aws

Developer Guide

AWS Backup



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

AWS Backup: Developer 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 AWS Backup? 1
Feature overview 1
Centralized backup management 1
Policy-based backup 1
Tag-based backup policies 2
Lifecycle management policies 2
Cross-Region backup 2
Cross-account management and cross-account backup 2
Auditing and reporting with AWS Backup Audit Manager 3
Incremental backups 3
Full AWS Backup management 4
Backup activity monitoring 4
Secure your data in backup vaults 5
AWS Backup and native service backups 5
Getting started 5
Supported AWS resources and applications 6
Pricing 7
AWS Backup feature availability 8
Features available for all supported resources 8
Feature availability by resource 8
Feature availability by AWS Region 14
Supported services by AWS Region 21
How it works 27
Working with supported AWS services 27
Working with Amazon S3 data 28
Working with VMware virtual machines 28
Working with Amazon DynamoDB 29
Working with Amazon FSx file systems 30
Working with Amazon EC2 30
Working with Amazon EFS 32
Working with Amazon EBS 32
Working with Amazon RDS and Aurora 33
Working with Aurora DSQL 34
Working with AWS BackInt 34

Working with AWS Storage Gateway ... 34
Working with Amazon DocumentDB ... 35
Working with Amazon Neptune ... 35
Working with Amazon Redshift and Amazon Redshift Serverless ... 35
Working with Amazon Timestream ... 35
Working with AWS Organizations ... 36
Working with CloudFormation ... 36
Working with AWS BackInt, AWS Systems Manager for SAP, and SAP HANA ... 36
Working with Amazon EKS ... 36
Working with Amazon GuardDuty ... 37
How AWS services back up their own resources ... 37
Metering, costs, and billing ... 38
AWS Backup pricing ... 7
AWS Backup billing ... 38
Cost allocation tags ... 39
AWS Backup Audit Manager pricing ... 39
Amazon Aurora pricing ... 39
Blogs, videos, tutorials, and other resources ... 39
Getting started ... 43
Prerequisites ... 43
Service Opt-in ... 44
Backup plans ... 46
Create a backup plan ... 47
Create backup plans using the AWS Backup console ... 47
Create backup plans using the AWS CLI ... 49
Backup plan options and configuration ... 50
CloudFormation templates for backup plans ... 59
Delete a backup plan ... 64
Update a backup plan ... 65
Understanding backup plan summary ... 65
Summary components ... 65
Scheduled backup runs preview ... 66
Feature compatibility validation ... 67
Accessing backup plan summaries ... 68
Opt in and assign resources ... 68
Opt in to services, then assign resources ... 68

Service opt-in ... 69
Resource assignment ... 69
Assign resources through console ... 70
Assign resources with AWS CLI ... 72
Assign resources with CloudFormation ... 80
Resource assignment quotas ... 83
Backup vaults ... 84
Backup vault creation and deletion ... 84
Required permissions ... 85
Creating a backup vault (console) ... 86
Creating a backup vault (programmatically) ... 87
Backup vault name ... 87
AWS KMS encryption key ... 87
Backup vault tags ... 87
Delete a vault ... 88
Logically air-gapped vault ... 88
Overview of logically air-gapped vaults ... 88
Use case for logically air-gapped vaults ... 90
Compare and contrast with a standard backup vault ... 91
Create a logically air-gapped vault ... 93
View logically air-gapped vault details ... 96
Creating backups in a logically air-gapped vault ... 97
Share a logically air-gapped vault ... 100
Restore a backup from a logically air-gapped vault ... 103
Delete a logically air-gapped vault ... 104
Additional programmatic options for logically air-gapped vaults ... 104
Understanding encryption key types for logically air-gapped vaults ... 105
Usage of service-owned key ... 113
Considerations for security auto-remediation ... 113
Troubleshoot a logically air-gapped vault issue ... 114
Primary backups to logically air-gapped vaults ... 117
Multi-party approval for logically air-gapped vaults ... 133
Vault access policies ... 149
Deny access to a resource type in a backup vault ... 150
Deny access to a backup vault ... 151
Deny access to delete recovery points in a backup vault ... 152

Vault Lock ... 154
Vault lock modes ... 155
Vault lock benefits ... 155
Lock a backup vault using the console ... 156
Lock a backup vault programmatically ... 157
Review a backup vault for its AWS Backup Vault Lock configuration ... 159
Vault lock removal during grace time (Compliance mode) ... 160
AWS account closure with a locked vault ... 161
Additional security considerations ... 161
Backup creation, maintenance, and restore ... 163
On-demand backups ... 164
Continuous backups and PITR ... 166
Continuous backup and PITR considerations ... 167
Supported services ... 168
Finding a continuous backup ... 175
Restoring a continuous backup ... 176
Stopping or deleting continuous backups ... 176
Copying continuous backups ... 177
Changing your retention period ... 178
Removing the only continuous backup rule from a backup plan ... 179
Backup creation by resource type ... 179
Creating automatic backups ... 179
Creating on-demand backups ... 179
Backup job statuses ... 180
Incremental backups ... 180
Access to source resources ... 181
CloudFormation stack backups ... 182
Amazon Aurora DSQL backups ... 188
Advanced DynamoDB backup ... 189
Amazon EBS backups ... 194
Amazon RDS backups ... 195
Amazon Redshift backups ... 200
Amazon Redshift Serverless backups ... 203
Amazon EKS ... 206
SAP HANA backup on Amazon EC2 ... 214
Amazon S3 backups ... 225

Amazon Timestream backups ... 242
Virtual machine backups ... 245
Create Windows VSS backups ... 287
Backup and tag copy ... 289
Copy job retry ... 291
Copy job concurrency ... 291
Cross-Region backup ... 292
Cross-account backup ... 295
Copy tags onto backups ... 309
Add or remove tags on existing recovery points ... 310
Required permissions ... 310
Add or remove tags using the console ... 310
Add or remove tags using the AWS CLI ... 311
Backup deletion ... 311
Deleting backups manually ... 313
Troubleshooting manual deletions ... 314
Backup and tag edits ... 315
Backup search ... 316
Overview ... 316
Use cases for backup indexes and search ... 316
Access ... 316
Process Flow ... 317
Backup indexes ... 318
Searches ... 321
Search results ... 324
Export search results to an S3 bucket ... 325
Cost considerations and best practices ... 326
Restore from search ... 326
Backup tiering ... 326
Overview ... 326
Tiering Configurations ... 327
Creating tiering configurations ... 328
Managing tiering configurations ... 330
How S3 backup tiering configurations apply ... 332
Cost structure and monitoring ... 332
Configuration Examples ... 333

Supported features and limitations 335
Troubleshooting 335
Access points 336
Overview 336
How backup access points work 337
Prerequisites 338
Creating a backup access point 339
Accessing backup data 340
Managing access points 341
Access point policies 342
Supported S3 operations 344
Recovery point lifecycle protection 345
Access points in logically air-gapped vaults 346
Monitoring and notifications 347
Supported features and limitations 347
Troubleshooting 348
Restore by resource type 349
How to restore 349
Non-destructive restores 350
Restore testing 350
Copy tags during a restore 351
Restore job statuses 354
Aurora restore 355
Aurora DSQL restore 359
CloudFormation restore 366
DocumentDB restore 368
DynamoDB restore 371
EBS restore 373
EC2 restore 379
EFS restore 384
EKS restore 390
FSx restore 408
Neptune restore 418
RDS restore 420
Redshift restore 422
Redshift Serverless restore 427

SAP HANA restore ... 431
S3 restore ... 440
Storage Gateway restore ... 446
Timestream restore ... 448
VM restore ... 451
Restore testing ... 458
Overview ... 459
Compare with restores ... 459
Plan management ... 461
Create testing plan ... 462
Update testing plan ... 467
View testing plans ... 468
View testing jobs ... 469
Delete plan ... 470
Audit testing ... 471
Quotas and parameters ... 471
Troubleshooting ... 472
Inferred metadata ... 474
Restore testing validation ... 485
Stop a backup job ... 488
View existing backups ... 488
Listing backups by protected resource in the console ... 488
Listing backups by backup vault in the console ... 489
Listing backups programmatically ... 489
AWS Backup Audit Manager ... 490
Working with audit frameworks ... 491
Choosing your controls ... 492
Turning on resource tracking ... 495
Creating frameworks using the AWS Backup console ... 501
Creating frameworks using the AWS Backup API ... 502
Viewing framework compliance status ... 517
Finding non-compliant resources ... 518
Updating audit frameworks ... 519
Deleting audit frameworks ... 519
Working with audit reports ... 519
Enhanced job report context ... 520

Choosing your report template ... 522
Creating report plans ... 533
Creating report plans using the AWS Backup API ... 538
Creating on-demand reports ... 540
Viewing audit reports ... 541
Updating report plans ... 542
Deleting report plans ... 542
Using CloudFormation to deploy AWS Backup Audit Manager resources ... 542
Turn on resource tracking ... 501
Deploy default controls ... 548
Exempt IAM roles from control evaluation ... 549
Create a report plan ... 550
Using AWS Backup Audit Manager with AWS Audit Manager ... 551
Controls and remediation ... 551
Backup resources are included in at least one backup plan ... 552
Backup plan minimum frequency and minimum retention ... 553
Vaults prevent manual deletion of recovery points ... 553
Recovery points are encrypted ... 554
Minimum retention established for recovery point ... 554
Cross-Region backup copy is scheduled ... 555
Cross-account backup copy is scheduled ... 556
Resources are in a backup plan with an AWS Backup Vault Lock ... 556
Last recovery point was created ... 557
Restore time for resources meet target ... 558
Resources in a logically air-gapped vault ... 559
Manage multiple accounts with AWS Organizations ... 560
Cross-account management overview ... 560
Cross-account management settings ... 561
Creating a management account in Organizations ... 562
Enabling cross-account management ... 562
Backup policies ... 563
Delegated administrator ... 564
Prerequisites ... 565
Register a member account as a delegated administrator account ... 566
Deregister a member account ... 567
Delegate AWS Backup policies through AWS Organizations ... 568

Monitoring activities in multiple AWS accounts ... 568
Resource opt-in rules ... 569
Defining policies, policy syntax, and policy inheritance ... 569
AWS Backup and CloudFormation ... 570
In general ... 570
Deploying a backup vault, backup plan, and resource assignment with CloudFormation ... 570
Deploying backup plans with CloudFormation ... 570
Deploying AWS Backup Audit Manager frameworks and report plans with CloudFormation ... 571
Using CloudFormation with AWS Organizations ... 571
Learning more ... 571
Network ... 572
Endpoints ... 572
AWS PrivateLink ... 572
Considerations for Amazon VPC endpoints ... 572
Create an AWS Backup VPC endpoint ... 573
Use a VPC endpoint ... 574
Creating a VPC endpoint policy ... 574
Security ... 576
Compliance validation ... 577
Data protection ... 577
Encryption for backups in AWS Backup ... 578
Virtual machine hypervisor credential encryption ... 589
Identity and access management ... 591
Authentication ... 591
Access control ... 593
IAM service roles ... 602
Managed policies ... 605
Using service-linked roles ... 694
Cross-service confused deputy prevention ... 702
Infrastructure security ... 703
Integrity ... 704
AWS Backup data integrity goal ... 704
AWS Backup data integrity implementation ... 704
Objective confirmation and audit of AWS Backup data integrity ... 705
Legal holds ... 705
Legal hold overview ... 705

Multiple legal holds ... 706
Create a legal hold ... 706
View legal holds ... 707
View recovery points in legal holds ... 710
Release a legal hold ... 711
Malware protection ... 712
Integration with Amazon GuardDuty ... 713
How to use malware scanning ... 713
Access ... 713
Incremental vs full scans ... 714
Monitoring your malware scans ... 715
Understanding scan results ... 716
Troubleshooting scan failures ... 718
Metering ... 718
Quotas ... 718
Console and CLI usage steps for malware scan types ... 719
Resilience ... 723
Quotas ... 724
Backup ... 724
Backup index and search quotas ... 726
Policy quotas ... 727
Malware scanning quotas ... 727
Amazon Timestream resource quotas ... 728
AWS Backup Audit Manager quotas ... 728
Restore testing plan quotas ... 729
AWS Backup gateway quotas ... 731
Amazon EKS quotas ... 731
Logically air-gapped vault quotas ... 732
Related quotas ... 732
Monitoring AWS Backup ... 734
Console dashboards ... 734
Overview ... 735
Jobs dashboard ... 735
Problematic reasons ... 736
Dashboard data with AWS CLI ... 741
Monitoring events using EventBridge ... 742

Backup Job events ... 743
Backup Plan events ... 748
Backup Vault events ... 750
Copy Job events ... 752
Recovery Point events ... 755
Region Settings events ... 757
Restore Job events ... 757
Recovery point indexing events ... 761
Malware scan Job events ... 763
Backup access point events ... 770
AWS Backup metrics with Amazon CloudWatch ... 772
CloudWatch Dashboard ... 772
Metrics with CloudWatch ... 774
Logging AWS Backup API calls with CloudTrail ... 778
AWS Backup events in CloudTrail ... 779
Understanding AWS Backup log file entries ... 780
Logging cross-account management events ... 784
Notification options with AWS Backup ... 788
User Notifications and AWS Backup ... 789
Amazon SNS and AWS Backup events ... 789
Troubleshooting AWS Backup ... 802
Troubleshooting general issues ... 802
Troubleshoot creating resources ... 803
Troubleshooting deleting resources ... 804
Troubleshooting restoring resources ... 804
Troubleshooting formatting errors ... 805
Common troubleshooting questions ... 805
Additional important links ... 806
AWS Backup API ... 808
FAQs on supported protocols ... 809
What are the AWS JSON and Smithy RPC v2 CBOR protocols, and how do they differ? ... 809
How do I get started with the Smithy RPC v2 CBOR protocol? ... 809
What are the risks of enabling the Smithy RPC v2 CBOR protocol? ... 810
What if my solution does not support Smithy RPC v2 CBOR? ... 810
What AWS Regions support the Smithy RPC v2 CBOR protocol? ... 810

What latency improvements can I expect when upgrading to SDK versions that support Smithy RPC v2 CBOR? 810
Will the AWS JSON protocol be deprecated? 810
Where can I find more information about Smithy RPC v2 CBOR? 811
Document history 812

What is AWS Backup?

AWS Backup is a fully-managed service that makes it easy to centralize and automate data protection across AWS services, in the cloud, and on premises. Using this service, you can configure backup policies and monitor activity for your AWS resources in one place. It allows you to automate and consolidate backup tasks that were previously performed service-by-service, and removes the need to create custom scripts and manual processes. With a few clicks in the AWS Backup console, you can automate your data protection policies and schedules.

AWS Backup does not govern backups you take in your AWS environment outside of AWS Backup. Therefore, if you want a centralized, end-to-end solution for business and regulatory compliance requirements, start using AWS Backup today.

Feature overview

AWS Backup provides many features and capabilities, including the following.

Centralized backup management

AWS Backup provides a centralized backup console, a set of backup APIs, and the AWS Command Line Interface (AWS CLI) to manage backups across the AWS services that your applications use. With AWS Backup, you can centrally manage backup policies that meet your backup requirements. You can then apply them to your AWS resources across AWS services, enabling you to back up your application data in a consistent and compliant manner. The AWS Backup centralized backup console offers a consolidated view of your backups and backup activity logs, making it easier to audit your backups and ensure compliance.

Policy-based backup

With AWS Backup, you can create backup policies known as *backup plans*. Use these backup plans to define your backup requirements and then apply them to the AWS resources that you want to protect across the AWS services that you use. You can create separate backup plans that each meet specific business and regulatory compliance requirements. This helps ensure that each AWS resource is backed up according to your requirements. Backup plans make it easy to enforce your backup strategy across your organization and across your applications in a scalable manner.

For all the configuration options for backup plans, see Backup plan options and configuration.

Tag-based backup policies

You can use AWS Backup to apply backup plans to your AWS resources in a wide variety of ways, including tagging them. Tagging makes it easier to implement your backup strategy across all your applications and to ensure that all your AWS resources are backed up and protected. AWS tags are a great way to organize and classify your AWS resources. Integration with AWS tags enables you to quickly apply a backup plan to a group of AWS resources, so that they are backed up in a consistent and compliant manner.

For all the ways you can assign your resources to backup plans, see Select AWS services to backup.

Lifecycle management policies

AWS Backup enables you to meet compliance requirements while minimizing backup storage costs by storing backups in a low-cost cold storage tier. You can configure lifecycle policies that automatically transition backups from warm storage to cold storage according to a schedule that you define.

For a list of resources which can be transitioned to cold storage, see Feature availability by resource. For steps to turn on cold storage in your backup plan, see Lifecycle and storage tiers.

Cross-Region backup

Using AWS Backup, you can copy backups to multiple different AWS Regions on demand or automatically as part of a scheduled backup plan. Cross-Region backup is particularly valuable if you have business continuity or compliance requirements to store backups a minimum distance away from your production data. For more information, see Creating backup copies across AWS Regions.

Cross-account management and cross-account backup

You can use AWS Backup to manage your backups across all AWS accounts inside your AWS Organizations structure. With cross-account management, you can automatically use backup policies to apply backup plans across the AWS accounts within your organization. This makes compliance and data protection efficient at scale and reduces operational overhead. It also helps eliminate manually duplicating backup plans across individual accounts. For more information, see Managing AWS Backup resources across multiple AWS accounts.

You can also copy backups to multiple different AWS accounts inside your AWS Organizations management structure. This way, you can "fan in" backups to a single repository account, then "fan out" backups for greater resilience. Creating backup copies across AWS accounts.

Before you can use the cross-account management and cross-account backup features, you must have an existing organization structure configured in AWS Organizations. An *organizational unit* (OU) is a group of accounts that can be managed as a single entity. AWS Organizations is a list of accounts that can be grouped into organizational units and managed as a single entity.

Auditing and reporting with AWS Backup Audit Manager

AWS Backup Audit Manager helps you simplify data governance and compliance management of your backups across AWS. AWS Backup Audit Manager provides built-in, customizable controls that you can align with your organizational requirements. You can also use these controls to automatically track your backup activities and resources.

AWS Backup Audit Manager can help you locate specific activities and resources that are not yet compliant with the controls that you defined. It also generates daily reports that you can use to demonstrate evidence of compliance with your controls over time.

To include your backup compliance alongside your overall compliance posture, you can automatically import AWS Backup Audit Manager findings into AWS Audit Manager.

Incremental backups

AWS Backup efficiently stores your periodic backups incrementally for supported resource types. The first backup of an AWS resource backs up a full copy of your data. For each successive incremental backup, only the changes to your AWS resources are backed up. Incremental backups enable you to benefit from the data protection of frequent backups while minimizing storage costs.

> **Important**
>
> Not all resource types support incremental backups. For resources that do not support incremental backups, each backup is a full copy, which can result in higher storage costs. For a list of which resources support incremental backups, see Feature availability by resource.

For more information on behaviors in vaults, see Incremental backups.

Full AWS Backup management

Some resource types support full AWS Backup management. The benefits of full AWS Backup management include:

- **Independent encryption**. AWS Backup automatically encrypts your backups with the KMS key of your AWS Backup vault, instead of using the same encryption key as your source resource. This increases your layers of defense. See Encryption for backups in AWS Backup for more information.
- **awsbackup Amazon Resource Names (ARNs)**. Backup ARNs begin with `arn:aws:backup` instead of `arn:aws:source-resource`. This allows you to create access policies that apply specifically to backups and not the source resources. See Access control for more information.
- **Centralized backup billing and Cost Explorer cost allocation tags.**. Charges for AWS Backup (including storage, data transfers, restores, and early deletion) appear under "Backup" in your Amazon Web Services bill, instead of appearing under each supported resource. You can also use Cost Explorer cost allocation tags to track and optimize your backup costs. See Metering, costs, and billing for AWS Backup for more information.

