Amazon FSx for Windows File Server Windows User Guide*
- Using Amazon FSx with Amazon WorkSpaces Applications
- File gateways in the *AWS Storage Gateway User Guide*

When a folder contains thousands of files, WorkSpaces Applications might take a long time to display the list of files.

WorkSpaces Applications uses API calls to retrieve the content of folders that are stored in WorkSpaces Applications persistent storage. There is a limit to the number of items that an API call can retrieve each time the call runs. For this reason, if WorkSpaces Applications must retrieve

thousands of files in a single folder, it might take more time to display the list of all the files than it would to display the list of files in a folder that contains fewer files.

To resolve this issue, if you have thousands of files in one folder, we recommend that you divide this content into groups of fewer files and store each group in a different folder. Doing so reduces the number of API calls that are required to display the list of files in each folder.

Troubleshooting Notification Codes

The following are notification codes and resolution steps for notifications that you might see when you set up and use Amazon WorkSpaces Applications. These notifications can be found in the **Notifications** tab in the WorkSpaces Applications console, after selecting an image builder or fleet. You can also get fleet notifications by using the WorkSpaces Applications API operation DescribeFleets or the describe-fleets CLI command.

Active Directory Internal Service

Follow these steps if you receive an internal service error when you set up and use Active Directory with Amazon WorkSpaces Applications.

INTERNAL_SERVICE_ERROR

Message: The user name or password is incorrect.

Resolution: This error might occur when the computer object that was created in the Microsoft Active Directory domain for the resource was deleted or disabled. You can resolve this error by enabling the computer object in the Active Directory domain, and then starting the resource again. You might also need to reset the computer object account in the Active Directory domain. If you continue to encounter this error, contact AWS Support. For more information, see AWS Support Center.

Active Directory Domain Join

The following are notification codes and resolution steps for issues with domain join that you might encounter when you set up and use Active Directory with Amazon WorkSpaces Applications.

DOMAIN_JOIN_ERROR_ACCESS_DENIED

Message: Access is denied.

Resolution: The service account specified in the directory configuration does not have permissions to create the computer object or reuse an existing one. Validate the permissions and start the image builder or fleet. For more information, see Granting Permissions to Create and Manage Active Directory Computer Objects.

DOMAIN_JOIN_ERROR_LOGON_FAILURE

Message: The username or password is incorrect.

Resolution: The service account specified in the directory configuration has an invalid username or password. Update the configuration and re-create the image builder or fleet that had the error.

DOMAIN_JOIN_NERR_PASSWORD_EXPIRED

Message: The password of this user has expired.

Resolution: The password for the service account specified in the WorkSpaces Applications directory configuration has expired. Change the password for the service account in your Active Directory domain, update the configuration, and then re-create the image builder or fleet that had the error.

DOMAIN_JOIN_ERROR_DS_MACHINE_ACCOUNT_QUOTA_EXCEEDED

Message: Your computer could not be joined to the domain. You have exceeded the maximum number of computer accounts you are allowed to create in this domain. Contact your system administrator to have this limit reset or increased.

Resolution: The service account specified on the directory configuration does not have permissions to create the computer object or reuse an existing one. Validate the permissions and start the image builder or fleet. For more information, see Granting Permissions to Create and Manage Active Directory Computer Objects.

DOMAIN_JOIN_ERROR_INVALID_PARAMETER

Message: A parameter is incorrect. This error is returned if the `LpName` parameter is NULL or the `NameType` parameter is specified as `NetSetupUnknown` or an unknown nametype.

Resolution: This error can occur when the distinguished name for the OU is incorrect. Validate the OU and try again. If you continue to encounter this error, contact AWS Support. For more information, see AWS Support Center.

DOMAIN_JOIN_ERROR_MORE_DATA

Message: More data is available.

Resolution: This error can occur when the distinguished name for the OU is incorrect. Validate the OU and try again. If you continue to encounter this error, contact AWS Support. For more information, see AWS Support Center.

DOMAIN_JOIN_ERROR_NO_SUCH_DOMAIN

Message: The specified domain either does not exist or could not be contacted.

Resolution: The streaming instance was unable to contact your Active Directory domain. To ensure network connectivity, confirm your VPC, subnet, and security group settings. For more information, see My WorkSpaces Applications streaming instances aren't joining the Active Directory domain.

DOMAIN_JOIN_NERR_WORKSTATION_NOT_STARTED

Message: The Workstation service has not been started.

Resolution: An error occurred starting the Workstation service. Ensure that the service is enabled in your image. If you continue to encounter this error, contact AWS Support. For more information, see AWS Support Center.

DOMAIN_JOIN_ERROR_NOT_SUPPORTED

Message: The request is not supported. This error is returned if a remote computer was specified in the `lpServer` parameter and this call is not supported on the remote computer.

Resolution: Contact AWS Support for assistance. For more information, see AWS Support Center.

DOMAIN_JOIN_ERROR_FILE_NOT_FOUND

Message: The system cannot find the file specified.

Resolution: This error occurs when an invalid organizational unit (OU) distinguished name is provided. The distinguished name must start with **OU=**. Validate the OU distinguished name and try again. For more information, see Finding the Organizational Unit Distinguished Name.

DOMAIN_JOIN_INTERNAL_SERVICE_ERROR

Message: The account already exists.

Resolution: This error can occur in the following scenarios:

- If the issue isn't permissions-related, check the Netdom logs for errors and make sure that you provided the correct OU.

- The service account specified in the directory configuration does not have permissions to create the computer object or reuse an existing one. If this is the case, validate the permissions and start the image builder or fleet. For more information, see Granting Permissions to Create and Manage Active Directory Computer Objects.
- After WorkSpaces Applications creates the computer object, it is moved from the OU in which it was created. In this case, the first image builder or fleet is created successfully, but any new image builder or fleet that uses the computer object fails. When Active Directory searches for the computer object in the specified OU and detects that an object with the same name exists elsewhere in the domain, the domain join is not successful.
- The name of the OU specified in the WorkSpaces Applications Directory Config includes spaces before or after the commas in the directory configuration. In this case, when a fleet or image builder attempts to rejoin the Active Directory domain, WorkSpaces Applications cannot cycle the computer objects correctly and the domain rejoin does not succeed. To resolve this issue for a fleet, do the following:

  1. Stop the fleet.
  2. Edit the Active Directory domain settings for the fleet to remove the Directory Config and Directory OU to which the fleet is joined. For more information, see Step 3: Create a Domain-Joined Fleet.
  3. Update the WorkSpaces Applications Directory Config to specify an OU that doesn't contain spaces. For more information, see Step 1: Create a Directory Config Object.
  4. Edit the Active Directory domain settings for the fleet to specify the Directory Config with the updated Directory OU.

  To resolve this issue for an image builder, do the following:

  1. Delete the image builder.
  2. Update the WorkSpaces Applications Directory Config to specify an OU that doesn't contain spaces. For more information, see Step 1: Create a Directory Config Object.
  3. Create a new image builder and specify the Directory Config with the updated Directory OU. For more information, see Launch an Image Builder to Install and Configure Streaming Applications.

Image Internal Service

If you receive an internal service error after you use managed WorkSpaces Applications image updates to initiate an image update, follow these steps.

INTERNAL_SERVICE_ERROR

Message: WorkSpaces Applications could not update image *image-name*. Failed to update/install/configure/disable <software name>. Check your source image and try again. If this problem persists, contact AWS Support.

Resolution: This error can occur when there is an issue with the source image. Try to update the image again.

If updating again doesn't work, make sure that you're using the latest version of SSM Agent. For version information, see the section called “Base Image and Managed Image Update Release Notes”. For installation information, see Manually install SSM Agent on EC2 instances for Windows Server.

If the error continues to occur, launch an image builder from the image. For more information, see Launch an Image Builder to Install and Configure Streaming Applications. If you can't launch image builder from the image, there is another issue with the image that needs to be resolved before you can use managed AppStream 2.0 image updates to update the image. If you continue to encounter this error, contact AWS Support. For more information, see AWS Support Center.

Session Provisioning

The following are notification codes and resolution steps for issues with session provisioning that you might encounter when your end users try to provision the streaming session.

> **Note**
>
> "X" below equals the number of sessions that encountered a given error code.

USER_PROFILE_MOUNTING_FAILURE

Message: X session(s) encountered user profile mounting failures.

Resolution: To troubleshoot this issue, check if any user profiles have been corrupted or if any third party processes on instance are interfering with user profile mounting. If you continue to encounter this error, contact AWS Support. For more information, see AWS Support Center.

USER_PROFILE_DOWNLOADING_FAILURE

Message: X session(s) encountered user profile downloading failures.

Resolution: To troubleshoot this issue, check your network configuration. If you continue to encounter this error, contact AWS Support. For more information, see AWS Support Center.

HOME_FOLDER_MOUNTING_FAILURE

Message: X session(s) encountered home folder mounting failures.

Resolution: To troubleshoot this issue, check your network configuration. If you continue to encounter this error, contact AWS Support. For more information, see AWS Support Center.

ONE_DRIVE_MOUNTING_FAILURE

Message: X session(s) encountered OneDrive mounting failures.

Resolution: To troubleshoot this issue, check your network configuration and ensure the user has valid Microsoft 365 credentials. If you continue to encounter this error, contact AWS Support. For more information, see AWS Support Center.

GOOGLE_DRIVE_MOUNTING_FAILURE

Message: X session(s) encountered Google Drive mounting failures.

Resolution: To troubleshoot this issue, check your network configuration and ensure the user has valid Google Workspace credentials. If you continue to encounter this error, contact AWS Support. For more information, see AWS Support Center.

Amazon WorkSpaces Applications Service Quotas

WorkSpaces Applications provides different resources that you can use. WorkSpaces Applications resources include stacks, fleets, images, and image builders. When you create your Amazon Web Services account, we set default quotas (also referred to as limits) on the number of resources that you can create, and on the number of users who can use the WorkSpaces Applications service.

To request a quota increase, you can use the Service Quotas console at https://console.aws.amazon.com/servicequotas/. For more information, see Requesting a quota increase in the *Service Quotas User Guide*.

The following table lists the default quotas for each WorkSpaces Applications resource and for users in the WorkSpaces Applications user pool. The actual quotas for your account could be higher or lower, depending on when you created your account.

> **Note**
>
> Graphics Pro instances are no longer available after 10/31/2025 due to End of Life of hardware supporting Graphics Pro instance types. WorkSpaces Applications does not accept limit increase requests for Graphics Pro instances.
> Graphics Design instances will no longer be available from AWS after 12/31/2025 due to End of Life of hardware supporting Graphics Design instance types. WorkSpaces Applications does not accept limit increase requests for Graphics Design instances.

Name	Default	Adjustable
Stacks	10	Yes
Fleets	10	Yes
Compute-optimized fleet instances *	• stream.compute.large: 5 • stream.compute.xlarge: 2 • stream.compute.2xlarge: 0 • stream.compute.4xlarge: 0 • stream.compute.8xlarge: 0	Yes

Name	Default	Adjustable
Graphics fleet instances *	• stream.graphics-design.large: 3 • stream.graphics-design.xlarge: 3 • stream.graphics-design.2xlarge: 3 • stream.graphics-design.4xlarge: 0 • stream.graphics.g4dn.xlarge: 0 • stream.graphics.g4dn.2xlarge: 0 • stream.graphics.g4dn.4xlarge: 0 • stream.graphics.g4dn.8xlarge: 0 • stream.graphics.g4dn.12xlarge: 0 • stream.graphics.g4dn.16xlarge: 0 • stream.graphics.g5.xlarge: 0 • stream.graphics.g5.2xlarge: 0 • stream.graphics.g5.4xlarge: 0 • stream.graphics.g5.8xlarge: 0 • stream.graphics.g5.12xlarge: 0 • stream.graphics.g5.16xlarge: 0 • stream.graphics.g5.24xlarge: 0 • stream.graphics.g6.xlarge: 0 • stream.graphics.g6.2xlarge: 0 • stream.graphics.g6.4xlarge: 0 • stream.graphics.g6.8xlarge: 0 • stream.graphics.g6.16xlarge: 0 • stream.graphics.g6.12xlarge: 0 • stream.graphics.g6.24xlarge: 0 • stream.graphics.gr6.4xlarge: 0 • stream.graphics.gr6.8xlarge: 0 • stream.graphics.g6f.large: 0 • stream.graphics.g6f.xlarge: 0 • stream.graphics.g6f.2xlarge: 0	Yes

Name	Default	Adjustable
	• stream.graphics.g6f.4xlarge: 0 • stream.graphics.gr6f.4xlarge: 0 • stream.graphics.g7.2xlarge: 0 • stream.graphics.g7.4xlarge: 0 • stream.graphics.g7.8xlarge: 0 • stream.graphics.g7.12xlarge: 0 • stream.graphics.g7.24xlarge: 0 • stream.graphics.g7.48xlarge: 0	
Memory-optimized fleet instances *	• stream.memory.large: 5 • stream.memory.xlarge: 2 • stream.memory.2xlarge: 0 • stream.memory.4xlarge: 0 • stream.memory.8xlarge: 0 • stream.memory.z1d.large: 5 • stream.memory.z1d.xlarge: 2 • stream.memory.z1d.2xlarge: 0 • stream.memory.z1d.3xlarge: 0 • stream.memory.z1d.6xlarge: 0 • stream.memory.z1d.12xlarge: 0	Yes
Standard fleet instances *	• stream.standard.small: 50 • stream.standard.medium: 50 • stream.standard.large: 50 • stream.standard.xlarge: 10 • stream.standard.2xlarge: 10	Yes

Name	Default	Adjustable
General purpose fleet instances*	• t3/t3a: small/medium/large/xlarge/2xlarge: 50/50/10/10/10 • m6i/m7i: large/xlarge/2xlarge+: 5/2/0 • m7i-flex: large/xlarge/2xlarge+: 5/2/0 • m6a/m7a: large/xlarge/2xlarge+: 5/2/0	Yes
Compute optimize fleet instances*	• c6i/c7i: large/xlarge/2xlarge+: 5/2/0 • c6a/c7a: large/xlarge/2xlarge+: 5/2/0 • c7i-flex: large/xlarge/2xlarge+: 5/2/0	Yes
Memory optimize fleet instances*	• r6i/r7i: large/xlarge/2xlarge+: 5/2/0 • r6a/r7a: large/xlarge/2xlarge+: 5/2/0	Yes
Accelerated fleet instances*	• G4dn/G5/G6/G6f/Gr6/Gr6f/G6e/G7: All sizes: 0	Yes
Image builders (total)	10	Yes
Images	10	Yes
Compute-optimized image builder instances	• stream.compute.large: 3 • stream.compute.xlarge: 3 • stream.compute.2xlarge: 0 • stream.compute.4xlarge: 0 • stream.compute.8xlarge: 0	Yes

Name	Default	Adjustable
Graphics image builder instances	• stream.graphics-design.large: 1 • stream.graphics-design.xlarge: 1 • stream.graphics-design.2xlarge: 1 • stream.graphics-design.4xlarge: 0 • stream.graphics.g4dn.xlarge: 0 • stream.graphics.g4dn.2xlarge: 0 • stream.graphics.g4dn.4xlarge: 0 • stream.graphics.g4dn.8xlarge: 0 • stream.graphics.g4dn.12xlarge: 0 • stream.graphics.g4dn.16xlarge: 0 • stream.graphics.g5.xlarge: 0 • stream.graphics.g5.2xlarge: 0 • stream.graphics.g5.4xlarge: 0 • stream.graphics.g5.8xlarge: 0 • stream.graphics.g5.12xlarge: 0 • stream.graphics.g5.16xlarge: 0 • stream.graphics.g5.24xlarge: 0 • stream.graphics.g6.xlarge: 0 • stream.graphics.g6.2xlarge: 0 • stream.graphics.g6.4xlarge: 0 • stream.graphics.g6.8xlarge: 0 • stream.graphics.g6.16xlarge: 0 • stream.graphics.g6.12xlarge: 0 • stream.graphics.g6.24xlarge: 0 • stream.graphics.gr6.4xlarge: 0 • stream.graphics.gr6.8xlarge: 0 • stream.graphics.g6f.large: 0 • stream.graphics.g6f.xlarge: 0 • stream.graphics.g6f.2xlarge: 0	Yes

Name	Default	Adjustable
	• stream.graphics.g6f.4xlarge: 0 • stream.graphics.gr6f.4xlarge: 0 • stream.graphics.g7.2xlarge: 0 • stream.graphics.g7.4xlarge: 0 • stream.graphics.g7.8xlarge: 0 • stream.graphics.g7.12xlarge: 0 • stream.graphics.g7.24xlarge: 0 • stream.graphics.g7.48xlarge: 0	
Memory-optimized image builder instances	• stream.memory.large: 3 • stream.memory.xlarge: 3 • stream.memory.2xlarge: 0 • stream.memory.4xlarge: 0 • stream.memory.8xlarge: 0 • stream.memory.z1d.large: 3 • stream.memory.z1d.xlarge: 3 • stream.memory.z1d.2xlarge: 0 • stream.memory.z1d.3xlarge: 0 • stream.memory.z1d.6xlarge: 0 • stream.memory.z1d.12xlarge: 0	Yes
Standard image builder instances	• stream.standard.small: 5 • stream.standard.medium: 5 • stream.standard.large: 5 • stream.standard.xlarge: 3 • stream.standard.2xlarge: 3	Yes

Name	Default	Adjustable
General purpose image builder instances	• t3/t3a: small/medium/large/xlarge/2xlarge: 5/5/5/3/3 • m6i/m7i: large/xlarge/2xlarge+: 3/3/0 • m7i-flex: large/xlarge/2xlarge+: 3/3/0 • m6a/m7a: large/xlarge/2xlarge+: 3/3/0	Yes
Compute optimize image builder instances	• c6i/c7i: large/xlarge/2xlarge+: 3/3/0 • c6a/c7a: large/xlarge/2xlarge+: 3/3/0 • c7i-flex: large/xlarge/2xlarge+: 3/3/0	Yes
Memory optimize image builder instances	• r6i/r7i: large/xlarge/2xlarge+: 3/3/0 • r6a/r7a: large/xlarge/2xlarge+: 3/3/0	Yes
Accelerated image builder instances	• G4dn/G5/G6/G6f/Gr6/Gr6f/G6e/G7: All sizes: 0	Yes
BYOL Standard fleet instances	• stream.standard.medium: 50 • stream.standard.large: 20 • stream.standard.xlarge: 20 • stream.standard.2xlarge: 10	Yes

Name	Default	Adjustable
BYOL Compute-optimized fleet instances	• stream.compute.large: 20 • stream.compute.xlarge: 20 • stream.compute.2xlarge: 10 • stream.compute.4xlarge: 10 • stream.compute.8xlarge: 10	Yes
BYOL Memory-optimized fleet instances	• stream.memory.large: 20 • stream.memory.xlarge: 20 • stream.memory.2xlarge: 10 • stream.memory.4xlarge: 10 • stream.memory.8xlarge: 10	Yes
BYOL Graphics fleet instances	• stream.graphics.g6.xlarge/2xlarge/4xlarge/8xlarge/12xlarge/16xlarge: 0 • stream.graphics.g6f.large/xlarge/2xlarge/4xlarge: 0 • stream.graphics.gr6.4xlarge/8xlarge: 0 • stream.graphics.gr6f.4xlarge: 0 • stream.graphics.g7.2xlarge/4xlarge/8xlarge/12xlarge/24xlarge/48xlarge: 0	Yes
BYOL Standard Image Builders	• stream.standard.medium: 10 • stream.standard.large: 5 • stream.standard.xlarge: 5 • stream.standard.2xlarge: 2	Yes

Name	Default	Adjustable
BYOL Compute-optimized Image Builders	• stream.compute.large: 5 • stream.compute.xlarge: 5 • stream.compute.2xlarge: 2 • stream.compute.4xlarge: 2 • stream.compute.8xlarge: 2	Yes
BYOL Memory-optimized Image Builders	• stream.memory.large: 5 • stream.memory.xlarge: 5 • stream.memory.2xlarge: 2 • stream.memory.4xlarge: 2 • stream.memory.8xlarge: 2	Yes
BYOL Graphics Image Builders	• stream.graphics.g6.xlarge/2xlarge/4xlarge/8xlarge/12xlarge/16xlarge: 0 • stream.graphics.g6f.large/xlarge/2xlarge/4xlarge: 0 • stream.graphics.gr6.4xlarge/8xlarge: 0 • stream.graphics.gr6f.4xlarge: 0 • stream.graphics.g7.2xlarge/4xlarge/8xlarge/12xlarge/24xlarge/48xlarge: 0	Yes
Number of AWS accounts an image can be shared with	100	Yes
Concurrent image copies per destination Region	2	Yes
Concurrent image updates	5	Yes
Users in the user pool	50	Yes

Name	Default	Adjustable
Max concurrent sessions for Elastic fleets	Ubuntu Pro 24.04 LTS • stream.standard.small: 10 • stream.standard.medium: 10 • stream.standard.large: 5 • stream.standard.xlarge: 2 • stream.standard.2xlarge: 2 Windows Server 2019 • stream.standard.small: 10 • stream.standard.medium: 10 • stream.standard.large: 5 • stream.standard.xlarge: 2 • stream.standard.2xlarge: 2	Yes
App block builders (total)	10	Yes
Max number of app block builders	• stream.standard.small: 1 • stream.standard.medium: 1 • stream.standard.large: 1 • stream.standard.xlarge: 1 • stream.standard.2xlarge: 1	Yes

* WorkSpaces Applications instance type and size quotas are per AWS account per AWS Region. If you have multiple fleets in the same Region that use the same instance type and size, the total number of instances in all fleets in that Region must be less than or equal to the applicable quota. To determine which instance types are available in which Regions or Availability Zones, see *Pricing by AWS Region – Always-On, On-Demand, app block builders, and image builder instances* in WorkSpaces Applications Pricing.

For fleets that have **Default Internet Access** enabled, the quota is 100 fleet instances. If your deployment must support more than 100 concurrent users, use the NAT gateway configuration instead. For more information about enabling internet access for a fleet, see Internet Access.

Guidance for WorkSpaces Applications Users

If you are an WorkSpaces Applications administrator, you can provide your users with the guidance in this section to help them get started with using WorkSpaces Applications.

If you are a user who now has access to WorkSpaces Applications, the topics in this section will help you use WorkSpaces Applications for application streaming. With WorkSpaces Applications, your administrator makes your applications available for you to access remotely, so that you don't have to install the applications on your own device. To access your applications, connect to WorkSpaces Applications and start an application streaming session.

Contents

- WorkSpaces Applications Access Methods and Clients
- File Storage Options
- Configure Regional Settings

WorkSpaces Applications Access Methods and Clients

You can connect to WorkSpaces Applications by using a web browser or the WorkSpaces Applications client for Windows.

Contents

- Web Browser Access
- WorkSpaces Applications Client Application for Windows
- WorkSpaces Applications Client Application for macOS

Web Browser Access

The following information helps you use a web browser to connect to WorkSpaces Applications and stream applications.

Contents

- Requirements
- Setup

- Connect to WorkSpaces Applications
- WorkSpaces Applications Web Browser Access (Version 2)
- Monitors and Display Resolution
- USB Devices
- Touchscreen Devices
- Function Keys
- Remap the Mac Option and Command Keys
- Video and Audio Conferencing
- Drawing Tablets
- Relative Mouse Offset
- Troubleshooting

Requirements

You can connect to WorkSpaces Applications from any location by using an HTML5-capable web browser. Supported browsers include the following:

- Google Chrome
- Mozilla Firefox
- Safari
- Microsoft Edge

WorkSpaces Applications supports the three most recent major versions of all supported browsers. Users accessing the web client with older browser versions will receive a notification recommending an update to ensure optimal performance.

> **Note**
>
> Only the Google Chrome or Mozilla Firefox browsers are supported for use with drawing tablets during WorkSpaces Applications streaming sessions. Webcam redirection for video and audio conferencing is supported on Chromium-based web browsers, including Google Chrome and Microsoft Edge.

Setup

No browser extensions or plugins are required to use WorkSpaces Applications in a web browser.

Connect to WorkSpaces Applications

Follow these steps to connect to WorkSpaces Applications and start an application streaming session.

1. If your administrator requires you to sign in first through your organization's sign-in page, complete the tasks in this step.

   If your administrator doesn't require you to sign in through your organization's sign-in page, skip the tasks in this step and proceed to step 2.

   a. Navigate to your organizational sign-in page and enter your domain credentials when prompted.

   b. After you sign in, you are redirected to a page that displays one or more applications that are available for your WorkSpaces Applications streaming session. **Desktop View** is also available, if enabled by your administrator.

   c. Choose an application or, if available, **Desktop View**.

2. If your administrator doesn't require you to sign in first through your organization's sign-in page, do either of the following:

   - If this is the first time that you've used WorkSpaces Applications and you receive a welcome email that notifies you to start accessing your applications using WorkSpaces Applications:

     1. Open the email, and then select the **Login page** link.
     2. Enter your email address and the temporary password that was provided in the email, and then choose **Log in**.
     3. When prompted, enter a new password, confirm it, and then choose **Set Password**.
     4. After a few moments, the WorkSpaces Applications portal opens, displaying one or more applications that are available for your WorkSpaces Applications streaming session. **Desktop View** is also available, if enabled by your administrator.
     5. Choose an application or, if available, **Desktop View**.

   - If this isn't the first time that you've used WorkSpaces Applications and your administrator has provided you with the web address (URL) for the WorkSpaces Applications portal:

1. Enter the URL provided by your administrator to navigate to the WorkSpaces Applications portal.
2. Enter your password when prompted, and choose **Connect**.
3. After a few moments, the WorkSpaces Applications portal opens, displaying one or more applications that are available for your WorkSpaces Applications streaming session. **Desktop View** is also available, if enabled by your administrator.

WorkSpaces Applications Web Browser Access (Version 2)

WorkSpaces Applications web browser access version 2 offers an enhanced end user experience, including menu options that are easily discoverable and textual guidance for end users. No new menu items have been added, and all configuration choices that were available in the previous version are still available in the new one. These setting options have been reorganized as a result of usability testing conducted by the WorkSpaces Applications team.

End users can access an enhanced WorkSpaces Applications toolbar, plus the following features available under **My files**:

- Download or delete multiple files. Select the file(s), choose **Actions**, and then choose **Delete** or **Download**.
- Upload a folder using drag and drop.
- Sort the files based on **Name**, **Last modified date**, and **Size**.
- Wrap the file name column to accommodate long file names. Choose the small gear icon in the top-right corner, **Wrap lines**, and **Confirm**.

End users can also access the following features:

- To switch between two visual modes (light and dark), choose **Preferences**, **General**, **Theme**, and **Light mode** or **Dark mode**.
- To take focus away from the streaming session, and put it on the first element of the toolbar to enable keyboard-based usage during sessions, use the keyboard shortcut **`ctrl + alt + shift + F11`**.