To see which resource types are eligible for full AWS Backup management, see Feature availability by resource.

Backup activity monitoring

AWS Backup provides a dashboard that makes it simple to audit backup and restore activity across AWS services. With just a few clicks on the AWS Backup console, you can view the status of recent backup jobs. You can also restore jobs across AWS services to ensure that your AWS resources are properly protected.

AWS Backup integrates with Amazon CloudWatch and Amazon EventBridge. CloudWatch allows you to track metrics and create alarms. EventBridge allows you to view and monitor AWS Backup events. For more information, see Monitoring AWS Backup events using EventBridge and Monitoring AWS Backup metrics with CloudWatch.

AWS Backup integrates with AWS CloudTrail. CloudTrail gives you a consolidated view of backup activity logs that make it quick and easy to audit how your resources are backed up. AWS Backup also integrates with Amazon Simple Notification Service (Amazon SNS), providing you with backup activity notifications, such as when a backup succeeds or a restore has been initiated. For more

information, see Logging AWS Backup API calls with CloudTrail and Amazon SNS and AWS Backup events.

Secure your data in backup vaults

The content of each AWS Backup backup is immutable, meaning that no one can alter that content. AWS Backup further secures your backups in backup vaults, which separates them safely from their source instances. For example, your vault will retain your Amazon EC2 and Amazon EBS backups according to the lifecycle policy you choose, even if you delete the source Amazon EC2 instance and Amazon EBS volumes.

Backup vaults offer encryption and resource-based access policies that let you define who has access to your backups. You can define access policies for a backup vault that define who has access to the backups within that vault and what actions they can take. This provides a simple and secure way to control access to your backups across AWS services. To review AWS and customer managed policies for AWS Backup, see Managed policies for AWS Backup.

You can use AWS Backup Vault Lock to prevent anyone (including you) from deleting backups or altering their retention period. AWS Backup Vault Lock helps you enforce a *write-once-read-many* (WORM) model and add another layer of defense to your defense in depth. To get started, see AWS Backup Vault Lock.

AWS Backup and native service backups

The relationship between AWS Backup and native service backup mechanisms varies by service. In some cases, AWS Backup uses the same underlying backup infrastructure as the native service (for example, Amazon RDS continuous backups). In other cases, AWS Backup operates as a separate backup layer (for example, DynamoDB). Understanding this relationship helps you avoid unintended duplicate costs. Consult the service-specific pages in this guide for details on how AWS Backup interacts with each service's native backup capabilities.

Getting started

To learn more about AWS Backup, we recommend that you start with Getting started with AWS Backup.

Supported AWS resources and applications

The following are AWS resources and third-party applications that you can back up and restore using AWS Backup. For more information, see the section called “AWS Backup feature availability”.

Service	Supported resource types
Amazon Elastic Compute Cloud (Amazon EC2)	Amazon EC2 instances backed by Amazon EBS volumes
Amazon Simple Storage Service (Amazon S3)	Amazon S3 data
Amazon Elastic Block Store (Amazon EBS)	Amazon EBS volumes
Amazon DynamoDB	Amazon DynamoDB tables
Amazon Relational Database Service (Amazon RDS)	Amazon RDS database instances (including all database engines); Multi-Availability Zone clusters
Amazon Aurora DSQL	Single-Region clusters and multi-Region clusters
Amazon Aurora	Aurora clusters
Amazon Elastic File System (Amazon EFS)	Amazon EFS file systems
FSx for Lustre	FSx for Lustre file systems
FSx for Windows File Server	FSx for Windows File Server file systems
Amazon FSx for NetApp ONTAP	FSx for ONTAP file systems
Amazon FSx for OpenZFS	FSx for OpenZFS file systems

Service	Supported resource types
AWS Storage Gateway (Volume Gateway)	AWS Storage Gateway volumes
Amazon DocumentDB	Amazon DocumentDB instance-based clusters
Amazon Neptune	Amazon Neptune clusters
Amazon Redshift	Amazon Redshift clusters
Amazon Redshift Serverless	Amazon Redshift namespaces
Amazon Timestream for LiveAnalytics	Amazon Timestream for LiveAnalytics tables. (AWS Backup does not support Amazon Timestream for InfluxDB. For more information, see Amazon Timestream backups.)
VMware Cloud™ on AWS	VMware Cloud™ virtual machines on AWS
VMware Cloud™ on AWS Outposts	VMware Cloud™ virtual machines on AWS Outposts
CloudFormation	CloudFormation stacks
SAP HANA databases	SAP HANA databases on Amazon EC2 instances
Amazon Elastic Kubernetes Service (Amazon EKS)	Amazon EKS clusters and persistent storage backups

Pricing

With AWS Backup, you pay for backup storage, data restored, restore testing, cross-Region data transfer, and AWS Backup Audit Manager. For more information, see AWS Backup Pricing.

AWS Backup feature availability

AWS Backup features are offered according to resource and AWS Region. The following sections and tables can help you determine feature availability.

Contents

- Features available for all supported resources
- Feature availability by resource
- Feature availability by AWS Region
- Supported services by AWS Region

Features available for all supported resources

AWS Backup offers the following features for its supported AWS services, as well as for supported third-party applications. Support for a feature or service should not be assumed unless explicitly mentioned.

- Automated backup schedules and retention management
- Centralized backup monitoring
- Encrypted backups
- Incremental backups
- Cross-account management with AWS Organizations
- Automated backup audits and reports with AWS Backup Audit Manager
- Write-once, read-many (WORM) with AWS Backup Vault Lock

Feature availability by resource

To use AWS Backup with a supported AWS service in a particular Region, the service must be available in the Region. To determine service availability in a Region, view the service endpoints in the *AWS General Reference*.

For information on opt-in Regions and what resources and features are supported within, see Feature availability by AWS Region.

Note

If a resource type does not have a checkmark in the Cross-Region backup or Cross-account backup columns, then copy operations for that resource type are not supported in any scenario, including same-Region and same-account copies to a different vault. If the underlying AWS service provides its own native copy or replication feature, refer to that service's documentation.

AWS Backu supp	Cross Reg ion backu	Cross acc ount backu	AWS Backu Audit Mana	Incre al backu	Conti s backu and point in-time resto	Full mana t	Lifec to cold stora	Item-leve l resto [1]	Resto testir	Logic air-gapp d vault	Backu searc	Backu tierin	Malw Prote n	Access points
Amaz EC2	✓	✓	✓	✓					✓	✓			✓	
Amaz S3	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓
Amaz EBS	✓	✓	✓	✓			✓	✓	✓	✓ [14]	✓		✓	
Amaz RDS DB instar (singl AZ and Multi AZ) [12]	✓ [3]	✓ [3]	✓ [4]	✓	✓				✓					

AWS Backu supp	Cross Reg ion backu	Cross acc ount backu	AWS Backu Audit Mana	Incre al backu	Conti s backu and point in-time resto	Full mana t	Lifec to cold stora	Item-leve l resto [1]	Resto testir	Logic air-gapp d vault	Backu searc	Backu tierin	Malw Prote n	Access points
Amaz RDS cluste [13]	✓ [3]	✓ [3]	✓ [4]	✓					✓					
Amaz Auror	✓ [3]	✓ [3]	✓	✓ [6]	✓				✓	✓				
Amaz Auror DSQL	✓	✓				✓	✓			✓				
Amaz EFS	✓	✓	✓	✓		✓	✓	✓	✓	✓				
FSx for Lustr	✓	✓	✓	✓					✓	✓				
FSx for Wind File Serve	✓	✓	✓	✓					✓ [8]	✓				
FSx for ONTA	✓	✓	✓	✓					✓	✓				

AWS Backu supp	Cross Regi ion backu	Cross acco unt backu	AWS Backu Audit Mana	Incre al backu	Conti s backu and point in-time resto	Full mana t	Lifec to cold stora	Item-level resto [1]	Resto testir	Logic air-gapp d vault	Backu searc	Backu tierin	Malw Prote n	Access points
FSx for Open	✓	✓	✓	✓					✓	✓				
AWS Stora Gatev	✓	✓	✓	✓						✓				
Amaz Docu B	✓ [3]	✓ [3]							✓	✓				
Amaz Neptu	✓ [3]	✓ [3]							✓	✓ [9]				
Amaz Redsh								✓	✓					
Amaz Redsh Serve s								✓						
Amaz Time m	✓	✓		✓		✓	✓			✓				
Wind VSS	✓	✓	✓	✓						✓				

AWS Backu supp	Cross Reg ion backu	Cross acc ount backu	AWS Backu Audit Mana	Incre al backu	Conti s backu and point in-time resto	Full mana t	Lifec to cold stora	Item-leve l resto [1]	Resto testir	Logic air-gapp d vault	Backu searc	Backu tierin	Malw Prote n	Access points
Virtu mach	✓	✓	✓	✓		✓	✓	✓		✓				
AWS Cloud ation	✓ [5]		✓	✓ [5]			✓							
Amaz Dyna			✓						✓					
Dyna with AWS Backu adva featu	✓	✓	✓			✓	✓		✓	✓				
SAP HANA datab on Amaz EC2 insta	✓	✓		✓ [6]	✓	✓	✓							
Amaz EKS[10]	✓	✓				✓		✓		✓[11]				

Some resource types have both continuous backup capability and cross-Region and cross-account copy available. When a cross-Region or cross-account copy of a continuous backup is made, the copied recovery point (backup) becomes a snapshot (periodic) backup. PITR (Point-in-Time Restore) is not available for these copies.

- Amazon RDS and Amazon S3 support cross-account and cross-Region copy from incremental backups. Amazon RDS also supports simultaneous cross-Region and cross-account snapshot copying in a single action.
- Amazon Aurora and SAP HANA on Amazon EC2 instances support cross-account and cross-Region copy from full backups. Amazon Aurora also supports simultaneous cross-Region and cross-account snapshot copying in a single action.

[1] The "item" in an item-level restore varies depending on the supported resource. For example, a file system item is a file or directory, whereas an S3 item is an S3 object. A VMware item is a disk. For more information, see the Restore a backup by resource type section for the supported resource.

[3] Amazon RDS, Aurora, DocumentDB, and Neptune now support cross-Region and cross-account snapshot copying in a single action. RDS multi availability zone (Multi-AZ) database instances can be copied, but Multi-AZ clusters do not currently support any copy operations. See Cross-Region copy considerations with specific resources for further information.

[4] See RDS multi-availability zone backups for Regions where Backup Audit Manager support is available.

[5] In CloudFormation stack backups, nested resources retain their source resource features. However, resources within the stack do not retain Point-in-Time Restore (PITR) functionality (such as Amazon S3 and Amazon RDS).

[6] Snapshots are full and incremental backup is offered through PITR.

[7] Amazon FSx for OpenZFS Multi-AZ (multi-availability zone) file systems can only be restored from the Amazon FSx console or the API request `createFileSystemFromBackup`.

[8] Is supported in a restore test if FSx for Windows File Server uses AWS managed active directory

[9] Is not currently available in Asia Pacific (Jakarta) Region

[10] Features for Persistent Storage Backups (e.g. Amazon EBS) taken as part of Amazon EKS Backups will reflect the feature support for the respective resource types in this table.

[11] Is not currently available in Middle East (Bahrain) and Middle East (UAE)

[12] AWS Backup supports continuous backups and point-in-time recovery for all Amazon RDS instances and Aurora that are supported by the native Amazon RDS service. Amazon RDS DB instances include both single-AZ and Multi-AZ instance deployments.

[13] AWS Backup does not support continuous backups or point-in-time recovery for Amazon RDS Multi-AZ clusters.

[14] Amazon EBS backups in a logically air-gapped vault can transition to the archive tier (Amazon EBS Snapshots Archive).

Feature availability by AWS Region

AWS Backup is available in all the following AWS Regions. AWS Backup features are available in all these Regions unless otherwise noted in the following table.

Some Regions require account opt-in, as noted in the following table. Some feature availability is determined by whether opt-in is required or not required. For more information, see AWS Regions your account can use in the *AWS Account Management Reference Guide*.

Considerations for opt-in Regions:

- **Cross-Region** copy of FSx for ONTAP, FSx for Lustre, FSx for Windows File Server, and FSx for OpenZFS is **not supported** in Middle East (Bahrain) and Middle East (UAE).
- **Cross-Region** and **Cross-account** copy of FSx for ONTAP is **not supported** in Asia Pacific (New Zealand), China (Beijing), and China (Ningxia).
- **Cross-account** copy is **not supported** for CloudFormation and Timestream in Regions where opt-in is required.

Considerations and limitations for cross-account management in opt-in Regions:

- Cross-account management in AWS Regions where opt-in is required includes cross-account monitoring and access to backup policies; delegated administrator accounts can launch policies but do not have access to the monitoring functions.
- Both management accounts and their child accounts can be opted into AWS Organizations. If a child account is opted into cross-account management prior to its management account being opted into cross-account management, there will be a delay (up to 24 hours) before cross-account monitoring will show job statuses across the organization.

AWS Backup suppor	Opt-in	Cross-Region backup copy	Cross-account management	Cross-account backup copy	AWS Backup Audit Manager and Jobs dashbo	Restore testing	Logical air-gapped vault	Backup search	Backup tiering	Malware Protection	Access points
US East (N. Virginia	Not require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
US East (Ohio)	Not require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
US West (N. Californ a)	Not require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
US West (Orego	Not require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Africa (Cape Town)	Require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Asia Pacific (Hong Kong)	Require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Asia Pacific	Require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

AWS Backup support	Opt-in	Cross-Region backup copy	Cross-account management	Cross-account backup copy	AWS Backup Audit Manager and Jobs dashboard	Restore testing	Logically air-gapped vault	Backup search	Backup tiering	Malware Protection	Access points
(Hyderabad)											
Asia Pacific (Jakarta)	Required	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Asia Pacific (Malaysia)	Required	✓	✓	✓		✓	✓	✓	✓	✓	
Asia Pacific (Melbourne)	Required	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Asia Pacific (Mumbai)	Not required	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Asia Pacific (New Zealand)	Required	✓		✓		✓	✓		✓		
Asia Pacific (Osaka)	Not required	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

AWS Backup suppor	Opt-in	Cross-Region backup copy	Cross-account management	Cross-account backup copy	AWS Backup Audit Manager and Jobs dashbo	Restore testing	Logical air-gapped vault	Backup search	Backup tiering	Malware Protection	Access points
Asia Pacific (Seoul)	Not require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Asia Pacific (Singapore)	Not require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Asia Pacific (Sydney)	Not require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Asia Pacific (Taipei)	Require	✓		✓		✓	✓		✓	✓	
Asia Pacific (Thailand)	Require	✓		✓		✓	✓	✓	✓	✓	
Asia Pacific (Tokyo)	Not require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Canada (Central)	Not require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

AWS Backup support	Opt-in	Cross-Region backup copy	Cross-account management	Cross-account backup copy	AWS Backup Audit Manager and Jobs dashboard	Restore testing	Logically air-gapped vault	Backup search	Backup tiering	Malware Protection	Access points
Canada West (Calgar	Require	✓	✓	✓		✓	✓	✓	✓	✓	
China (Beijing[1]	AWS in China	✓	✓	✓					✓		
China (Ningxi[1]	AWS in China	✓	✓	✓					✓		
Europe (Frankfurt)	Not require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Europe (Irelanc	Not require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Europe (Londo	Not require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Europe (Milan)	Require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Europe (Paris)	Not require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Europe (Spain)	Require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

AWS Backup suppor	Opt-in	Cross-Region backup copy	Cross-account management	Cross-account backup copy	AWS Backup Audit Manager and Jobs dashbo	Restore testing	Logical air-gapped vault	Backup search	Backup tiering	Malware Protection	Access points
Europe (Stockholm)	Not require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Europe (Zurich	Require	✓	✓	✓	✓	✓	✓	✓	✓	✓	
Israel (Tel Aviv)	Require	✓	✓	✓		✓	✓	✓	✓	✓	✓
Mexico (Centra	Require	✓		✓		✓	✓	✓	✓	✓	
Middle East (Bahrai	Require	✓	✓	✓	✓	✓	✓	✓	✓	✓	
Middle East (UAE)	Require	✓	✓	✓	✓	✓	✓	✓	✓	✓	
South Americ (São Paulo)	Not require	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

AWS Backup support	Opt-in	Cross-Region backup copy	Cross-account management	Cross-account backup copy	AWS Backup Audit Manager and Jobs dashboard	Restore testing	Logical air-gapped vault	Backup search	Backup tiering	Malware Protection	Access points
AWS GovCloud (US-East)	AWS GovCloud (US)	✓	✓	✓[3]	✓			✓	✓	✓	
AWS GovCloud (US-West)	AWS GovCloud (US)	✓	✓	✓[3]	✓			✓	✓	✓	
AWS European Sovereign Cloud (Germany)	AWS European Sovereign Cloud (Germany)		✓	✓					✓		

[1]China (Beijing) and China (Ningxia) support cross-Region copy only between these two Regions. Cross-account backup copy and cross-account management are supported in both China Regions for accounts that belong to the same organization in AWS Organizations. AWS Backup does not support cross-Region copy between the China Regions and Regions outside the AWS China partition.

[2]Jobs dashboard and AWS Backup Audit Manager organizational reporting, and Jobs dashboard aggregation are only available in Regions that support cross-account management and AWS Backup Audit Manager.

[3]AWS GovCloud (US-East) and AWS GovCloud (US-West) support cross-Region copy from one of these two Regions to the other. Cross-Region copy is not supported from these Regions to other Regions, or from other Regions into these Regions.

Supported services by AWS Region

AWS Backup is available for these resource types in all Regions in which AWS Backup and the listed resource operates:

- Aurora
- AWS CloudFormation
- DynamoDB
- DynamoDB with AWS Backup advanced features
- Amazon EBS
- Amazon EC2
- Amazon EFS
- Amazon Redshift
- Amazon RDS
- Amazon Redshift Serverless
- Amazon DocumentDB

> **Note**
>
> AWS Backup cannot be deployed or operated locally on AWS Outposts infrastructure. All backup operations and storage are performed in the associated AWS Region, and local backup retention on Outposts is not supported.

Regio and servic	Auror DSQL [2]	Amaz FSx	Amaz RDS Multi-AZ	SAP HANA on EC2 instan	Amaz S3	Storag Gatew	Amaz Times m	VMwa and Backu gatew	Amaz EKS	Amaz Neptu	Amaz Docun B	Amazon Redshift Serverles s
US East (N. Virgini	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Regio and servic	Auror DSQL [2]	Amaz FSx	Amaz RDS Multi-AZ	SAP HANA on EC2 instan	Amaz S3	Storag Gatew	Amaz Times m	VMwa and Backu gatew	Amaz EKS	Amaz Neptu	Amaz Docun B	Amazon Redshift Serverles s
US East (Ohio)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
US West (N. Califor a)		✓		✓	✓	✓		✓	✓	✓		✓
US West (Orego	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Africa (Cape Town)		✓	✓	✓	✓	✓		✓	✓	✓	✓	✓
Asia Pacific (Hong Kong)	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓
Asia Pacific (Hyder d)		✓	✓		✓	✓			✓		✓	✓
Asia Pacific (Jakar		✓	✓		✓	✓			✓	✓	✓	✓

Regio and servic	Auror DSQL [2]	Amaz FSx	Amaz RDS Multi-AZ	Amaz SAP HANA on EC2 instan	Amaz S3	Storag Gatew	Amaz Times m	VMwa and Backu gatew	Amaz EKS	Amaz Neptu	Amaz Docun B	Amazon Redshift Serverles s
Asia Pacific (Malay)		✓	✓		✓	✓			✓	✓	✓	✓
Asia Pacific (Melbo e)	✓	Windo Lustre ONTAI	✓		✓	✓			✓	✓	✓	✓
Asia Pacific (Mumb	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Asia Pacific (New Zealar		✓			✓	✓			✓		✓	✓
Asia Pacific (Osaka	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓
Asia Pacific (Seoul	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓
Asia Pacific (Singa e)	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓

Regio and servic	Auror DSQL [2]	Amaz FSx	Amaz RDS Multi-AZ	SAP HANA on EC2 instan	Amaz S3	Storag Gatew	Amaz Times m	VMwa and Backu gatew	Amaz EKS	Amaz Neptu	Amaz Docun B	Amazon Redshift Serverles s
Asia Pacific (Sydne	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Asia Pacific (Taipe		✓			✓	✓			✓			✓
Asia Pacific (Thaila)		✓			✓	✓			✓		✓	✓
Asia Pacific (Tokyo	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Canad (Centr	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓
Canad West (Calga	✓	✓	✓		✓	✓			✓	✓		
China (Beijin		Windo Lustre	✓		✓ [1]	✓				✓	✓	✓
China (Ningx		Windo Lustre	✓		✓ [1]	✓				✓	✓	✓
Europe (Frank t)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Regio and servic	Auror DSQL [2]	Amaz FSx	Amaz RDS Multi-AZ	SAP HANA on EC2 instan	Amaz S3	Storag Gatew	Amaz Times m	VMwa and Backu gatew	Amaz EKS	Amaz Neptu	Amaz Docun B	Amazon Redshift Serverles s
Europ (Irelan	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Europ (Londo	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓
Europ (Milan		✓	✓	✓	✓	✓		✓	✓		✓	✓
Europ (Paris)	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓
Europ (Spain	✓	✓	✓		✓	✓			✓	✓	✓	✓
Europ (Stock m)	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓
Europ (Zurich		✓	✓		✓	✓			✓	✓	✓	✓
Israel (Tel Aviv)		✓	✓		✓	✓			✓	✓	✓	✓
Mexic (Centr		✓			✓	✓			✓		✓	✓
Middl East (Bahra		✓	✓	✓	✓	✓		✓	✓	✓		

Regio and servic	Auror DSQL [2]	Amaz FSx	Amaz RDS Multi-AZ	SAP HANA on EC2 instan	Amaz S3	Storag Gatew	Amaz Times m	VMwa and Backu gatew	Amaz EKS	Amaz Neptu	Amaz Docun B	Amazon Redshift Serverles s
Middle East (UAE)		✓	✓		✓	✓			✓	✓	✓	✓
South Ameri (São Paulo)	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓
AWS GovCl (US-West)		Windo Lustre ONTAI			✓	✓	✓	✓	✓	✓	✓	
AWS GovCl (US-East)		Windo Lustre ONTAI			✓	✓		✓	✓	✓		
AWS Europ Sovere Cloud (Germ					✓				✓			

A check under Amazon FSx indicates that FSx for Windows File Server, FSx for Lustre, FSx for ONTAP, and FSx for OpenZFS are all supported in that Region by AWS Backup; otherwise, the supported configurations will be listed.

[1] Cross-account copy is not supported.

[2] For multi-Region backup and restore supported Regions, see Amazon Aurora DSQL backups.

AWS Backup: How it works

AWS Backup is a fully managed backup service that makes it easy to centralize and automate the backing up of data across AWS services. With AWS Backup, you can create backup policies called *backup plans*. You can use these plans to define your backup requirements, such as how frequently to back up your data and how long to retain those backups.

AWS Backup lets you apply backup plans to your AWS resources by simply tagging them. AWS Backup then automatically backs up your AWS resources according to the backup plan that you defined.

The following sections describe how AWS Backup works, its implementation details, and security considerations.

Topics

- How AWS Backup works with supported AWS services
- Metering, costs, and billing for AWS Backup
- AWS Backup blogs, videos, tutorials, and other resources

How AWS Backup works with supported AWS services

Some AWS Backup-supported AWS services offer their own, stand-alone backup features. Those features are available to you independent of whether you use AWS Backup. However, the backups other AWS services create are not available for central governance through AWS Backup.

To configure AWS Backup to centrally manage data protection for all your supported services, you must opt in to managing that service with AWS Backup, create an on-demand backup or schedule backups using a backup plan, and store your backups in backup vaults.

See Assigning resources for information on choosing which services (resource types) you want to back up.

Topics

- Working with Amazon S3 data
- Working with VMware virtual machines
- Working with Amazon DynamoDB

- Working with Amazon FSx file systems
- Working with Amazon EC2
- Working with Amazon EFS
- Working with Amazon EBS
- Working with Amazon RDS and Aurora
- Working with Aurora DSQL
- Working with AWS BackInt
- Working with AWS Storage Gateway
- Working with Amazon DocumentDB
- Working with Amazon Neptune
- Working with Amazon Redshift and Amazon Redshift Serverless
- Working with Amazon Timestream
- Working with AWS Organizations
- Working with CloudFormation
- Working with AWS BackInt, AWS Systems Manager for SAP, and SAP HANA
- Working with Amazon EKS
- Working with Amazon GuardDuty
- How AWS services back up their own resources

Working with Amazon S3 data

AWS Backup offers fully-managed backup and restore for Amazon S3 backups. To learn more, see Amazon S3 backups.

- How to back up resources: Getting started with AWS Backup
- How to restore Amazon S3 data using AWS Backup: Restore S3 data using AWS Backup

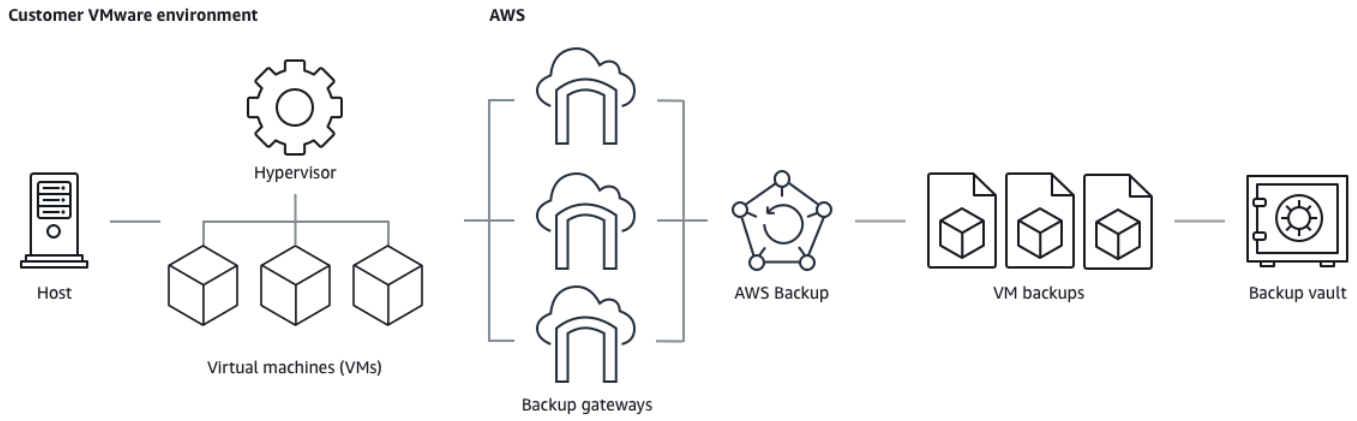
For detailed information about S3 data, see the Amazon S3 documentation.

Working with VMware virtual machines

AWS Backup supports centralized and automated data protection for on-premises VMware virtual machines (VMs) along with VMs in the VMware Cloud™ (VMC) on AWS. You can back up from your

on premises and VMC virtual machines to AWS Backup. Then, you can restore from AWS Backup to either on premises or VMC.

Backup gateway is downloadable AWS Backup software that you deploy to your VMware infrastructure to connect your VMware VMs to AWS Backup. The gateway connects to your VM management server to discover your VMs, encrypt data, and efficiently transfer data to AWS Backup. The following diagram illustrates how Backup gateway connects to your VMs:



- How to back up resources: Virtual machine backups
- How to restore VM resources: Restore a virtual machine using AWS Backup

Working with Amazon DynamoDB

AWS Backup supports backing up and restoring Amazon DynamoDB tables. DynamoDB is a fully-managed NoSQL database service that provides fast and predictable performance with seamless scalability.

Since its launch, AWS Backup has always supported DynamoDB. Starting November 2021, AWS Backup also introduced advanced features for DynamoDB backups. Those advanced features include copying your backups across AWS Regions and accounts, tiering backups to cold storage, and using tags for permissions and cost management.

New AWS Backup customers onboarding after November 2021 will have advanced DynamoDB backup features enabled by default.

We recommend all existing AWS Backup customers enable advanced features for DynamoDB. There is no difference in warm backup storage pricing after you enable advanced features, and you can

save money by tiering backups to cold storage and optimize your costs by using cost allocation tags.

For a full list of advanced features and how to enable them, see Advanced DynamoDB backup.

- How to back up resources: Getting started with AWS Backup
- How to restore DynamoDB resources: Restore a Amazon DynamoDB table

For detailed information about DynamoDB, see What is Amazon DynamoDB? in the *Amazon DynamoDB Developer Guide*.

Working with Amazon FSx file systems

AWS Backup supports backing up and restoring Amazon FSx file systems. Amazon FSx provides fully managed third-party file systems with the native compatibility and feature sets for workloads. AWS Backup uses the built-in backup functionality of Amazon FSx. So backups taken from the AWS Backup console have the same level of file system consistency and performance, and the same restore options as backups that are taken through the Amazon FSx console.

If you use AWS Backup to manage these backups, you gain additional functionality, such as unlimited retention options, and the ability to create scheduled backups as frequently as every hour. In addition, AWS Backup retains your backups even after the source file system is deleted. This protects against accidental or malicious deletion.

Use AWS Backup to protect Amazon FSx file systems if you want to configure backup policies and monitor backup tasks from a central backup console that also extends support for other AWS services.

- How to back up resources: Getting started with AWS Backup
- How to restore Amazon FSx resources: Restore an FSx file system

For detailed information about Amazon FSx file systems, see the Amazon FSx documentation.

Working with Amazon EC2

AWS Backup supports Amazon EC2 instances.

- How to back up resources: Getting started with AWS Backup
- How to restore Amazon EC2 resources: Restore an Amazon EC2 instance

You can schedule or perform on-demand backup jobs that include entire EC2 instances, including its Amazon EBS volumes. Therefore, you can restore an entire Amazon EC2 instance from a single recovery point, including the root volume, data volumes, and some instance configuration settings, such as the instance type and key pair.

You can also back up and restore your VSS-enabled Microsoft Windows applications. You can schedule application-consistent backups, define lifecycle policies, and perform consistent restores as part of an on-demand backup or a scheduled backup plan. For more information, see Create Windows VSS backups.

AWS Backup does not reboot your EC2 instances at any time.

Images and snapshots

When backing up an Amazon EC2 instance, AWS Backup takes a snapshot of the root Amazon EBS storage volume, the launch configurations, and all associated EBS volumes. AWS Backup stores certain configuration parameters of the EC2 instance, including instance type, security groups, Amazon VPC, monitoring configuration, and tags. The backup data is stored as an Amazon EBS volume-backed Amazon Machine Image (AMI).

If the instance was launched from an AMI from AWS Marketplace, the instance has a product code. An AMI created from the instance also has a product code. With Amazon EC2, you can't copy an AMI with a product code to another account. Therefore, AWS Backup also has this restriction.

If you delete an Amazon Machine Image (AMI) or Amazon EBS snapshot that is managed by AWS Backup using AWS Backup and you have the Amazon EC2 recycle bin configured, the image or snapshot might incur charges per the Amazon EC2 recycle bin policy. Snapshots and images in the Amazon EC2 recycle bin are no longer managed by AWS Backup and will not be managed by AWS Backup policies if you restore them from the recycle bin.

> **Note**
>
> AWS Backup periodically identifies orphaned AMIs - AMIs created by AWS Backup that are no longer associated with a recovery point - and places them as expired recovery points in the account's Default vault. You can then delete these expired recovery points, which also deletes the associated AMI and snapshots. If the Default vault does not exist in the Region, create one so that orphaned AMIs can be placed there as expired recovery points. You must then delete these recovery points manually.

AWS Backup managed Amazon EBS snapshots and snapshots associated with a AWS Backup managed Amazon EC2 AMI which have Amazon EBS Snapshot Lock applied may not be deleted as part of the recovery point lifecycle if the snapshot lock duration exceeds the backup lifecycle. Instead, these recovery points will have the status of `EXPIRED`. These recovery points can be deleted manually if you choose to first remove the Amazon EBS snapshot lock.

AWS Backup can encrypt EBS snapshots associated with an Amazon EC2 backup. This is similar to how it encrypts EBS snapshots. AWS Backup uses the same encryption applied on the underlying EBS volumes when creating a snapshot of the Amazon EC2 AMI, and the configuration parameters of the original instance are persisted in the restore metadata.

A snapshot derives its encryption from the volume, and the same encryption is applied to the corresponding snapshots. EBS snapshots of a copied AMI are always encrypted. If you specify a KMS key during the copy, the specified key is applied. If you don't specify a KMS key, a default KMS key is applied.

For more information, see Amazon EC2 instances in the *Amazon EC2 User Guide* and Amazon EBS encryption in the *Amazon EBS User Guide*.

Working with Amazon EFS

AWS Backup supports Amazon Elastic File System (Amazon EFS).

- How to back up resources: Getting started with AWS Backup
- How to restore Amazon EFS resources: Restore an Amazon EFS file system

The Amazon EFS automatic backup vault `aws/efs/automatic-backup-vault` is reserved for those automatic backups only.

This vault should not be used to create cross-account copies or as a destination for backups created by other non-automated backup plans. If you use it as a destination for other backup plans, you will receive an "insufficient privileges" error.

For detailed information about Amazon EFS file systems, see What is Amazon Elastic File System? in the *Amazon Elastic File System User Guide*.

Working with Amazon EBS

AWS Backup supports Amazon Elastic Block Store (Amazon EBS) volumes.

AWS Backup managed Amazon EBS snapshots and snapshots associated with a AWS Backup managed Amazon EC2 AMI which have Amazon EBS Snapshot Lock applied may not be deleted as part of the recovery point lifecycle if the snapshot lock duration exceeds the backup lifecycle. Instead, these recovery points will have the status of `EXPIRED`. These recovery points can be deleted manually if you choose to first remove the Amazon EBS snapshot lock.

- How to back up resources: Getting started with AWS Backup
- How to restore Amazon EBS volumes: Restore an Amazon EBS volume

You can also learn more using the following tutorial: Amazon EBS Backup and Restore Using AWS Backup.

For more information, see Amazon EBS volumes in the *Amazon EBS User Guide*.

Working with Amazon RDS and Aurora

AWS Backup supports Amazon RDS database engines and Aurora clusters.

- How to back up resources: Getting started with AWS Backup
- Amazon Relational Database Service backups
- How to restore Amazon RDS resources: Restore an RDS database
- How to restore Aurora clusters: Restoring an Amazon Aurora cluster

You can also learn by trying the following how-to guide: Amazon RDS Backup and Restore Using AWS Backup.

For more information about Amazon RDS, see What is Amazon Relational Database Service? in the *Amazon RDS User Guide*.

For detailed information about Aurora, see What is Amazon Aurora? in the *Amazon Aurora User Guide*.

If you initiate a backup job from the Amazon RDS console, this can conflict with an Aurora clusters backup job, causing the error Backup job expired before completion. If this occurs, configure a longer backup window in AWS Backup.

AWS does not charge for Aurora snapshots stored inside a backup vault as long as Aurora has automated backups enabled and the retention period for Aurora automated backups is more than the retention period of Aurora snapshots. Any snapshots within the backup vault will be

charged if the snapshots' database is deleted (deletions may occur accidentally or during blue/green deployment).

Large snapshots and frequent backups from a deleted database could result in significant storage charges. Visit the AWS Backup calculator to estimate potential AWS Backup charges.

Working with Aurora DSQL

AWS Backup supports Aurora DSQL clusters for centralized backup and restore operations. Aurora DSQL is a serverless, distributed SQL database that provides ACID transactions with strong consistency and snapshot isolation across multiple AWS Regions.

Aurora DSQL supports both single-Region and multi-Region cluster configurations. Multi-Region clusters include linked cluster Regions that handle read and write operations, and witness Regions that participate in consensus without serving client traffic.

- How to back up resources: Getting started with AWS Backup
- How to restore Aurora DSQL clusters: Amazon Aurora DSQL restore

For detailed information about Aurora DSQL, see What is Aurora DSQL? in the *Aurora DSQL User Guide*.

When backing up multi-Region Aurora DSQL clusters, AWS Backup coordinates backup operations across all linked cluster Regions while leveraging witness Regions for transaction consistency. This ensures that your backup captures a consistent state across the entire distributed database system.

Working with AWS BackInt

AWS Backup works with AWS Backint to support SAP HANA database backup and restore on Amazon EC2 instances.

- Instructions to backup and restore SAP HANA resources: SAP HANA Amazon EC2 Instances backup and restore
- Set up AWS Backint Agent: AWS Backint Agent for SAP HANA

Working with AWS Storage Gateway

AWS Backup supports Storage Gateway Volume Gateway. You can also restore Amazon EBS snapshots as Storage Gateway volumes.

- How to back up resources: Getting started with AWS Backup
- How to restore Storage Gateway resources: Restore a Storage Gateway volume.

Working with Amazon DocumentDB

AWS Backup supports Amazon DocumentDB clusters.

- How to back up resources: Getting started with AWS Backup
- How to restore Amazon DocumentDB resources: Restoring a DocumentDB cluster.

Working with Amazon Neptune

AWS Backup supports Amazon Neptune clusters.

- How to back up resources: Getting started with AWS Backup
- How to restore Amazon Neptune clusters: Restore a Neptune cluster.

Working with Amazon Redshift and Amazon Redshift Serverless

AWS Backup supports Amazon Redshift provisioned clusters and Redshift Serverless namespaces.

- How to backup Amazon Redshift provisioned clusters.
- How to backup Redshift Serverless data warehouses.
- How to restore Amazon Redshift.
- How to restore Redshift Serverless.

Working with Amazon Timestream

AWS Backup supports Amazon Timestream tables.

- How to backup Timestream tables.
- How to restore Timestream tables.

Working with AWS Organizations

AWS Backup works with AWS Organizations to simplify cross-account monitoring and management

- Create a management account in Organizations.
- Turn on cross-account management.
- Designate delegated administrator accounts and delegate policies.

Working with CloudFormation

AWS Backup support CloudFormation templates and application stacks

- CloudFormation stack backups

Working with AWS BackInt, AWS Systems Manager for SAP, and SAP HANA

AWS Backup works with AWS BackInt and with SSM for SAP to support SAP HANA backup and restore functions.

- SAP HANA databases on Amazon EC2 instances backup
- Get started with AWS Systems Manager for SAP
- AWS Backint Agent for SAP HANA

Working with Amazon EKS

AWS Backup supports backups of Amazon EKS clusters, including Kubernetes cluster state and persistent storage attached to the EKS cluster via a persistent volume claim (EBS volumes, EFS file systems, and S3 buckets).

An Amazon EKS backup will create a composite recovery point, where a child recovery point will be for each resource backed up.