Monitors and Display Resolution

WorkSpaces Applications supports the use of multiple monitors during streaming sessions, including monitors that have different resolutions. To help ensure an optimal streaming experience, we recommend that you set the display scale for your monitors to 100 percent if you use multiple monitors.

You can use dual monitors for application streaming sessions that are started on the following web browsers:

- Google Chrome
- Mozilla Firefox
- Safari
- Microsoft Edge

For browser-based streaming sessions on dual monitors, a maximum display resolution of 4K (4096x2160) is supported per monitor.

> **Important**
>
> Higher resolution monitors require significantly more compute capacity and encoding resources to stream content effectively, with actual performance depending on your display configuration (resolution and number of monitors) and the compute instance being used. If you experience suboptimal performance on 4K monitors, we recommend reducing your display resolution to improve responsiveness.
> Ultrawide monitors with resolutions exceeding 4096 pixels in either dimension (e.g., 5120x2160) will display black bars on the sides, as the maximum supported resolution is limited to 4096 pixels per dimension.

USB Devices

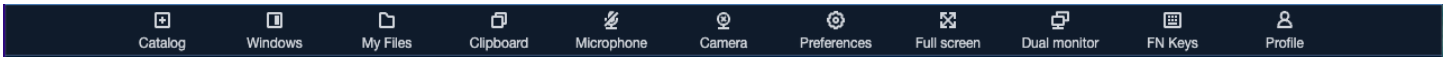
USB devices are not supported for browser-based WorkSpaces Applications streaming sessions. To use your USB devices with applications streamed through WorkSpaces Applications, you must use the WorkSpaces Applications client. For more information, see WorkSpaces Applications Client Application for Windows.

Touchscreen Devices

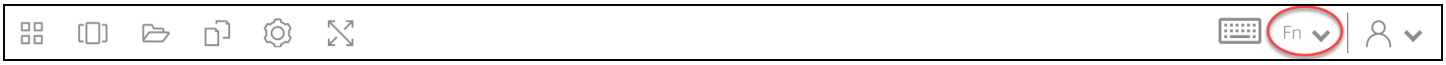
WorkSpaces Applications supports gestures on touch-enabled iPads, Android tablets, and Windows devices. Examples of supported touch gestures include long-tap to right-click, swipe to scroll, pinch to zoom, and two-finger rotation for supporting applications.

> **Note**
>
> Touchscreen devices with a screen size of less than 8 inches are not supported.

To display the on-screen keyboard on an iPad or Android tablet, tap the keyboard icon on the WorkSpaces Applications toolbar. The keyboard icon turns blue, and you can use the on-screen keyboard to input text in the streaming application. Tap the keyboard icon again to hide the on-screen keyboard.



Tap the Fn icon to display a row of Windows-specific keys and keyboard shortcuts.



For touch-enabled devices, the *remote keyboard*, which is displayed when you tap the keyboard icon on the WorkSpaces Applications toolbar, is different than the *local keyboard*, the on-screen keyboard that a touch-enabled device automatically displays when you tap inside an input control in a locally running application. During WorkSpaces Applications streaming sessions, you can use the remote keyboard to input text into streaming applications only. You can display or hide the remote keyboard only by tapping the keyboard icon on the WorkSpaces Applications toolbar. A blue keyboard icon on the WorkSpaces Applications toolbar indicates that the remote keyboard is active.

You can use the local keyboard to input text into elements of the WorkSpaces Applications web portal, including the **My Files** dialog box. However, you can't use this keyboard to input text into streaming applications. Also, you can't display or hide it by using the keyboard icon on the WorkSpaces Applications toolbar.

> **Note**
>
> To display the on-screen keyboard on a Windows computer, tap the keyboard icon in the Windows system tray. If the keyboard icon doesn't appear in the Windows system tray,

switch to Windows tablet mode. Tap the keyboard icon in the Windows system tray again to hide the on-screen keyboard.

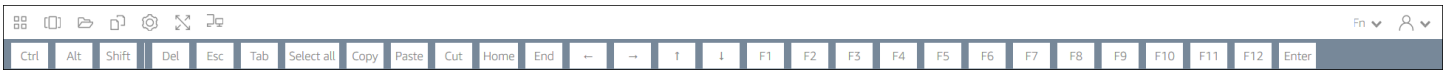
For more information about function keys, see the next section.

Function Keys

You can use keyboard shortcuts during WorkSpaces Applications streaming sessions to enter special keystrokes or key combinations. To display a row of Windows-specific keys and keyboard shortcuts during your streaming session, choose the Fn icon (or the FN Keys on WorkSpaces Applications web browser access v2). The Fn icon is displayed in the WorkSpaces Applications toolbar in the top right of your session window.



Following is an example of how Windows-specific keys and keyboard shortcuts are displayed when you choose the Fn icon (or the FN Keys on WorkSpaces Applications web browser access v2). If not all keys are displayed, you can scroll to the right or left on the shortcut toolbar to display more keys.



To use a key combination that includes the Windows Control key, choose the Ctrl key on the shortcut toolbar, and then type any key on the shortcut toolbar (or, if you are using a touch-enabled device, the on-screen keyboard). Choosing the Ctrl key changes the color to blue. In this case, any other key that you select is interpreted as a key combination that includes the Control key.



Choose the Ctrl key again to release it. For example, to use the keyboard shortcut Ctrl + F, choose the Ctrl key on the shortcut toolbar, and then type the f key. Choose the Ctrl key on the shortcut toolbar again to release the Control key. To use shortcuts that include the Alt or Shift keys, choose the Alt key or the Shift key on the shortcut toolbar in the same way. You can use the Shift key on the shortcut toolbar only for keyboard shortcuts. If you are using a touch-enabled device, this key doesn't affect the capitalization of keys that you type on the on-screen keyboard.

Remap the Mac Option and Command Keys

When you use a device that runs macOS or Mac OS X to connect to WorkSpaces Applications, you can remap the Mac Option and Command keys on your keyboard.

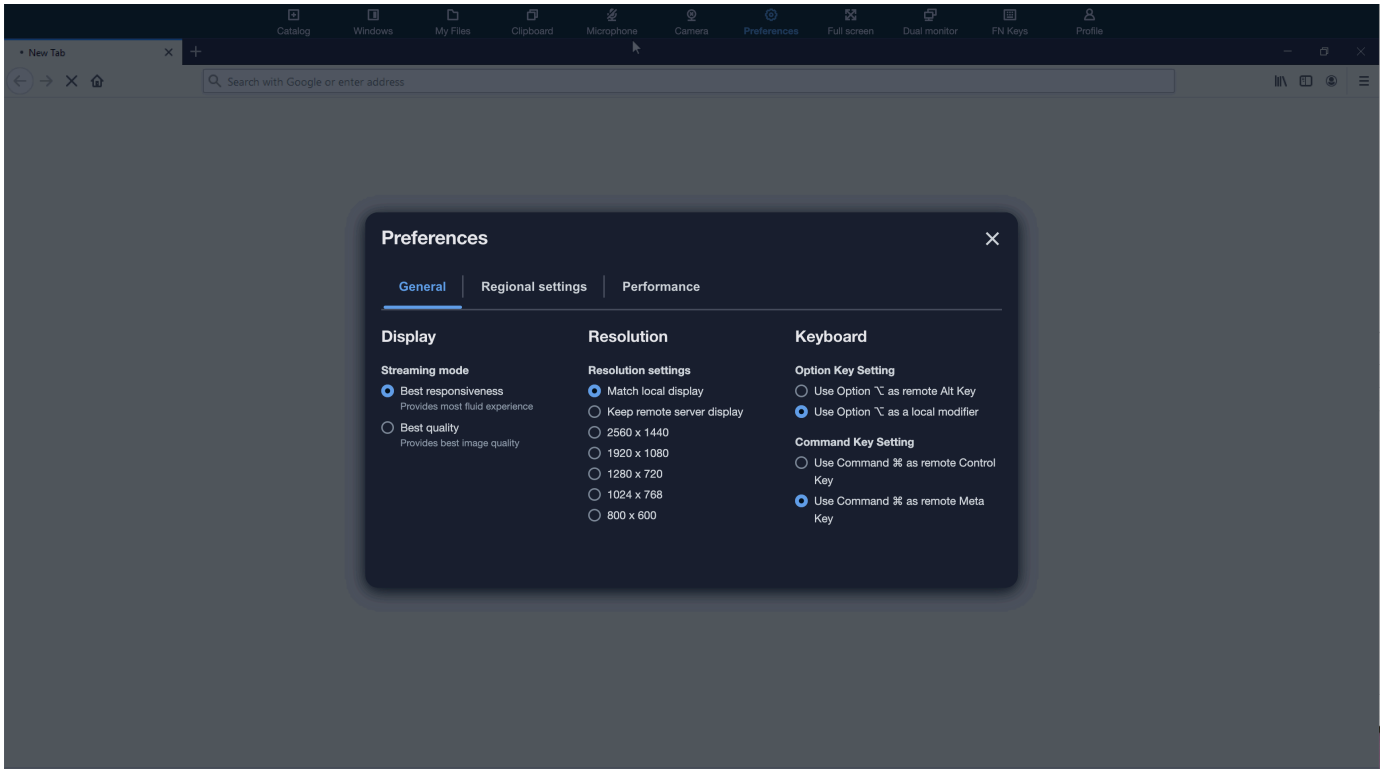
A *modifier key* modifies the action of another key when you use both keys together. You can use a modifier key with another key to perform a task such as printing. A *Meta key* is a special type of modifier key. You can use a Meta key to temporarily change the function of another key when you use both keys together.

You can remap this Mac key	To this key during a streaming session
Option key 	• Remote Alt key • Local modifier key
Command key 	• Remote Control key • Remote Meta key

Follow these steps to remap the Mac Option and Command keys during an WorkSpaces Applications streaming session.

To remap the Mac Option and Command keys

1. Use a web browser to connect to WorkSpaces Applications.
2. In the top left of the WorkSpaces Applications toolbar, choose the **Settings** icon, and choose **Keyboard Settings**.
3. Choose the options that correspond to the keys that you want to remap.



Follow these steps to remap the Mac Option and Command keys on WorkSpaces Applications web browser access v2.

To remap the Mac Option and Command keys on WorkSpaces Applications web browser access v2

1. Use a web browser to connect to WorkSpaces Applications.
2. From the top menu of the WorkSpaces Applications toolbar, choose the **Preferences** menu.
3. Choose **General**, **Keyboard**, and the options that correspond to the keys that you want to remap.

Video and Audio Conferencing

WorkSpaces Applications real-time audio-video (AV) redirects your local webcam video and microphone audio input to WorkSpaces Applications streaming sessions. That way, you can use your local devices for video and audio conferencing within your WorkSpaces Applications streaming session.

To use a local webcam and microphone within an WorkSpaces Applications streaming session

1. Connect to WorkSpaces Applications from a Chromium-based web browser, including Google Chrome and Microsoft Edge.

   > **Note**
   >
   > Most popular HTML5 compatible browsers support audio input in an WorkSpaces Applications session, including Chrome, Edge, and Firefox.

   > **Note**
   >
   > If your web browser doesn't support video or audio input, the options won't appear on the WorkSpaces Applications toolbar.

2. Configure your web browser's camera and microphone permissions to set default devices and allow access to WorkSpaces Applications.

   > **Note**
   >
   > For information about how to configure Google Chrome, see Use your camera & microphone.

3. In the top-left of the WorkSpaces Applications toolbar, choose the Settings icon, and then choose **Enable webcam**. For WorkSpaces Applications web browser access v2, choose the **Camera** option from the WorkSpaces Applications toolbar (which turns the option blue).

   > **Note**
   >
   > If the microphone or webcam icons don't appear in the **Settings** menu, contact your WorkSpaces Applications administrator. Your web browser might not support video or audio input, or your administrator might need to perform additional configuration tasks. For more information, see the section called "Real-Time Audio-Video".

4. Depending on your web browser settings, you might be prompted to allow the camera to be used by your web browser. Choose **Allow** to enable your camera.

5. In the top-left of the WorkSpaces Applications toolbar, choose the Settings icon, and then choose Enable microphone. For WorkSpaces Applications web browser access v2, choose the **Microphone** option from the WorkSpaces Applications toolbar (which turns the option blue).
6. Depending on your web browser settings, you might be prompted to the microphone to be used by your web browser. Choose **Allow** to enable your microphone.

> **Note**
>
> If you have more than one webcam or microphone and want to change the devices that you use for streaming within an WorkSpaces Applications session, you must clear your web browser settings for the WorkSpaces Applications website URL and configure default devices. Then, refresh your browser or start a new session for the changes to take effect, and repeat the above steps to enable your webcam and microphone.

Drawing Tablets

Drawing tablets, also known as pen tablets, are computer input devices that let you draw with a stylus (pen). With WorkSpaces Applications, you can connect a drawing tablet, such as a Wacom drawing tablet, to your local computer and use the tablet with your streaming applications.

Following are requirements and considerations for using drawing tablets with your streaming applications.

- To use this feature, you must connect to WorkSpaces Applications through the Google Chrome or Mozilla Firefox browsers only, or by using the WorkSpaces Applications client.
- The applications that you stream must support Windows Ink technology. For more information, see Pen interactions and Windows Ink in Windows apps.
- Depending on the streaming applications that you use, your drawing tablet might require USB redirection to function as expected. This is because some applications, such as GIMP, require USB redirection to support pressure sensitivity. If this is the case for your streaming applications, you must connect to WorkSpaces Applications by using the WorkSpaces Applications client, and share the drawing tablet with your streaming session. For information about how to share USB devices with your streaming session, see USB Devices.
- This feature is not supported on Chromebooks.

To get started with using a drawing tablet during your application streaming sessions, connect your drawing tablet to your local computer with USB, share the device with WorkSpaces Applications if required for pressure sensitivity detection, and then start an WorkSpaces Applications streaming session. You can use a supported web browser or the WorkSpaces Applications client, if it is installed, to start a streaming session.

Relative Mouse Offset

By default, during a streaming session, WorkSpaces Applications transmits information about mouse movements by using absolute coordinates and rendering the mouse movements locally. For graphics-intensive applications, such as computer-aided design (CAD)/computer-aided manufacturing (CAM) software or video games, mouse performance improves when relative mouse mode is enabled. Relative mouse mode uses relative coordinates, which represent how far the mouse moved since the last frame, rather than the absolute x-y coordinate values within a window or screen. When you enable relative mouse mode, WorkSpaces Applications renders the mouse movements remotely.

You can enable this feature during an WorkSpaces Applications streaming session by doing either of the following:

- Windows: Pressing Ctrl+Shift+F8
- Mac: Pressing Control+Fn+Shift+F8

Troubleshooting

If issues occur when you use WorkSpaces Applications, your WorkSpaces Applications session ID can help your administrator with troubleshooting. This section describes how to find the session ID.

The session ID is created when you request a streaming session. The session ID, and other information used by WorkSpaces Applications, is stored in the session storage location for your browser. You can use the developer tools that are available for your browser interface to find this location.

For information about the developer tools that are available for common web browsers, see the following resources:

- Apple Safari Developer Help: Storage tab
- View And Edit Session Storage With Chrome DevTools

- Firefox Developer Tools: Local Storage / Session Storage
- Microsoft Edge (Chromium) Developer Tools
- Microsoft Edge (EdgeHTML) Developer Tools

After you locate the developer tools for your browser, search for the session storage for the WorkSpaces Applications website. The domain for the website is **https://appstream2.<*aws-region*>.aws.amazon.com**. Expand the domain, and choose **sessionStorage.as2SessionData**. The session ID is stored in the key **sessionId**.

WorkSpaces Applications Client Application for Windows

The following information helps you use the WorkSpaces Applications client for Windows to connect to WorkSpaces Applications and stream applications.

Contents

- Features
- Requirements
- Setup for Windows
- Connect to WorkSpaces Applications on Windows Client
- Monitors
- USB Devices
- Local File Access
- Printer Redirection
- Video and Audio Conferencing
- Drawing Tablets
- Relative Mouse Offset
- Logging
- Troubleshooting
- WorkSpaces Applications Windows Client Release Notes

Features

The WorkSpaces Applications client for Windows is an application that you install on your Windows PC. This application provides additional capabilities that are not available when you access

WorkSpaces Applications by using a web browser. For example, the WorkSpaces Applications client lets you do the following:

- Use more than two monitors or 4K resolution.
- Use your USB devices with applications streamed through WorkSpaces Applications.
- Access your local drives and folders during your streaming sessions.
- Redirect print jobs from your streaming application to a printer that is connected to your local computer.
- Use your local webcam for video and audio conferencing within your streaming sessions.
- Use keyboard shortcuts during your streaming sessions.
- Interact with your remote streaming applications in much the same way as you interact with locally installed applications.

Requirements

The WorkSpaces Applications client for Windows must be installed on a computer that meets the following requirements:

- Operating system — Windows 10 (32-bit or 64-bit), Windows 11 (64-bit)
- Microsoft Visual C++ 2019 Redistributable or later for WorkSpaces Applications client version 1.1.1066 and above. For information about the latest Visual C++ redistributable packages for Visual Studio 2015, 2017, and 2019, see The latest supported Visual C++ downloads in the Microsoft Support documentation.
- RAM — 2 GB minimum
- Hard drive space — 200 MB minimum

In addition, to install the WorkSpaces Applications USB driver for USB driver support, you must have local administrator rights on your PC.

Setup for Windows

> **Note**
>
> The Amazon AppStream 2.0 client has been renamed to the WorkSpaces Applications client. It is the same application. To stream your applications, download and use the WorkSpaces Applications client.

Follow these steps to install the client.

1. On the PC where you want to install the WorkSpaces Applications client, download the WorkSpaces Applications client for Windows application from WorkSpaces Applications supported clients.
2. Navigate to the location where you downloaded the application .exe file, and then double-click the file to begin the installation.

   > **Important**
   >
   > Contact your network administrator if nothing happens when you double-click the file or if an error message is displayed. Your organization might be using antivirus software that prevents the WorkSpaces Applications client installation program from running.

3. The installation wizard displays links to the AWS Customer Agreement, AWS Service Terms, and the AWS Privacy Notice, and third-party notices. Review this information, and then choose **Next**.
4. On the **Client Diagnostics** page, to enable the WorkSpaces Applications client to automatically upload device logs to help with troubleshooting issues, keep **Client logging** selected, and then choose **Next**.
5. On the **Optional Components** page, to enable your USB devices to be used with streaming applications, select the **WorkSpaces Applications Client USB Driver** check box, and then choose **Finish**.
6. If the **WorkSpaces Applications USB driver** wizard setup wizard opens, choose **Install**.
7. If prompted by **User Account Control** to choose whether to allow the app to make changes to your device, choose **Yes**.
8. When a message notifies you that the USB driver installation is complete, choose **Close**.

The WorkSpaces Applications sign-in page opens. For information about how to connect to WorkSpaces Applications and start an application streaming session, see Connect to WorkSpaces Applications.

Connect to WorkSpaces Applications on Windows Client

After the WorkSpaces Applications client for Windows is installed on your PC, you can use it to connect to WorkSpaces Applications.

Topics

- WorkSpaces Applications Client Connection Modes
- Connect to WorkSpaces Applications
- How to Switch WorkSpaces Applications Connection Modes

WorkSpaces Applications Client Connection Modes

The WorkSpaces Applications client provides two connection modes: *Native application mode* and *classic mode*. The connection mode that you choose determines the options that are available to you during application streaming, and how your streaming applications function and display. In addition, **Desktop view** is also available, if your administrator has enabled it.

Native application mode

Native application mode lets you work with remote streaming applications in much the same way that you work with applications that are installed on your local PC.

When you connect to WorkSpaces Applications in native application mode, the WorkSpaces Applications Application Launcher window opens and displays the list of applications that are available for you to stream. When you open a streaming application in this mode, the WorkSpaces Applications Application Launcher window remains open, and the application opens in its own window. During your streaming session, the remote streaming application functions in much the same way as a locally installed application. The application icon is displayed in the taskbar of your local PC, just as the icons do for your local applications. Unlike the icons for your local applications, the icons for your streaming applications in native application mode include the WorkSpaces Applications logo.

During your WorkSpaces Applications streaming session, you can switch quickly between your locally installed applications and your remote streaming applications by clicking the taskbar

icon of the remote or local application you want to work with. You can also switch WorkSpaces Applications connection modes. If you want to work in classic mode instead, you can switch from native application mode to classic mode.

Classic mode

When you use classic application mode, you work with remote streaming applications in the WorkSpaces Applications session window. If your administrator has made more than one application available to you, you can open multiple applications during your session. All applications that you open are displayed in the same WorkSpaces Applications session window.

When you connect to WorkSpaces Applications in classic mode, the WorkSpaces Applications Application Launcher window opens and displays the list of applications that are available for you to stream. When you open a streaming application in this mode, the Application Launcher window closes, and the application opens in the WorkSpaces Applications session window.

If your administrator has not disabled native application mode, you can switch from classic mode to native application mode. For more information, see How to Switch WorkSpaces Applications Connection Modes.

Desktop view

When you connect to WorkSpaces Applications and choose **Desktop view**, WorkSpaces Applications provides a standard Windows desktop view for your streaming session. The icons of applications that are available for you to stream appear on the Windows desktop. In addition, the WorkSpaces Applications toolbar, which enables you to configure settings for your streaming session, appears in the top left area of your streaming session window.

Connect to WorkSpaces Applications

Follow these steps to connect to WorkSpaces Applications and start an application streaming session.

1. If your administrator requires you to sign in first through your organization's sign-in page, complete the tasks in this step, then proceed to step 3.

   If your administrator doesn't require you to sign in through your organization's sign-in page, skip the tasks in this step and proceed to step 2.

   a. Navigate to your organizational sign-in page and enter your domain credentials when prompted.

b. After you sign in, you are redirected to the WorkSpaces Applications portal, which displays one or more applications that are available for your WorkSpaces Applications streaming session. **Desktop View** is also available, if enabled by your administrator.

c. Choose an application or, if available, **Desktop View**.

2. If your administrator doesn't require you to sign in first through your organization's sign-in page, complete the following steps to start the WorkSpaces Applications client:

a. On your local computer where the WorkSpaces Applications client is installed, in the lower left of your screen, choose the Windows search icon on the taskbar and enter **`AppStream`** in the Search box.

b. In the search results, select **Amazon AppStream** to start the WorkSpaces Applications client.

c. On the WorkSpaces Applications client sign-in page, you can choose whether to use the client in native application mode or classic mode.

- To use native application mode, keep the **Start in native application mode** check box selected.
- To use classic mode, clear the **Start in native application mode** check box.

d. Do either of the following:

- If the client sign-in page is prepopulated with a web address (URL), choose **Connect**.
- If the client sign-in page is not prepopulated with a URL, enter the URL that your WorkSpaces Applications administrator provided for WorkSpaces Applications, and then choose **Connect**. If you don't know the URL, contact your administrator.

e. After a few moments, the WorkSpaces Applications portal opens, displaying one or more applications that are available for your WorkSpaces Applications streaming session. **Desktop View** is also available, if enabled by your administrator.

f. Choose an application or, if available, **Desktop View**.

3. Depending on the authentication settings that your WorkSpaces Applications administrator enabled, after you choose an application or **Desktop View**, you might be prompted to enter your Active Directory domain credentials to sign in to your WorkSpaces Applications session. If this is the case, do one of the following:

- If your organization has enabled password authentication, enter your Active Directory domain password, and then choose **Password sign in**.

- If your organization has enabled smart card authentication, select **Choose a smart card**, follow the instructions to choose your smart card certificate and enter your smart card PIN, and then choose **Smart card sign in**.
- If your organization has enabled both authentication methods, either enter your Active Directory domain password and choose **Password sign in**, or select **Choose a smart card**, and follow the instructions to complete the smart card sign-in.

How to Switch WorkSpaces Applications Connection Modes

If your administrator has not disabled native application mode for your streaming sessions, you can switch between native application mode and classic mode.

To switch from native application mode to classic mode

1. In the upper left of the WorkSpaces Applications Application Launcher window, choose the **Settings** icon, and then choose **Switch to classic mode**.
2. When you switch to classic mode, the Application Launcher window closes and the WorkSpaces Applications session window opens. Any application that you are streaming in native application mode opens in the WorkSpaces Applications session window.

Follow these steps to switch from classic mode to native application mode.

To switch from classic mode to native application mode

1. In the upper left of the WorkSpaces Applications session window, choose the **Settings** icon, and then choose **Switch to native application mode**.
2. When you switch from classic mode back to native application mode, the WorkSpaces Applications session window closes and the WorkSpaces Applications Application Launcher window opens. Any application that you are streaming in classic mode opens in a separate window.

Monitors

Monitors and Display Resolution

WorkSpaces Applications supports the use of multiple monitors during streaming sessions, including monitors that have different resolutions. To help ensure an optimal streaming experience,

we recommend that you set the display scale for your monitors to 100 percent if you use multiple monitors.

The WorkSpaces Applications client supports multiple monitors with the following display resolutions:

- Multiple monitors (up to 4K resolution) — Up to 4 monitors with a maximum display resolution of 4096x2160 pixels per monitor

> **Note**
>
> If you are connected to an WorkSpaces Applications streaming session using native application mode, you can use monitors with up to 4K resolution. If you use higher-resolution monitors, the WorkSpaces Applications client falls back to classic mode. In this case, the WorkSpaces Applications classic mode streaming view occupies 4K of the screen, and the remaining portion of the screen is black.

> **Important**
>
> Higher resolution monitors require significantly more compute capacity and encoding resources to stream content effectively, with actual performance depending on your display configuration (resolution and number of monitors) and the compute instance being used. If you experience suboptimal performance on 4K monitors, we recommend reducing your display resolution to improve responsiveness.
> Ultrawide monitors with resolutions exceeding 4096 pixels in either dimension (e.g., 5120x2160) will display black bars on the sides, as the maximum supported resolution is limited to 4096 pixels per dimension.

Using Multiple Monitors

When using multiple monitors, you can choose from the following options:

- Extend full-screen across a *single* monitor
- Extend full-screen across *all* monitors
- Extend full-screen across *selected* monitors

Extending full-screen across a single monitor

You can extend full screen only on the current monitor if multiple monitors are connected to your local computer. To enable this feature, complete the following steps:

1. On the toolbar at the top of the window, choose the Full Screen (crossed arrows) icon.
2. From the drop-down menu, choose **Across a single monitor**.

Extending full-screen across all monitors

You can extend the display for a session across all monitors at full screen resolution. The extended display matches your physical display layout and screen resolutions. For example, three monitors are connected to your local computer. The server extends the display for a session across all three monitors and matches the specific screen resolutions of your display.

To enable this feature, complete the following steps:

1. On the toolbar at the top of the window, choose the Full Screen (crossed arrows) icon.
2. From the drop-down menu, choose **Across all monitors**.