- How to back up resources: Getting started with AWS Backup
- How to backup Amazon EKS clusters: Amazon EKS backups
- How to restore Amazon EKS clusters: Restore an Amazon EKS cluster

For detailed information about Amazon EKS, see What is Amazon EKS? in the *Amazon EKS User Guide*.

Working with Amazon GuardDuty

AWS Backup integrates with Amazon GuardDuty to provide automated malware scanning of your recovery points. This integration helps you detect potential malware in your EC2, EBS, and S3 backups.

- **Automated scanning** - Configure backup plans to automatically scan recovery points as they are created
- **On-demand scanning** - Manually initiate scans of specific recovery points using the StartScanJob API
- **Incremental scanning** - Scan only changed data since the last scan to reduce costs
- **Scan results** - View scan status and results in the AWS Backup console and APIs

To enable malware scanning, you must configure IAM roles that allow GuardDuty to read your recovery points and AWS Backup to initiate scans. For more information, see the section called "Managed policies".

How AWS services back up their own resources

You might refer to the technical documentation for a specific AWS service's backup and restore process, particularly when, during a restore, you need to configure a new instance of that AWS service. The following is a list of documentation:

- Amazon EC2 Related Services
- Using AWS Backup with Amazon EFS
- Backup and restore for DynamoDB
- Amazon EBS Snapshots
- Backing Up and Restoring Amazon RDS DB Instances
  - Overview of Backing Up and Restoring an Aurora DB Cluster
- Using AWS Backup with FSx for Windows File Server
- Using AWS Backup with FSx for Lustre
- Backing up your volumes
- Backing Up and Restoring in Amazon DocumentDB

- Backing Up and Restoring an Amazon Neptune Cluster

Metering, costs, and billing for AWS Backup

AWS Backup pricing

Current AWS Backup prices are available at AWS Backup pricing.

> **Important**
>
> To avoid additional charges, configure your retention policy with a warm storage duration of **at least one week**.
> For example, assume you take daily backups and retain them for one day. Further, assume that your protected resources are so large it takes the entire day to complete your backup. AWS Backup implements your retention period of one day and removes your backup from warm storage when your backup job completes. The next day, AWS Backup cannot create an incremental backup because you have no backup in warm storage. Since this retention period did not follow best practices, you run the risk and expense of creating a full backup every day.
> Contact AWS Support for further assistance.

AWS Backup billing

When a resource type supports full AWS Backup management, charges for AWS Backup activity (including storage, data transfers, restores, and early deletion) appear in the "Backup" section of your Amazon Web Services bill. For a list of services that support full AWS Backup management, see the Full AWS Backup management section in the Feature availability by resource table.

When a resource type does not support full AWS Backup management, some of your AWS Backup activity such as storage costs for your backups, have billing reflected by the respective AWS service. However, if you are using a Logically air-gapped vault, all storage or data transfer charges occur under AWS Backup.

Copy Jobs

You will only be charged once a recovery point has been created in the destination vault. There is no charge when a copy job fails and no recovery point is created.

When you configure cross account and cross Region copies, you will see the following billing depending on which resources you copy. For fully managed resources, the source account will see the data transfer charges, and for non fully managed resources, the destination account will see the data transfer charges.

Cost allocation tags

You can use cost allocation tags to track and optimize AWS Backup costs on a detailed level, and view and filter those tags using AWS Cost Explorer.

To use cost allocation tags, see Automating backups and optimizing backup costs for Amazon EFS using AWS Backup and Using Cost Allocation Tags.

AWS Backup Audit Manager pricing

AWS Backup Audit Manager charges for usage based on the number of control evaluations. A control evaluation is the evaluation of one resource against one control. Control evaluation charges appear on your AWS Backup bill. For current control evaluation pricing, see AWS Backup pricing.

To use AWS Backup Audit Manager controls, you must enable AWS Config recording to track your backup activity. AWS Config charges for each configuration item recorded, and these charges appear on your AWS Config bill. For current configuration item recorded pricing, see AWS Config pricing.

Amazon Aurora pricing

During the configured retention period for Aurora continuous backups (up to 35 days), snapshots do not incur a storage charge. Snapshots retained past this window are charged as full backups.

AWS Backup blogs, videos, tutorials, and other resources

For more information about AWS Backup, see the following:

- Lower your Amazon S3 backup costs with AWS Backup S3 tiering by Stuart Lupton and Sravan Rachiraju (November 2025)
- Scan backups for malware with Amazon GuardDuty Malware Protection for AWS Backup by Mehak Mann and Samreen Taj K P (November 2025)
- Secure EKS clusters with the new support for Amazon EKS in AWS Backup by Veliswa Boya (November 2025)

- Encrypt AWS Backup logically air-gapped vaults with customer-managed keys by Desiree Brunner, Mani Manasa Mylavarapu, and Sabith Venkitachalapathy (November 2025)
- AWS Backup adds new Multi-party approval for logically air-gapped vaults by Veliswa Boya (June 2025)
- Streamline search and item level recovery with AWS Backup. With Bisman Sethi (December 2024).
- Building cyber resiliency with AWS Backup logically air-gapped vault. With Sushmitha Srinivasa Murthy and Sabith Venkitachalapathy (August 2024).
- Attach an Amazon EC2 key pair to an AWS Backup restore of a VMware virtual machine. With Olumuyiwa Koya and Kenie Ogunsemowo (August 2024).
- Streamline and automate compliance monitoring and reporting with AWS Backup Audit Manager. With Sabith Venkitachalapathy, Glenn Chia, and Mark Rowland (June 2024).
- Application-consistent backup for Windows application on Amazon EC2 with AWS Backup. With Bhaskar Mazumdar and Karthikeyan KM (May 2024).
- Optimizing AWS Backup costs. With Kenneth Hui (April 2024).
- Designing a resilient and cost-effective backup strategy for Amazon S3. With Mojgan Toth and Harish Mandhadi (March 2024).
- Automate the delivery of AWS Backup Audit Manager reports via email. With Ezekiel Oyerinde, Charles Meruwoma, and Sri Gudavalli.(January 2024)
- Streamlining Point-in-Time Recovery (PITR) for Aurora with AWS Backup. With Enrique Ramirez (January 2024)
- Backup and restore on-premises VMware virtual machines using AWS Backup. With Olumuyiwa Koya and Ezekiel Oyerinde (June 2022).
- Using AWS Backup to protect Amazon Aurora databases. With Chris Hendon, Brandon Rubadou, and Thomas Liddle (May 2022).
- Protecting encrypted Amazon RDS instances with cross-account and cross-Region backups. With Evan Peck and Sabith Venkitachalapathy (May 2022).
- Automate and improve your security posture using AWS Backup and AWS PrivateLink. With Bilal Alam (Apr. 2022).
- Obtain aggregated daily cross-account multi-Region AWS Backup reporting. With Wali Akbari and Sabith Venkitachalapathy (Feb. 2022).
- Automate visibility of backup findings using AWS Backup and AWS Security Hub CSPM. With Kanishk Mahajan (Jan. 2022).

- Top 10 security best practices for securing backups in AWS. With Ibukun Oyewumi (Jan. 2022).
- Optimizing SAS Grid on AWS with FSx for Lustre (and optimizing disaster recovery using AWS Backup). With Matt Saeger and Shea Lutton (Jan. 2022).
- Centralizing data protection and compliance in Amazon Neptune with AWS Backup. With Brian O'Keefe (Nov. 2021).
- Manage backup and restore of Amazon DocumentDB (with MongoDB compatibility) with AWS Backup. With Karthik Vijayraghavan (Nov. 2021).
- Simplify auditing your data protection policies with AWS Backup Audit Manager. With Jordan Bjorkman and Harshitha Putta (Nov. 2021).
- Enhance the security posture of your backups with AWS Backup Vault Lock. With Rolland Miller (Oct. 2021).
- How to retain resource tags in AWS Backup restore jobs. With Ibukun Oyewumi, Amee Shah, and Sabith Venkitachalapathy (Sep. 2021).
- Managing access to backups using service control policies with AWS Backup. With Sabith Venkitachalapathy and Ibukun Oyewumi (Aug. 2021).
- Automate centralized backup at scale across AWS services using AWS Backup. With Ibukun Oyewumi and Sabith Venkitachalapathy (Jul. 2021).
- Blog: How to simplify Microsoft SQL Server backup using AWS Backup and VSS. With Siavash Irani and Sepehr Samiei (Jul. 2021).
- Automate data recovery validation with AWS Backup. With Mahanth Jayadeva (Jun. 2021).
- Configuring notifications to monitor AWS Backup jobs. With Virgil Ennes (Jun. 2021).
- Automating backups and optimizing backup costs for Amazon EFS using AWS Backup. With Prachi Gupta and Rohit Verma (Jun. 2021).
- Manage Amazon EFS backup costs: AWS Backup support for cost allocation tags. With Aditya Maruvada (May 2021).
- Create and share encrypted backups across accounts and Regions using AWS Backup. With Prachi Gupta (May 2021).
- AWS Backup is now FedRAMP High approved for your compliance and data protection needs. With Andy Grimes (May 2021).
- ZS Associates enhances backup efficiency with AWS Backup. With Mitesh Naik, Hiranand Mulchandani, and Sushant Jadhav (May 2021).
- Tutorial: Amazon EBS Backup and Restore using AWS Backup. With Fathima Kamal (Apr. 2021).
- Video Tutorial: Managing Cross-Region Copies of Backups. With David DeLuca (Apr. 2021).

- Delete multiple AWS Backup recovery points using AWS Tools for PowerShell. With Sherif Talaat (Apr. 2021).
- Cross-region and cross-account backups for Amazon FSx using AWS Backup. With Adam Hunter and Fathima Kamal (Apr. 2021).
- Amazon CloudWatch Events and Metrics for AWS Backup. With Rolland Miller (Mar. 2021).
- Tutorial: Amazon Relational Database Service (RDS) Backup and Restore using AWS Backup. With Fathima Kamal (Mar. 2021).
- Point-in-time recovery and continuous backup for Amazon RDS with AWS Backup. With Kelly Griffin (Mar. 2021).
- Automate AWS Backup with AWS Service Catalog. with John Husemoller (Jan. 2021).
- Secure data recovery with cross-account backup and Cross-Region copy using AWS Backup. With Cher Simon (Jan. 2021).
- AWS re:Invent recap: Data protection and compliance with AWS Backup. With Nancy Wang (Dec. 2020).
- AWS Backup provides centralized data protection across your AWS resources. With Nancy Wang (Nov. 2020).
- Tech Talk: Data protection at scale with AWS Backup. With Kareem Behairy (Sep. 2020).
- Centralized cross-account management with cross-Region copy using AWS Backup. With Cher Simon (Sep. 2020).
- Video Tutorial: Managing backups at scale in your AWS Organizations using AWS Backup. With Ildar Sharafeev (Jul. 2020).
- Managing backups at scale in your AWS Organizations using AWS Backup. With Nancy Wang, Avi Drabkin, Ganesh Sundaresan, and Vikas Shah (Jun. 2020).
- Recover Amazon EFS files and folders with AWS Backup. With Abrar Hussain and Gurudath Pai (May 2020).
- Scheduling automated backups using Amazon EFS and AWS Backup. With Rob Barnes (Dec. 2019).
- re:Invent Recording: AWS re:Invent 2019: Deep dive on AWS Backup ft. Rackspace. With Nancy Wang and Jason Pavao (Dec. 2019).
- Protecting your data with AWS Backup. With Anthony Fiore (Jul. 2019).
- Marketing Video: Introducing AWS Backup. Jan. 2019.
- Video: Introduction to AWS Backup. With AWS Training and Certification.

Getting started with AWS Backup

This tutorial shows you the generic steps for using AWS Backup features and functionality. As with any part of this technical documentation, you should follow along with the AWS Management Console in the other window.

Prerequisites

Before you begin, ensure that you have the following:

- An AWS account. For more information, see Create an AWS account.
- At least one resource supported by AWS Backup. See the list of supported AWS resources and third-party applications.

  When new AWS services become available, enable AWS Backup to use those services.
- If you intend to include multiple resource types, such as Amazon EBS volumes and Amazon S3 buckets, in the same backup plan, confirm those resources are in the same AWS Region.

To configure the AWS services to use with AWS Backup

1. Sign in to the AWS Management Console, and open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Settings**.
3. On the **Service opt-in** page, choose **Configure resources**.
4. On the **Configure resources** page, use the toggle switches to enable or disable the services that are used with AWS Backup. Choose **Confirm** when your services are configured. Make sure that the AWS service you're opting in is available in your AWS Region.

See for additional information. The AWS Backup console allows a user to assign a resource type to a backup plan; this will be included even if the opt-in is not enabled for that particular service.

To complete this tutorial, you can use your AWS account root user to sign in to the AWS Management Console. However, AWS Identity and Access Management (IAM) recommends that you not use the AWS account root user. Instead, create an administrator in your account and use those credentials to manage resources in your account.

The AWS Backup console provides different options to back up your resources. You can create a backup on-demand, schedule and configure how you want the resource backed up, or configure resources to back up automatically when the resource is created.

Service Opt-in

The AWS Backup console has two ways to include resource types in a backup plan: explicitly assign the resource type in a backup plan or include all resources. See the points below to understand how these selections work with service opt ins.

- If resource assignments are only based on tags, then service opt-in settings are applied.
- If a resource type is explicitly assigned to a backup plan, it will be included in the backup even if the opt-in is not enabled for that particular service. This does not apply to Aurora, Neptune, and Amazon DocumentDB. For these services to be included, the opt-in must be enabled.
- If both resource type and tags are specified in a resource assignment, the specified resource types are filtered first, then tags further filter those resources.

  Service opt-in settings are ignored for most resource types. However Aurora, Neptune, and Amazon DocumentDB require service opt-in.
- For Amazon FSx for NetApp ONTAP, when using tag-based resource selection, apply tags to individual volumes instead of the whole file system.

Opt-in choices apply to the specific account and AWS Region. When an account uses AWS Backup (creates a backup vault or backup plan) in a Region, the account automatically is opted into all resource types supported by AWS Backup in the Region at that time. Supported services added to that Region at a later date will not be automatically included in a backup plan. You can choose to opt into those resource types once they become supported.

As AWS Backup supports more and more AWS services and third-party applications, you might need to revisit this step to opt in to those newly-supported resources.

AWS Backup does not govern or manage backups taken in AWS environments other than AWS Backup.

> **Note**
>
> For more information on how service opt-ins work for management and delegated administrator accounts, see Managing AWS Backup resources across multiple AWS accounts.

To opt in to use AWS Backup to protect all supported resource types

1. Sign in to the AWS Management Console, and open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the left navigation pane, choose **Settings**.
3. Under **Service opt-in**, choose **Configure resources**.
4. Opt in to all AWS Backup-supported **Resources** by moving all the toggles to the right.
5. Choose **Confirm**.

Backup plans

In AWS Backup, a *backup plan* is a policy expression that defines when and how you want to back up your AWS resources, such as Amazon DynamoDB tables or Amazon Elastic File System (Amazon EFS) file systems. You can assign resources to backup plans, and AWS Backup automatically backs up and retains backups for those resources according to the backup plan. You can create multiple backup plans if you have workloads with different backup requirements. By default, backup windows are optimized by AWS Backup. You can customize the backup window in the console or programmatically.

AWS Backup efficiently stores your periodic backups incrementally for supported resource types. The first backup of an AWS resource backs up a full copy of your data. For each successive incremental backup, only the changes to your AWS resources are backed up. Incremental backups enable you to benefit from the data protection of frequent backups while minimizing storage costs.

> **Important**
>
> Not all resource types support incremental backups. For resources that do not support incremental backups, each backup is a full copy, which can result in higher storage costs. For a list of which resources support incremental backups, see Feature availability by resource.

AWS Backup also seamlessly manages your backup plan's lifecycle based on your retention settings, which allows you to restore when needed.

The following sections provide the basics of managing your backup strategy in AWS Backup.

Topics

- Create a backup plan
- Understanding backup plan summary
- Select AWS services to backup

Create a backup plan

You can create a backup plan using the AWS Backup console, API, CLI, SDK, or an AWS CloudFormation template.

Topics

- Create backup plans using the AWS Backup console
- Create backup plans using the AWS CLI
- Backup plan options and configuration
- CloudFormation templates for backup plans
- Delete a backup plan
- Update a backup plan

Create backup plans using the AWS Backup console

Open the AWS Backup console at https://console.aws.amazon.com/backup. From the dashboard, choose **Manage Backup** plans. Or, using the navigation pane, choose **Backup plans** and choose **Create Backup** plan.

Start options

You have three choices for your new backup plan:

- Create a backup plan based on an existing plan
- Build a new plan
- Create a backup plan using the AWS CLI

In this procedure, we build a new plan. Each part of the configuration has a link to an expanded section further on the page to where you can navigate for more detail.

1. Enter a plan name in **Backup plan name**.

   If you try to create a backup plan that is identical to an existing plan, you receive an `AlreadyExistsException` error.

2. Optionally, you can add tags to your backup plan.

3. **Backup rule configuration:** In the backup rule configuration section, you will set the backup schedule, window, and lifecycle.
4. **Schedule:**
   a. Enter a **backup rule name** in the text field.
   b. In the backup vault menu, choose **Default** or choose **Create new Backup vault** to create a vault.
   c. In the backup frequency menu, choose how often you want this plan to create a backup.
5. **Backup window:**
   a. **Start time** defaults to 12:30 AM (00:30 in 24hr time) in your system's local timezone.
   b. **Start within** defaults to 8 hours. You can change this to specify a window of time for the backup to start.
   c. **Complete within** defaults to 7 days. Ensure that there is enough time for the backup up to complete even if the job starts at the end of the start window.
6. **Continuous backups and point-in-time recovery (PITR):** You can select **Enable continuous backups for point-in-time recovery (PITR)**. To verify which resources are supported for this type of backup, see the Feature availability by resource matrix.
7. **Lifecycle**
   a. **Cold storage:** Select this box to let eligible resource types transition to cold storage in accordance with the timetable you specify in the total retention period. To use cold storage, you must have a total retention period of 90 days or greater. However please note the following. Some services support incremental backups. For incremental backups, you must have at least one warm full backup. AWS Backup recommends that you set your lifecycle settings to not move your backup to cold storage until after at least 8 days. If the full backup is transitioned to cold storage too soon (for example, a transition to cold storage after 1 day), AWS Backup will create another warm full backup.
   b. **Cold storage for Amazon EBS** is Amazon EBS Snapshots Archive. Snapshots transitioned to archive storage tier will display in the console as cold tier. If cold storage is enabled, and if your backup frequency is monthly or less often, you can have your backup plan transition EBS snapshots.
   c. The **total retention period** is the number of days that you store your resource in AWS Backup. It is the total number of days of warm storage plus cold storage.

8. (*Optional*) You can opt in to have a backup index created with each periodic backup of a supported resource type (continuous backups will have daily indexes created). Only recovery points (backups) that have an associated index can be included in a backup search.

   For example, each time your backup plan creates an S3 backup, you can have a backup index for that backup created, also. This will allow that particular backup to be included in a future search.

   Place a check next to the resource type(s) for which you want to have indexes created.

9. (*Optional*) Enable malware scanning to automatically scan backups after they are created. When configuring malware protection, specify which resource types to scan (Amazon EC2, Amazon EBS, Amazon S3, or all supported resources) and the scanning types (full or incremental). Malware scanning applies only to your selected resource types. For example, if your backup plan includes both Amazon S3 and Amazon EC2 resources, but you enable malware scanning only for Amazon EC2, the service will scan only your EC2 backups. For each backup rule, you can configure which scanning type to use. The schedule of the backup rule will determine how frequently the scanning type takes place.

   > **Important**
   >
   > Before enabling malware protection, ensure your backup role and scanner role have the required permissions. For more information, see the permissions documentation.