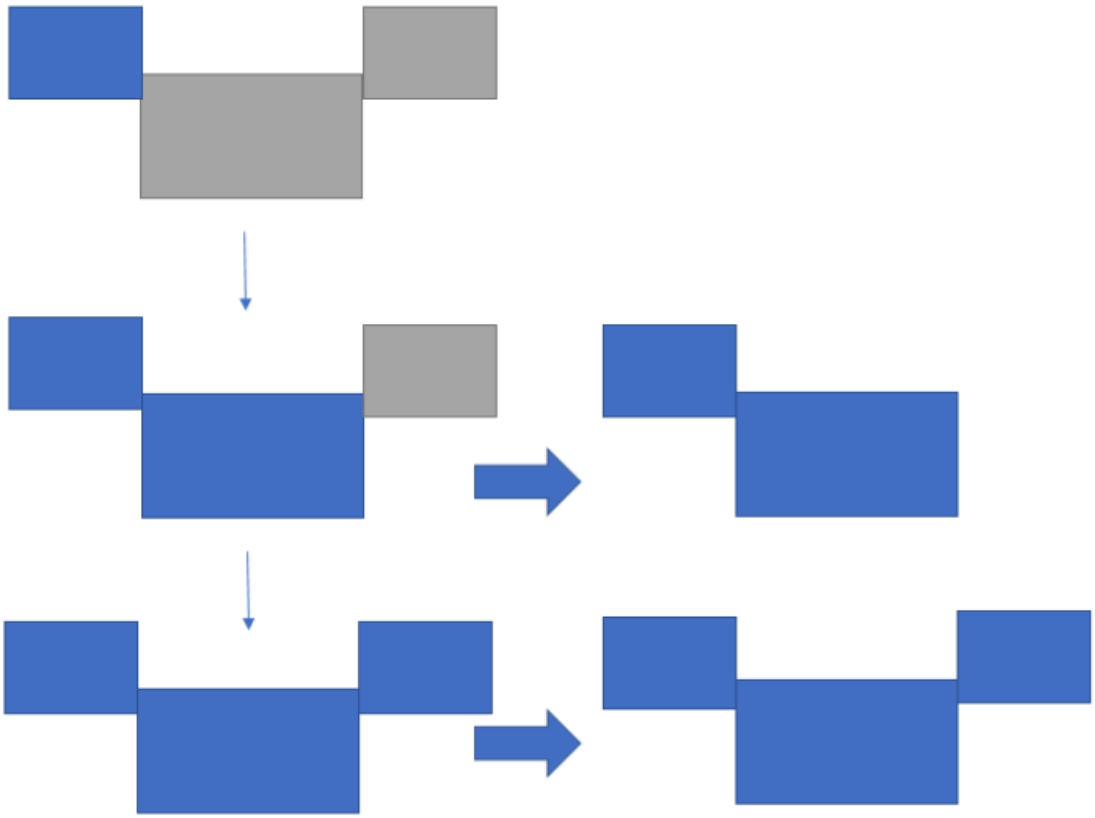
Extending full-screen across selected monitors

If there are three or more monitors connected, WorkSpaces Applications can also extend full-screen across a selection of those available monitors. If your selected monitors cannot go full screen, an error message will appear and you will need to perform the procedure again. Selected monitors must be set adjacent, or sharing a side with each other, in your display setting.

The following are examples of adjacent monitor placement. If your monitors are not set adjacent in your Windows display configuration, you must exit WorkSpaces Applications and change your Display settings on your local machine.

> **Note**
>
> The blue boxes are WorkSpaces Applications-enabled monitors, and the gray boxes are other monitors.

Examples of adjacent monitor placement



Examples of nonadjacent monitor placement



To enable this feature, complete the following steps:

1. On the toolbar at the top of the window, choose the Full Screen (crossed arrows) icon.
2. From the drop-down menu, choose **Across selected monitors**.
3. The **Across selected monitors** window appears, displaying your current monitor layout. Select which monitors you want DCV to be displayed full screen, and choose **Apply**.

USB Devices

With certain exceptions, USB redirection is required for the WorkSpaces Applications client to support USB devices. When USB redirection is required for a device, you must share the device with WorkSpaces Applications every time you start a new streaming session.

Topics

- How to Use a Smart Card During a Streaming Session
- How to Share a USB Device with WorkSpaces Applications

How to Use a Smart Card During a Streaming Session

Depending on the authentication settings that your administrator has enabled, you might need to use a smart card for authentication during an WorkSpaces Applications streaming session. For example, if you open a browser during your streaming session and navigate to an internal organizational site that requires smart card authentication, you must enter your smart card credentials.

By default, smart card redirection is enabled for WorkSpaces Applications streaming sessions, which means that you can use the smart card reader that is attached to your local computer without sharing it with WorkSpaces Applications. During your streaming session, your smart card reader and smart card are available for you to use with local applications, as well as with streaming applications.

If your administrator has disabled smart card redirection, you must share your smart card reader with WorkSpaces Applications. For more information, see the next section.

How to Share a USB Device with WorkSpaces Applications

If you are using a drawing tablet, USB redirection might not be required to use it with WorkSpaces Applications. However, if you are streaming an application such as the Gnu Image Manipulation Program (GIMP), which requires USB redirection to support pressure sensitivity, you must share your drawing tablet with WorkSpaces Applications. For information about drawing tablets, see Drawing Tablets.

To share a USB device with WorkSpaces Applications

1. Use the WorkSpaces Applications client to start a streaming session.

2. In the top left area, choose the **Settings** icon, and then choose **USB Devices**.
3. If your USB device is connected to your computer, the USB device name appears in the dialog box. If your USB device is not detected, contact your WorkSpaces Applications administrator for assistance.
4. Switch the **Share** toggle key next to the name of the USB device that you want to share with the streaming session.

   Your USB device is now available for use with your streaming applications.

   > **Important**
   >
   > USB devices can't be simultaneously used between local and remote applications. So after you share a USB device with a streaming session, you can't use it with applications on your local computer. To use your USB device on your local computer, switch the **Share** toggle key next to the name of the USB device that you want to use locally. This disables sharing with the streaming session.

5. You can also enable your USB device to automatically connect when a new streaming session starts. To do so, select the option next to the toggle key for the USB device that you want to connect. After you enable this option, when your next streaming session starts, the USB device is connected automatically.

Local File Access

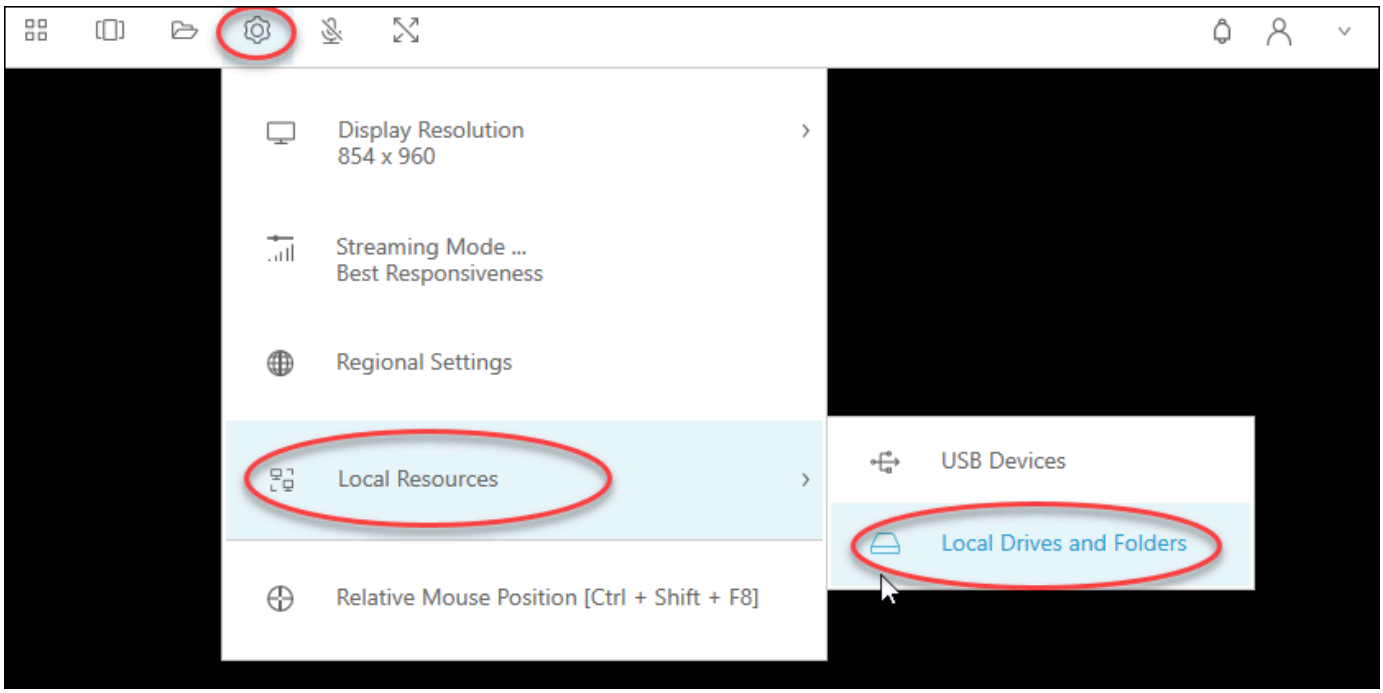
WorkSpaces Applications file redirection lets you access files on your local computer from your WorkSpaces Applications streaming session. To use file redirection, open the WorkSpaces Applications client, connect to a streaming session, and choose the drives and folders that you want to share. After you share a local drive or folder, you can access all files in the shared drive or folder from your streaming session. You can stop sharing local drives and folders at any time.

> **Important**
>
> To use WorkSpaces Applications file redirection, you must have the WorkSpaces Applications client installed on your local computer. File redirection is not available when you connect to WorkSpaces Applications by using a web browser.

> **Note**
>
> If your image uses an WorkSpaces Applications agent released on or after September 25, 2025, the drive for file system redirection will be available as a removable disk instead of a network drive.

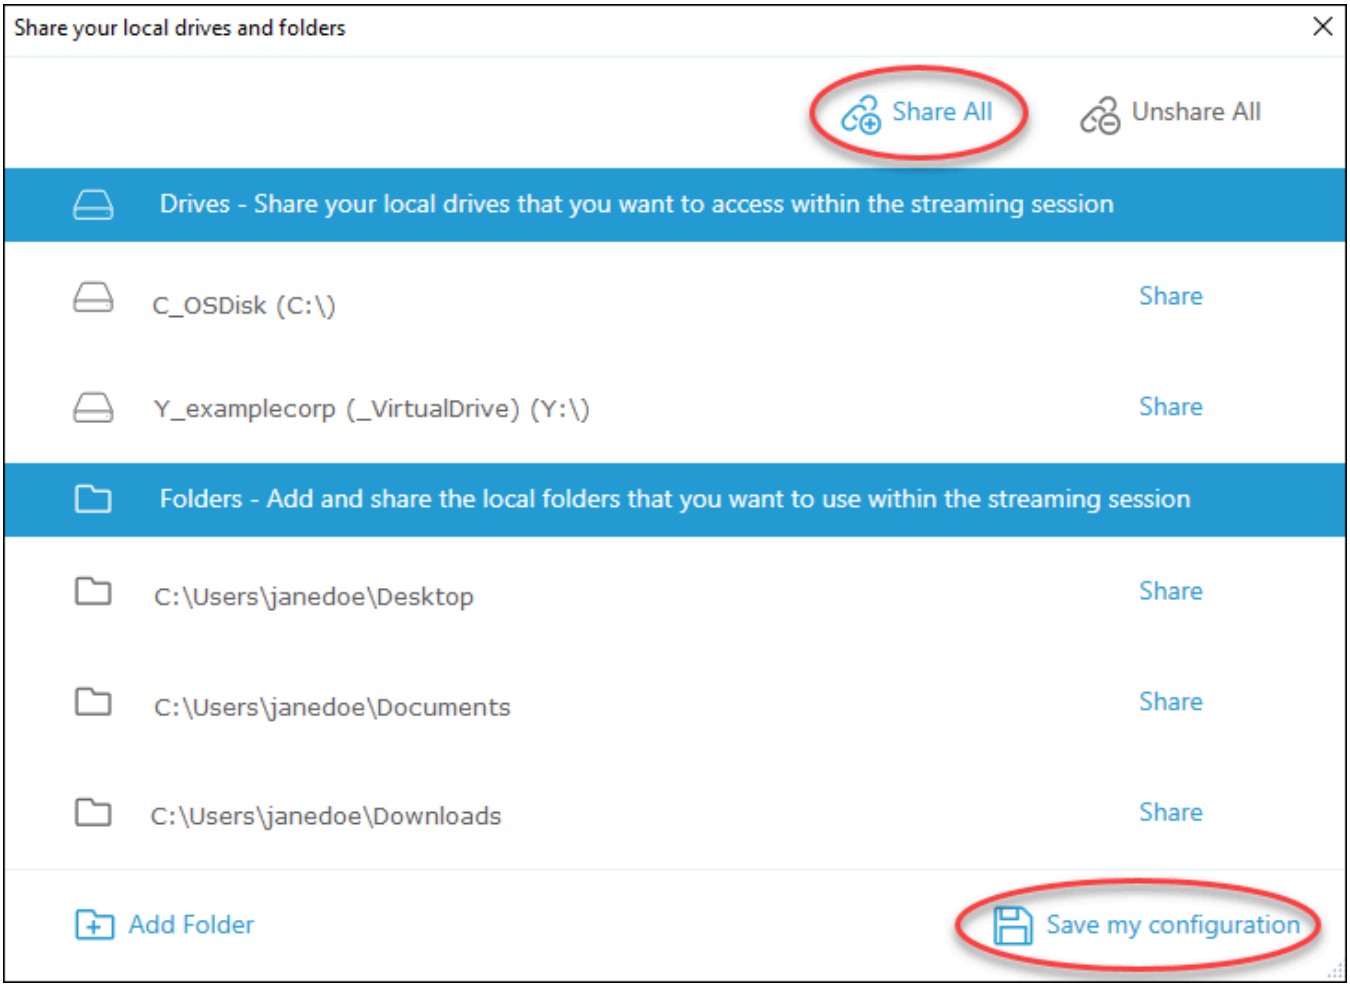
To share local drives and folders

1. Open the WorkSpaces Applications client and connect to a streaming session.
2. In your WorkSpaces Applications session, in the top left area, choose the **Settings** icon, and then choose **Local Resources**, **Local Drives and Folders**.

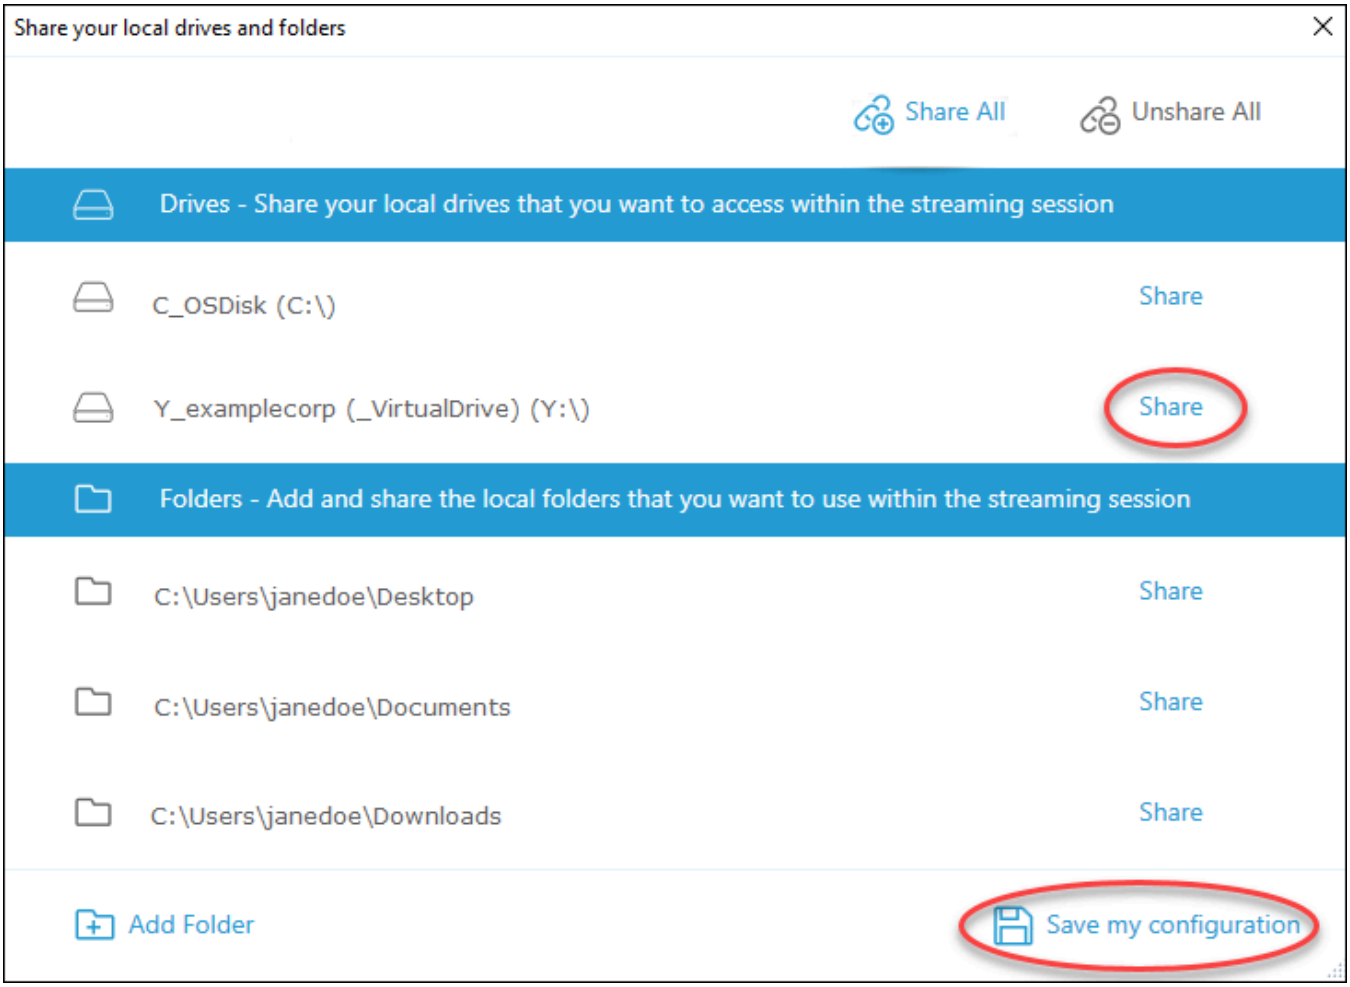


The **Share your local drives and folders** dialog box displays the drives and folders that your administrator has made available for you to share. You can share all or specific drives and folders, or just one. You can also add your own drives and folders. To share drives and folders, do one of the following:

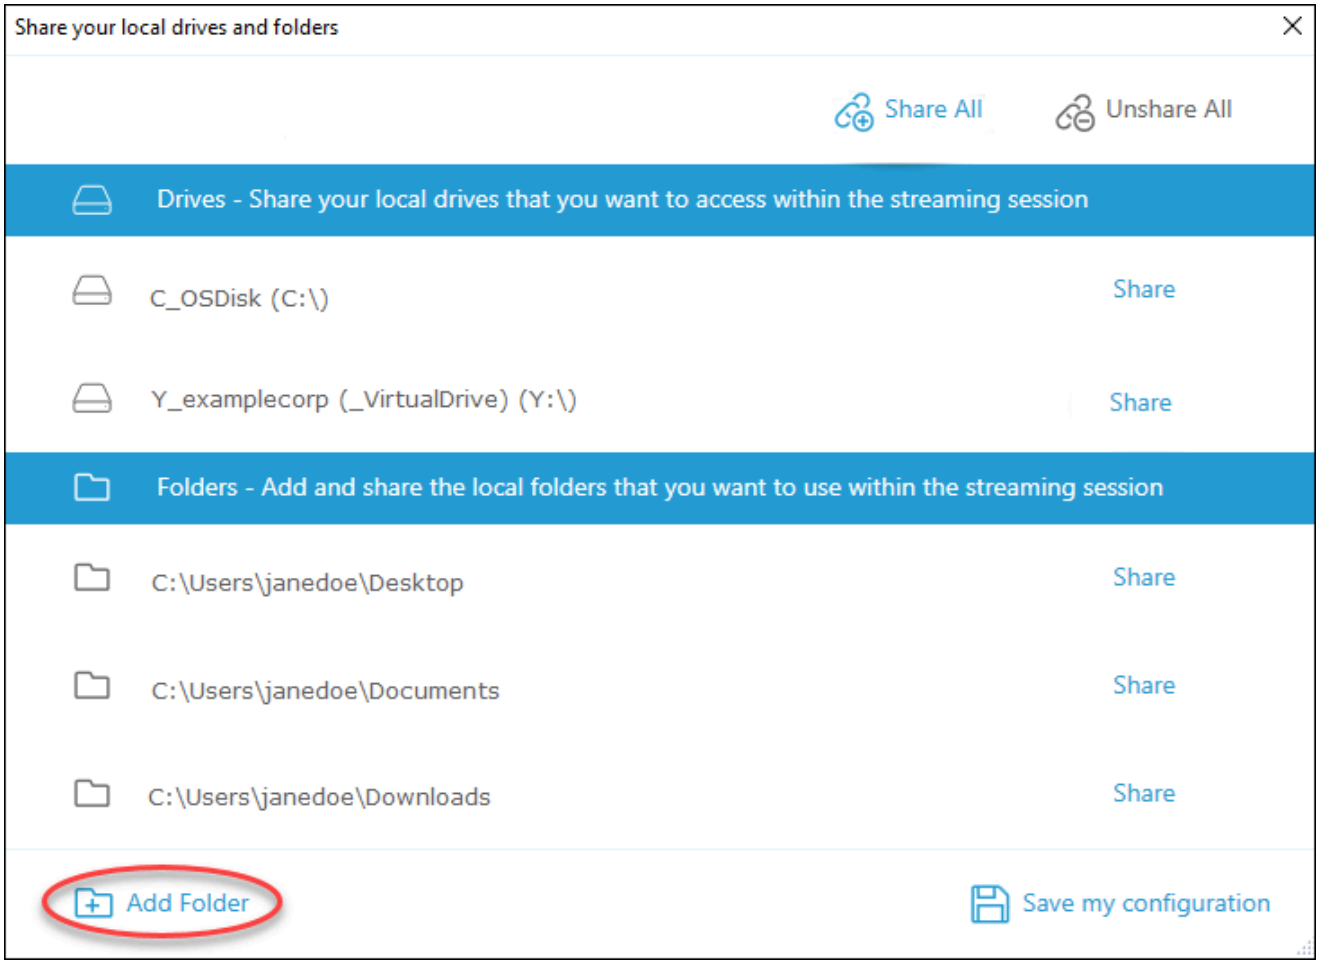
- To share all local drives and folders displayed in the **Share your local drives and folders** dialog box, choose **Share All**. To apply your changes to future streaming sessions, choose **Save my configuration**.



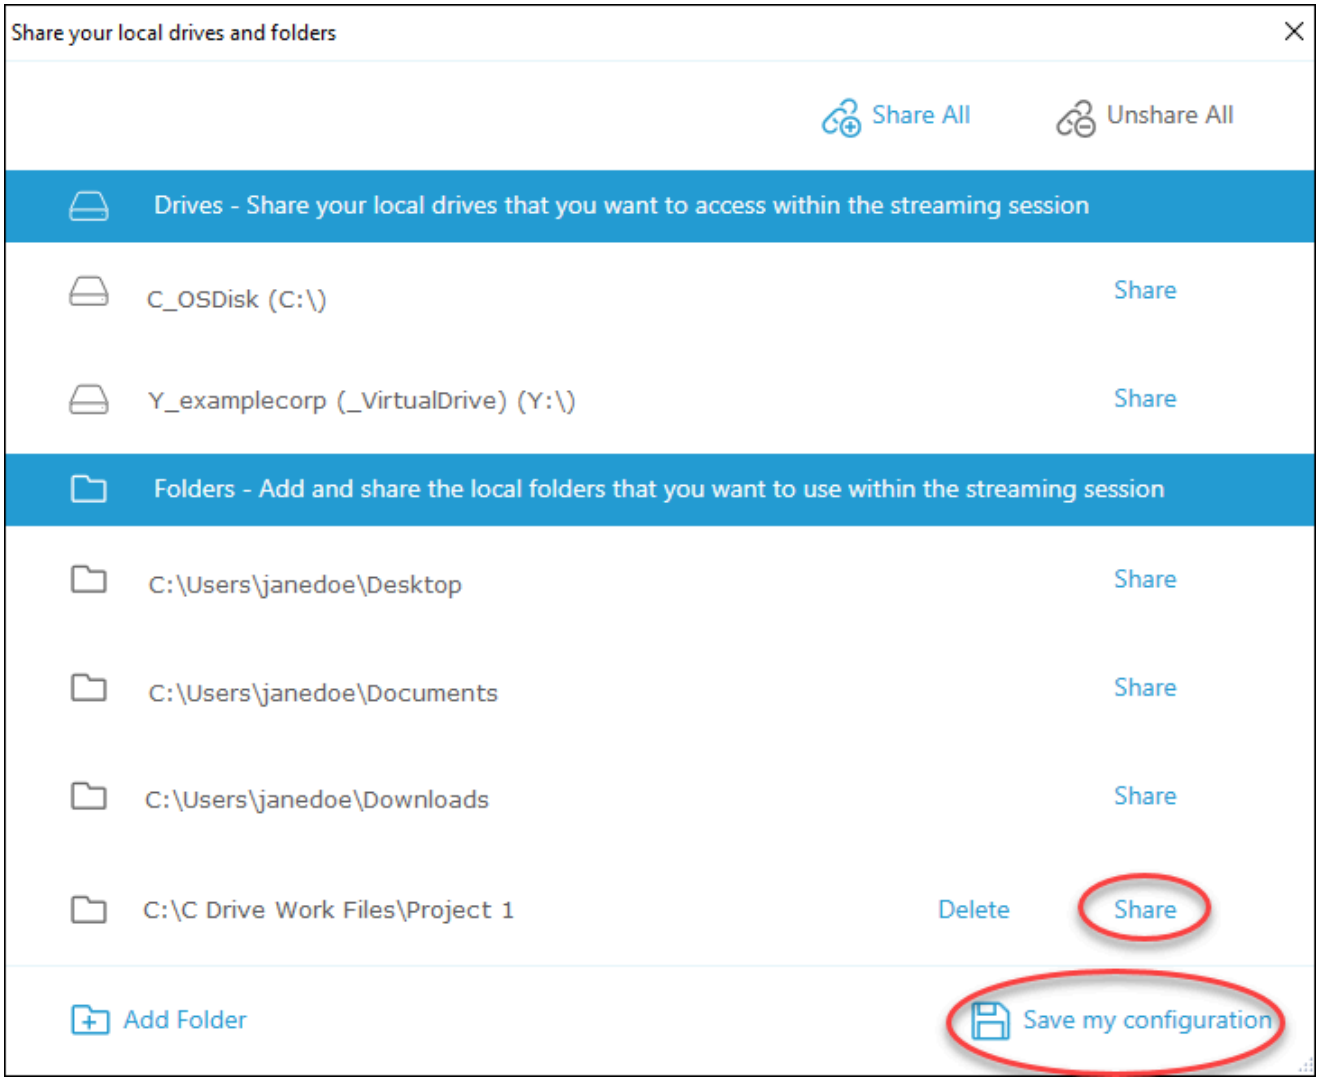
- To share a specific local drive or folder, select the drive or folder that you want to access, and choose **Share**, **Save my configuration**. To share another local drive or folder, repeat these steps as needed.