10. (*Optional*) Use **Copy to destination** to create a cross-Region copy of eligible resources if you want to store a copy of a backup in a different AWS Region.
11. (*Optional*) Tags added to recovery points.
12. When all sections are set to your specifications, choose **Save Backup rule**.

Create backup plans using the AWS CLI

You can also define your backup plan in a JSON document and provide it using the AWS Backup console or AWS CLI. The following JSON document contains a sample backup plan that creates a daily backup at 1:00 Pacific time (the local time adjusts to daylight, standard, or summer time conditions if applicable). It automatically deletes a backup after one year.

```
{
  "BackupPlan":{
```

```
    "BackupPlanName":"test-plan",
    "Rules":[
      {
        "RuleName":"test-rule",
        "TargetBackupVaultName":"test-vault",
        "ScheduleExpression":"cron(0 1 ? * * *)",
        "ScheduleExpressionTimezone":"America/Los_Angeles",
        "StartWindowMinutes":integer, // Value is in minutes
        "CompletionWindowMinutes":integer, // Value is in minutes
        "IndexActions": [
               {
                  "ResourceTypes": [ "string" ]
               }
            ],
        "Lifecycle":{
          "DeleteAfterDays":integer, // Value is in days
        }
      }
    ]
  }
}
```

You can store your JSON document with a name you choose. The following CLI command shows `create-backup-plan` with a JSON named `test-backup-plan.json`:

```
aws backup create-backup-plan --cli-input-json file://PATH-TO-FILE/test-backup-plan.json
```

Note that while some systems number the days of the week from 0 to 6, we number them from 1 to 7. For more information, see Cron and rate expressions. For more information about timezones, see TimeZone in the *Amazon Location Service API reference*.

Backup plan options and configuration

When you define a backup plan in the AWS Backup console, you configure the following options:

Backup plan name

You must provide a name for your backup plan. Names are limited to 50 characters, including alphanumeric characters, dashes, underscores, and periods.

Backup rules

Backup plans are composed of one or more backup rules. To add backup rules to a backup plan, or to edit existing rules in a backup plan:

1. From the AWS Backup console, in the left navigation pane, choose **Backup plans**.
2. Under **Backup plan name**, select a backup plan.
3. Under the **Backup rules** section:
   - To add a backup rule, choose **Add backup rule**.
   - To edit an existing backup rule, select a rule, then choose **Edit**.

> **Note**
>
> If you have a backup plan with multiple rules and the time frames of the two rules overlap, AWS Backup optimizes the backup and takes a backup for the rule with the longer retention time. The optimization takes into account the full start window, not just when the daily backup is taken.

Each backup rule consists of the following elements.

Backup rule name

Backup rule names are case sensitive. They must contain from 1 to 50 alphanumeric characters or hyphens.

Backup frequency

The backup frequency determines how often AWS Backup creates a snapshot backup. Using the console, you can choose a frequency of every hour, 12 hours, daily, weekly, or monthly. You can also create a cron expression that creates snapshot backups as frequently as hourly. Using the AWS Backup CLI, you can schedule snapshot backups as frequently as hourly.

If you select weekly, you can specify which days of the week you want backups to be taken. If you select monthly, you can choose a specific day of the month.

You can also check the **Enable continuous backups for supported resources** checkbox to create a point-in-time restore (PITR)-enabled continuous backup rule. Unlike snapshot backups, continuous

backups allow you to perform point-in-time restore. To learn more about continuous backups, see Point-in-Time Recovery.

Backup window

Backup windows consist of the time that the backup window begins and the duration of the window in hours. Backup jobs are started within this window. The default settings in the console are:

- **12:30 AM** local to your system's timezone (0:30 in 24-hour systems)
- **Start within** 8 hours
- **Complete within** 7 days

  (**complete within** parameter does not apply to Amazon FSx resources)

You can customize the backup frequency and backup window start time using a cron expression. To see the six fields of AWS cron expressions, see Cron and rate expressions in the *Amazon EventBridge User Guide*. Two examples of AWS cron expressions are `15 * ? * * *` (take a backup every hour at 15 minutes past the hour) and `0 12 * * ? *` (take a backup every day at 12 noon UTC). For a table of examples, click the preceding link and scroll down the page.

AWS Backup evaluates cron expressions between 00:00 and 23:59. If you create a backup rule for "every 12 hours" but provide a start time of later than 11:59, it will only run once per day.

Backup plans in a timezone that observes daylight savings time might be impacted by the time shift forward. You can switch to UTC or create a manual backup on the day that time shifts forward. For more information see Daylight savings time on EventBridge Scheduler.

Continuous backups and point-in-time restore (PITR) reference the changes recorded over a period of time; therefore, they cannot be scheduled with a time or cron expression.

In general, AWS database services cannot start backups 1 hour before or during their maintenance window and Amazon FSx cannot start backups 3 hours before or during their maintenance window or automatic backup window (Amazon Aurora is exempt from this maintenance window restriction). Snapshot backups scheduled during those times will fail. An exception occurs when you opt in to using AWS Backup for both snapshot and continuous backups for a supported service. AWS Backup will schedule backup windows automatically to avoid conflicts. See Point-in-Time Recovery for a list of supported services and instructions on how to use AWS Backup to take continuous backups.

Overlapping backup rules

On occasion, a backup plan might contain multiple, overlapping rules. When the start windows of different rules overlap, AWS Backup retains the backup under the rule with the longer retention period. For example, consider a backup plan with two rules:

1. Backup hourly, with a 1-hour start window, and retain for 1 day.
2. Backup every 12 hours, with an 8-hour start window, and retain for 1 week.

After 24 hours, the second rule creates two backups (because it has the longer retention period). The first rule creates eight backups (because the second rule's 8-hour start window prevented more hourly backups from running). Specifically:

During this Start Window	This Rule Creates 1 Backup
Midnight to 8AM	12 hours
8 to 9	Hourly
9 to 10	Hourly
10 to 11	Hourly
11 to Noon	Hourly
Noon to 8PM	12 hours
8 to 9	Hourly
9 to 10	Hourly
10 to 11	Hourly
11 to Midnight	Hourly

During the start window, the backup job status remains in `CREATED` status until it has successfully begun or until the start window time has run out. If within the start window time AWS Backup receives an error that allows the job to be retried, AWS Backup will automatically retry to begin the job at least every 10 minutes until the backup successfully begins (the job status changes to

RUNNING) or until the job status changes to EXPIRED (which is expected to occur when the start window time is over).

Lifecycle and storage tiers

Backups are stored for the number of days you specify, known as the backup *lifecycle*. Backups can be restored until the end of their lifecycle.

This is set as the **total retention period** in the lifecycle section of backup rule configuration in the AWS Backup console.

If you use AWS CLI, this is set using the parameter `DeleteAfterDays`. The retention period for snapshots can range between 1 day and 100 years (or indefinitely if you don't enter one), while the retention period for continuous backups can range from 1 day to 35 days. The creation date of a backup is the date the backup job started, not the date it completed. If your backup job doesn't complete on the same date it started, use the date on which it began to help calculate retention periods.

Backups are maintained in a storage tier. Each tier incurs a different cost for storage and for restore, as outlined by AWS Backup pricing. Every backup is created and is stored in warm storage. Depending on how long you choose to store your backup, you may wish to transition your backup to a lower-cost tier called cold storage. Feature availability by resource displays which resources have this optional feature.

Console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Create or edit a backup plan.
3. In the lifecycle section of backup rule configuration, check the box **Move backups from warm to cold storage**.
4. (*optional*) If Amazon EBS is one of the resources you back up and your backup frequency is monthly or less frequent, you can transition them to cold tier using EBS snapshot archival.
5. Input a value (in days) that you want your backups to remain in warm storage. AWS Backup recommends at least 8 days.
6. Input a value (in days) for the total retention period. The difference between total retention period and time in warm storage will be the amount of days the backups remain in cold storage.

AWS CLI

1. Use `create-backup-plan` or `update-backup-plan`.
2.
3. Include the Boolean parameter `OptInToArchiveForSupportedResources` for EBS resources.
4. Include the parameter `MoveToColdStorageAfterdays`.
5. Use the parameter `DeleteAfterDays`. This value must be 90 (days) plus the value you input for `MoveToColdStorageAfterDays`.

Cold storage is currently available for the following resource types:

Resource type	Incremental or Full backup in cold storage
AWS CloudFormation	Incremental
DynamoDB with advanced features	Full; no Incremental backups in any tier
Amazon EBS (using EBS Snapshot Archive)	Full; Incremental backups will become Full after transition.
Amazon EFS	Incremental
SAP HANA databases running on Amazon EC2 instances	Incremental
Amazon Timestream	Incremental
VMware virtual machines	Incremental

Once you have enabled transition to cold storage through the console or command line, the following conditions are true for backups in cold storage (or archive):

- Backups transitioned must be stored in cold storage for a minimum of 90 days, in addition to the time in warm storage. AWS Backup requires the retention to be set for 90 days longer than the “transition to cold after days” setting. You can't change the “transition to cold after days” setting after a backup has been transitioned to cold.

- Some services support incremental backups. For incremental backups, you must have at least one warm full backup. AWS Backup recommends that you set your lifecycle settings to not move your backup to cold storage until after at least 8 days. If the full backup is transitioned to cold storage too soon (for example, a transition to cold storage after 1 day), AWS Backup will create another warm full backup.
- For resource types that support incremental backups, AWS Backup transitions data from warm to cold storage if the transitioned data is no longer referenced by warm backups. Data in backups retained in cold storage that is only referenced by other cold backups is billed at cold storage tier prices. Other backups continue at warm storage tier pricing.

Backup vault

A backup vault is a container to organize your backups in. Backups created by a backup rule are organized in the backup vault that you specify in the backup rule. You can use backup vaults to set the AWS Key Management Service (AWS KMS) encryption key that is used to encrypt backups in the backup vault and to control access to the backups in the backup vault. You can also add tags to backup vaults to help you organize them. If you don't want to use the default vault, you can create your own. For step-by-step instructions for creating a backup vault, see Backup vault creation and deletion.

Copy to Regions

As part of your backup plan, you can optionally create a backup copy in the same or another AWS Region. These copies can be made either in the same account or another account. For more information about backup copies, see Creating backup copies across AWS Regions.

When you define a backup copy, you configure the following options:

Destination Region

The destination Region for the backup copy.

(Advanced Settings) Backup vault

The destination backup vault for the copy.

(Advanced Settings) IAM Role

The IAM role that AWS Backup uses when creating the copy. The role must also have AWS Backup listed as a trusted entity, which enables AWS Backup to assume the role. If you choose **Default**

and the AWS Backup default role is not present in your account, a role is created for you with the correct permissions.

(Advanced Settings) Lifecycle

Specifies when to transition the backup copy to cold storage and when to expire (delete) the copy. Backups transitioned to cold storage must be stored in cold storage for a minimum of 90 days. You can't change this value after a copy has transitioned to cold storage.

Expire specifies the number of days after creation that the copy is deleted. This must be greater than 90 days beyond the **Transition to cold storage** value.

If the value for `Lifecycle:DeleteAfterDays` (shown as **Expire** in the console) is not specified in the copy settings, the copy will follow the lifecycle settings of the backup from which it is copied.

Tags added to recovery points

The tags that you list here are automatically added to backups when they are created.

Tags added to backup plans

These tags are associated with the backup plan itself to help you organize and track your backup plan.

Advanced backup settings

Advanced backup settings allow you to configure resource-specific backup options for different AWS services.

Amazon EC2 advanced backup settings

Enables application consistent backups for third-party applications that are running on Amazon EC2 instances. Currently, AWS Backup supports Windows VSS backups. AWS Backup excludes specific Amazon EC2 instance types from Windows VSS backups. For more information, see Create Windows VSS backups.

Amazon S3 advanced backup settings

AWS Backup provides advanced settings to control what metadata is included in your Amazon S3 backups. You can optionally exclude Access Control Lists (ACLs) and object tags from your backups, which aligns with Amazon S3 best practices of using bucket-level permissions instead of object-level ACLs.

For detailed information on configuring Amazon S3 backup options for ACLs and object tags, see Advanced Amazon S3 backup settings.

Important

When you exclude ACLs from backups, objects restored without ACLs will use the destination bucket's default ownership settings. The destination bucket must have appropriate object ownership configuration.

Note

When a continuous backup job is running for an Amazon S3 bucket and you initiate a snapshot backup job, the snapshot will use the same ACL and object tag settings as the continuous backup, regardless of the settings specified for the snapshot job.

Malware scanning

AWS Backup integrates with Amazon GuardDuty to provide automated malware scanning of your recovery points. When you enable malware scanning in your backup plan, AWS Backup automatically scans your backups for malware and provides scan results to help you make informed decisions about restoring your data.

To configure malware scanning for your backup plan:

1. Create an IAM role that trusts `malware-protection.guardduty.amazonaws.com` and attach the AWS managed policy `AWSBackupGuardDutyRolePolicyForScans`.
2. Attach the AWS managed policy `AWSBackupServiceRolePolicyForScans` to your backup selection's IAM role.
3. In your backup plan configuration, add scanning settings that specify:
   - The scan service (GuardDuty)
   - The resource types to scan (Amazon EC2, Amazon EBS, Amazon S3)
   - The IAM role ARN for GuardDuty to assume
4. Configure scan actions in your backup rules to specify:
   - The scan service (GuardDuty)

- The scan type (incremental or full scan)

For more information about the managed policies, see the section called "AWSBackupGuardDutyRolePolicyForScans" and the section called "AWSBackupServiceRolePolicyForScans".

CloudFormation templates for backup plans

We provide three sample CloudFormation templates for your reference. The first template creates a simple backup plan. The second template enables VSS backups in a backup plan. The third template enables Amazon GuardDuty Malware Protection scanning in a backup plan.

> **Note**
>
> If you are using the default service role, replace *backup-policy-name* with AWSBackupServiceRolePolicyForBackup.

```
Description: backup plan template to back up all resources daily at 5am UTC, and tag
 all recovery points with backup:daily.

Resources:
  KMSKey:
    Type: AWS::KMS::Key
    Properties:
      Description: "Encryption key for daily"
      EnableKeyRotation: True
      Enabled: True
      KeyPolicy:
        Version: "2012-10-17"
        Statement:
          - Effect: Allow
            Principal:
              "AWS": { "Fn::Sub": "arn:${AWS::Partition}:iam::${AWS::AccountId}:root" }
            Action:
              - kms:*
            Resource: "*"

  BackupVaultWithDailyBackups:
    Type: "AWS::Backup::BackupVault"
```

```
    Properties:
      BackupVaultName: "BackupVaultWithDailyBackups"
      EncryptionKeyArn: !GetAtt KMSKey.Arn

  BackupPlanWithDailyBackups:
    Type: "AWS::Backup::BackupPlan"
    Properties:
      BackupPlan:
        BackupPlanName: "BackupPlanWithDailyBackups"
        BackupPlanRule:
          - RuleName: "RuleForDailyBackups"
            TargetBackupVault: !Ref BackupVaultWithDailyBackups
            ScheduleExpression: "cron(0 5 ? * * *)"
    DependsOn: BackupVaultWithDailyBackups

  DDBTableWithDailyBackupTag:
    Type: "AWS::DynamoDB::Table"
    Properties:
      TableName: "TestTable"
      AttributeDefinitions:
        - AttributeName: "Album"
          AttributeType: "S"
      KeySchema:
        - AttributeName: "Album"
          KeyType: "HASH"
      ProvisionedThroughput:
        ReadCapacityUnits: "5"
        WriteCapacityUnits: "5"
      Tags:
        - Key: "backup"
          Value: "daily"

  BackupRole:
    Type: "AWS::IAM::Role"
    Properties:
      AssumeRolePolicyDocument:
        Version: "2012-10-17"
        Statement:
          - Effect: "Allow"
            Principal:
              Service:
                - "backup.amazonaws.com"
            Action:
              - "sts:AssumeRole"
```

```
      ManagedPolicyArns:
        - "arn:aws:iam::aws:policy/service-role/backup-policy-name"

  TagBasedBackupSelection:
    Type: "AWS::Backup::BackupSelection"
    Properties:
      BackupSelection:
        SelectionName: "TagBasedBackupSelection"
        IamRoleArn: !GetAtt BackupRole.Arn
        ListOfTags:
          - ConditionType: "STRINGEQUALS"
            ConditionKey: "backup"
            ConditionValue: "daily"
      BackupPlanId: !Ref BackupPlanWithDailyBackups
    DependsOn: BackupPlanWithDailyBackups
```

```
Description: backup plan template to enable Windows VSS and add backup rule to take
 backup of assigned resources daily at 5am UTC.

Resources:
  KMSKey:
    Type: AWS::KMS::Key
    Properties:
      Description: "Encryption key for daily"
      EnableKeyRotation: True
      Enabled: True
      KeyPolicy:
        Version: "2012-10-17"
        Statement:
          - Effect: Allow
            Principal:
              "AWS": { "Fn::Sub": "arn:${AWS::Partition}:iam::${AWS::AccountId}:root" }
            Action:
              - kms:*
            Resource: "*"

  BackupVaultWithDailyBackups:
    Type: "AWS::Backup::BackupVault"
    Properties:
      BackupVaultName: "BackupVaultWithDailyBackups"
      EncryptionKeyArn: !GetAtt KMSKey.Arn

  BackupPlanWithDailyBackups:
```

```
    Type: "AWS::Backup::BackupPlan"
    Properties:
      BackupPlan:
        BackupPlanName: "BackupPlanWithDailyBackups"
        AdvancedBackupSettings:
          - ResourceType: EC2
            BackupOptions:
              WindowsVSS: enabled
        BackupPlanRule:
          - RuleName: "RuleForDailyBackups"
            TargetBackupVault: !Ref BackupVaultWithDailyBackups
            ScheduleExpression: "cron(0 5 ? * * *)"

    DependsOn: BackupVaultWithDailyBackups
```

```
Description: Backup plan template with Amazon GuardDuty Malware Protection scanning
 enabled.

Resources:
  BackupVault:
    Type: "AWS::Backup::BackupVault"
    Properties:
      BackupVaultName: "MalwareScanBackupVault"

  BackupPlanWithMalwareScanning:
    Type: "AWS::Backup::BackupPlan"
    Properties:
      BackupPlan:
        BackupPlanName: "BackupPlanWithMalwareScanning"
        BackupPlanRule:
          - RuleName: "DailyBackupWithIncrementalScan"
            TargetBackupVault: !Ref BackupVault
            ScheduleExpression: "cron(0 5 ? * * *)"
            Lifecycle:
              DeleteAfterDays: 35
            ScanActions:
              - MalwareScanner: GUARDDUTY
                ScanMode: INCREMENTAL_SCAN
          - RuleName: "MonthlyBackupWithFullScan"
            TargetBackupVault: !Ref BackupVault
            ScheduleExpression: "cron(0 5 1 * ? *)"
            Lifecycle:
              DeleteAfterDays: 365
```

```
          ScanActions:
            - MalwareScanner: GUARDDUTY
              ScanMode: FULL_SCAN
      ScanSettings:
        - MalwareScanner: GUARDDUTY
          ResourceTypes:
            - EBS
          ScannerRoleArn: !GetAtt ScannerRole.Arn
  DependsOn: BackupVault

ScannerRole:
  Type: "AWS::IAM::Role"
  Properties:
    AssumeRolePolicyDocument:
      Version: "2012-10-17"
      Statement:
        - Effect: "Allow"
          Principal:
            Service:
              - "malware-protection.guardduty.amazonaws.com"
          Action:
            - "sts:AssumeRole"
    ManagedPolicyArns:
      - "arn:aws:iam::aws:policy/AWSBackupGuardDutyRolePolicyForScans"

BackupRole:
  Type: "AWS::IAM::Role"
  Properties:
    AssumeRolePolicyDocument:
      Version: "2012-10-17"
      Statement:
        - Effect: "Allow"
          Principal:
            Service:
              - "backup.amazonaws.com"
          Action:
            - "sts:AssumeRole"
    ManagedPolicyArns:
      - "arn:aws:iam::aws:policy/service-role/backup-policy-name"
      - "arn:aws:iam::aws:policy/AWSBackupServiceRolePolicyForScans"

TagBasedBackupSelection:
  Type: "AWS::Backup::BackupSelection"
  Properties:
```

```
      BackupSelection:
        SelectionName: "MalwareScanSelection"
        IamRoleArn: !GetAtt BackupRole.Arn
        ListOfTags:
          - ConditionType: "STRINGEQUALS"
            ConditionKey: "backup"
            ConditionValue: "true"
      BackupPlanId: !Ref BackupPlanWithMalwareScanning
    DependsOn: BackupPlanWithMalwareScanning
```

Delete a backup plan

You can delete a backup plan only after all associated selections of resources have been deleted. These selections are also known as *resource assignments*. If these have not been deleted prior to deletion of the backup plan, the console will display the error: "Related backup plan selections must be deleted prior to backup plan deletion." Use the console or use `DeleteBackupSelection`.

Deleting a backup plan deletes the current version of the plan. The current and previous versions, if any, still exist, but they are no longer listed on the console under **Backup plans**.

> **Note**
>
> When a backup plan is deleted, existing backups are not deleted. To remove existing backups, delete them from the backup vault using the steps in Deleting backups.

To delete a backup plan using the AWS Backup console

1. Sign in to the AWS Management Console, and open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane on the left, choose **Backup plans**.
3. Choose your backup plan in the list.
4. Select any resource assignments that are associated with the backup plan.
5. Choose **Delete**.

Update a backup plan

After creating a backup plan, you can edit the plan—for example, you can add tags, or you can add, edit, or delete backup rules. Any changes that you make to a backup plan have no effect on existing backups created by the backup plan. The changes apply only to backups that are created in the future.

For example, when you update the retention period in a backup rule, the retention period of backups created before you made the update remain the same. Any backups that are created by that rule going forward reflect the updated retention period.

To edit a backup plan using the AWS Backup console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Backup plans**.
3. Under the second pane, **Backup plans**, existing back plans are displayed. Select the underlined link in the column **Backup plan name** to see details of the chosen backup plan.
4. You can edit a backup rule, view resource assignments, view backup jobs, manage tags, or change Windows VSS settings.
5. To update a backup rule, select the name of the backup rule.

   Select **Manage tags** to add or delete tags.

   Select **Edit** next to **Advanced backup settings** to turn Windows VSS on or off.
6. Change the setting(s) you prefer, and then select **Save**.

Understanding backup plan summary

The backup plan summary feature provides a comprehensive view of your backup plan configuration to help you validate that your backup strategy meets your requirements. The summary includes three key components that address common customer challenges when managing complex backup configurations across multiple AWS services.

Summary components

Each backup plan summary contains the following information:

- **Next 10 scheduled backup runs**: Shows when your backup rules will execute, including merged rules and proper timezone handling
- **Feature compatibility notes**: Information about which services support point-in-time recovery (PITR) and cold tier storage if these are part of your backup plan rule.

Scheduled backup runs preview

The scheduled runs preview shows up to 10 future backup executions based on your backup plan rules. This preview helps you understand when backups will occur and how multiple rules interact with each other.

How multiple rules are merged

When you have multiple backup rules in a plan, AWS Backup applies merging logic to determine the final backup behavior:

- **Overlapping backup windows**: It is best practice not to have two rules with conflicting rules. When two rules would execute at the same time, AWS Backup selects one.
- **PITR and snapshot combinations**: PITR applies to supported resources, while snapshots apply to all resources in the selection
- **Multiple time zones**: Each rule respects its configured timezone for execution scheduling
- **Cross-region copy actions**: Copy actions are preserved and displayed with their destination information. Note that if a copy rule is attached to a conflicting schedule rule, the copy rule will only run if the schedule rule has run.

Example Rule merging example

Consider a backup plan with two rules:

- Rule 1: Daily snapshot at 8am PDT, 35-day retention
- Rule 2: Weekly snapshot at 8am PDT, 90-day retention

When both rules would execute on the same day (such as Sunday), the preview shows Rule 2 because it has longer retention (90 days vs 35 days).

Execution types

The scheduled runs preview displays different execution types based on your rule configuration:

- `CONTINUOUS`: Continuous backups for supported services
- `SNAPSHOTS`: Snapshot (periodic) backups for all resources in the selection
- `CONTINUOUS_AND_SNAPSHOTS`: Point-in-time-restore (PITR) for resource types which support it; snapshot (periodic) for all other resource types.

Feature compatibility validation

The summary includes compatibility reminders to help you understand which AWS services support advanced backup features. This prevents configuration errors and unexpected costs.

Point-in-time recovery support

When you enable continuous backups in a rule, the summary reminds you that PITR is only supported for:

- Aurora databases
- Amazon RDS databases
- SAP HANA databases
- Amazon S3 buckets

Cold tier storage support

When you configure cold tier transitions, the summary shows which services support this feature:

- Amazon DynamoDB tables
- Amazon Elastic File System file systems
- SAP HANA databases
- Amazon Timestream databases
- VMware virtual machines

Accessing backup plan summaries

You can view backup plan summaries in the AWS Backup console after creating or updating a backup plan. The summary appears in the plan details page and provides immediate feedback about your configuration.

To view a backup plan summary

1. Open the AWS Backup console at https://console.aws.amazon.com/backup/.
2. In the navigation pane, choose **Backup plans**.
3. Choose the backup plan you want to review.
4. The backup plan summary appears in the plan details page. Expand the summary to see scheduled runs and compatibility notes.

You can also retrieve backup plan information using the AWS CLI:

```
aws backup get-backup-plan --backup-plan-id 012345678
```

Select AWS services to backup

Opt in to services, then assign resources

AWS Backup works with many different AWS services. Before you decide which services to include in backup plan, use the AWS Backup console or AWS CLI to opt in to using AWS Backup to work with those services.

Then, in each backup plan, specify in the console or through CLI which resource types to include in that plan.

For example, you can opt in to all services which AWS Backup supports, then include only Amazon S3 buckets and Aurora clusters in a backup plan.

Topics

- AWS Backup service opt-in
- Backup plan resource assignment
- Assign resources using the AWS Backup console
- Assign resources with AWS CLI

- Assign AWS Backup resources through CloudFormation
- Backup plan resource assignments quotas

AWS Backup service opt-in

Service opt-in through the AWS Backup console

To configure the AWS services to use with AWS Backup

1. Sign in to the AWS Management Console, and open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Settings**.
3. On the **Service opt-in** page, choose **Configure resources**.
4. On the **Configure resources** page, use the toggle switches to enable or disable the services that are used with AWS Backup. Choose **Confirm** when your services are configured. Make sure that the AWS service you're opting in is available in your AWS Region.

Service opt-in through AWS CLI

Use the `aws backup update-region-settings` command to change the services (resource types) your account or organization will use AWS Backup to orchestrate backup creation. Use the `aws backup describe-region-settings` command to determine which services you have opted into in a specific Region.

Backup plan resource assignment

Through the AWS Backup console or through AWS CLI, the resource assignment in your backup plan specifies which resources AWS Backup will include. AWS Backup provides both simple default settings and fine-grained controls to assign resources.

You can assign resources in the following ways:

- Explicitly assign resource types to the backup plan
- Include all resources (AWS Backup will then scan for all supported resource types)
- Use tags to include or exclude resources

If you only use tags for resource assignment, then the service opt-in settings will still apply.

You can further refine the resource assignment using conditions and tags. There are some limits on the number of ARNs, conditions, and tags that can be used in a single resource assignment.

Resource selection through CLI is based on service names and resource types. See Assign resources with AWS CLI for considerations about resource election.

Assign resources using the AWS Backup console

To navigate to the Assign resources page:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Choose **Backup plans**.
3. Choose **Create Backup plan**.
4. Select any template in the **Choose template** dropdown list, then choose **Create plan**.
5. Type in a **Backup plan name**.
6. Choose **Create plan**.
7. Choose **Assign resources**.

To begin your resource assignment, in the General section:

1. Type in a **Resource assignment name**.
2. Choose the **Default role** or **Choose an IAM role**.

> **Note**
>
> If you choose an IAM role, verify that it has permission to back up all the resources you are about assign. If your role encounters a resource that it doesn't have permission to back up, your backup plan will fail.

> **Note**
>
> If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

To assign your resources, in the **Assign resources** section, choose one of the two options under **Define resource selection**:

- **Include all resource types**. This option configures your backup plan to protect all current and future AWS Backup-supported resources assigned to your backup plan. Use this option to quickly and easily protect your data estate.

  When you choose this option, you can optionally **Refine selection using tags** as the next step.

- **Include specific resource types**. When you choose this option, you must **Select specific resource types** with the following steps:

  1. Using the **Select resource types** dropdown menu, assign one or more resource types.

     Once you finish, AWS Backup presents you the list of resource types you selected and its default setting, which is to protect all resources for each selected resource type.

  2. Optionally, if you want to exclude specific resources from a resource type you selected:

     1. Use the **Choose resources** dropdown menu and deselect the default option.
     2. Select the specific resources to assign to your backup plan.

  3. Optionally, you can **Exclude specific resource IDs from the selected resource types**. Use this option if you want to exclude one or a few resources out of many, because doing so might be faster than selecting many resources during the previous step. You must include a resource type before you can exclude resources from that resource type. Exclude a resource ID using the following steps:

     1. Under **Exclude specific resource IDs from the selected resource types**, choose one or more of the resource types that you included using **Select resource types**.
     2. For each resource type, use the **Choose resources** menu to select one or more resources to exclude.

In addition to your previous choices, you can make even more granular selections using the optional **Refine selection using tags** feature. This feature allows you to refine your current selection to include a subset of your resources using tags.

Tags are key-value pairs that you can assign to specific resources to help you identify, organize, and filter your resources. Tags are case sensitive. For more information about tags, see Tagging your AWS resources.

When you refine your selection using two or more tags, the effect is an AND condition. For example, if you refine your selection using two tags, `env: prod` and `role: application`, you only assign resources with BOTH tags to your backup plan.

To refine your selection using tags:

1. Under **Refine selection using tags**, choose a **Key** from the list.
2. Choose a **Condition for value** from the list.
   - *Value* refers to the next input, the value of your key-value pair.
   - **Condition** can be `Equals`, `Contains`, `Begins with`, or `Ends with`, or their inverse: `Does not equal`, `Does not contain`, `Does not begin with`, or `Does not end with`.
3. Choose a **Value** from the list.

   > **Important**
   >
   > Wildcard conditions such as "Begins with" or "Does not begin with" do not match spaces. For example, `prod*` matches `prod_server` but not `prod server`. To match values with spaces, include each space explicitly in the pattern (for example, `prod *` matches `prod server`). We recommend avoiding spaces in tag values used with wildcard conditions.

4. To further refine using another tag, choose **Add tag**.

Assign resources with AWS CLI

Filter by services or resource types

Resource selection is based on service names and resource types. The method of resource selection determines whether a resource is included in the backup. This inclusion depends on service names, resource types, and opt-in settings.

Selection by service name

When you specify only a service name in the resource selection, the backup inclusion depends on the opt-in setting for the underlying resource types. For example, with `arn:aws:ec2:*`, EC2 instances will be included in the backup only if the opt-in setting for the EC2 resource type is enabled.

Selection by resource type

If you specify the resource selection directly with the resource type, it will be included in the backup regardless of the opt-in setting for that particular service. For example, with `arn:aws:ec2:::instance/*`, EC2 instances will be backed up regardless of the opt-in setting.

Shared resource types

When multiple resources share the same resource type, you need to enable opt-in settings for specific resource types to initiate backups.

Example

Aurora and RDS Clusters share the ARN format: `arn:aws:rds:::cluster:*`. To backup Aurora databases, you must enable the opt-in setting for Aurora.

FSx and FSx for OpenZFS share the ARN format `arn:aws:fsx:::file-system/*`. Enable the respective opt-in settings to backup these file systems.

Use a JSON to define backup plan resource assignment

You can define a resource assignment in a JSON document.

You can specify conditions, tags, or resources to define what will be included in your backup plan. For more information to help you determine which parameters to include, see `BackupSelection`.

This sample resource assignment assigns all Amazon EC2 instances to the backup plan *BACKUP-PLAN-ID*:

```
{
  "BackupPlanId":"BACKUP-PLAN-ID",
  "BackupSelection":{
    "SelectionName":"resources-list-selection",
    "IamRoleArn":"arn:aws:iam::ACCOUNT-ID:role/IAM-ROLE-ARN",
    "Resources":[
      "arn:aws:ec2:*:*:instance/*"
    ]
  }
}
```

Assuming this JSON is stored as `backup-selection.json`, you can assign these resources to your backup plan using the following CLI command:

```
aws backup create-backup-selection --cli-input-json file://PATH-TO-FILE/backup-selection.json
```

The following are example resource assignments, along with the corresponding JSON document. To make this table easier for you to read, the examples omit the fields `"BackupPlanId"`, `"SelectionName"`, and `"IamRoleArn"`. The wildcard `*` represents zero or more non-whitespace characters.

> **Important**
>
> The `*` wildcard does not match whitespace characters such as spaces. If your tag values contain spaces, a single `*` will not match the entire value. For example, `prod*` will match `prod_server` but will not match `prod server`. To match a value with spaces, each space must be explicitly included in the pattern (for example, `prod *` matches `prod server`). We recommend avoiding spaces in tag values used with wildcard conditions.

Example: Select all resources in my account

```
{
  "BackupSelection":{
    "Resources":[
      "*"
    ]
  }
}
```

Example: Select all resources in my account, but exclude EBS volumes

```
{
  "BackupSelection":{
    "Resources":[
      "*"
    ],
    "NotResources":[
      "arn:aws:ec2:*:*:volume/*"
    ]
  }
}
```

Example: Select all resources tagged with "backup":"true", but exclude EBS volumes

```
{
  "BackupSelection":{
    "Resources":[
      "*"
    ],
    "NotResources":[
      "arn:aws:ec2:*:*:volume/*"
    ],
    "Conditions":{
      "StringEquals":[
        {
          "ConditionKey":"aws:ResourceTag/backup",
          "ConditionValue":"true"
        }
      ]
    }
  }
}
```

Important

RDS, Aurora, Neptune, and DocumentDB ARNs start with `arn:aws:rds:`. Refine your selection with tags and conditional operators if you don't intend to include all those types.

Example: Select all EBS volumes and RDS DB instances tagged with both "backup":"true" and "stage":"prod"

The Boolean arithmetic is similar to that in IAM policies, with those in "Resources" combined using a Boolean OR and those in `"Conditions"` combined with a Boolean AND.

The `"Resources"` expression `"arn:aws:rds:*:*:db:*"` only selects RDS DB instances because there are no corresponding Aurora, Neptune, or DocumentDB resources.

```
{
  "BackupSelection":{
    "Resources":[
      "arn:aws:ec2:*:*:volume/*",
      "arn:aws:rds:*:*:db:*"
```

```
    ],
    "Conditions":{
      "StringEquals":[
        {
          "ConditionKey":"aws:ResourceTag/backup",
          "ConditionValue":"true"
        },
        {
          "ConditionKey":"aws:ResourceTag/stage",
          "ConditionValue":"prod"
        }
      ]
    }
  }
}
```

Example: Select all EBS volumes and RDS instances tagged with "backup":"true" but not "stage":"test"

```
{
  "BackupSelection":{
    "Resources":[
      "arn:aws:ec2:*:*:volume/*",
      "arn:aws:rds:*:*:db:*"
    ],
    "Conditions":{
      "StringEquals":[
        {
          "ConditionKey":"aws:ResourceTag/backup",
          "ConditionValue":"true"
        }
      ],
      "StringNotEquals":[
        {
          "ConditionKey":"aws:ResourceTag/stage",
          "ConditionValue":"test"
        }
      ]
    }
  }
}
```

Example: Select all resources tagged with "key1" and a value which begins with "include" but not with "key2" and value that contains the word "exclude"

You can use the wildcard character at the start, end, and middle of a string. Note the use of the wildcard character (*) in `include*` and `*exclude*` in the example above. You can also use the wildcard character in the middle of a string as shown in the previous example, `arn:aws:rds:*:*:db:*`.

```
{
  "BackupSelection":{
    "Resources":[
      "*"
    ],
    "Conditions":{
      "StringLike":[
        {
          "ConditionKey":"aws:ResourceTag/key1",
          "ConditionValue":"include*"
        }
      ],
      "StringNotLike":[
        {
          "ConditionKey":"aws:ResourceTag/key2",
          "ConditionValue":"*exclude*"
        }
      ]
    }
  }
}
```

Example: Select all resources tagged with "backup":"true" except FSx file systems and RDS, Aurora, Neptune, and DocumentDB resources

Items in `NotResources` are combined using the Boolean OR.

```
{
  "BackupSelection":{
    "Resources":[
      "*"
    ],
    "NotResources":[
      "arn:aws:fsx:*",
```

```
      "arn:aws:rds:*"
    ],
    "Conditions":{
      "StringEquals":[
        {
          "ConditionKey":"aws:ResourceTag/backup",
          "ConditionValue":"true"
        }
      ]
    }
  }
}
```

Example: Select all resources tagged with a tag "backup" and any value

```
{
  "BackupSelection":{
    "Resources":[
      "*"
    ],
    "Conditions":{
      "StringLike":[
        {
          "ConditionKey":"aws:ResourceTag/backup",
          "ConditionValue":"*"
        }
      ]
    }
  }
}
```

Example: Select all FSx file systems, the Aurora cluster "my-aurora-cluster", and all resources tagged with "backup":"true", except for resources tagged with "stage":"test"

```
{
  "BackupSelection":{
    "Resources":[
      "arn:aws:fsx:*",
      "arn:aws:rds:*:*:cluster:my-aurora-cluster"
    ],
    "ListOfTags":[
      {
        "ConditionType":"StringEquals",
```

```
        "ConditionKey":"backup",
        "ConditionValue":"true"
      }
    ],
    "Conditions":{
      "StringNotEquals":[
        {
          "ConditionKey":"aws:ResourceTag/stage",
          "ConditionValue":"test"
        }
      ]
    }
  }
}
```

Example: Select all resources tagged with tag "backup":"true" except for EBS volumes tagged with "stage":"test"

Use two CLI commands to create two selections to select this group of resources. The first selection applies to all resources except for EBS volumes. The second selection applies to EBS volumes.

```
{
  "BackupSelection":{
    "Resources":[
      "*"
    ],
    "NotResources":[
      "arn:aws:ec2:*:*:volume/*"
    ],
    "Conditions":{
      "StringEquals":[
        {
          "ConditionKey":"aws:ResourceTag/backup",
          "ConditionValue":"true"
        }
      ]
    }
  }
}
```

```
{
  "BackupSelection":{
    "Resources":[
```

```
      "arn:aws:ec2:*:*:volume/*"
    ],
    "Conditions":{
      "StringEquals":[
        {
          "ConditionKey":"aws:ResourceTag/backup",
          "ConditionValue":"true"
        }
      ],
      "StringNotEquals":[
        {
          "ConditionKey":"aws:ResourceTag/stage",
          "ConditionValue":"test"
        }
      ]
    }
  }
}
```