- If the local drive or folder that you want to share is not displayed, you can add it. For example, your administrator might make your entire local C Drive available for you to share. However, you might only need to access a specific folder on that drive. In this case, you can add the folder that you need and share only that folder. To choose a folder, do the following:
  - In the **Share your local drives and folders** dialog box, choose **Add Folder**.



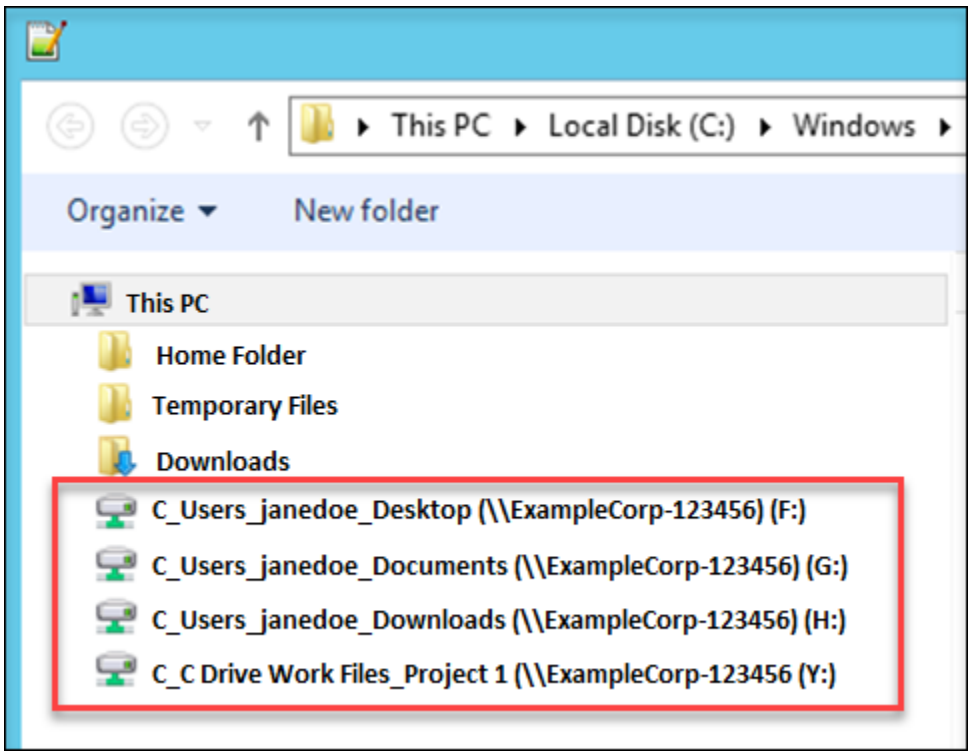
- Browse to the folder that you want to share, and choose **OK**.
- The folder that you selected is now available to share. Select the folder, and choose **Share**, **Save my configuration**. To add another local drive or folder, repeat these steps as needed.



After you share a local drive or folder, perform the following steps to access files in the shared drive or folder from your streaming session.

To access files in a shared local drive or folder

1. Open the WorkSpaces Applications client and connect to a streaming session.
2. In your WorkSpaces Applications session, open the application that you want to use.
3. From your application interface, choose **File Open**, and browse to the file that you want to access. The following screenshot shows how shared local drives and folders appear in the Notepad++ browse dialog box for Jane Doe when she browses for a file.



In the browse dialog box, the corresponding paths for her shared drives and folders are shown in the red box. The paths appear with backslashes replaced by underscores. At the end of each path is the name of Jane's computer, ExampleCorp-123456, and a drive letter.

4. When you're done working with the file, use the **File Save** or **File Save As** command to save it to the location that you want.

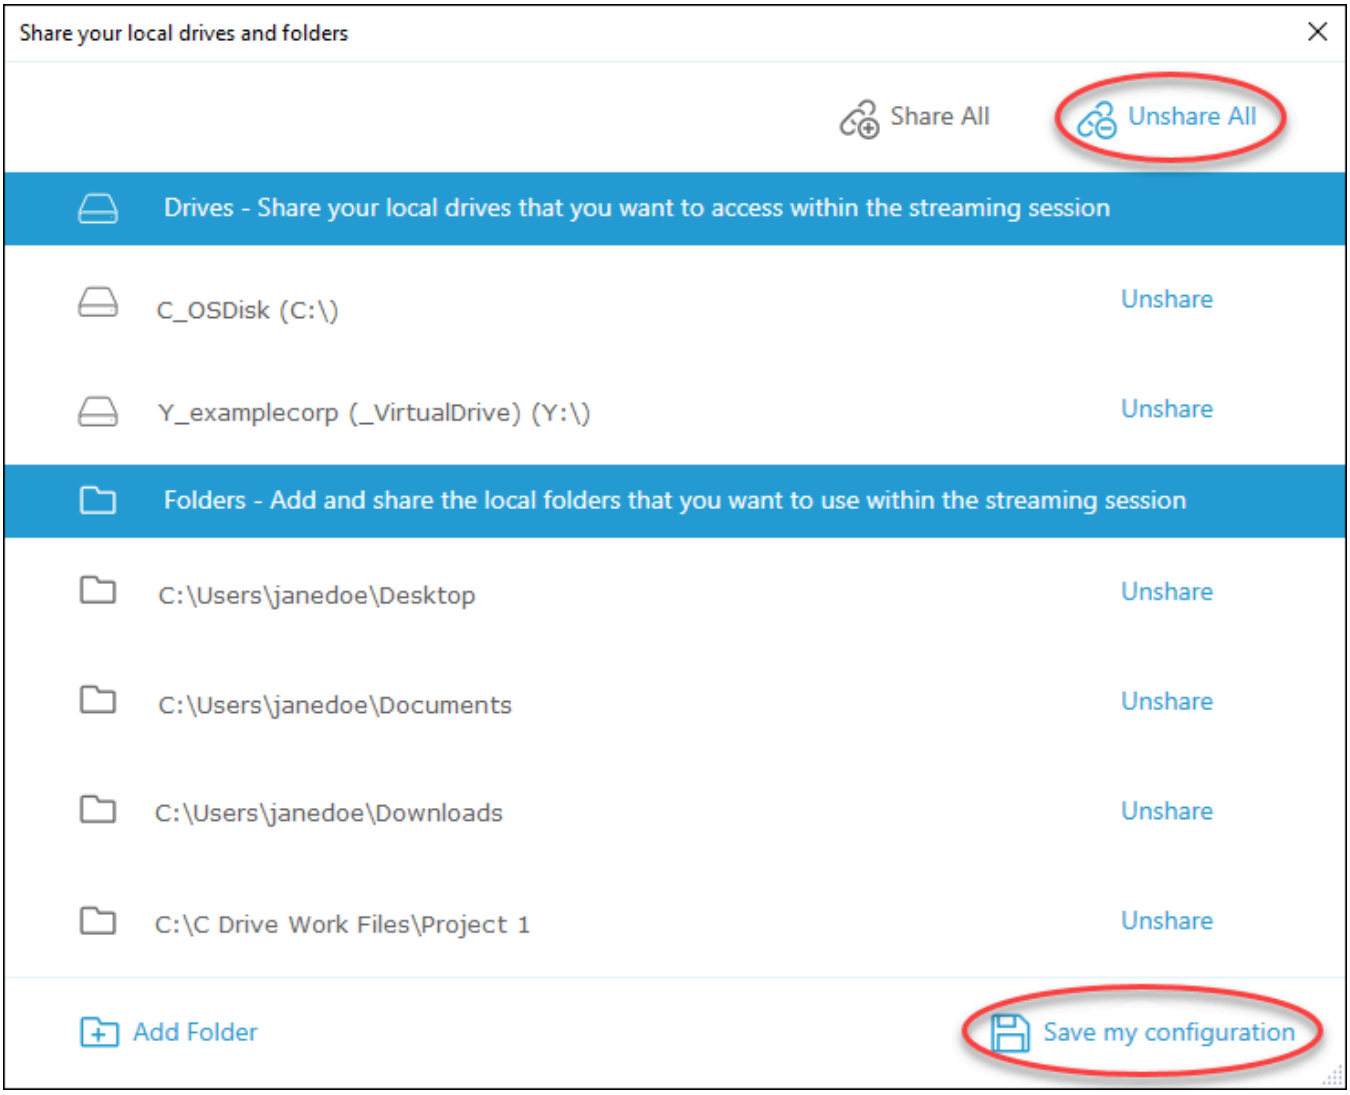
If you want to stop sharing a local drive or folder, perform the following steps.

To stop sharing local drives and folders

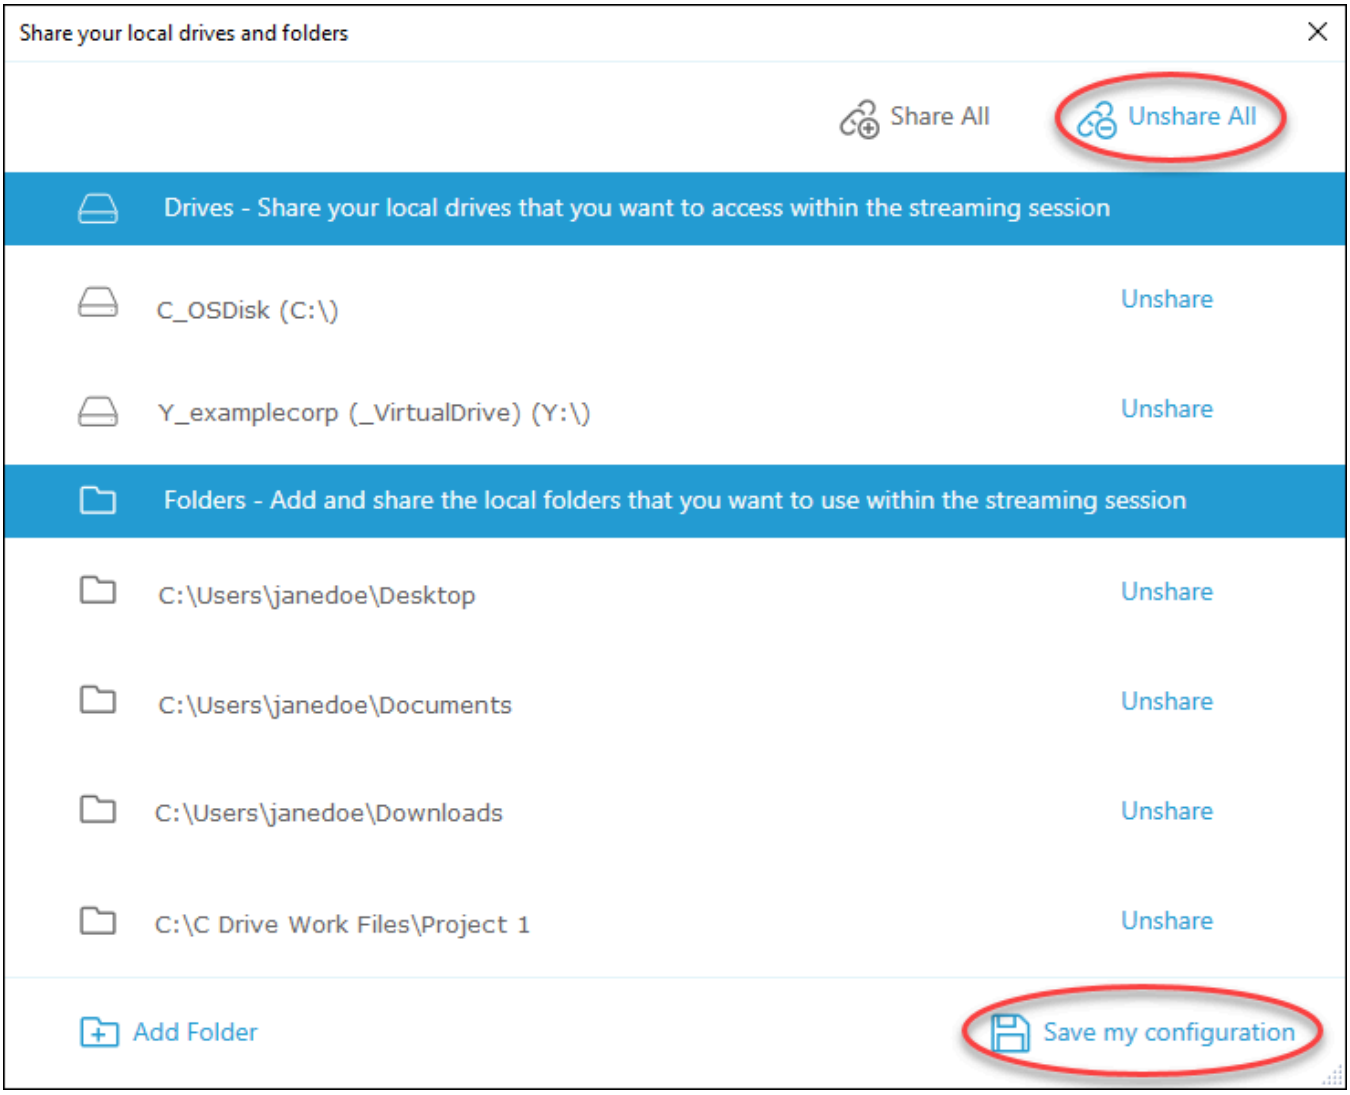
1. Open the WorkSpaces Applications client and connect to a streaming session.
2. In your WorkSpaces Applications session, in the top left area, choose the **Settings** icon, and then choose **Local Resources**, **Local Drives and Folders**.

   The **Share your local drives and folders** dialog box displays the drives and folders that your administrator has made available for you to share, and any that you added, if applicable. To stop sharing one or more local drives and folders, do either of the following:

   - To stop sharing all shared local drives and folders, choose **Unshare All**, **Save my configuration**.



- To stop sharing a specific shared local drive or folder, select the drive or folder, and choose **Unshare**, **Save my configuration**. To stop sharing another local drive or folder, repeat these steps as needed.



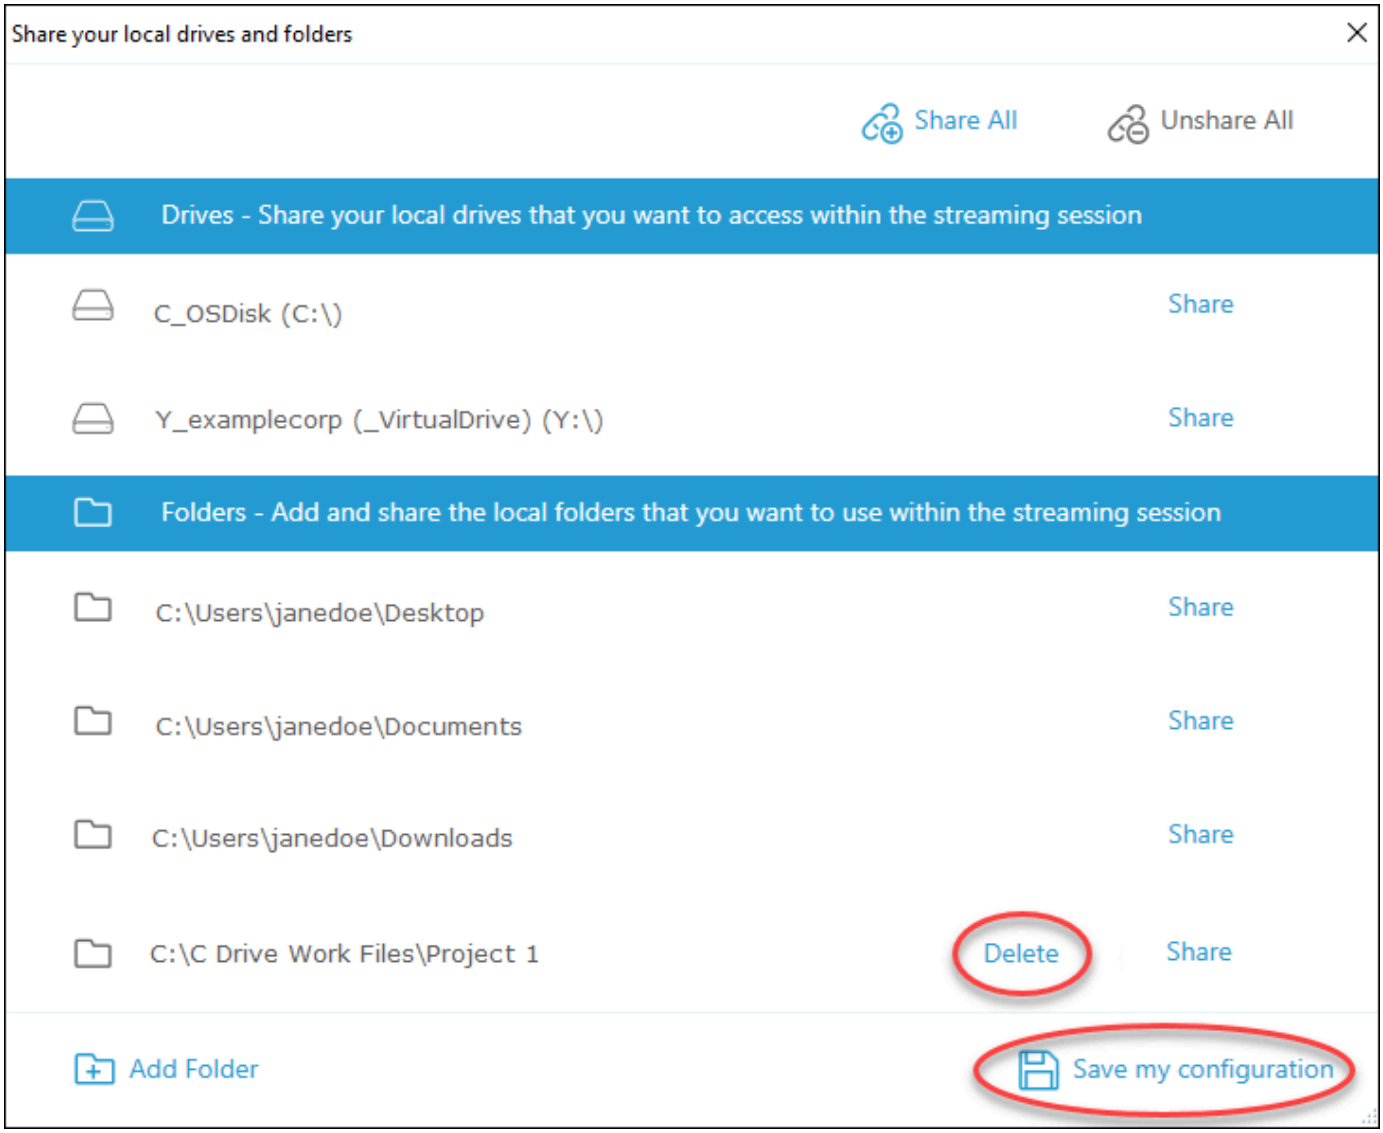
You can delete local drives and folders that you add to the **Share your local drives and folders** dialog box. However, you can't delete local drives or folders that your administrator has made available for you to share. Also, if you have already shared a local drive or folder, you must stop sharing it before you can delete it.

To delete local drives and folders

1. Open the WorkSpaces Applications client and connect to a streaming session.
2. In your WorkSpaces Applications session, in the top left area, choose the **Settings** icon, and then choose **Local Resources**, **Local Drives and Folders**.

The **Share your local drives and folders** dialog box displays the drives and folders that your administrator has made available for you to share. If you added any drives or folders, they are also displayed.

3. Select the local drive or folder that you want to delete, and then choose **Delete**, **Save my configuration**.



Printer Redirection

WorkSpaces Applications local printer redirection lets you access printers that are connected to your local computer from your WorkSpaces Applications streaming session. That way, you can redirect print jobs from your streaming application to a local printer, or to a network printer that you have mapped.

> **Important**
>
> To use WorkSpaces Applications printer redirection, you must have the WorkSpaces Applications client installed on your local computer, and you must use the client to connect to a streaming session. Printer redirection is not available when you connect to WorkSpaces Applications by using a web browser.

To redirect a print job to a local printer

1. Open the WorkSpaces Applications client and connect to a streaming session.
2. In your streaming application, choose **File**, **Print Now**.
3. In the top-right area of the WorkSpaces Applications session window, select the new notification that appears next to the notification icon.
4. In the **Notifications** dialog box, choose the **Print Jobs** tab.
5. On the **Print Jobs** tab, choose **Print**.
6. The **Print** dialog box for your streaming application opens.
7. In the **Print** dialog box, a list of available local printers is displayed. Choose the local printer that you want to use, and then proceed with printing.

Video and Audio Conferencing

WorkSpaces Applications real-time audio-video (AV) redirects your local webcam video input to WorkSpaces Applications streaming sessions. That way, you can use your local devices for video and audio conferencing within your WorkSpaces Applications streaming session.

To use a local webcam and microphone within an WorkSpaces Applications streaming session

1. Open the WorkSpaces Applications client and connect to a streaming session.
2. In the WorkSpaces Applications toolbar in the top left of your session window, do either of the following:

   - If the video icon has a diagonal line through it (as shown in the following screenshot), this indicates that the WorkSpaces Applications real-time AV feature is available for use but no webcams are attached to your streaming session. Choose the video icon to attach one or more webcams.

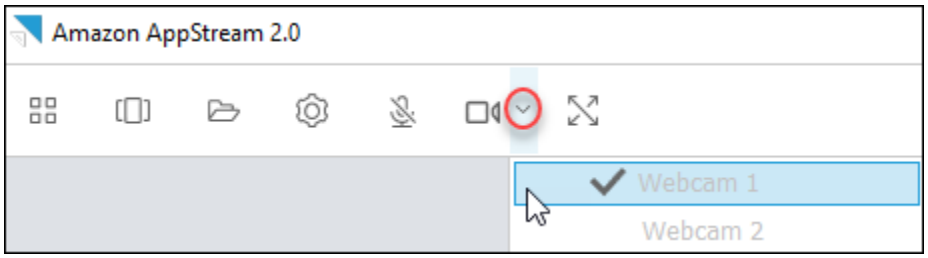


- If the video icon does not have a diagonal line through it (as shown in the following screenshot), one or more webcams are already attached to your streaming session. Skip this step and proceed to the next step.

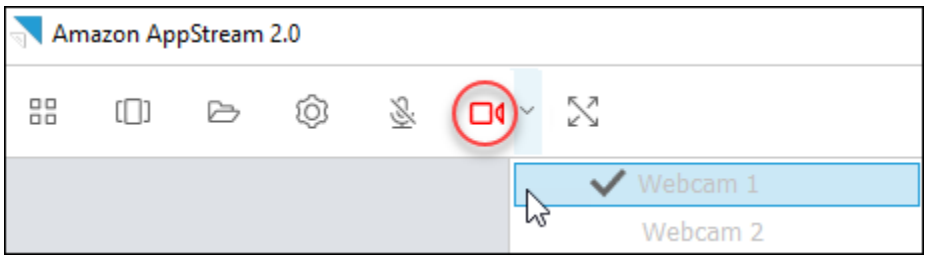


> **Note**
>
> If the video icon doesn't display in the WorkSpaces Applications toolbar, contact your WorkSpaces Applications administrator. Your administrator might need to perform additional configuration tasks, as described in Real-Time Audio-Video.

3. To display the names of the webcams that are attached to your streaming session, choose the downward arrow next to the video icon. If you have more than one webcam (for example, if you have a USB webcam that is connected to your laptop and a built-in webcam), a check mark appears next to the name of the webcam that is selected for use for video conferencing within your streaming session.



4. To use the selected webcam for video conferencing within your WorkSpaces Applications streaming session, start the video conferencing application that you want to use. When the webcam is active (being used for video conferencing within your streaming session), the video icon is red.



5. To enable the microphone, choose the microphone icon.

> **Note**
>
> If you have more than one webcam and want to change the one that you use for streaming within an WorkSpaces Applications session, you must first detach your webcams from the session. For more information, see the next procedure.

To change the local webcam to use within an WorkSpaces Applications streaming session

1. Within your WorkSpaces Applications streaming session, in the WorkSpaces Applications toolbar in the top left of your session window, do either of the following:

   - If the video icon does not have a diagonal line through it (as shown in the following screenshot), this indicates that the WorkSpaces Applications real-time AV feature is available for use and that webcams are still attached to your streaming session. Choose the video icon to detach the webcams.

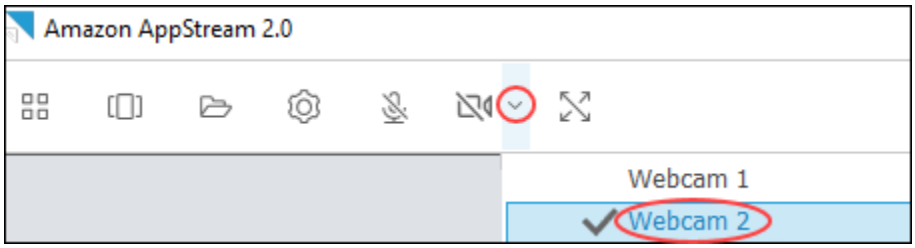
     

   - If the video icon has a diagonal line through it (as shown in the following screenshot), your webcams are already detached from your streaming session. Skip this step and proceed to the next step.

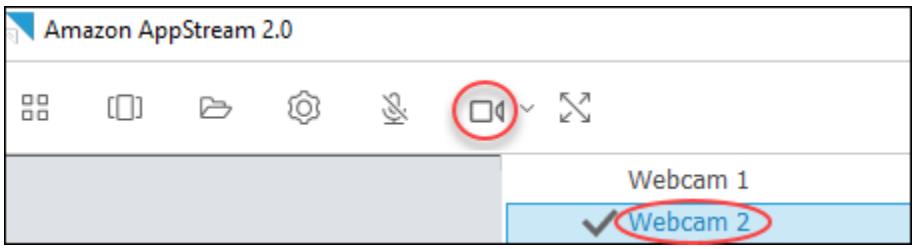
     

2. Display the names of your webcams by choosing the downward arrow next to the video icon, then select the name of the webcam that you want to use.

> **Note**
>
> You must select the name of the webcam you want to use. If you select the check mark next to the name of the webcam you want to use, the webcam won't change.



3. Choose the video icon to reattach the webcams to your WorkSpaces Applications streaming session.



Drawing Tablets

Drawing tablets, also known as pen tablets, are computer input devices that let you draw with a stylus (pen). With WorkSpaces Applications, you can connect a drawing tablet, such as a Wacom drawing tablet, to your local computer and use the tablet with your streaming applications.

Following are requirements and considerations for using drawing tablets with your streaming applications.

- To use this feature, you must connect to WorkSpaces Applications by using the WorkSpaces Applications client, or through the Google Chrome or Mozilla Firefox browsers only.
- The applications that you stream must support Windows Ink technology. For more information, see Pen interactions and Windows Ink in Windows apps.
- Depending on the streaming applications that you use, your drawing tablet might require USB redirection to function as expected. This is because some applications, such as GIMP, require USB redirection to support pressure sensitivity. If this is the case for your streaming applications, you

must connect to WorkSpaces Applications by using the WorkSpaces Applications client and share the drawing tablet with your streaming session.

- This feature is not supported on Chromebooks.

To get started with using a drawing tablet during your application streaming sessions, connect your drawing tablet to your local computer with USB, share the device with WorkSpaces Applications if required for pressure sensitivity detection, and then start an WorkSpaces Applications streaming session. You can use the WorkSpaces Applications client or a supported web browser to start a streaming session.

Relative Mouse Offset

By default, during a streaming session, WorkSpaces Applications transmits information about mouse movements by using absolute coordinates and rendering the mouse movements locally. For graphics-intensive applications, such as computer-aided design (CAD)/computer-aided manufacturing (CAM) software or video games, mouse performance improves when relative mouse mode is enabled. Relative mouse mode uses relative coordinates, which represent how far the mouse moved since the last frame, rather than the absolute x-y coordinate values within a window or screen. When you enable relative mouse mode, WorkSpaces Applications renders the mouse movements remotely.

You can enable this feature during an WorkSpaces Applications streaming session in either of the following ways:

- Pressing Ctrl+Shift+F8
- Choosing **Relative Mouse Position [Ctrl+Shift+F8]** from the **Settings** menu on the WorkSpaces Applications toolbar in the top left area of your streaming session window. This method works when you use classic mode or **Desktop View**.

Logging

To help with troubleshooting if an issue with the WorkSpaces Applications client occurs, you can enable diagnostic logging. The log files that are sent to WorkSpaces Applications (AWS) include detailed information about your device and connection to the AWS network. You can enable automatic log uploads so that these files are sent to WorkSpaces Applications (AWS) automatically. You can also upload log files on an as-needed basis, before or during an WorkSpaces Applications streaming session.

Automatic logging

You can enable automatic logging when you install the WorkSpaces Applications client. For information about how to enable automatic logging when you install the WorkSpaces Applications client, see step 5 in Setup for Windows.

On-demand logging

If an issue occurs during an WorkSpaces Applications streaming session, you can also send log files on an as-needed basis. If an issue occurs that causes the WorkSpaces Applications client to stop responding, a notification prompts you to choose whether to send an error report and the associated log files to WorkSpaces Applications (AWS).

The following procedures describe how to send log files before you sign in to an WorkSpaces Applications streaming session and during an WorkSpaces Applications streaming session.

To send log files before an WorkSpaces Applications streaming session

1. On your local PC where the WorkSpaces Applications client is installed, in the lower left of your screen, choose the Windows search icon on the taskbar, and enter **`AppStream`** in the Search box.
2. In the search results, select **Amazon AppStream** to start the WorkSpaces Applications client.
3. At the bottom of the WorkSpaces Applications sign-in page, choose the **Send Diagnostic Logs** link.
4. To continue connecting to WorkSpaces Applications, if your WorkSpaces Applications administrator has provided you with a web address (URL) to use to connect to WorkSpaces Applications for application streaming, enter the URL, and then choose **Connect**.

To send log files during an WorkSpaces Applications streaming session

1. If you are not already connected to WorkSpaces Applications and streaming an application, use the WorkSpaces Applications client to start a streaming session.
2. In the upper right of the WorkSpaces Applications session window, choose the **Profiles** icon, and then choose **Send Diagnostic Logs**.

Troubleshooting

If issues occur when you use the WorkSpaces Applications client for Windows, your WorkSpaces Applications client ID and version number can help your administrator with troubleshooting. The following sections describe how to find the client ID and client version number.

How to find the WorkSpaces Applications client ID

The WorkSpaces Applications client ID uniquely identifies your device. This ID is created when you install the WorkSpaces Applications client for Windows. To find your client ID, open the WorkSpaces Applications client. On the bottom left of the client sign-in page, choose the **Client Options** link. The client ID is displayed at the top of the **WorkSpaces Applications Client Options** dialog box. You can use your mouse to select the client ID, and then copy it to your clipboard by using your keyboard or mouse.

How to find the WorkSpaces Applications client version number

WorkSpaces Applications periodically releases new client versions to add features and functionality or resolve issues. To find the version of the WorkSpaces Applications client that you have installed, open the WorkSpaces Applications client. On the bottom of the client sign-in page, choose the **About Amazon WorkSpaces Applications** link. The client version is displayed below the Amazon WorkSpaces Applications logo.

WorkSpaces Applications Windows Client Release Notes

The WorkSpaces Applications client is a native application that is designed for users who require additional functionality during their WorkSpaces Applications streaming sessions. The following table describes the latest updates that are available in released versions of the WorkSpaces Applications client.

For more information about the client, see Provide Access Through the WorkSpaces Applications Client.

Client version	Release date	Changes
1.5.2.1096	11-11-2026	• Fixed an issue that could prevent connecting to streaming URLs on custom domains.

Client version	Release date	Changes
1.5.0.1071	10-15-2026	• Updated the client to use a Windows Installer (MSI) package, enabling per-device installation for all users. The client is no longer distributed as an EXE installer; if you deploy the client using automation, update your scripts to use the MSI package. • Updated the DCV SDK. • Other bug fixes and enhancements.
1.4.0	09-25-2026	• Re-enabled support for honoring printer preferences when printing from a session, which had been temporarily disabled in version 1.3.1992. • Made improvements to reduce client startup time. • Updated the DCV SDK. • Other bug fixes and enhancements.

Client version	Release date	Changes
1.3.1992	09-03-2026	• Temporarily disabled support for honoring printer preferenc es, an enhancement introduced in client version 1.3.1954, due to an issue causing delays with printer detection.
1.3.1978	08-30-2026	• Fixed an issue introduced in client version 1.3.1954 that prevented redirecte d printers and file transfers from working properly in prolonged sessions. • Updated the Enterpris e Deployment Tool to use a Windows Installer (MSI) package, enabling per-device client installation for all users.

Client version	Release date	Changes
1.3.1954	08-18-2026	• Added support for persisting the state of the full screen or maximized client window across sessions. • Added support for honoring printer preferences when printing from a session. • Fixed an issue where expired authentic ation tokens could cause connection failures during session reconnection. • Fixed focus-related issues with native application windows. • Fixed an issue that caused the file storage dialog to freeze during loading under certain conditions. • Updated the Chromium Embedded Framework (CEF) to version 149.0.60. • Updated the DCV SDK. • Other bug fixes and enhancements.

Client version	Release date	Changes
1.2.1830	05-29-2026	• Updated the title of native application mode windows to start with the application name first, improving usability. • Fixed an issue that caused a white screen to appear during reconnection under certain conditions in native application mode. • Fixed an issue that prevented downloadi ng files with multibyte characters (such as Japanese, Korean, or Chinese characters) in the file name. • Updated the Chromium Embedded Framework (CEF) to version 146.0.10. • Updated the DCV SDK. • Other bug fixes and enhancements.

Client version	Release date	Changes
1.2.1792	04-13-2026	• Made an improvement so that the client automatically attempts to reconnect your session after temporary network interruptions. • Fixed an issue with certificate-based authentication to better handle certificate errors. • Other bug fixes and enhancements.
1.2.1778	03-31-2026	• Added support for automatically minimizing the application catalog window when using a streaming URL with a pre-configured application ID. • Updated the Chromium Embedded Framework (CEF) to version 145.0.26. • Updated the DCV SDK. • Other bug fixes and enhancements.

Client version	Release date	Changes
1.2.1733	02-12-2026	• Added streaming support for Ubuntu-based Elastic fleets. • Added support for the client to consume the centrally-managed time zone redirection setting from the enable-timezone-redirection parameter on the DCV server. • Added a configurable indicator to provide real-time visual feedback on streaming data loss. • Fixed an issue with incorrect monitor resolutions being applied when the display scaling factor was set to higher than 100% in a multi-display configuration. • Fixed a timing issue that caused streaming session launches to result in a black screen in rare conditions. • Fixed an issue to show the application display name during

Client version	Release date	Changes
		the application launch process. • Fixed an issue that showed an incorrect session expiration error during user pool streaming sessions. • Updated the DCV SDK. • Updated the Chromium Embedded Framework (CEF) to version 140.1.140. • Other bug fixes and enhancements.
1.2.1610	11-14-2025	• Fixed an issue introduced in client version 1.2.1581 that disabled smart card redirection functiona lity.
1.2.1581	11-06-2025	• Launching the new Amazon WorkSpaces Applications client • Adds support for streaming over IPv6 • Additional bug fixes and improvements
1.1.1532	09-18-2025	• Includes bug fixes and improvements

Client version	Release date	Changes
1.1.1490	08-11-2025	• Includes bug fixes and improvements
1.1.1458	06-10-2025	• Upgrades the embedded Chromium browser to version 135.0.170 • Includes bug fixes and improvements
1.1.1440	06-06-2025	• Reverts the embedded Chromium browser to version 131.3.50 • Includes bug fixes and improvements
1.1.1437	05-21-2025	• Upgrades the embedded Chromium browser to version 135.0.170 • Includes bug fixes and improvements
1.1.1423	03-31-2025	• Upgrades the embedded Chromium browser to version 131.3.50 • Includes bug fixes and improvements
1.1.1414	01-16-2025	• Adds support for automatic time zone redirection • Includes bug fixes and improvements

Client version	Release date	Changes
1.1.1408	12-19-2024	• Includes bug fixes and improvements
1.1.1403	12-12-2024	• Adds support to save user preferences between streaming sessions • Includes bug fixes and improvements
1.1.1395	11-18-2024	• Upgrades the embedded Chromium browser to version 129.0.110 • Includes bug fixes and improvements
1.1.1360	08-01-2024	• Adds support for extending full-scre en across selected monitors • Adds support to stream Red Hat Enterprise Linux images • Upgrades the embedded Chromium browser to version 125.0.210 • Includes bug fixes and improvements
1.1.1332	07-03-2024	• Includes bug fixes and improvements

Client version	Release date	Changes
1.1.1326	06-17-2024	• Improves the user experience for the IdP-initiated SSO workflow by automatically opening the client after user sign-in with the system browser • Other bug fixes and improvements
1.1.1303	04-03-2024	• Includes bug fixes and improvements
1.1.1300	03-28-2024	• Added support for launching the WorkSpaces Applications client from IdP-initiated streaming sessions • Added support for new relay state regional endpoints • Upgrades the embedded Chromium browser to version 121.3.70 • Includes bug fixes and improvements
1.1.1259	02-08-2024	• Includes bug fixes and improvements

Client version	Release date	Changes
1.1.1246	01-18-2024	• Includes improved accessibility features • Includes bug fixes and improvements • Upgrades the embedded Chromium browser to version 119.4.30
1.1.1228	11-01-2023	• Includes bug fixes and improvements • Upgrades the embedded Chromium browser to version 114.1.120
1.1.1183	06-22-2023	• Includes bug fixes and improvements • Upgrades the embedded Chromium browser to version 111.2.20
1.1.1159	05-09-2023	• Includes bug fixes and improvements
1.1.1130	02-09-2023	• Upgrades the embedded Chromium browser to version 108.4.130
1.1.1118	11-07-2022	• Upgrades the embedded Chromium browser to version 106.0.26

Client version	Release date	Changes
1.1.1099	10-13-2022	• Includes bug fixes and improvements
1.1.1066	08-17-2022	• Upgrades the embedded Chromium browser to version 102.0.9. Microsoft Visual C++ 2019 Redistributable must be installed as a prerequisite.
1.1.1025	06-29-2022	• Adds support for UDP streaming. For more information, see Amazon WorkSpaces Applications enables UDP streaming for Windows native client.
1.1.421	05-19-2022	• Includes bug fixes
1.1.414	04-26-2022	• Includes bug fixes and UI improvements
1.1.398	02-23-2022	• Includes bug fixes
1.1.394	02-08-2022	• Upgrades the embedded Chromium browser to version 97
1.1.386	12-20-2021	• Upgrades the embedded Chromium browser to version 94.4 • Includes bug fixes

Client version	Release date	Changes
1.1.360	11-15-2021	• Adds support for Linux application streaming • Adds support for Elastic fleets. For more information, see Amazon AppStream 2.0 launches Elastic fleets. • Fixes a bug with the Japanese keyboard
1.1.333	09-08-2021	• Bug fixes for the embedded Chromium browser
1.1.319	08-16-2021	• Resolves an issue with the caps lock, number lock, and scroll lock keys • Resolves an issue for the domain join sign-in experience
1.1.304	08-02-2021	• Upgrades the embedded Chromium browser to version 91 • Updated USB driver to include important fixes

Client version	Release date	Changes
1.1.294	04-26-2021	• Resolves an issue with SAML 2.0 authentic ation • Resolves a client stability issue with Windows 7 • Resolves an issue with folder sharing on client reconnection
1.1.285	03-08-2021	• Includes fixes that improve compatibility with antivirus software
1.1.257	12-28-2020	• Adds support for real-time audio-video (AV) • Adds support for using a smart card for Windows sign in to Active Directory -joined streaming instances and in-sessio n authentication for streaming applications • Resolves an issue that causes Microsoft Excel sheets to lose focus during streaming sessions

Client version	Release date	Changes
1.1.195	08-18-2020	• Improves the experience of sharing local drives and folders that belong to cloud-based persistent storage solutions such as OneDrive when file redirection is used during streaming sessions • Upgrades the embedded Chromium browser to version 81 • Resolves `AS2TrustedDomain` DNS TXT record lookup failures for domains specified in the `AS2TrustedDomains` list. These failures may occur with some URI schemes. For more information, see Create the AS2TrustedDomains DNS TXT Record to Enable Your Domain for the WorkSpaces Applications Client Without Registry Changes • Resolves an intermittent issue that causes the client to stop

Client version	Release date	Changes
		functioning when audio is enabled
1.1.179	07-08-2020	• Adds support for local printer redirection • Resolves an issue with concurrent HTTP connections that prevents streaming with some proxy settings • Resolves an issue that causes file downloads for files greater than a few gigabytes to stop, and then fail • Resolves an issue that causes subsequent connection attempts to WorkSpaces Applications to fail if users sign in and connect to WorkSpaces Applications over SAML, disconnect from the session without closing the WorkSpaces Applications client, and then try to start a new WorkSpaces Applications streaming session

Client version	Release date	Changes
1.1.160	04-28-2020	• Resolves an issue that prevents the application catalog page from opening on a Windows PC that has .NET Framework version 4.7.1 or earlier installed • Resolves an intermitt ent issue that causes the client to stop responding when users close the client application

Client version	Release date	Changes
1.1.156	04-22-2020	• Adds support for defining trusted subdomains for user connections in a DNS TXT record • Adds support for on-demand diagnostic log and minidump uploads • Adds support for displaying custom branding for users who stream in native application mode **Note** Users who have this version of the WorkSpaces Applications client installed must have .NET Framework version 4.7.2 or later installed on the same PC. For a list of the .NET Framework versions available for download, see Download .NET Framework.

Client version	Release date	Changes
1.1.137	03-08-2020	• Reverts the updates in version 1.1.136
1.1.136	03-05-2020	• Adds support for defining trusted subdomains for user connections in a DNS TXT record
1.1.129	02-28-2020	• Adds support for native application mode • Improves the user interface for the DCV Printer experience • Resolves an issue with using Surface Pro Pen with streaming applications • Resolves an issue with downloading files with file names that have international characters
1.0.525	12-12-2019	• Resolves a DPI issue that causes the mouse cursor to point to the wrong location when a user clicks on an application during a streaming session

Client version	Release date	Changes
1.0.511	10-16-2019	• Adds support for up to 4 monitors with a maximum display resolution of 4096x2160 pixels per monitor • Adds support for up to 2 monitors with a maximum display resolution of 4096x2160 pixels per monitor on the Graphics Design and Graphics Pro instance types • Adds support for seamless user connections to streaming sessions that were started using custom uniform resource identifier (URI) redirects • Adds support for adding trusted domains for start URLs

Client version	Release date	Changes
1.0.499	09-26-2019	• Resolves an issue with client-side hardware rendering • Resolves an issue with the client not working correctly when Bluetooth headsets are connected to the local computer
1.0.480	08-20-2019	• Adds support for WorkSpaces Applications file system redirection
1.0.467	07-29-2019	• Includes fixes and enhancements to ensure compatibility with updates made to WorkSpaces Applications portal endpoints

Client version	Release date	Changes
1.0.407	05-16-2019	• Adds support for configuring the amount of time that users can be idle (inactive) before they are disconnected from their streaming session. For more information, see "Create a Fleet" in Create an Amazon WorkSpaces Applications Fleet and Stack. • Resolves an issue with the “session alert” window appearing when a SAML 2.0 session has expired • Includes bug fixes for printing a document to a print server

Client version	Release date	Changes
1.0.375	03-07-2019	• Adds touch screen support on Windows PCs • Adds support for automatically connecting USB devices when a new streaming session starts • Adds support for running session scripts • Adds support for delivering virtualized applications using the WorkSpaces Applicati ons dynamic applicati on framework APIs

Client version	Release date	Changes
1.0.320	01-19-2019	• Adds multi monitor support for Graphics Design instances • Adds support for client display scaling factors greater than 100 percent • Adds support for WorkSpaces Applications regional settings • Adds support for the WorkSpaces Applications user pool • Adds support for honoring client-side proxy settings
1.0.247	11-20-2018	Initial release

WorkSpaces Applications Client Application for macOS

The following information helps you use the WorkSpaces Applications client for macOS to connect to WorkSpaces Applications and stream applications.

Contents

- Requirements
- Setup and installation for macOS
- Connect to WorkSpaces Applications on macOS client
- Monitors
- Video and Audio Conferencing
- Relative Mouse Offset
- Remap the Windows Logo Key or Command Key

- Remember My Settings
- Printer Redirection
- Disconnect and End Session
- Troubleshooting
- WorkSpaces Applications macOS Client Release Notes

Requirements

The WorkSpaces Applications client for macOS must be installed on a computer that meets the following requirements:

- Operating system — macOS 13 (Ventura), macOS 14 (Sonoma), or macOS 15 (Sequoia)
- Hard drive space — 200 MB minimum

Setup and installation for macOS

Follow these steps to download and install the client application.

1. On your macOS device, open Amazon WorkSpaces Applications Downloads, and choose the macOS link.
2. Download and install the application.
3. Verify that the WorkSpaces Applications client application icon appears on your Mac Launchpad, or check in `/Users/username/Applications/` or `~/Applications`.

Connect to WorkSpaces Applications on macOS client

After the WorkSpaces Applications client for macOS is installed on your PC, you can use it to connect to WorkSpaces Applications.

Topics

- WorkSpaces Applications macOS Client Connection Mode
- Connect to WorkSpaces Applications

WorkSpaces Applications macOS Client Connection Mode

The WorkSpaces Applications macOS client supports two connection modes: *Classic mode* and *Desktop view*. Your administrator will set up the connection mode for you.

Classic mode

When you use classic application mode, you work with remote streaming applications in the WorkSpaces Applications session window. If your administrator has made more than one application available to you, you can open multiple applications during your session. All applications that you open are displayed in the same WorkSpaces Applications session window.

When you connect to WorkSpaces Applications in classic mode, the WorkSpaces Applications Application Launcher window opens and displays the list of applications that are available for you to stream. When you open a streaming application in this mode, the Application Launcher window closes, and the application opens in the WorkSpaces Applications session window.

Desktop view

When you connect to WorkSpaces Applications and choose **Desktop view**, WorkSpaces Applications provides a standard Windows desktop view for your streaming session. The icons of applications that are available for you to stream appear on the Windows desktop. In addition, the WorkSpaces Applications toolbar, which enables you to configure settings for your streaming session, appears in the top left area of your streaming session window.

Connect to WorkSpaces Applications

Follow these steps to connect to WorkSpaces Applications and start an application streaming session.

1. If your administrator requires you to sign in first through your organization's sign-in page, complete the tasks in this step, then proceed to step 3.

   If your administrator doesn't require you to sign in through your organization's sign-in page, skip the tasks in this step and proceed to step 2.

   a. Navigate to your organizational sign-in page and enter your domain credentials when prompted.
   b. After you sign in, you are redirected to the WorkSpaces Applications Application Manager catalog page, which displays one or more applications that are available for your

WorkSpaces Applications streaming session. **Desktop View** is also available, if enabled by your administrator.

c. Choose an application or, if available, **Desktop View**.

2. If your administrator doesn't require you to sign in first through your organization's sign-in page, complete the following steps to start the WorkSpaces Applications client:

   a. On your local computer where the WorkSpaces Applications client is installed, choose **Amazon WorkSpaces Applications** to start the WorkSpaces Applications client.

   b. Do either of the following:

      - If the client sign-in page is prepopulated with a web address (URL), choose **Connect**. You will be redirected to the system's default web browser for authentication before transitioning into the streaming session in the macOS client.
      - If the client sign-in page is not prepopulated with a URL, enter the URL that your WorkSpaces Applications administrator provided for WorkSpaces Applications, and then choose **Connect**. You might also be redirected to the system's default web browser for authentication before transitioning into the streaming session in the macOS client. If you don't know the URL, contact your administrator.

   c. After a few moments, the WorkSpaces Applications Application Manager catalog page opens, displaying one or more applications that are available for your WorkSpaces Applications streaming session. **Desktop View** is also available, if enabled by your administrator.

   d. Choose an application or, if available, **Desktop View**.

3. Depending on the authentication settings that your WorkSpaces Applications administrator enabled, after you choose an application or **Desktop View**, you might be prompted to enter your Active Directory domain credentials to sign in to your WorkSpaces Applications session. If this is the case, enter your Active Directory domain password, and then choose **Password sign in**.

Monitors

Monitors and Display Resolution

The WorkSpaces Applications client supports multiple monitors with the following display resolutions:

- Multiple monitors (up to 4K resolution) — Up to 4 monitors with a maximum display resolution of 4096x2160 pixels per monitor

> **Important**
>
> Higher resolution monitors require significantly more compute capacity and encoding resources to stream content effectively, with actual performance depending on your display configuration (resolution and number of monitors) and the compute instance being used. If you experience suboptimal performance on 4K monitors, we recommend reducing your display resolution to improve responsiveness.
> Ultrawide monitors with resolutions exceeding 4096 pixels in either dimension (e.g., 5120x2160) will display black bars on the sides, as the maximum supported resolution is limited to 4096 pixels per dimension.

If you prefer a fixed resolution, which does not change even when the client window is resized, choose **Settings**, **Display**, **Display Resolution**, and specify the desired resolution. To re-enable automatic resizing, choose **Adapt automatically**.

Using Multiple Monitors

When using multiple monitors, you can choose from the following options:

- Extend full-screen across a *single* monitor
- Extend full-screen across *all* monitors
- Extend full-screen across *selected* monitors

Extending full-screen across a single monitor

You can extend full screen only on the current monitor if multiple monitors are connected to your local computer. To enable this feature, complete the following steps:

1. On the toolbar at the top of the window, choose the Full Screen (crossed arrows) icon.
2. From the drop-down menu, choose **Full screen current monitor**.

Extending full-screen across all monitors

You can extend the display for a session across all monitors at full screen resolution. The extended display matches your physical display layout and screen resolutions. For example, three monitors are connected to your local computer. The server extends the display for a session across all three monitors and matches the specific screen resolutions of your display.

To enable this feature, complete the following steps:

1. On the toolbar at the top of the window, choose the Full Screen (crossed arrows) icon.
2. From the drop-down menu, choose **Full screen all monitors**.

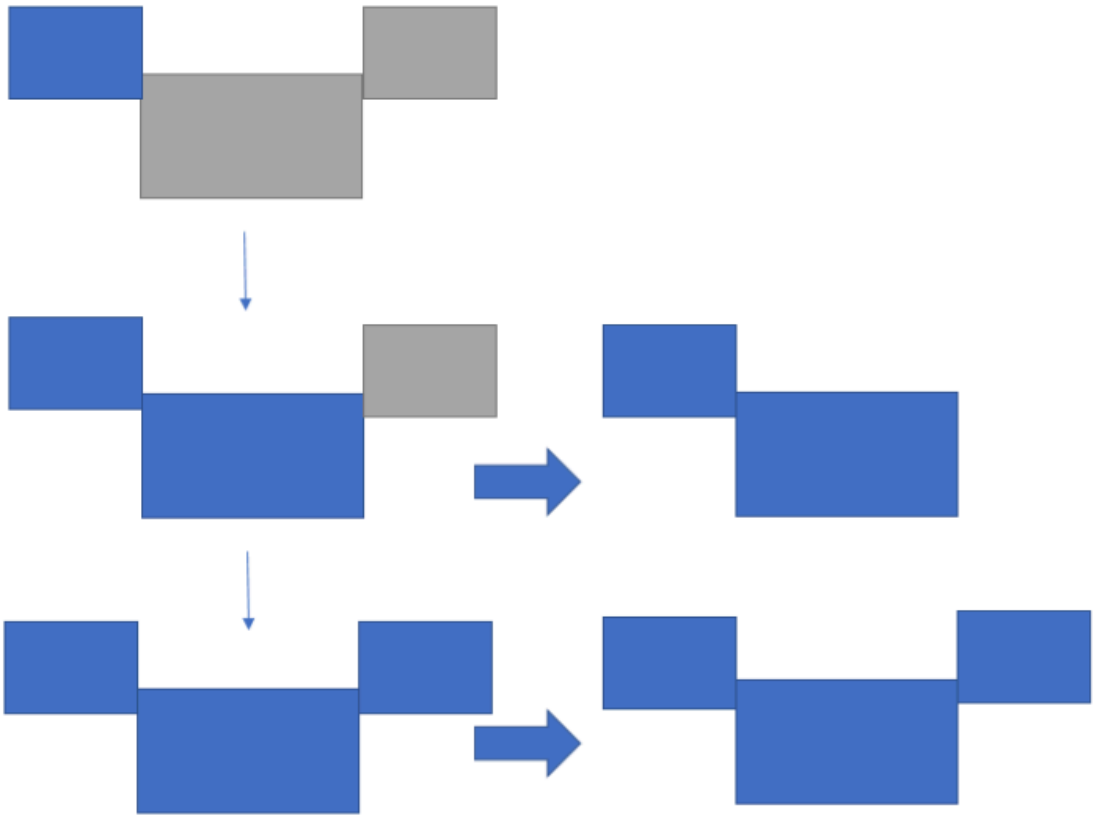
Extending full-screen across selected monitors

If there are three or more monitors connected, WorkSpaces Applications can also extend full-screen across a selection of those available monitors. If your selected monitors cannot go full screen, an error message will appear and you will need to perform the procedure again. Selected monitors must be set adjacent, or sharing a side with each other, in your display setting.