Assign AWS Backup resources through CloudFormation

This end-to-end CloudFormation template creates a resource assignment, a backup plan, and a destination backup vault:

- A backup vault named *`CloudFormationTestBackupVault`*.
- A backup plan named *`CloudFormationTestBackupPlan`*. This plan will run two contains two backup rules, both of which take backups daily at 12 noon UTC and retain them for 210 days.
- A resource selection named *`BackupSelectionName`*.
- • The resource assignment backs up the following resources:
    - Any resource tagged with the key-value pair `backupplan:dsi-sandbox-daily`.
    - Any resource tagged with the value `prod` or values beginning with `prod/`.
  - The resource assignment does not back up the following resources:
    - Any RDS, Aurora, Neptune, or DocumentDB cluster.
    - Any resource tagged with the value `test` or values beginning with `test/`.

```
Description: "Template that creates Backup Selection and its dependencies"
Parameters:
  BackupVaultName:
```

```
    Type: String
    Default: "CloudFormationTestBackupVault"
  BackupPlanName:
    Type: String
    Default: "CloudFormationTestBackupPlan"
  BackupSelectionName:
    Type: String
    Default: "CloudFormationTestBackupSelection"
  BackupPlanTagValue:
    Type: String
    Default: "test-value-1"
  RuleName1:
    Type: String
    Default: "TestRule1"
  RuleName2:
    Type: String
    Default: "TestRule2"
  ScheduleExpression:
    Type: String
    Default: "cron(0 12 * * ? *)"
  StartWindowMinutes:
    Type: Number
    Default: 60
  CompletionWindowMinutes:
    Type: Number
    Default: 120
  RecoveryPointTagValue:
    Type: String
    Default: "test-recovery-point-value"
  MoveToColdStorageAfterDays:
    Type: Number
    Default: 120
  DeleteAfterDays:
    Type: Number
    Default: 210
Resources:
  CloudFormationTestBackupVault:
    Type: "AWS::Backup::BackupVault"
    Properties:
      BackupVaultName: !Ref BackupVaultName
  BasicBackupPlan:
    Type: "AWS::Backup::BackupPlan"
    Properties:
      BackupPlan:
```

```
        BackupPlanName: !Ref BackupPlanName
        BackupPlanRule:
          - RuleName: !Ref RuleName1
            TargetBackupVault: !Ref BackupVaultName
            ScheduleExpression: !Ref ScheduleExpression
            StartWindowMinutes: !Ref StartWindowMinutes
            CompletionWindowMinutes: !Ref CompletionWindowMinutes
            RecoveryPointTags:
              test-recovery-point-key-1: !Ref RecoveryPointTagValue
            Lifecycle:
              MoveToColdStorageAfterDays: !Ref MoveToColdStorageAfterDays
              DeleteAfterDays: !Ref DeleteAfterDays
          - RuleName: !Ref RuleName2
            TargetBackupVault: !Ref BackupVaultName
            ScheduleExpression: !Ref ScheduleExpression
            StartWindowMinutes: !Ref StartWindowMinutes
            CompletionWindowMinutes: !Ref CompletionWindowMinutes
            RecoveryPointTags:
              test-recovery-point-key-1: !Ref RecoveryPointTagValue
            Lifecycle:
              MoveToColdStorageAfterDays: !Ref MoveToColdStorageAfterDays
              DeleteAfterDays: !Ref DeleteAfterDays
      BackupPlanTags:
        test-key-1: !Ref BackupPlanTagValue
    DependsOn: CloudFormationTestBackupVault

  TestRole:
    Type: "AWS::IAM::Role"
    Properties:
      AssumeRolePolicyDocument:
        Version: "2012-10-17"
        Statement:
          - Effect: "Allow"
            Principal:
              Service:
                - "backup.amazonaws.com"
            Action:
              - "sts:AssumeRole"
      ManagedPolicyArns:
        - !Sub "arn:${AWS::Partition}:iam::aws:policy/service-
role/AWSBackupServiceRolePolicyForBackup"
  BasicBackupSelection:
    Type: 'AWS::Backup::BackupSelection'
    Properties:
```

```
      BackupPlanId: !Ref BasicBackupPlan
      BackupSelection:
        SelectionName: !Ref BackupSelectionName
        IamRoleArn: !GetAtt TestRole.Arn
        ListOfTags:
          - ConditionType: STRINGEQUALS
            ConditionKey: backupplan
            ConditionValue: dsi-sandbox-daily
        NotResources:
          - 'arn:aws:rds:*:*:cluster:*'
        Conditions:
          StringEquals:
            - ConditionKey: 'aws:ResourceTag/path'
              ConditionValue: prod
          StringNotEquals:
            - ConditionKey: 'aws:ResourceTag/path'
              ConditionValue: test
          StringLike:
            - ConditionKey: 'aws:ResourceTag/path'
              ConditionValue: prod/*
          StringNotLike:
            - ConditionKey: 'aws:ResourceTag/path'
              ConditionValue: test/*
```

Backup plan resource assignments quotas

The following quotas apply to a single resource assignment:

- 500 Amazon Resource Names (ARNs) without wildcards
- 30 ARNs with wildcard expressions
- 30 conditions
- 30 tags per resource assignment (and an unlimited number of resources per tag)

Backup vaults

In AWS Backup, a *backup vault* is a container that stores and organizes your backups.

When creating a backup vault, you must specify the AWS Key Management Service (AWS KMS) encryption key that encrypts some of the backups placed in this vault. Encryption for other backups is managed by their source AWS services. For more information about encryption, see the chart in Encryption for backups in AWS.

The following sections provide an overview of how to manage your backup vaults in AWS Backup.

Topics

- Backup vault creation and deletion
- Logically air-gapped vault
- Vault access policies
- AWS Backup Vault Lock

Backup vault creation and deletion

You must create at least one vault before creating a backup plan or starting a backup job.

When you first use the **Backup Vaults** page of the AWS Backup console in an AWS Region, the console automatically creates a default vault for the Region.

> **Note**
>
> If a Default vault exists in a Region, AWS Backup periodically identifies orphaned AMIs - AMIs created by AWS Backup that are no longer associated with a recovery point - and places them as expired recovery points in the Default vault. You can then delete these expired recovery points from the Default vault, which also deletes the associated AMI and snapshots. If the Default vault does not exist in the Region, create one so that orphaned AMIs can be placed there as expired recovery points. You must then delete these recovery points manually.

However, if you use AWS Backup through the AWS CLI, AWS SDK, or CloudFormation, a default vault is not created. You must create your own vault.

Required permissions

You must have the following permissions to create a backup vault using AWS Backup.

JSON

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Sid": "AllowKMSOperations",
      "Effect": "Allow",
      "Action": [
        "kms:CreateGrant",
        "kms:DescribeKey",
        "kms:RetireGrant",
        "kms:Decrypt",
        "kms:GenerateDataKey"
      ],
      "Resource": "arn:aws:kms:us-
east-1:444455556666:key/1234abcd-12ab-34cd-56ef-1234567890ab"
    },
    {
      "Sid": "AllowCreateBackupVault",
      "Effect": "Allow",
      "Action": [
        "backup:CreateBackupVault"
      ],
      "Resource": "arn:aws:backup:us-east-1:444455556666:backup-vault:*"
    },
    {
      "Sid": "AllowMountCapsule",
      "Effect": "Allow",
      "Action": [
        "backup-storage:MountCapsule"
      ],
      "Resource": "*"
    }
  ]
}
```

Creating a backup vault (console)

Instead of using the default backup vault that is automatically created for you on the AWS Backup console, you can create specific backup vaults to save and organize groups of backups in the same vault.

To create a backup vault

1. On the AWS Backup console, in the navigation pane, choose **Backup vaults**.

   > **Note**
   >
   > If the navigation pane is not visible on the left side, you can open it by choosing the menu icon in the upper-left corner of the AWS Backup console.

2. Choose **Create backup vault**.
3. Enter a name for your backup vault. You can name your vault to reflect what you will store in it, or to make it easier to search for the backups you need. For example, you could name it **`FinancialBackups`**.
4. Select an AWS Key Management Service (AWS KMS) key. You can use either a key that you already created, or select the default AWS Backup KMS key.

   > **Note**
   >
   > The AWS KMS key that is specified here applies only to backups of services that support AWS Backup independent encryption. To see the list of resources types that support AWS Backup independent encryption, see the "Full management" section of the Feature availability by resource table.

5. Optionally, add tags that will help you search for and identify your backup vault. For example, you could add a **`BackupType:Financial`** tag.
6. Choose **Create Backup vault**.
7. In the navigation pane, choose **Backup vaults**, and verify that your backup vault has been added.

> **Note**
>
> You can now edit a backup rule in one of your backup plans to store backups created by that rule in the backup vault you just created.

Creating a backup vault (programmatically)

The following AWS Command Line Interface command creates a backup vault:

```
aws backup create-backup-vault --backup-vault-name test-vault
```

You can also specify the following configurations for a backup vault.

Backup vault name

Backup vault names are case sensitive. They must contain from 2 to 50 alphanumeric characters, hyphens, or underscores.

AWS KMS encryption key

The AWS KMS encryption key protects your backups in this backup vault. By default, AWS Backup creates a KMS key with the alias `aws/backup` for you. You can choose that key or choose any other key in your account (cross-account KMS keys can be used via CLI).

You can create a new encryption key by following the Creating Keys procedure in the *AWS Key Management Service Developer Guide*.

After you create a backup vault and set the AWS KMS encryption key, you can no longer edit the key for that backup vault.

The encryption key that is specified in an AWS Backup vault applies to the backups of certain resource types. For more information about backup encryption, see Encryption for backups in AWS Backup in the Security section. Backups of all other resource types are backed up using the key that is used to encrypt the source resource.

Backup vault tags

These tags are associated with the backup vault to help you organize and track your backup vaults.

Delete a vault

To guard against accidental or malicious mass deletion, you can delete a backup vault in AWS Backup only after you delete (or your backup plan lifecycles) all the recovery points in your backup vault. To delete your recovery points manually, see Deleting backups.

When you delete a backup vault, update your backup plans to point to new backup vaults. A backup plan that points to a deleted backup vault will cause the backup creation to fail.

You can't delete the default backup vault or the Amazon EFS automatic backup vault using the AWS Management Console. You can delete a default backup vault using the AWS CLI if it is empty. However, if you open the AWS Backup console and select that Region, the console recreates the default backup vault. You can delete unused snapshots in the Amazon EFS automatic backup vault.

To delete a backup vault using the AWS Backup console

1. Sign in to the AWS Management Console, and open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Backup vaults**.
3. Choose the name of the backup vault to open its details page.
4. Choose and delete any backups that are associated with the backup vault.
5. Choose **Delete vault**. When prompted for confirmation, enter the vault name and then choose **Delete Backup vault**.

Logically air-gapped vault

Overview of logically air-gapped vaults

AWS Backup offers a secondary type of vault which can store backups in a container with additional security features. A **logically air-gapped vault** is a specialized vault which provides increased security beyond a standard backup vault, as well as the ability to share vault access to other accounts so that recovery time objectives (RTOs) can be faster and more flexible in case of an incident that requires rapid restoration of resources.

Logically air-gapped vaults come equipped with additional protection features; each vault is encrypted with either an AWS owned key (default) or optionally with a customer-managed KMS key, and each vault is equipped with AWS Backup Vault Lock's compliance mode. The encryption

key type information is visible through AWS Backup APIs and console for transparency and compliance reporting.

You can integrate your logically air-gapped vaults with Multi-party approval (MPA) to enable recovery of backups in the vaults even if the vault-owning account is inaccessible, which helps to maintain business continuity. Further more, you can choose to integrate with AWS Resource Access Manager (RAM) to share a logically air-gapped vault with other AWS accounts (including accounts in other organizations) so that the backups stored within the vault can be restored from an account with which the vault is shared, if needed for data loss recovery or restore testing. As part of this added security, a logically air-gapped vault stores its backups in an AWS Backup service owned account (which results in backups shown as shared outside your organization in modify attribute items in AWS CloudTrail logs).

In the scenario where your logically air-gapped vault owning account is closed (maliciously or otherwise), you can still access backups in the vault (restore or copy them) via MPA until the post-closure period ends. After the post-closure period expires, the backups are no longer accessible. During the post-closure period, you can reference the AWS Account Management documentation to regain control of your account while working on recovery.

For greater resiliency, we recommend creating cross-Region copies in logically air-gapped vaults in same or separate accounts. However, if you want to reduce storage costs by maintaining just a single copy, you can use Primary backups to logically air-gapped vaults, after onboarding to AWS MPA.

You can view the storage pricing for backups of supported services in a logically air-gapped vault on the AWS Backup pricing page.

See Feature availability by resource for resource types you can copy to a logically air-gapped vault.

Topics

- Use case for logically air-gapped vaults
- Compare and contrast with a standard backup vault
- Create a logically air-gapped vault
- View logically air-gapped vault details
- Creating backups in a logically air-gapped vault
- Share a logically air-gapped vault
- Restore a backup from a logically air-gapped vault

- Delete a logically air-gapped vault
- Additional programmatic options for logically air-gapped vaults
- Understanding encryption key types for logically air-gapped vaults
- Usage of service-owned key
- Considerations for security auto-remediation
- Troubleshoot a logically air-gapped vault issue
- Primary backups to logically air-gapped vaults
- Multi-party approval for logically air-gapped vaults

Use case for logically air-gapped vaults

A logically air-gapped vault is a secondary vault that serves as part of a data protection strategy. This vault can help enhance your organization's retention strategy and recovery when you desire a vault for your backups that

- Is automatically set with a vault lock in compliance mode
- By default offers encryption with an AWS owned key. Optionally you can provide a customer managed key
- Contains backups which, through AWS RAM or MPA, can be shared with and restored from a different account than the one that created the backup

Considerations and limitations

- Unencrypted Amazon Aurora, Amazon DocumentDB, and Amazon Neptune clusters are not supported for logically air-gapped vault, as they do not support encryption of unencrypted DB cluster snapshots.
- Amazon EC2 offers EC2 Allowed AMIs. If this setting is enabled in your account, add the alias `aws-backup-vault` to your allowlist.

  If this alias is not included, copy operations from a logically air-gapped vault to a backup vault and restore operations of EC2 instances from a logically air-gapped vault will fail with an error message such as "Source AMI ami-xxxxxx not found in Region."
- The ARN (Amazon Resource Name) of a recovery point stored in a logically air-gapped vault will have `backup` in place of the underlying resource type. For example, if the original ARN begins with `arn:aws:ec2:`*`region`*`::image/ami-*` , then the ARN of the recovery point in

the logically air-gapped vault will be `arn:aws:backup:region:account-id:recovery-point:*`.

You can use the CLI command `list-recovery-points-by-backup-vault` to determine the ARN.

- After an Amazon EBS backup in a logically air-gapped vault transitions to the archive tier, you can't copy it to another vault or AWS Region. **This includes copies from a restore access backup vault through Multi-party approval.** To keep a copy in another account or Region, copy the backup while it is still in warm storage.

Compare and contrast with a standard backup vault

A **backup vault** is the primary and standard type of vault used in AWS Backup. Each backup is stored in a backup vault when the backup is created. You can assign resource-based policies to manage backups stored in the vault, such as the lifecycle of backups stored within the vault.

A **logically air-gapped vault** is a specialized vault with additional security and flexible sharing for faster recovery time (RTO). This vault stores primary backups or copies of backups that were initially created and stored within a standard backup vault.

Backup vaults are encrypted with a key, a security mechanism that limits access to intended users. These keys can be customer managed or AWS managed. See Copy encryption for encryption behavior during copy jobs, including copying into a logically air-gapped vault.

Additionally, a backup vault can have additional security through a vault lock; logically air-gapped vaults come equipped by a vault lock in compliance mode.

Similar to backup vaults, logically air-gapped vaults also support restricted tags for Amazon EC2 backups.

Feature	Backup vault	Logically air-gapped vault
AWS Backup Audit Manager	You can use AWS Backup Audit Manager Controls and remediation to monitor your backup vaults.	Ensure a backup of a specific resource is stored in at least one logically air-gapped vault on a schedule you determine , in addition to controls available to standard vaults.

Feature	Backup vault	Logically air-gapped vault
Billing	Storage and data transfer charges for resources fully managed by AWS Backup occur under "AWS Backup". Other resource type storage and data transfer charges will occur under their respective services. For example, Amazon EBS backups will show under "Amazon EBS"; Amazon S3 backups will show under "AWS Backup".	All billing charges from these vaults (storage or data transfer) occur under "AWS Backup".
Regions	Available in all Regions in which AWS Backup operates	Available in most Regions supported by AWS Backup. Not currently available in China (Beijing), China (Ningxia), AWS GovCloud (US-East), or AWS GovCloud (US-West).
Resources	Can store copies of backups for most resource types that support cross-account copy.	See the logically air-gapped vault column in Feature availability by resource for resources that can be copied to this vault.
Restore	Backups can be restored by the same account to which the vault belongs.	Backups can be restored by a different account than the one to which the vault belongs if the vault is shared with that separate account.

Feature	Backup vault	Logically air-gapped vault
Security	Can optionally be encrypted with a key (customer managed or AWS managed) Can optionally use a vault lock in compliance or governance mode	Can be encrypted with an AWS owned key or a customer managed key Is always locked with a vault lock in compliance mode Encryption key type informati on is preserved and visible when vaults are shared through AWS RAM or MPA
Sharing	Access can be managed through policies and AWS Organizations Not compatible with AWS RAM	Can optionally be shared across accounts using AWS RAM

Create a logically air-gapped vault

You can create a logically air-gapped vault either through the AWS Backup console or through a combination of AWS Backup and AWS RAM CLI commands.

Each logically air-gapped comes equipped with a vault lock in compliance mode. See AWS Backup Vault Lock to help determine the retention period values most appropriate for your operation

Console

Create a logically air-gapped vault from the console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, select **Vaults**.
3. Both types of vaults will be displayed. Select **Create new vault**.

4. Enter a name for your backup vault. You can name your vault to reflect what you will store in it, or to make it easier to search for the backups you need. For example, you could name it `FinancialBackups`.
5. Select the radio button for **Logically air-gapped vault**.
6. *(Optional)* Choose an encryption key. You can select a customer-managed KMS key for additional control over encryption, or use the default AWS-owned key (recommended).
7. Set the **Minimum retention period**.

   This value (in days, months, or years) is the shortest amount of time a backup can be retained in this vault. Backups with retention periods shorter than this value cannot be copied to this vault.

   The minimum value allowed is 7 days. Values for months and years meet this minimum.
8. Set the **Maximum retention period**.

   This value (in days, months, or years) is the longest amount of time a backup can be retained in this vault. Backups with retention periods greater than this value cannot be copied to this vault.
9. *(Optional)* Set the **Encryption key**.

   Specify the key to use with your vault. You can choose an **AWS owned key (managed by AWS Backup)** or enter the ARN for a **Customer managed key** that preferably belongs to a different account to which you have access. AWS Backup recommends using an AWS owned key.
10. *(Optional)* Add tags that will help you search for and identify your logically air-gapped vault. For example, you could add a `BackupType:Financial` tag.
11. Select **Create vault**.
12. Review the settings. If all settings show as you intended, select **Create logically air-gapped vault**.
13. The console will take you to the details page of your new vault. Verify the vault details are as expected.
14. Select **Vaults** to view vaults in your account. Your logically air-gapped vault will be displayed. The KMS key will be available approximately 1 to 3 minutes after the vault creation. Refresh the page to see the associated key. Once the key is visible, the vault is in an available state and can be used.

AWS CLI

Create a logically air-gapped vault from CLI

You can use AWS CLI to programmatically carry out operations for logically air-gapped vaults. Each CLI is specific to the AWS service in which it originates. Commands related to sharing are prepended with `aws ram`; all other commands should be prepended with `aws backup`.

Use the CLI command `create-logically-air-gapped-backup-vault`, modified with the following parameters:

```
aws backup create-logically-air-gapped-backup-vault
--region us-east-1 // optional
--backup-vault-name sampleName // required
--min-retention-days 7 // required Value must be an integer 7 or greater
--max-retention-days 35 // required
--encryption-key-arn arn:aws:kms:us-
east-1:123456789012:key/12345678-1234-1234-1234-123456789012 // optional
--creator-request-id 123456789012-34567-8901 // optional
```

The optional `--encryption-key-arn` parameter allows you to specify a customer-managed KMS key for vault encryption. If not provided, the vault will use an AWS-owned key.

Example CLI command to create a logically air-gapped vault:

```
aws backup create-logically-air-gapped-backup-vault
--region us-east-1
--backup-vault-name sampleName
--min-retention-days 7
--max-retention-days 35
--creator-request-id 123456789012-34567-8901 // optional
```

Example CLI command to create a logically air-gapped vault with customer-managed encryption:

```
aws backup create-logically-air-gapped-backup-vault
--region us-east-1
--backup-vault-name sampleName
--min-retention-days 7
--max-retention-days 35
--encryption-key-arn arn:aws:kms:us-
east-1:123456789012:key/12345678-1234-1234-1234-123456789012
```

```
--creator-request-id 123456789012-34567-8901 // optional
```

See CreateLogicallyAirGappedBackupVault API response elements for information after the create operation. If the operation was successful, the new logically air-gapped vault will have the VaultState of `CREATING`.

Once the creation is complete and the KMS encrypted key has been assigned, the VaultState will transition to `AVAILABLE`. Once available, the vault can be used. `VaultState` can be retrieved by calling `DescribeBackupVault` or `ListBackupVaults`.

View logically air-gapped vault details

You can see the vault details such as summary, the recovery points, the protected resources, account sharing, access policy, and tags through the AWS Backup console or the AWS Backup CLI.

Console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Select **Vaults** from the left-hand navigation.
3. Below the descriptions of vaults will be three lists, **Vaults created by this account**, **Vaults shared through RAM**, and **Vaults accessible through Multi-party approval**. Select the desired tab to view the vaults.
4. Under **Vault name**, click on the name of the vault to open the details page. You can see the summary, the recovery points, the protected resources, account sharing, access policy, and tag details.

   Details display depending on account type: Accounts which own a vault can view account sharing; accounts which do not own a vault will not be able to view account sharing. For shared vaults, the encryption key type (AWS-owned or customer-managed KMS key) is displayed in the vault summary.

AWS CLI

View details of a logically air-gapped vault through CLI

The CLI command `describe-backup-vault` can be used to obtain details about a vault. Parameter `backup-vault-name` is required; `region` is optional.