The following are examples of adjacent monitor placement. If your monitors are not set adjacent in your display configuration, you must exit WorkSpaces Applications and change your Display settings on your local machine.

> **Note**
>
> The blue boxes are WorkSpaces Applications-enabled monitors, and the gray boxes are other monitors.

Examples of adjacent monitor placement



Examples of nonadjacent monitor placement



To enable this feature, complete the following steps:

1. On the toolbar at the top of the window, choose the Full Screen (crossed arrows) icon.
2. From the drop-down menu, choose **Full screen selected monitors**.
3. The **Full screen selected monitors** window appears, displaying your current monitor layout. Select which monitors you want DCV to be displayed full screen, and choose **Apply**.

Video and Audio Conferencing

WorkSpaces Applications real-time audio-video (AV) redirects your local webcam video input to WorkSpaces Applications streaming sessions. That way, you can use your local devices for video and audio conferencing within your WorkSpaces Applications streaming session.

To use a webcam or a microphone on WorkSpaces Applications, choose **Settings** on the system menu and **Audio & Video** to enable or disable the microphone and webcam. If you have more than one webcam (for example, if you have a USB webcam that is connected to your laptop and a built-in webcam), you can also select one webcam from **Active Webcam**.

If you have selected **Show Toolbar**, you can also enable or disable the microphone and webcam by choosing the microphone or video icon. If you have more than one webcam, you can select the down arrow next to the video icon, and select one webcam to use.

> **Note**
>
> If the video icon doesn't display in the WorkSpaces Applications toolbar, contact your WorkSpaces Applications administrator. Your administrator might need to perform additional configuration tasks, as described in the section called “Real-Time Audio-Video”.

Relative Mouse Offset

By default, during a streaming session, WorkSpaces Applications transmits information about mouse movements by using absolute coordinates and rendering the mouse movements locally. For graphics-intensive applications, such as computer-aided design (CAD)/computer-aided manufacturing (CAM) software or video games, mouse performance improves when relative mouse mode is enabled. Relative mouse mode uses relative coordinates, which represent how far the mouse moved since the last frame, rather than the absolute x-y coordinate values within a window or screen. When you enable relative mouse mode, WorkSpaces Applications renders the mouse movements remotely.

You can enable this feature during an WorkSpaces Applications streaming session in either of the following ways:

- Pressing Ctrl+Shift+Fn+F8
- Choosing **Enable relative mouse** from the **Settings** and enabling it.

Remap the Windows Logo Key or Command Key

You can remap the Mac Option and Command keys on your keyboard.

A *modifier* key modifies the action of another key when you use both keys together. You can use a modifier key with another key to perform a task such as printing. A *Meta* key is a special type of modifier key. You can use a Meta key to temporarily change the function of another key when you use both keys together.

To remap the Mac Option and Command keys, choose **Settings** and **Keyboard & Mouse**.

You can remap the Option key to the following keys during a streaming session:

- Remote Alt key
- Local modifier key

You can remap the Command key to the following keys during a streaming session:

- Remote Control key
- Meta key

Remember My Settings

The WorkSpaces Applications macOS client application can save the preferences you configured in **Settings** for future sessions, except for **Audio & Video** settings. If you want to remember your audio (microphone) and video (webcam) settings for future sessions, choose **Settings**, **Audio & Video**, and **Remember Audio & Video (enabled/disabled) for future sessions**. When you enable this function, your audio (microphone) and video (webcam) settings are saved and persist across sessions when you access the same stack from the macOS client on the same device.

Printer Redirection

WorkSpaces Applications local printer redirection lets you access printers that are connected to your local computer from your WorkSpaces Applications streaming session. That way, you can redirect print jobs from your streaming application to a local printer, or to a network printer that you have mapped.

> **Important**
>
> To use WorkSpaces Applications printer redirection, you must have the WorkSpaces Applications client installed on your local computer, and you must use the client to connect to a streaming session. Printer redirection is not available when you connect to WorkSpaces Applications by using a web browser.

To redirect a print job to a local printer

1. Open the WorkSpaces Applications client and connect to a streaming session.
2. In your streaming application, choose **File**, **Print Now**.
3. The **Print** dialog box for your streaming application opens.
4. In the **Print** dialog box, a list of available local printers is displayed. Choose the local printer that you want to use, and then proceed with printing.

Disconnect and End Session

To disconnect the streaming session, choose one of the following options:

- On the WorkSpaces Applications toolbar, choose **Disconnect**.
- On your Mac, on the menu bar at the top of the screen, choose **Amazon WorkSpaces Applications** and **Disconnect**.

You can reconnect to the previous streaming session after a disconnection within a timeout time interval. The amount of time that a streaming session remains active after you disconnect is configured by your administrator.

To end the current session, choose one of the following options:

- On the WorkSpaces Applications toolbar, choose **End Session**.
- On your Mac, on the menu bar at the top of the screen, choose **Amazon WorkSpaces Applications** and **End Session**.

When you end the session, you are prompted to save any open documents, and you are immediately disconnected from the streaming instance.

Troubleshooting

Use the following steps to enable diagnostic log uploads and determine your client version and client ID.

Enable Diagnostic Log Uploads

To troubleshoot issues with the WorkSpaces Applications client, you can enable diagnostic logging. The log files that are sent to WorkSpaces Applications include detailed information about your device and connection to the AWS network. You can enable diagnostic log uploads before or during WorkSpaces Applications streaming sessions, so these files are sent to WorkSpaces Applications automatically. As a best practice, we recommend that you enable log upload to help the WorkSpaces Applications team troubleshoot issues.

To enable file logging, follow these steps:

1. Choose **AppStream 2.0** from the system menu bar, or navigate to the top-right corner of the **Connect** page.
2. Choose **Client Options** and **Client automatic logging**.

Collect Logs for WorkSpaces Applications Client for macOS

WorkSpaces Applications logs can be used by your administrator to identify and troubleshoot configuration issues. They can also help enable AWS Support to diagnose and troubleshoot cases. To collect and share the logs, choose from the following options:

- Option 1: Open a terminal and enter **`open ~/Library/Containers/com.amazon.appstreamclient/Data/logs`**
- Option 2: Open **Finder**, and choose **Users**, **User_Name**, **Library**, **Containers**, **Appstream**, **Data**, and **logs**
- Option 3: Open **Finder**, and from the top-left system menu bar, choose **Go** and **Go to folder**. Enter **`~/Library/Containers/com.amazon.appstreamclient/Data/logs`**

Determine Client Version and Client ID

If issues occur when you use the WorkSpaces Applications client for macOS, your WorkSpaces Applications version number and client ID can help your administrator and AWS support team with troubleshooting. To find the version of the WorkSpaces Applications client that you have

installed, open the WorkSpaces Applications client. On the system menu bar, choose **Amazon WorkSpaces Applications** and **About Amazon AppStream 2.0**. The client version is displayed below the Amazon WorkSpaces Applications logo.

To find the client ID of the WorkSpaces Applications client that you have installed, choose **Amazon WorkSpaces Applications** on the system menu bar, or navigate to the top-right corner of the **Connect** page and choose **Client Option**.

WorkSpaces Applications macOS Client Release Notes

The following table describes the latest updates that are available in released versions of the macOS WorkSpaces Applications client.

Client version	Release date	Changes
1.5.2.1118	10-30-2026	• Fixed an issue in application view mode where application windows could not be moved or resized. • Updated the DCV SDK.
1.5.1.1066	08-25-2026	• Fixed an issue where expired authentic ation tokens could cause connection failures during session reconnection. • Updated the DCV SDK. • Other bug fixes and enhancements.
1.5.0.1035	07-09-2026	• Fixed an issue that caused the streaming session to be positione d incorrectly when entering full screen on a non-primary display,

Client version	Release date	Changes
		when multiple displays are connected. • Improved the connection reliabili ty during network interruption scenarios and in proxy environments. • Updated the DCV SDK. • Other bug fixes and enhancements.
1.4.0.1012	04-28-2026	• Added support for URL redirection, enabling administr ators to configure certain websites to be redirected from the streaming session to the local device's web browser. • Fixed an issue that was preventing copying and pasting from the clipboard between the local client device and the streaming session in some scenarios. • Updated the DCV SDK. • Other bug fixes and enhancements.

Client version	Release date	Changes
1.3.0.962	02-10-2026	• Added support for user pools, an alternati ve method for access management that provides a persistent portal for users. • Added streaming support for Ubuntu-ba sed Elastic fleets. • Improved the SAML login flow to support your preferred browser, including Apple Safari, Google Chrome, Mozilla Firefox, and Microsoft Edge browsers. • Fixed an issue to sort applications in the catalog by application display name. • Fixed an issue that was preventing a streaming session from being redirected from the web browser to the client if the client application was not previously open. • Updated the DCV SDK. • Other bug fixes and enhancements.

Client version	Release date	Changes
1.2.0	11-06-2025	• Launching the new Amazon WorkSpaces Applications client • Adds support for streaming over IPv6 • Additional bug fixes and improvements
1.1.0	06-02-2025	• Accessibility fixes and improvements • Support for dynamic application providers • Support for client customizations
1.0.1	03-25-2025	• Bug fixes and improvements • Support for certificate-based authentication & multi-stack access for SAML instances
1.0.0	12-19-2024	• Initial release

File Storage Options

If your WorkSpaces Applications administrator has enabled it, you can use one or more of the following storage options for your files and folders during application streaming sessions.

- Home folders
- Google Drive
- OneDrive for Business
- Custom shared folders

Note

Google Drive and OneDrive for Business are currently not supported for Linux-based streaming instances.

Use Home Folders

If your WorkSpaces Applications administrator has enabled this file storage option, when you are signed in to an WorkSpaces Applications streaming session, you can use your home folder. You can do the following with your home folder:

- Open and edit files and folders that you store in your home folder. Content that is stored in your home folder cannot be accessed by other users.
- Upload and download files between your local computer and your home folder. WorkSpaces Applications continuously checks for the most recently modified files and folders and backs them up to your home folder.
- When you are working in an application, you can access files and folders that are stored in your home folder. Choose **File Open** from the application and browse to the file or folder that you want to open. To save your changes in a file to your home folder, choose **File Save** from the application interface, and browse to the location in your home folder where you want to save the file.
- You can also access your home folder by choosing **My Files** from the web view session toolbar.

Warning

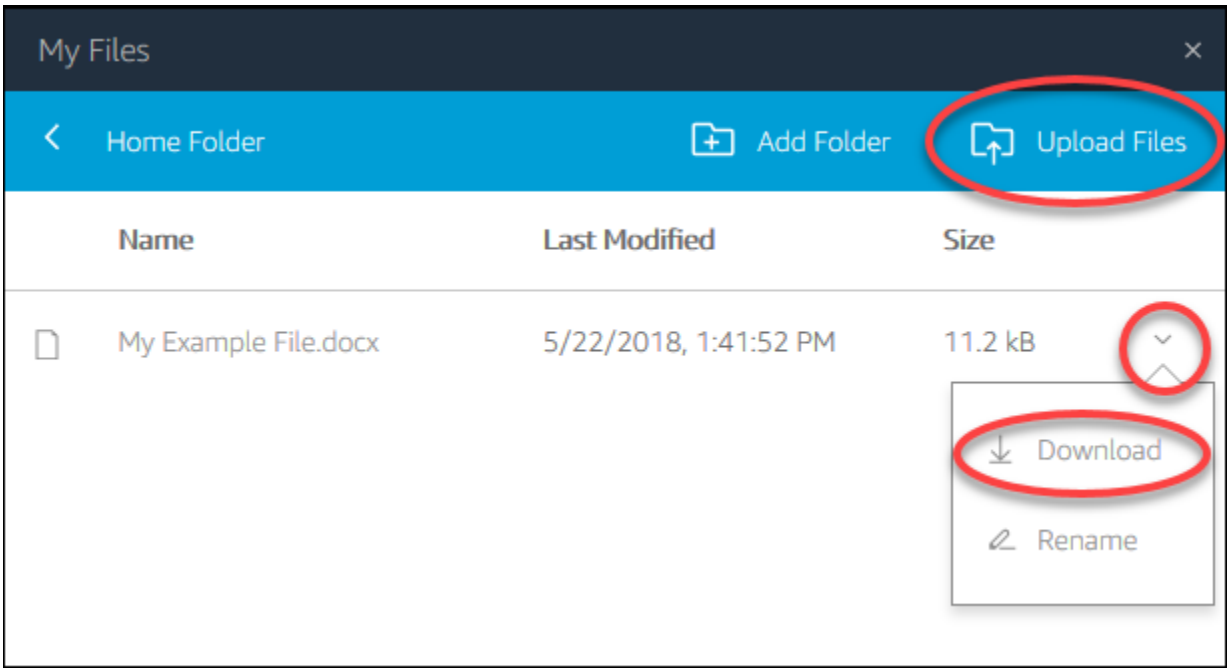
Files larger than 5 GB will not be persisted between WorkSpaces Applications sessions.

Note

If your home folder doesn't appear, view your home folder files by browsing to the following directory in File Explorer: C:\Users\PhotonUser\My Files\Home Folder.

To upload and download files between your local computer and your home folder

1. In the top left of the WorkSpaces Applications toolbar, choose the **My Files** icon.
2. Navigate to an existing folder, or choose **Add Folder** to create a folder.
3. When the folder that you want is displayed, do one of the following:
   - To upload a file to the folder, select the file that you want to upload, and choose **Upload**.
   - To download a file from the folder, select the file that you want to download, choose the down arrow to the right of the file name, and choose **Download**.



Use Google Drive

> **Note**
>
> Amazon WorkSpaces Applications's use and transfer to any other app of information received from Google APIs will adhere to Google API Services User Data Policy, including the Limited Use requirements.

If your WorkSpaces Applications administrator has enabled this file storage option, you can add your Google Drive account to WorkSpaces Applications. After you add your account and you sign in to an WorkSpaces Applications streaming session, you can do the following in Google Drive:

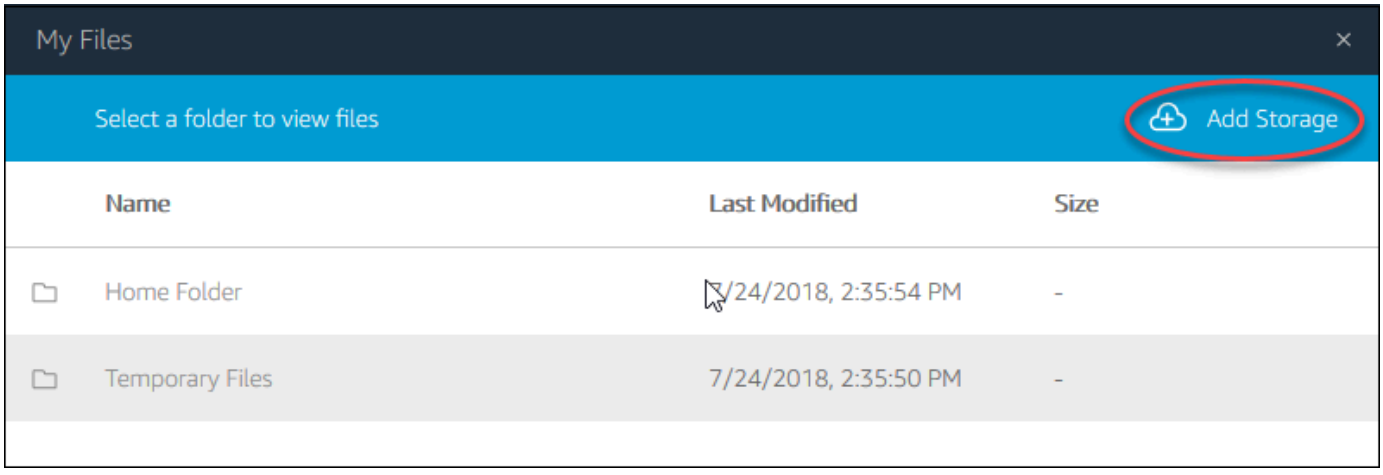
> **Note**
>
> Google Drive is currently not supported for Linux-based streaming instances.

- Open and edit files and folders that you store in Google Drive. Other users cannot access your content unless you choose to share it.
- Upload and download files between your local computer and Google Drive. Any changes that you make to your files and folders in Google Drive during a streaming session are automatically backed up and synchronized. They are available to you when you sign in to your Google Drive account and access Google Drive outside of your streaming session.
- When you are working in an application, you can access your files and folders that are stored in Google Drive. Choose **File**, **Open** from the application interface and browse to the file or folder that you want to open. To save your changes in a file to your Google Drive, choose **File**, **Save** from the application and browse to the location in Google Drive where you want to save the file.
- You can also access Google Drive by choosing **My Files** from top left of the WorkSpaces Applications toolbar.

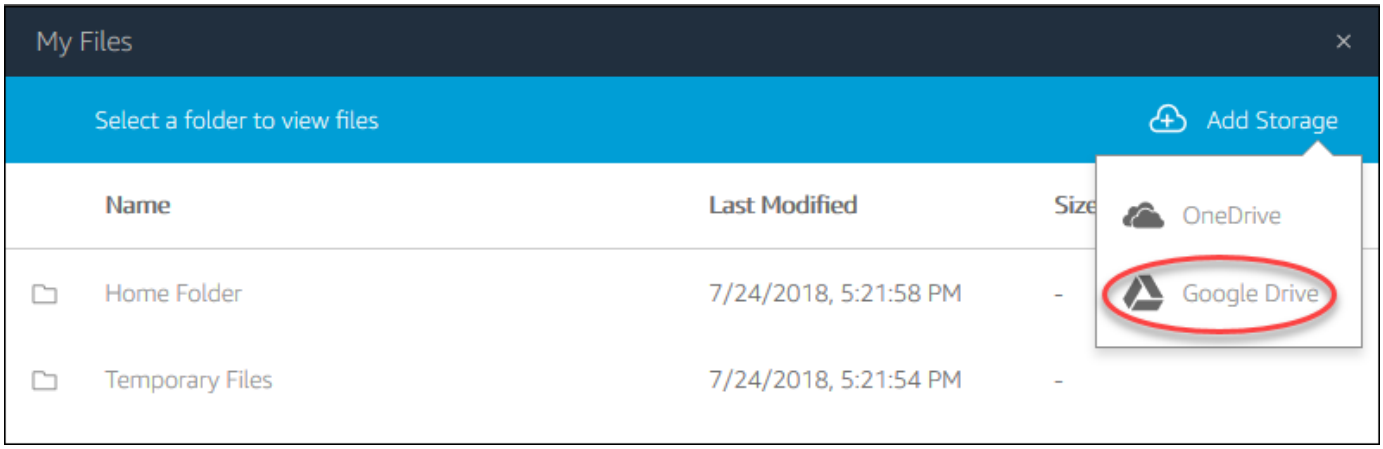
To add your Google Drive account to WorkSpaces Applications

To access your Google Drive during WorkSpaces Applications streaming sessions, you must first add your Google Drive account to WorkSpaces Applications.

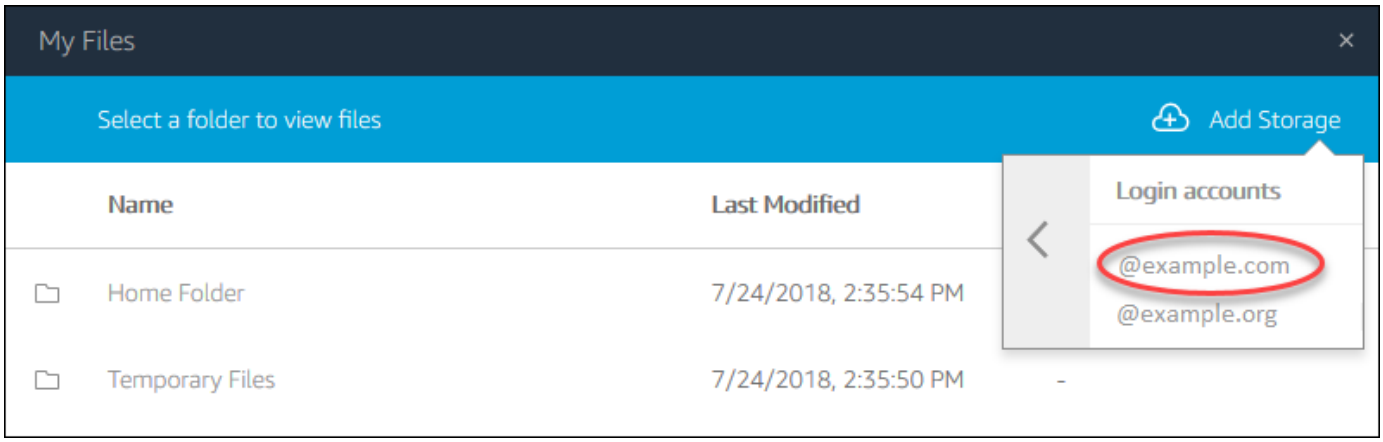
1. In the top left of the WorkSpaces Applications toolbar, choose the **My Files** icon.
2. In the **My Files** dialog box, choose **Add Storage**.



3. Choose **Google Drive**.

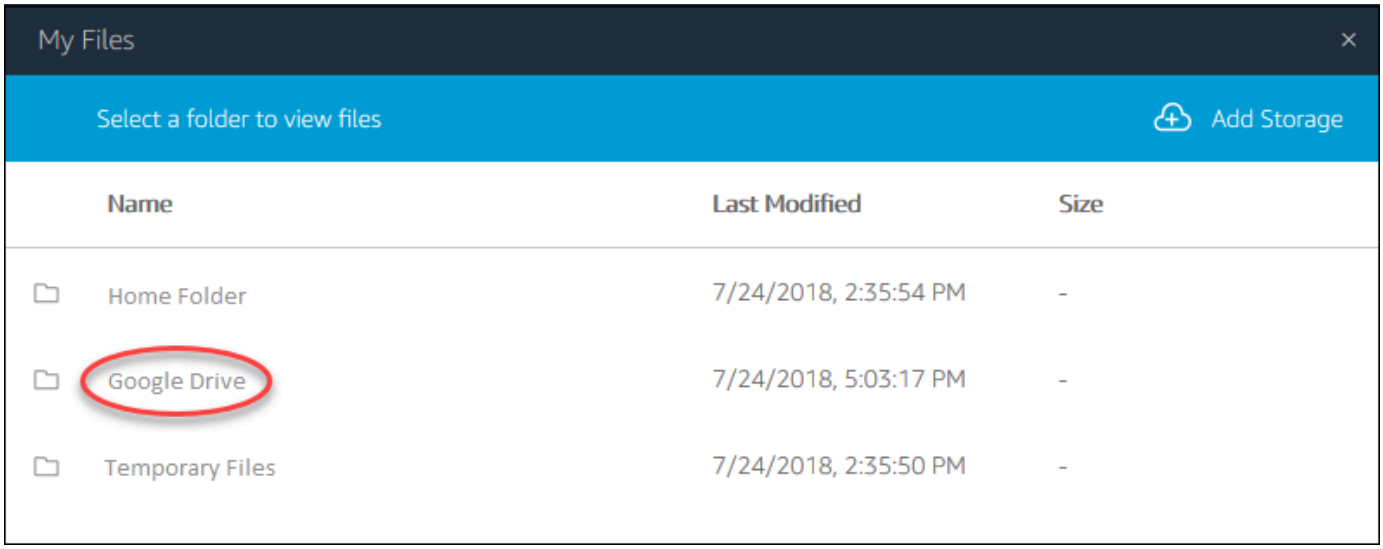


4. Choose the domain for your Google Drive account.



5. The **Sign in with Google** dialog box is displayed. Enter the sign-in credentials for your Google Drive account when prompted.

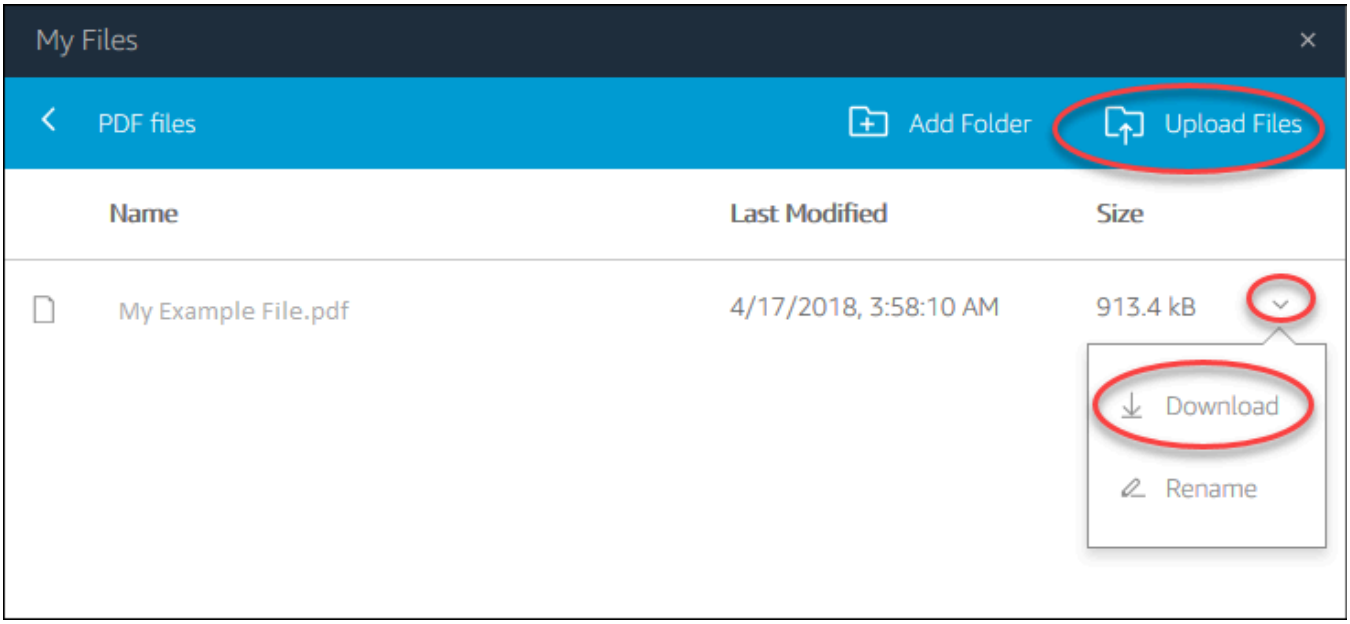
   After your Google Drive account is added to WorkSpaces Applications, your Google Drive folder is displayed in **My Files**.



6. To work with your files and folders in Google Drive, choose the **Google Drive** folder and browse to the file or folder you want. If you do not want to work with files in Google Drive during this streaming session, close the **My Files** dialog box.

To upload and download files between your local computer and your Google Drive

1. In the top left of the WorkSpaces Applications toolbar, choose the **My Files** icon.
2. In the **My Files** dialog box, choose **Google Drive**.
3. Navigate to an existing folder, or choose **Add Folder** to create a folder.
4. When the folder that you want is displayed, do one of the following:
   - To upload a file to the folder, select the file that you want to upload, and choose **Upload**.
   - To download a file from the folder, select the file that you want to download, choose the down arrow to the right of the file name, and choose **Download**.



Use OneDrive for Business

> **Note**
>
> OneDrive for Business is currently not supported for Linux-based streaming instances.

If your WorkSpaces Applications administrator has enabled this file storage option, you can add your OneDrive account to WorkSpaces Applications. After you add your account and sign in to an WorkSpaces Applications streaming session, you can do the following in OneDrive:

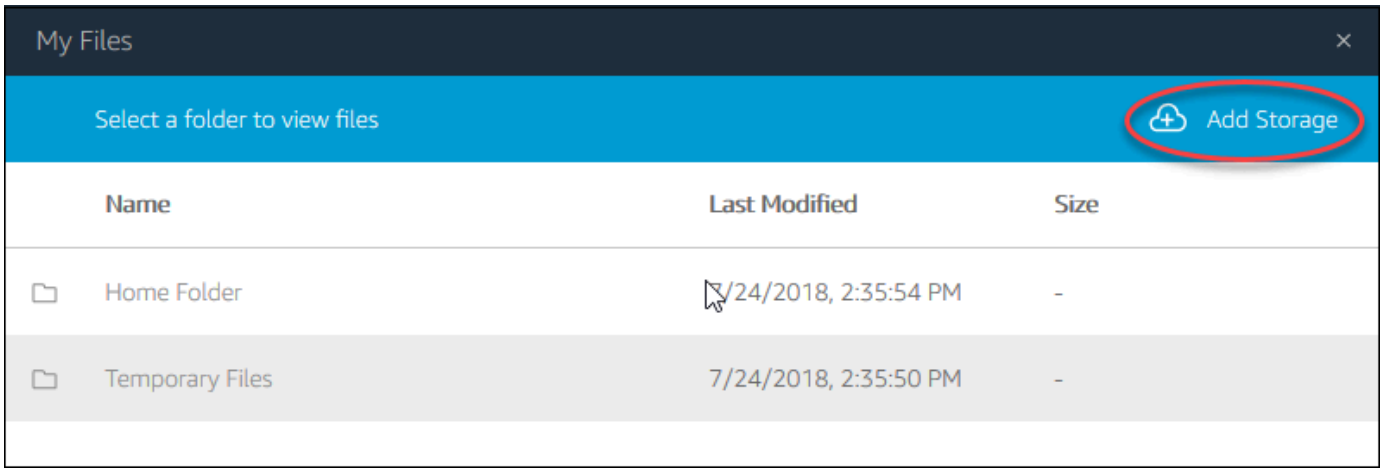
- Open and edit files and folders that you store in OneDrive. Other users cannot access your content unless you choose to share it.
- Upload and download files between your local computer and OneDrive. Any changes that you make to your files and folders in OneDrive during a streaming session are backed up and synchronized automatically. They are available to you when you sign in to your OneDrive account and access OneDrive outside of your streaming session.
- When you are working in an application, you can access your files and folders that are stored in OneDrive. Choose **File**, **Open** from the application interface and browse to the file or folder that you want to open. To save your changes in a file to OneDrive, choose **File**, **Save** from the application and browse to the location in OneDrive where you want to save the file.

- You can also access OneDrive by choosing **My Files** from the top left of the WorkSpaces Applications toolbar.

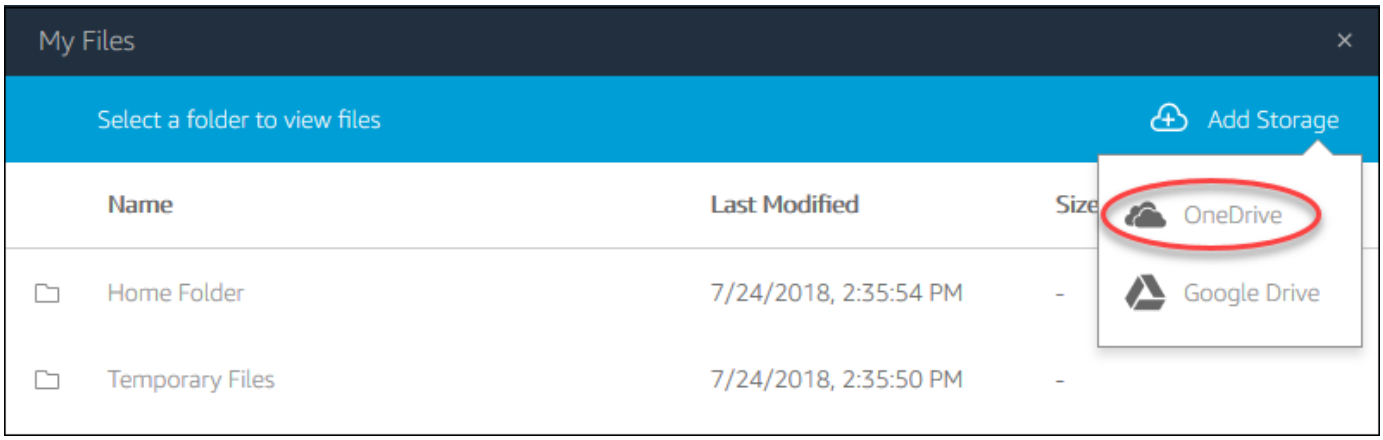
To add your OneDrive account to WorkSpaces Applications

To access your OneDrive during WorkSpaces Applications streaming sessions, you must first add your OneDrive account to WorkSpaces Applications.

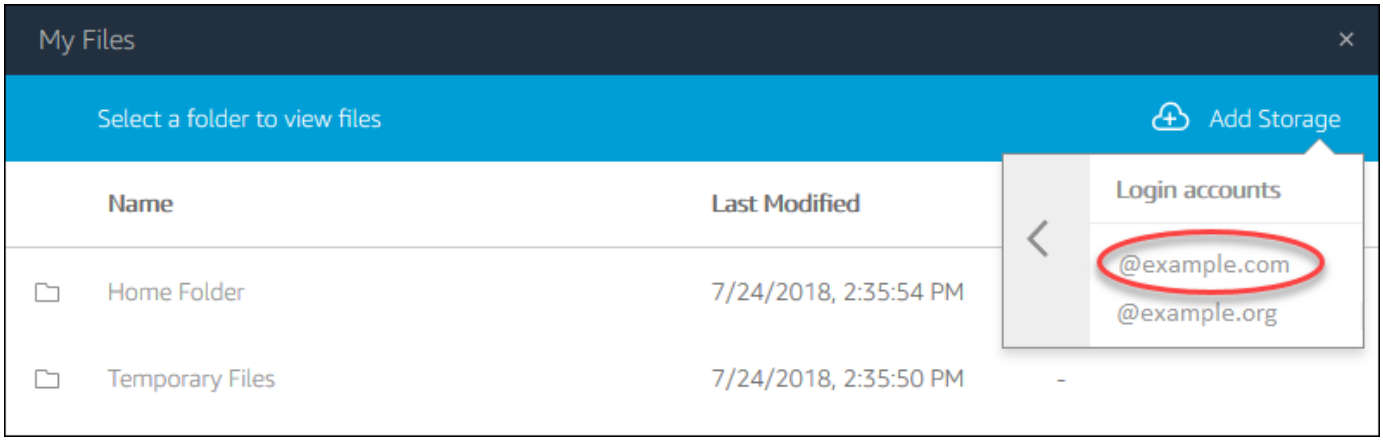
1. In the top left of the WorkSpaces Applications toolbar, choose the **My Files** icon.
2. In the **My Files** dialog box, choose **Add Storage**.



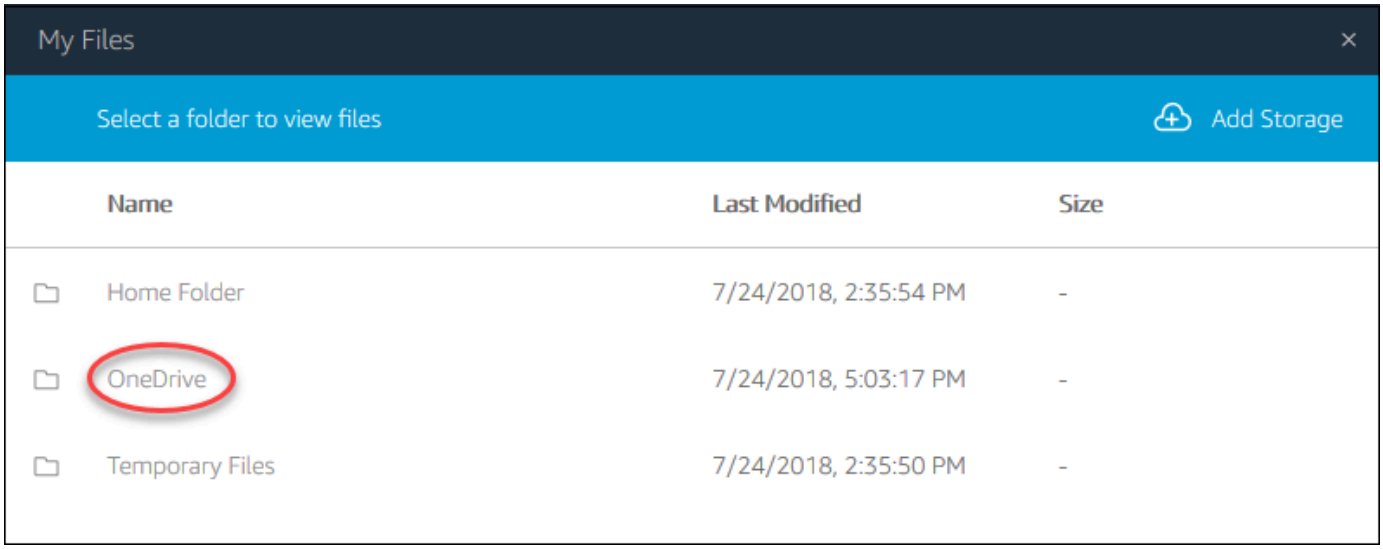
3. Choose **OneDrive**.



4. Under **Login accounts**, choose the domain for your OneDrive account.



5. In the **Sign in** dialog box, enter the sign-in credentials for your account.
6. After your OneDrive account is added to WorkSpaces Applications, your OneDrive folder is displayed in **My Files**.

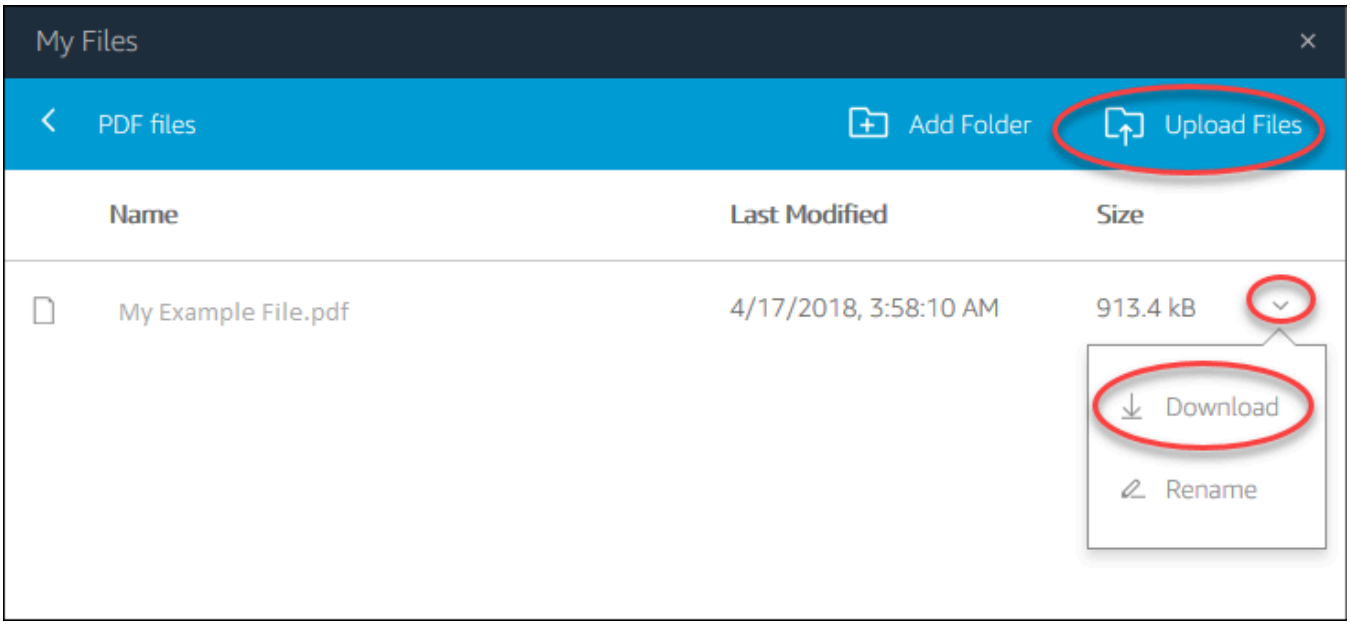


7. To work with your files and folders in OneDrive, choose the **OneDrive** folder and browse to the file or folder you want. If you do not want to work with files in OneDrive during this streaming session, close the **My Files** dialog box.

To upload and download files between your local computer and your OneDrive

1. In the top left of the WorkSpaces Applications toolbar, choose the **My Files** icon.
2. In the **My Files** dialog box, choose **OneDrive**.
3. Navigate to an existing folder, or choose **Add Folder** to create a folder.
4. When the folder is displayed, do one of the following:

- To upload a file to the folder, select the file that you want to upload, and choose **Upload**.
- To download a file from the folder, select the file that you want to download, choose the down arrow to the right of the file name, and choose **Download**.



To remove OneDrive permissions from WorkSpaces Applications

If you no longer want to use OneDrive during your WorkSpaces Applications streaming sessions, follow these steps to remove OneDrive permissions from WorkSpaces Applications.

> **Note**
>
> You can restore these permissions at any time during an WorkSpaces Applications streaming session.

1. Sign in to Office 365 or OneDrive for Business.
2. In the right pane, under **My accounts**, choose **My account**.
3. On the account dashboard page, in **App permissions**, choose **Change app permissions**.
4. On the **App permissions** page, under **Amazon WorkSpaces Applications**, choose **Revoke**.

Use Custom Shared Network Folders

If your WorkSpaces Applications administrator has enabled this file storage option, after you sign in to a streaming session, your administrator will have a custom shared folder configured and named for you. Contact your administrator for the name of the shared folder. Multiple users can access and collaborate within the shared custom folder.

You can do the following with custom shared folders:

- Open and edit files and folders that you store in OneDrive. Other users cannot access your content unless you choose to share it.
- Open and edit files and folders that you store in your custom shared folder. Content stored in it can be accessed by other users depending on the permissions configured by your administrator.
- Upload and download files between your local computer and your shared custom folder.
- When you are working in an application, you can access files and folders that are stored in your shared custom folder.

To upload and download files between your local computer and your shared custom folder

1. In the top left of the WorkSpaces Applications toolbar, choose the **My Files** icon.
2. Navigate to the shared custom folder. Contact your administrator for your folder name.
3. Choose the shared custom folder.
    - To upload a file to the folder, select the file that you want to upload, and choose **Upload**.
    - To download a file from the folder, select the file that you want to download, choose the down arrow to the right of the file name, and choose **Download**.

Configure Regional Settings

You can configure regional settings so that your WorkSpaces Applications Windows streaming sessions use settings that are specific to your location or language. Changes that you make during your streaming session are applied to future streaming sessions.

To configure regional settings for your Windows WorkSpaces Applications streaming sessions

1. In the top left of the WorkSpaces Applications toolbar, choose the **Settings** icon, and then choose **Regional settings**.

2. In the **Regional settings** dialog box, set the following options as needed. When you're done, choose **Save**.

   - **Time zone** — Determines the system time used by Windows and any applications that rely on the operating system time. Choose one of the following options:
     - To sync the time zone for your streaming session to match the time zone set on your device, choose **Set my time zone automatically based on my device**. This is enabled by default for both single-session and multi-session fleets, and can be disabled.
     - For single-session fleets, you can choose a specific time zone for your streaming session instead of using automatic redirection. To set a custom time zone, disable the **Set my time zone automatically based on my device option** in **Regional settings**, and choose a preferred time zone from the available list.
   - **Locale** (also known as culture) — Determines how Windows displays numbers, currency, time, and dates. WorkSpaces Applications supports the following locales: Chinese (Simplified and Traditional), Dutch, English, French, German, Italian, Japanese, Korean, Portuguese, Spanish, and Thai.
   - **Input method** — Determines the keystroke combinations that can be used to input characters in another language.

Configuring regional settings is not yet supported for Linux streaming sessions. However, you can switch between different input methods available in your streaming sessions with shortcut key combinations specified by your administrator. The default shortcut key combinations are “Super + Space” and “Shift + Super + Space”. “Super” is the “Windows” key on a Windows keyboard or the “Command” key on an Apple keyboard. Always check with your administrators for the shortcut keys they specified when creating the image. For example, in Tutorial: Enable Japanese Support for Your Linux Images, the shortcut key combinations have been changed to “Control + Space” and “Shift + Control + Space”.

Extension SDK Developer Guide for Amazon WorkSpaces Applications

Amazon WorkSpaces Applications uses Amazon DCV technology to provide secure, high-performance access to your applications. With the Amazon DCV Extension SDK, developers can customize WorkSpaces Applications experiences for end users, including the following actions:

- Facilitate custom hardware support.
- Enhance the usability of third-party applications in remote sessions. For example, you can add local audio termination for VoIP applications or local video playback for conferencing applications.
- Provide accessibility software such as screen readers with information about the remote session and applications running remotely.
- Allow security software to analyze the security posture of the local endpoint to allow conditional access policies.
- Perform arbitrary data transfers over an established remote session.

To get started with the Amazon DCV Extension SDK, see What is the Amazon DCV Extension SDK?. The SDK itself can be found in the Amazon DCV Extension SDK Github repository. In addition, integration examples of the SDK can be found in the Amazon DCV Extension SDK Samples Github repository.

Topics

- Extension SDK Prerequisites for Amazon WorkSpaces Applications
- Third-party vendor extensions for Amazon WorkSpaces Applications

Extension SDK Prerequisites for Amazon WorkSpaces Applications

Before you start working with the Amazon DCV Extension SDK, make sure that your WorkSpaces Applications client applications and your WorkSpaces Applications servers meet the following requirements.

Supported WorkSpaces Applications clients:

- WorkSpaces Applications Windows client version 1.1.1154 or above

> **Note**
>
> WorkSpaces Applications web access doesn't support the Amazon DCV Extension SDK.

Supported WorkSpaces Applications streaming instances:

- Windows Server 2016, 2019, 2022, and 2025
- WorkSpaces Applications agent version for Windows releaed on May 8, 2023 or later
- Managed WorkSpaces Applications image updates released on May 8, 2023 or later

Third-party vendor extensions for Amazon WorkSpaces Applications

AWS supports the Amazon DCV Extension SDK API within the WorkSpaces Applications host and client processes. However, please note that extensions developed by third-party Independent Software Vendors (ISVs) are not developed or maintained by AWS. Therefore, support for the extensions themselves, including their installation, configuration, troubleshooting, and updates, is the responsibility of the third-party vendor who developed the extension. If you have any issues or questions related to third-party extensions, please contact the relevant third-party vendor for support.

FAQs on supported protocols

The Amazon WorkSpaces Applications service is adding SDK support for the Smithy RPC v2 CBOR protocol in addition to the existing AWS JSON protocol. The Smithy RPC v2 CBOR protocol is more performant than AWS JSON for most workloads. AWS SDKs will prioritize the protocol that is the most performant for each language.

What are the AWS JSON and Smithy RPC v2 CBOR protocols, and how do they differ?

A wire protocol acts as a mediator between a client and the service. The AWS SDK handles the serialization (converting an object to the wire format) and deserialization (converting the wire format back to an object) between the client and server transparently for both protocols.

AWS JSON - JSON is one of the most widely used and accepted data formats for communication between systems. This service uses AWS JSON as a medium to communicate between an AWS SDK client and the server. An HTTP request for a service API operation accepts input in the form of JSON, the operation is executed, and the response is returned to the SDK client in the form of JSON.

Smithy RPC v2 CBOR - Smithy RPC v2 CBOR (Concise Binary Object Representation) is a protocol developed by Amazon for more efficient data transmission. It is a data format designed to keep message sizes compact and to allow the format to evolve without requiring version negotiation between systems. The key difference from AWS JSON is that AWS JSON uses human-readable text (like `{"name": "John"}`), while Smithy RPC v2 CBOR uses a binary format (sequences of 1s and 0s). This binary approach creates smaller data packages that travel faster over the network and reduces serialization time since computers process binary data directly.

How do I get started with the Smithy RPC v2 CBOR protocol?

To get started, use the latest AWS SDK/CLI version to use the Smithy RPC v2 CBOR protocol for this service. The SDK will automatically use the Smithy RPC v2 CBOR protocol when available. Upgrade your AWS SDK to the specified version or any subsequent version.

Starting 2026-07-23, the following AWS SDKs will use the Smithy RPC v2 CBOR protocol, while the other SDKs will continue using the AWS JSON protocol:

- AWS SDK for C++
- AWS SDK for Go V2
- AWS SDK for Java 2.x
- AWS SDK for .NET v4
- AWS Tools for PowerShell v5
- AWS SDK for Swift
- AWS SDK for Kotlin

What are the risks of enabling the Smithy RPC v2 CBOR protocol?

The added support of Smithy RPC v2 CBOR in the SDKs is expected to be fully transparent and effortless for customers. However, if you are using a custom implementation of the AWS SDK, or a combination of custom clients and the AWS SDK, that depends on AWS JSON-specific serialization behavior, it may be incompatible with the Smithy RPC v2 CBOR protocol.

What if my solution does not support Smithy RPC v2 CBOR?

To use the AWS JSON protocol, downgrade your SDK to a version released before 2026-07-23.

What AWS Regions support the Smithy RPC v2 CBOR protocol?

The Smithy RPC v2 CBOR protocol is supported in all AWS Regions where this service is available.

What latency improvements can I expect when upgrading to SDK versions that support Smithy RPC v2 CBOR?

Smithy RPC v2 CBOR provides up to 13% lower latency and up to 15% reduction in payload size compared to AWS JSON. The binary encoding eliminates the overhead of JSON text parsing and string escaping, which reduces both network transfer time and serialization cost.

Will the AWS JSON protocol be deprecated?

AWS JSON protocol will continue to be supported through older versions of SDKs. Also, some language-specific SDKs will continue using AWS JSON as it is more performant than Smithy RPC

v2 CBOR due to the available JSON implementation for the language being faster than its CBOR implementation.

Where can I find more information about Smithy RPC v2 CBOR?

You can find more information about the AWS JSON protocol at AWS JSON 1.0 protocol and AWS JSON 1.1 protocol in the Smithy documentation, and about Smithy RPC v2 CBOR at Smithy RPC v2 CBOR protocol in the Smithy documentation.

Document History for Amazon WorkSpaces Applications

- **API version:** 2016-12-01

The following table describes important additions to the WorkSpaces Applications service (including WorkSpaces Applications base image, WorkSpaces Applications agent, and WorkSpaces Applications client releases) and to the *Amazon WorkSpaces Applications Administration Guide* documentation from June 4, 2018 onward. We also update the documentation frequently to address the feedback that you send us.

For notification about these updates, you can subscribe to the Amazon WorkSpaces Applications RSS feed.

Change	Description	Date
Control storage connector re-authentication	Added guidance on controlling OneDrive for Business and Google Drive re-authentication at the stack level. You can configure a stack to force users to re-authenticate with their cloud storage provider at the start of each streaming session.	October 15, 2026
Unified graphics images	Adds the unified graphics base image (AppStream-Graphics-NV-*WinServerOperatingSystemVersion* -*MM-DD-YYYY*), a single graphics image that is compatible with the Graphics G4dn, G5, G6, Gr6, G6f, Gr6f, and G7 instance families. Available for Windows, Red	September 24, 2026

	Hat Enterprise Linux 8, and Rocky Linux 8.	
Unified graphics images	Added support for unified graphics images. A single graphics image works across the Graphics G4dn, G5, G6, Gr6, G6f, Gr6f, and G7 instance families, and you can change a fleet's instance type across these families.	September 24, 2026
MCP server OAuth authorization and simplified connection	The WorkSpaces Applications MCP server now supports OAuth 2.1 authorization through AWS Sign-In. You can also connect by providing the stack and user ID instead of generating a streaming URL, and use the `get_session_info` tool to retrieve session identifiers.	September 18, 2026
Agent access Region expansion	Agent access is now available in the Asia Pacific (Malaysia), Europe (Milan), Europe (Spain), and South America (São Paulo) Regions. For the MCP endpoint in each Region, see the WorkSpaces Applications MCP endpoints table.	September 17, 2026
Agent Toolkit for AWS guidance	Added guidance on using the Agent Toolkit for AWS and the AWS MCP Server to configure agent access and build your first agent.	September 17, 2026

Graphics G7 instance family	Added support for the Graphics G7 instance family (stream.graphics.g7.* instance types), powered by NVIDIA RTX PRO 4500 Blackwell Server Edition GPUs.	August 31, 2026
New agent access tools for launching applications	Added the launch_application and toggle_app_switche r tools to the WorkSpace s Applications MCP server. These tools let agents launch applications from the fleet catalog and switch between them.	August 21, 2026
Added FAQ page for supported protocols	Added an FAQ page with answers to common questions about the AWS JSON and Smithy RPC v2 CBOR wire protocols. For more informati on, see FAQs on supported protocols.	July 23, 2026
New agent access guidance	WorkSpaces Applications agent access public preview	May 1, 2026
Updated AmazonApp StreamServiceAccess	Updated the AmazonApp StreamServiceAccess managed policy	November 17, 2025
Updated AmazonApp StreamReadOnlyAccess	Updated the AmazonApp StreamReadOnlyAccess managed policy	October 22, 2025
License included apps	Stream Microsoft license included applications	October 14, 2025

New agent version	09-30-2025 version for AppStream 2.0 Agent	October 3, 2025
New client version	New client version 1.1.1532	September 18, 2025
New agent version	07-15-2025 version for AppStream 2.0 Agent	August 4, 2025
Default quotas	Updated the default quotas for Graphics fleet instances and Graphics image builder instances	July 17, 2025
Latest released images	Updated the latest released images	June 18, 2025
New version of the macOS client	Client version 1.1.0	June 2, 2025
New client version	New client version 1.1.1437	May 21, 2025
New allowed domain	New allowed domain for Europe (Paris)	May 6, 2025
New client version	New client version 1.1.1423	March 31, 2025
New version of the macOS client	Client version 1.0.1	March 26, 2025
Group policy settings	Added additional group policy settings	February 20, 2025
New client version	New client version 1.1.1414	January 16, 2025
New client version	New client version 1.1.1408	December 20, 2024
Rocky Linux 8 support	WorkSpaces Applications supports the Rocky Linux 8 operating system	December 19, 2024

Client application for macOS	Use the AppStream 2.0 client for macOS to connect to AppStream 2.0 and stream applications	December 19, 2024
New client version	New client version 1.1.1403	December 13, 2024
New agent version	10-31-2024 version for AppStream 2.0 Agent	November 15, 2024
New agent version	10-21-2024 version for AppStream 2.0 Agent	October 30, 2024
Web browser access version 2	Web browser access v2 is now the default experience.	October 4, 2024
Set my time zone automatically based on my device	Sync the time zone to match the time zone set on your device.	October 2, 2024
Heavy File Sync Mode for Home Folders	Enable Amazon Simple Storage Service Home Folders options for your organization	October 1, 2024
Application settings updates	Various updates to application settings	September 30, 2024
Cookie-based authentication	Take proactive measures to prevent cookie theft	September 25, 2024
Branding update	Choose your organization logo or favicon from your Amazon S3 buckets	September 23, 2024
New client version	New client version 1.1.1360	August 1, 2024
Red Hat Enterprise Linux base images	Added Red Hat Enterprise Linux base images	July 30, 2024
New client version	New client version 1.1.1332	July 3, 2024

Latest Windows base images	Windows base images released on 06-17-2024	July 2, 2024
New client version	New client version 1.1.1326	June 17, 2024
New agent version	05-21-2024 version for AppStream 2.0 Agent	May 30, 2024
Multi-session use case	Audio conferencing added to multi-session use cases	May 30, 2024
Latest base image	Base image released on 05-08-2024	May 30, 2024
Managed WorkSpaces Applications image updates	Managed WorkSpaces Applications image updates released on April 25, 2024	May 15, 2024
New agent version	04-15-2024 version for AppStream 2.0 Agent	April 26, 2024
Latest base image	Base image released on 03-24-2024	April 26, 2024
Cross-account PCA sharing	Grant permissions for other accounts to use a centralized CA	April 25, 2024
New client version	Client version 1.1.1303	April 4, 2024
New relay state region endpoints	Relay state endpoints for Windows client application version 1.1.1300 and later	April 1, 2024
New client version	Client version 1.1.1300	March 28, 2024
Latest base image	Base image released on 01-26-2024	February 16, 2024

Use session scripts on multi-session fleets	When using session scripts on multi-session fleets, there are additional requirements and considerations to ensure optimal performance and security.	February 15, 2024
New agent version	01-17-2024 version for AppStream 2.0 Agent	February 15, 2024
Clipboard update	You can choose **Copy to local device character limit** or **Paste to remote session character limit** or both to limit the amount of data that users can copy or paste when using the clipboard	February 15, 2024
New client version	Client version 1.1.1259	February 8, 2024
New client version	Client version 1.1.1246	January 18, 2024
Windows Server 2022 support	Added support for Windows Server 2022 Base	December 14, 2023
Web browser access version 2	AppStream 2.0 web browser access version 2 offers an enhanced end user experienc e, including menu options that are easily discoverable and textual guidance for end users	December 11, 2023
New client version	Client version 1.1.1228	November 1, 2023

Multi-session fleets	A multi-session fleet enables you to provision more than one user session on a single fleet instance	October 26, 2023
New agent version	08-22-2023 version for AppStream 2.0 Agent	August 25, 2023
New Instance Family	New Graphics G5 Instance Family	July 26, 2023
New Agent version	06-11-2023 version for AppStream 2.0 Agent	July 25, 2023
App block builder	An app block builder is a reusable resource that you can use to package your applications (or app block)	June 29, 2023
New client version	Client version 1.1.1183	June 26, 2023
New Agent version	05-30-2023 version for AppStream 2.0 Agent	June 15, 2023
Extension SDK Developer Guide	Amazon WorkSpaces Applications uses Amazon DCV technology to provide secure, high-performance access to your applications	May 26, 2023
New Agent version	05-08-2023 version for AppStream 2.0 Agent	May 12, 2023
New client version	Client version 1.1.1159	May 9, 2023
New Agent version	04-13-2023 version for AppStream 2.0 Agent	April 25, 2023

Latest base images	Latest Linux base images released on 03-15-2023	April 5, 2023
Documentation update: New region — AWS GovCloud (US-East)	Updated Amazon WorkSpaces Applications in the *AWS GovCloud (US) User Guide* and updated the relay state endpoint table in "Setting Up SAML" and other content as needed.	April 5, 2023
New Agent version	03-21-2023 version for AppStream 2.0 Agent	April 3, 2023
New region support	South America (São Paulo) is now supported	December 15, 2022
New client version	Client version 1.1.1118	November 7, 2022
Client features table	Added a table that compares the features that are supported by the different access types.	November 7, 2022
Certificate-based authentication	You can use certificate-based authentication with WorkSpaces Applications fleets joined to Microsoft Active Directory.	October 31, 2022
Latest base images	Latest Linux base images released on 10-05-2022	October 27, 2022
New agent version	Agent version 10-13-2022	October 24, 2022
New client version	Client version 1.1.1099	October 13, 2022
Webcam support for Linux	Enable and disable webcam for Linux-based images	October 5, 2022

Latest base images	Latest Linux base images released on 09-21-2022	October 3, 2022
New client version	Client version 1.1.1066	August 17, 2022
Latest base images	Latest Windows base images released on 07-12-2022	July 21, 2022
New Agent version	06-20-2022 version for AppStream 2.0 Agent	June 30, 2022
New client version	Client version 1.1.421	June 29, 2022
New region support	US East (Ohio) is now supported	June 28, 2022
New default quotas	Max concurrent sessions for Elastic fleets	May 31, 2022
New client version	Client version 1.1.414	April 27, 2022
Japanese support	Enable Japanese support for a Linux image	April 19, 2022
Session scripts for Elastic fleets	Configure and specify session scripts for Elastic fleets	April 14, 2022
Canada (Central) support	Canada (Central) region now supported	March 31, 2022
New Agent version	03-14-2022 version for AppStream 2.0 Agent	March 25, 2022
New agent version	Agent version 03-14-2022	March 19, 2022
Latest base images	Latest Windows base images released on 03-03-2022	March 14, 2022
Managed image update	Released base image 02-18-2022	March 3, 2022

New client version	Client version 1.1.398	February 23, 2022
New client version	Client version 1.1.394	February 8, 2022
Managed image update	AppStream 2.0 Agent version 12-20-2021	January 6, 2022
Attribute-based application entitlements	Application entitlements control access to specific applications within your WorkSpaces Applications stacks.	January 5, 2022
App blocks and applications	When using an Elastic fleet, you can create app blocks and applications.	November 19, 2021
Create Linux images	You can now create Linux-based Amazon WorkSpaces Applications images.	November 15, 2021
New agent version	Agent version 10-19-2021.	October 26, 2021
Managed image and base image updates	Various updates.	October 21, 2021
New client version	Client version 1.1.333	September 14, 2021
Managed image updates	Managed imaged updates for August 12, 2021.	August 23, 2021
New agent version	Agent version 08-02-2021.	August 18, 2021
WorkSpaces Applications client update: version 1.1.304	Upgrades the embedded Chromium browser to version 91	August 2, 2021
WorkSpaces Applications agent update	Resolves multiple issues	August 2, 2021

New client version	Client version 1.1.304.	August 2, 2021
Base image update	Base image update (07-19-2021).	July 23, 2021
WorkSpaces Applications agent update	Resolves multiple issues	July 1, 2021
Base image update	Base image update (06-01-2021).	June 10, 2021
Install AMD Driver on Graphics Design Instances	If you need to update the AMD driver on your Image Builder that is using a Graphics Design instance, you can either use the latest AppStream 2.0 Graphics Design base images, or download the AMD driver and install it on your Image Builder.	June 4, 2021
Documentation update: New Graphics Design base images	Added two new released images and updated other content as needed.	June 3, 2021
Documentation update: New WorkSpaces Applications agent version	Added the entry for the 05-17-2021 agent version and updated other content as needed.	May 26, 2021
WorkSpaces Applications agent update	Resolves multiple issues	May 17, 2021
WorkSpaces Applications client update: version 1.1.294	Resolves issues with SAML 2.0 authentication, client stability on Windows 7, and folder sharing on client reconnection	April 26, 2021

New client version	Client version 1.1.1154	April 25, 2021
Documentation update: Support for managed WorkSpaces Applications image updates	Created the *Update an Image by Using Managed WorkSpaces Applications Image Updates* section in "Administer Your Images" and updated other content as needed.	April 8, 2021
WorkSpaces Applications client update: version 1.1.285	Includes fixes that improve compatibility with antivirus software	March 8, 2021
WorkSpaces Applications agent update	Resolves multiple issues	March 4, 2021
New agent version	Agent version 02-21-2022	February 24, 2021
New client version	Client version 1.1.1130	February 9, 2021
Documentation update: Support for smart cards	Created the *Smart Cards* section in "System Requirements and Feature Support" and updated other content as needed.	January 12, 2021
WorkSpaces Applications agent update	Adds support for using smart cards for Windows sign in to streaming instances and in-session authentication	January 4, 2021
WorkSpaces Applications client update: version 1.1.257	Adds support for real-time audio video (AV) and smart card authentication and resolves an issue with Excel	December 28, 2020

WorkSpaces Applications base image update	Updates for Base, Graphics Design, Graphics G4dn, Graphics Pro: Includes Microsoft Windows updates up to December 9, 2020; AWS CLI version 1.18.138; and Amazon SSM Agent version 3.0.431.0	December 28, 2020
Documentation update: Support for real-time audio-video (AV)	Created the *Real-Time Audio-Video* section in "System Requirements and Feature Support" and updated other content as needed.	December 28, 2020
WorkSpaces Applications agent update	Resolves multiple issues	December 17, 2020
WorkSpaces Applications agent update	Resolves multiple issues	October 8, 2020
WorkSpaces Applications agent update	Resolves multiple issues	September 1, 2020
WorkSpaces Applications client update: version 1.1.195	Improves local drive and folder sharing when file redirection is used, and provides other enhancements and fixes.	August 18, 2020
Documentation update: Support for local printer redirection	Created the *Enable Local Printer Redirection* section in "System Requirements and Feature Support" and updated other content as needed.	August 7, 2020

WorkSpaces Applications agent update	Adds support for local printer redirection to the WorkSpaces Applications client and resolves multiple issues	July 30, 2020
WorkSpaces Applications base image update	Updates for Base, Graphics Design, Graphics G4dn, Graphics Pro: Includes Microsoft Windows updates up to June 9, 2020; AWS CLI version 1.18.86; and Amazon SSM Agent version 2.3.1319.0	July 16, 2020
WorkSpaces Applications client update: version 1.1.179	Adds support for local printer redirection, and provides other enhancements and fixes.	July 8, 2020
New Region: Mumbai	Updated the relay state endpoint table in "Setting Up SAML" and other content as needed.	July 8, 2020
Documentation update: Support for drawing tablets	Created the *Drawing Tablets* section in "System Requirements and Feature Support" and updated other content as needed.	June 26, 2020
WorkSpaces Applications agent update	Resolves multiple issues	May 27, 2020

WorkSpaces Applications client update: version 1.1.160	Resolves an issue that prevents the application catalog page from opening on a Windows PC that has .NET Framework version 4.7.1 or earlier installed; also resolves another intermittent issue	April 28, 2020
WorkSpaces Applications client update: version 1.1.156	Adds support for on-demand diagnostic log and minidump uploads, defining trusted subdomains for user connections in a DNS TXT record, and other enhancements	April 22, 2020
Documentation update: Support for on-demand logging	Created the *Automatic and On-Demand Diagnostic Log Uploads* section in "System Requirements and Feature Support" and updated other content as needed.	April 22, 2020
Documentation update: Support for defining trusted subdomains for user connections in a DNS TXT record	Created the *Create the AS2TrustedDomains DNS TXT Record to Enable Your Domain for the WorkSpaces Applications Client Without Registry Changes* section in "System Requirements and Feature Support" and updated other content as needed.	April 22, 2020
WorkSpaces Applications agent update	Resolves an issue that causes streaming sessions to fail; improves performance with IAM roles	April 20, 2020

WorkSpaces Applications base image update	Updates for Base, Graphics Design, Graphics G4dn, Graphics Pro (Windows Server 2019): Includes Microsoft Windows updates up to March 10, 2020; AWS CLI version 1.18.21; and Amazon SSM Agent version 2.3.930.0	April 18, 2020
WorkSpaces Applications base image update	Updates for Base, Graphics Design, Graphics Pro: Includes Microsoft Windows updates up to February 11, 2020; AWS CLI version 1.17.5; and Amazon SSM Agent version 2.3.842.0	March 18, 2020
WorkSpaces Applications base image update	Adds support for Graphics G4dn instances (Windows Server 2012 R2); also includes Microsoft Windows updates up to February 11, 2020; AWS CLI version 1.17.5; and Amazon SSM Agent version 2.3.842.0	March 16, 2020
WorkSpaces Applications client update: version 1.1.137	Reverts the updates in version 1.1.136	March 8, 2020
WorkSpaces Applications client update: version 1.1.136	Adds support for defining trusted subdomains for user connections in a DNS TXT record, and provides other enhancements and fixes	March 5, 2020

WorkSpaces Applications base image update	Adds support for Graphics g4dn instances (Windows Server 2016, Windows Server 2019); also includes Microsoft Windows updates up to February 11, 2020; AWS CLI version 1.17.5; and Amazon SSM Agent version 2.3.842.0	March 5, 2020
WorkSpaces Applications client update: version 1.1.129	Adds support for native application mode and provides other enhancements and fixes	February 28, 2020
Documentation update: Support for native application mode	Created the *Native applicati on mode* section in "System Requirements and Feature Support" and updated other content as needed.	February 28, 2020
WorkSpaces Applications agent update	Adds support for native application mode and the **Desktop** stream view	February 19, 2020
WorkSpaces Applications base image update	Updates for Graphics Design: Adds support for Windows Server 2019, with Microsoft Windows updates up to November 12, 2019	January 13, 2020
WorkSpaces Applications agent update	Resolves multiple issues	January 13, 2020
Documentation update: Enhanced the documentation for security in WorkSpaces Applications	Created "Security in Amazon WorkSpaces Applications" and updated other content as needed.	December 23, 2019

WorkSpaces Applications client update: version 1.0.525	Resolves a DPI issue that causes the mouse cursor to point to the wrong location when a user clicks on an application during a streaming session	December 12, 2019
WorkSpaces Applications base image update	Includes Microsoft Windows updates up to December 12, 2019; AWS CLI version 1.16.284; and Amazon SSM Agent version 2.3.760.0	December 12, 2019
Documentation update: Support for WorkSpaces Applications z1d-based instances	Updated "WorkSpaces Applications Instance Families"	November 21, 2019
WorkSpaces Applications agent update	WorkSpaces Applications assemblies are now signed, including executables and installer packages	November 13, 2019
Documentation update: Embedded WorkSpaces Applications streaming sessions	Created "Embed WorkSpaces Applications Streaming Sessions" and updated other content as needed.	November 1, 2019
WorkSpaces Applications client update: version 1.0.511	Adds support for up to 4 monitors and provides other enhancements	October 16, 2019

Documentation update: New region — AWS GovCloud (US-West)	Created Amazon WorkSpaces Applications in the *AWS GovCloud (US) User Guide* and updated the relay state endpoint table in "Setting Up SAML" and other content as needed.	October 9, 2019
WorkSpaces Applications agent update	Modifies the WorkSpaces Applications storage connector to no longer bypass the system proxy server	October 8, 2019
Documentation update: FIPS-compliant endpoints	Created "Protecting Data in Transit with FIPS Endpoints" and updated other content as needed.	October 7, 2019
WorkSpaces Applications client update: version 1.0.499	Resolves issues with client-side hardware rendering and with the client not working correctly when Bluetooth headsets are connected to the local computer	September 26, 2019
WorkSpaces Applications agent update	Resolves multiple issues	September 23, 2019

WorkSpaces Applications base image update	Updates for all Base and Graphics Pro instances, and for Graphics Design Windows Server 2012 R2 instances: Includes Microsoft Windows updates up to August 13th, 2019 and AWS CLI version 1.16.222. Additional updates for Base, Graphics Design, and Graphics Pro instances	September 18, 2019
Documentation update: Support for applying IAM roles to WorkSpaces Applications streaming instances	Created "Using an IAM Role to Grant Permissions to Applications and Scripts Running on WorkSpaces Applications Streaming Instances" and updated other content as needed.	September 9, 2019
WorkSpaces Applications base image update	Updates for Graphics Design instances: Includes Microsoft Windows updates up to August 13th, 2019; AWS CLI version 1.16.222; and AMD driver 24.20.13028.3002	September 5, 2019
WorkSpaces Applications agent update	Adds support for applying IAM roles to WorkSpaces Applications streaming instances	September 3, 2019
Documentation update: WorkSpaces Applications file system redirection	Created "Enable File System Redirection for Your WorkSpaces Applications Users" and updated other content as needed.	August 20, 2019

Documentation update: Interface VPC endpoints	Created "Creating and Streaming From Interface VPC Endpoints." Also created "Access WorkSpaces Applications API Operations and CLI Commands Through an Interface VPC Endpoint" and updated other content as needed.	August 19, 2019
WorkSpaces Applications client update: version 1.0.480	Adds support for WorkSpaces Applications file system redirection	August 14, 2019
WorkSpaces Applications agent update	Adds support for WorkSpaces Applications file system redirection	August 8, 2019
Documentation update: Programmatic WorkSpaces Applications image creation	Created "Create Your WorkSpaces Applications Image Programmatically" and updated other content as needed.	August 1, 2019
WorkSpaces Applications agent update	Adds support for creating WorkSpaces Applications images programmatically	July 26, 2019
Documentation update: Support for Windows Server 2016 and Windows Server 2019 base images	Updated "WorkSpaces Applications Base Image Version History" and other content as needed.	June 28, 2019
WorkSpaces Applications agent update	Adds support for Windows Server 2016 and Windows Server 2019 base images	June 19, 2019

WorkSpaces Applications base image update	Adds support for Windows Server 2016 and Windows Server 2019	June 10, 2019
WorkSpaces Applications base image update	Includes Microsoft Windows updates up to May 14, 2019	May 28, 2019
Documentation update: WorkSpaces Applications Usage reports	Created "WorkSpaces Applications Usage Reports" and updated other content as needed.	May 21, 2019
Documentation update: Support for disconnecting idle users	Updated "Create a Fleet" in "Create an WorkSpaces Applications Fleet and Stack."	May 17, 2019
WorkSpaces Applications client update: version 1.0.407	Adds support for configuring the amount of time that users can be idle (inactive) before they are disconnected from their streaming session	May 16, 2019
WorkSpaces Applications agent update	Adds support for configuring the amount of time that users can be idle (inactive) before they are disconnected from their streaming session. Also adds support for subscribing to WorkSpaces Applications usage reports.	May 7, 2019
WorkSpaces Applications base image update	Includes Microsoft Windows updates up to April 20, 2019; AWS CLI version 1.16.126; and NVIDIA Graphics Driver 412.16 for Graphics Pro instances	April 29, 2019

Documentation update: Logging WorkSpaces Applications API calls with AWS CloudTrail	Created "Logging WorkSpaces Applications API Calls with AWS CloudTrail."	April 25, 2019
Documentation update: HIPAA compliance	Created "WorkSpaces Applications Compliance."	March 28, 2019
Documentation update: Gesture support	Created "Touchscreen Devices" and updated other content as needed.	March 13, 2019
WorkSpaces Applications client update: version 1.0.375	Adds touch screen support on Windows PCs and support for: Automatically connecting USB devices when a new streaming session starts, running session scripts, and delivering virtualized applications using the WorkSpaces Applications dynamic application framework APIs	March 7, 2019
WorkSpaces Applications agent update	Adds support for gestures on touch-enabled iPads, Android tablets, and Windows devices	March 7, 2019
New Region: Seoul	Updated the relay state endpoint table in "Setting Up SAML" and other content as needed.	February 13, 2019
Documentation update: Session scripts	Created "Use Session Scripts to Manage Your WorkSpaces Applications Users' Streaming Experience" in "Images" and updated other content as needed.	January 27, 2019

WorkSpaces Applications base image update	Includes Microsoft Windows updates up to December 10, 2018; AWS CLI version 1.16.84; and NVIDIA Graphics Driver 391.58 for Graphics Pro instances	January 22, 2019
WorkSpaces Applications agent update	Adds support for using on-instance session scripts. Also adds support for adding tags to the following WorkSpaces Applications resource types during resource creation: image builders, images, fleets, and stacks.	January 22, 2019
WorkSpaces Applications client update: version 1.0.320	Adds support for WorkSpaces Applications dynamic application framework APIs, WorkSpaces Applications regional settings, the WorkSpaces Applications user pool, and provides other enhancements	January 19, 2019
Documentation update: Default regional settings	Created "Configure Default Regional Settings for Your WorkSpaces Applications Users" and updated other content as needed.	December 13, 2018
Documentation update: Dynamic application framework	Created "Manage App Entitlement" and updated other content as needed.	December 7, 2018

WorkSpaces Applications agent update	Adds support for using the WorkSpaces Applications dynamic application framework to build a dynamic app provider. Also adds support for using a Japanese keyboard with web clients that run Windows.	December 4, 2018
WorkSpaces Applications client update: version 1.0.247	Initial release	November 20, 2018
Documentation update: WorkSpaces Applications client	Created "The WorkSpaces Applications Client" and updated other content as needed.	November 20, 2018
WorkSpaces Applications agent update	Adds support for launching streaming sessions using the WorkSpaces Applications Windows client	November 14, 2018
Documentation update: Image sharing	Created "Administer Your WorkSpaces Applications Images" and updated other content as needed.	September 14, 2018
Documentation update: Application settings persistence	Created "Enable Application Settings Persistence for Your WorkSpaces Applications Users" and updated other content as needed.	September 5, 2018
WorkSpaces Applications agent update	Adds support for application settings persistence	August 29, 2018

Documentation update: OneDrive support	Created "Enable and Administer OneDrive for Your WorkSpaces Applications Users" and updated other content as needed.	July 31, 2018
WorkSpaces Applications agent update	Adds support for OneDrive persistent storage	July 26, 2018
WorkSpaces Applications agent update	Resolves an issue with optimizing images for application launch.	June 19, 2018
Documentation update: Regional settings	Created "Enable Regional Settings for Your WorkSpace s Applications Users" and updated other content as needed.	June 14, 2018
Documentation update: Default application and Windows settings	Added "Step 4: Create Default Application and Windows Settings" to "Tutorial: Create a Custom Image" and updated other content as needed.	June 14, 2018
WorkSpaces Applications base image update	Includes Microsoft Windows updates up to May 9, 2018 and Windows PowerShell 5.1.	June 12, 2018
WorkSpaces Applications agent update	Adds support for regional settings and default applicati on and Windows settings.	June 6, 2018
Documentation update: Google Drive support	Created "Enable and Administer Google Drive for Your WorkSpaces Applicati ons Users" and updated other content as needed.	June 4, 2018

Earlier Updates

The following table describes important additions to the WorkSpaces Applications service and the *Amazon WorkSpaces Applications Administration Guide* documentation before June 4, 2018.

> **Note**
>
> Individual WorkSpaces Applications base image and WorkSpaces Applications agent releases are not included in this table. For information about these releases, see WorkSpaces Applications Base Image and Managed Image Update Release Notes and WorkSpaces Applications Agent Release Notes.

Change	Description	Date
Administrative controls for data transfer	Updated the "Create a Stack" section in "Create WorkSpaces Applications Fleets and Stacks" and updated other content as needed	May 24, 2018
New region: Frankfurt	Updated the relay state endpoint table in "Setting Up SAML" and other content as needed	March 28, 2018
Custom branding	Created "Add Your Custom Branding to Amazon WorkSpaces Applications" and updated other content as needed.	March 26, 2018
Image copy	Updated "Tutorial: Create a Custom Image" and updated other content as needed.	February 23, 2018
New regions: Singapore and Sidney	Updated the relay state endpoint table in "Setting Up SAML" and other content as needed.	January 24, 2018
Resource tagging	Created "Tagging Your Amazon Amazon WorkSpaces Applications Resources" and updated other content as needed.	December 15, 2017

Change	Description	Date
Managed WorkSpaces Applications agent updates	Created "Amazon WorkSpaces Applications Agent Version History" and updated other content as needed.	December 7, 2017
On-Demand fleets	Added "Fleet Type" section to “Amazon WorkSpaces Applications Fleets and Stacks” and updated other content as needed.	September 19, 2017
Instance types	Created "Amazon WorkSpaces Applications Instance Families" and updated other content as needed.	July 25, 2017
Active Directory	Created "Using Active Directory with Amazon WorkSpaces Applications" and updated other content as needed.	July 24, 2017
User pool	Created "Manage Access Using the WorkSpaces Applications User Pool" and updated other content as needed.	June 15, 2017
Security groups	Added "Security Groups" section to "Network Settings for Amazon WorkSpaces Applications" and updated other content as needed.	May 26, 2017
Home folders	Created "Enable and Administer Home Folders for Your WorkSpaces Applications Users" and updated other content as needed.	May 18, 2017
Default internet access	Created "Network Settings for Amazon WorkSpaces Applications" and updated other content as needed.	April 21, 2017
Fleet automatic scaling	Created "Fleet Auto Scaling for Amazon WorkSpaces Applications" and updated other content as needed.	March 23, 2017

Change	Description	Date
Fleet management	Created "Amazon WorkSpaces Applications Fleets and Stacks" and updated other content as needed.	February 22, 2017
SAML 2.0 support	Created "Single Sign-on Access to WorkSpaces Applications Using SAML 2.0" and updated other content as needed.	February 15, 2017
Image builders	Created "WorkSpaces Applications Image Builders" and updated other content as needed.	January 19, 2017
Initial documentation release	Created the initial release of the Amazon WorkSpaces Applications Administration Guide.	December 01, 2016