```
aws backup describe-backup-vault
--region us-east-1
--backup-vault-name testvaultname
```

Example of response:

```
{
            "BackupVaultName": "LOG-AIR-GAP-VAULT-TEST",
            "BackupVaultArn": "arn:aws:backup:us-east-1:234567890123:backup-
vault:IAD-LAGV-01",
            "VaultType": "LOGICALLY_AIR_GAPPED_BACKUP_VAULT",
            "EncryptionKeyType": "AWS_OWNED_KMS_KEY",
            "CreationDate": "2024-07-25T16:05:23.554000-07:00",
            "NumberOfRecoveryPoints": 0,
            "Locked": true,
            "MinRetentionDays": 8,
            "MaxRetentionDays": 30,
            "LockDate": "2024-07-25T16:05:23.554000-07:00"
}
```

> **Note**
>
> The `VaultType` field is not included in the API response in regions where logically air-gapped vaults are not available.

Creating backups in a logically air-gapped vault

Logically air-gapped vaults can be a copy job destination target in a backup plan or a target for an on-demand copy job. It can also be used as a primary backup target. See Primary backups to logically air-gapped vaults.

Compatible encryption

A successful copy job from a backup vault to a logically air-gapped vault requires an encryption key that is determined by the resource type being copied.

When you create or copy a backup of a fully managed resource type, the source resource can be encrypted by a customer managed key or by an AWS managed key.

When you create or copy a backup of other resource types (ones not fully managed), the source must be encrypted with a customer managed key. AWS managed keys for not fully managed resources are not supported.

Create or copy backups to a logically air-gapped vault through a backup plan

You can copy a backup (recovery point) from a standard backup vault to a logically air-gapped vault by creating a new backup plan or updating an existing one in the AWS Backup console or through the AWS CLI commands `create-backup-plan` and `update-backup-plan`. You can also create backups directly in a logically air-gapped vault by using it as a primary target. See Primary backups to logically air-gapped vaults for more details.

You can copy a backup from one logically air-gapped vault to another logically air-gapped vault on-demand (this type of backup cannot be scheduled in a backup plan). You can copy a backup from a logically air-gapped vault to a standard backup vault as long as the copy is encrypted with a customer managed key.

Archive Amazon EBS backups in a logically air-gapped vault

You can transition Amazon EBS backups in a logically air-gapped vault to Amazon EBS Snapshots Archive, a lower-cost storage tier for backups that you rarely access. In a backup rule that runs monthly or less often, add a copy action with a logically air-gapped vault as the destination. In the copy lifecycle, select Archive Amazon EBS snapshots, then set the time in warm storage and the total retention period. If you use a logically air-gapped vault as the primary backup target, select Archive Amazon EBS snapshots in the backup rule lifecycle instead. The backup transitions to the archive tier after it is copied to the vault. The total retention period must be at least 90 days longer than the time in warm storage, and it must fall within the minimum and maximum retention period of the vault.

If the vault is encrypted with a customer managed key, update the key policy before you enable this option. For more information, see the section called “Key policy for CMK encrypted logically air-gapped vault creation”.

On-demand backup copy to a logically air-gapped vault

To create a one-time on-demand copy of a backup to a logically air-gapped vault, you can copy from a standard backup vault. Cross-Region or cross-account copies are available if the resource type supports the copy type.

Copy availability

A copy of a backup can be created from the account to which the vault belongs. Accounts with which the vault has been shared have the ability to view or a restore a backup, but not to create a copy.

Only resource types that support cross-Region or cross-account copy can be included.

Console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Select **Vaults** from the left-hand navigation.
3. In the vault detail page, all recovery points within that vault are displayed. Place a check mark next to the recovery point you wish to copy.
4. Select **Actions**, and then select **Copy** from the drop-down menu.
5. On the next screen, input the details of the destination.
    a. Specify the destination Region.
    b. Destination backup vault drop-down menu displays eligible destination vaults. Select one with the type `logically air-gapped vault`
6. Select **Copy** once all details are set to your preferences.

On the **Jobs** page in the console, you can select **Copy** jobs to see current copy jobs.

AWS CLI

Use start-copy-job to copy an existing backup in a backup vault to a logically air-gapped vault.

Sample CLI input:

```
aws backup start-copy-job
--region us-east-1
--recovery-point-arn arn:aws:resourcetype:region::snapshot/snap-12345678901234567
--source-backup-vault-name sourcevaultname
--destination-backup-vault-arn arn:aws:backup:us-east-1:123456789012:backup-
vault:destinationvaultname
--iam-role-arn arn:aws:iam::123456789012:role/service-role/servicerole
```

For more information, see Copying a backup, cross-Region backup, and Cross-account backup.

Share a logically air-gapped vault

You can use AWS Resource Access Manager (RAM) to share a logically air-gapped vault with other accounts you designate. When sharing vaults, the encryption key type information (AWS-owned or customer-managed KMS key) is preserved and visible to accounts with which the vault is shared.

A vault can be shared with individual AWS account IDs only. You can share with an account in your organization or with an account in another organization. The vault cannot be shared with an entire organization or with organizational units (OUs) — only individual account IDs are supported as share principals.

Only accounts with specific IAM privileges can share and manage the sharing of vaults.

To share using AWS RAM, ensure you have the following:

- Two or more accounts that can access AWS Backup
- Vault-owning account that intends to share has necessary RAM permissions. The permission `ram:CreateResourceShare` is necessary for this procedure. The policy `AWSResourceAccessManagerFullAccess` contains all needed RAM-related permissions:
  - `backup:DescribeBackupVault`
  - `backup:DescribeRecoveryPoint`
  - `backup:GetRecoveryPointRestoreMetadata`
  - `backup:ListProtectedResourcesByBackupVault`
  - `backup:ListRecoveryPointsByBackupVault`
  - `backup:ListTags`
  - `backup:StartRestoreJob`
- At least one logically air-gapped vault

Console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Select **Vaults** from the left-hand navigation.
3. Below the descriptions of vaults will be two lists, **Vaults owned by this account** and **Vaults shared with this account**. Vaults owned by the account are eligible to be shared.
4. Under **Vault name**, select the name of the logically air-gapped vault to open the details page.

5. The **Account sharing** pane shows with which accounts the vault is being shared.
6. To begin sharing with another account or to edit accounts already being shared, select **Manage sharing**.
7. The AWS RAM console opens when **Manage sharing** is selected. For steps to share a resource using AWS RAM, see Creating a resource share in AWS RAM in the *AWS RAM User Guide*.
8. The account invited to accept an invitation to receive a share has 12 hours to accept the invitation. See Accepting and rejecting resource share invitations in the *AWS RAM User Guide*.
9. If the sharing steps are completed and accepted, the vault summary page will show under **Account sharing = "Shared - see account sharing table below"**.

AWS CLI

AWS RAM uses the CLI command `create-resource-share`. The access to this command is only available to accounts with sufficient permissions. See Creating a resource share in AWS RAM for CLI steps.

Steps 1 through 4 are conducted with the account that owns the logically air-gapped vault. Steps 5 through 8 are conducted with the account with which the logically air-gapped vault will be shared.

1. Log into the owning account OR request a user at your organization with sufficient credentials for accessing the source account completes these steps.
   - If a resource share was previously created and you wish to add an additional resource to it, use CLI `associate-resource-share` instead with the ARN of the new vault.
2. Fetch credentials of a role with sufficient permissions to share via RAM. Input these into the CLI.
   - The permission `ram:CreateResourceShare` is necessary for this procedure. The policy AWSResourceAccessManagerFullAccess contains all RAM-related permissions.
3. Use create-resource-share.
   a. Include the ARN of the logically air-gapped vault.
   b. Example input:

```
aws ram create-resource-share
--name MyLogicallyAirGappedVault
--resource-arns arn:aws:backup:us-east-1:123456789012:backup-vault:test-
vault-1
--principals 123456789012
--region us-east-1
```

c. Example output:

```
{
   "resourceShare":{
      "resourceShareArn":"arn:aws:ram:us-east-1:123456789012:resource-
share/12345678-abcd-09876543",
      "name":"MyLogicallyAirGappedVault",
      "owningAccountId":"123456789012",
      "allowExternalPrincipals":true,
      "status":"ACTIVE",
      "creationTime":"2021-09-14T20:42:40.266000-07:00",
      "lastUpdatedTime":"2021-09-14T20:42:40.266000-07:00"
   }
}
```

4. Copy the resource share ARN in the output (which is needed for subsequent steps). Give the ARN to the operator of account you are inviting to receive the share.
5. Obtain the resource share ARN

   a. If you did not perform steps 1 through 4, obtain the resourceShareArn from whomever did.

   b. Example: `arn:aws:ram:us-east-1:123456789012:resource-share/12345678-abcd-09876543`

6. In the CLI, assume credentials of the recipient account.
7. Get resource share invitation with `get-resource-share-invitations`. For more information, see Accepting and rejecting invitations in the *AWS RAM User Guide*.
8. Accept the invitation in destination (recovery) account.

   - Use `accept-resource-share-invitation` (can also `reject-resource-share-invitation`).

You can use AWS RAM CLI commands to view shared items:

- Resources you have shared:

  ```
  aws ram list-resources --resource-owner SELF --resource-type
  backup:backup-vault --region us-east-1
  ```

- Show the principal:

  ```
  aws ram get-resource-share-associations --association-type PRINCIPAL
  --region us-east-1
  ```

- Resources shared by other accounts:

  ```
  aws ram list-resources --resource-owner OTHER-ACCOUNTS --resource-type
  backup:backup-vault --region us-east-1
  ```

Restore a backup from a logically air-gapped vault

You can restore a backup stored in a logically air-gapped vault from either the account that owns the vault or from any account with which the vault is shared.

See Restoring a backup for information on how to restore a recovery point through the AWS Backup console.

Once a backup has been shared from a logically air-gapped vault to your account, you can use `start-restore-job` to restore the backup.

A sample CLI input can include the following command and parameters:

```
aws backup start-restore-job
--recovery-point-arn arn:aws:backup:us-east-1:accountnumber:recovery-
point:RecoveryPointID
--metadata {\"availabilityzone\":\"us-east-1d\"}
--idempotency-token TokenNumber
--resource-type ResourceType
--iam-role arn:aws:iam::number:role/service-role/servicerole
--region us-east-1
```

Delete a logically air-gapped vault

See delete a vault. Vaults cannot be deleted if they still contain backups (recovery points). Ensure the vault is empty of backups before you initiate a delete operation.

> **Note**
>
> A restore access backup vault is a view of an underlying logically air-gapped vault and contains no recovery points of its own. You can delete it with `DeleteBackupVault` from the recovery account even when it is in FAILED state. Vault Lock (compliance mode) does not block this deletion.

Deletion of a vault also deletes the key associated with the vault seven days after the vault is deleted in accordance with key deletion policy.

The following sample CLI command `delete-backup-vault` can be used to delete a vault.

```
aws backup delete-backup-vault
--region us-east-1
--backup-vault-name testvaultname
```

Additional programmatic options for logically air-gapped vaults

The CLI command `list-backup-vaults` can be modified to list all the vaults owned by and present in the account:

```
aws backup list-backup-vaults
--region us-east-1
```

To list just the logically air-gapped vaults, add the parameter

```
--by-vault-type LOGICALLY_AIR_GAPPED_BACKUP_VAULT
```

Include the parameter `by-shared` to filter the returned list of vaults to show only shared logically air-gapped vaults. The response will include encryption key type information for each shared vault.

```
aws backup list-backup-vaults
--region us-east-1
```

```
--by-shared
```

Example response showing encryption key type information:

```
{
    "BackupVaultList": [
        {
            "BackupVaultName": "shared-logically air-gapped-vault",
            "BackupVaultArn": "arn:aws:backup:us-east-1:123456789012:backup-
vault:shared-logically air-gapped-vault",
            "VaultType": "LOGICALLY_AIR_GAPPED_BACKUP_VAULT",
            "EncryptionKeyType": "AWS_OWNED_KMS_KEY",
            "CreationDate": "2024-07-25T16:05:23.554000-07:00",
            "Locked": true,
            "MinRetentionDays": 7,
            "MaxRetentionDays": 30
        }
    ]
}
```

Note

The `VaultType` field is not included in the API response in regions where logically air-gapped vaults are not available.

Understanding encryption key types for logically air-gapped vaults

Logically air-gapped vaults support different encryption key types, and this information is visible through AWS Backup APIs and console. When vaults are shared through AWS RAM or MPA, the encryption key type information is preserved and made visible to accounts with which the vault is shared. This transparency helps you understand the encryption configuration of vaults and make informed decisions about backup and restore operations.

Encryption key type values

The `EncryptionKeyType` field can have the following values:

- `AWS_OWNED_KMS_KEY` - The vault is encrypted with an AWS-owned key. This is the default encryption method for logically air-gapped vaults when no customer-managed key is specified.

- `CUSTOMER_MANAGED_KMS_KEY` - The vault is encrypted with a customer-managed KMS key that you control. This option provides additional control over encryption keys and access policies.

> **Note**
>
> - AWS Backup recommends using AWS owned keys with logically air-gapped vaults.
> - If your organization policy requires using a customer managed key, AWS does not recommend using keys from the same account, except for testing. For production workloads, use a customer managed key from another account in a secondary organization dedicated to recovery as a best practice. You can reference the blog Encrypt AWS Backup logically air-gapped vaults with customer-managed keys to gather more insights into setting up CMK based logically air-gapped vaults.
> - You can only select an AWS KMS encryption key during vault creation. Once created, all backups contained in the vault will be encrypted with that key. You cannot change or migrate your vaults to use a different encryption key.

Key policy for CMK encrypted logically air-gapped vault creation

When creating a logically air-gapped vault with a customer managed key, you must apply the AWS-managed policy `AWSBackupFullAccess` to your account role. This policy includes `Allow` actions that enable AWS Backup to interact with AWS KMS for grant creation on KMS keys during backup, copy, and storage operations. Additionally, you must ensure your customer managed key (if used) policy includes specific required permissions.

- The CMK must be shared with the account where the logically air-gapped vault resides

```
{
    "Sid": "Allow use of the key to create a logically air-gapped vault",
    "Effect": "Allow",
    "Principal": {
        "AWS": "arn:aws:iam::[account-id]:role/TheRoleToAccessAccount"
    },
    "Action": [
        "kms:CreateGrant",
        "kms:DescribeKey"
    ],
```

```
    "Resource": "*",
   "Condition": {
        "StringLike": {
            "kms:ViaService": "backup.*.amazonaws.com"
        }
   }
}
```

Key policy for copy/restore

To prevent job failures, review your AWS KMS key policy to ensure it includes all required permissions and doesn't contain any deny statements that could block operations. The following conditions apply:

- For all copy scenarios, the CMKs must be shared with the source copy role

```
{
    "Sid": "Allow use of the key for copy",
    "Effect": "Allow",
    "Principal": {
        "AWS": "arn:aws:iam::[source-account-id]:role/service-role/
AWSBackupDefaultServiceRole"      //[Source copy role]
    },
    "Action": [
        "kms:Encrypt",
        "kms:Decrypt",
        "kms:ReEncrypt*",
        "kms:GenerateDataKey*",
        "kms:DescribeKey"
    ],
    "Resource": "*",
   "Condition": {
         "StringLike": {
            "kms:ViaService": "backup.*.amazonaws.com"
         }
   }
},
{
    "Sid": "Allow AWS Backup to create grant on the key for copy",
    "Effect": "Allow",
    "Principal": {
```

```
        "AWS": "arn:aws:iam::[source-account-id]:role/service-role/
AWSBackupDefaultServiceRole" //[Source copy role]
    },
    "Action": [
        "kms:CreateGrant"
    ],
    "Resource": "*",
    "Condition": {
        "Bool": {
            "kms:GrantIsForAWSResource": "true"
        },
        "StringLike": {
            "kms:ViaService": "backup.*.amazonaws.com"
        }
    }
}
```

- When copying from a CMK encrypted logically air-gapped vault to a backup vault, the CMK must also be shared with the destination account SLR

```
{
    "Sid": "Allow use of the key for copy from a CMK encrypted logically air-gapped
 vault to normal backup vault",
    "Effect": "Allow",
    "Principal": {
        "AWS": ["arn:aws:iam::[source-account-id]:role/service-role/
AWSBackupDefaultServiceRole",      //[Source copy role]
                "arn:aws:iam::[destination-account-id]:role/aws-service-role/
backup.amazonaws.com/AWSServiceRoleForBackup"], //[Destination SLR]
    },
    "Action": [
        "kms:Encrypt",
        "kms:Decrypt",
        "kms:ReEncrypt*",
        "kms:GenerateDataKey*",
        "kms:DescribeKey"
    ],
    "Resource": "*"
},
{
    "Sid": "Allow AWS Backup to create grant on the key for copy",
    "Effect": "Allow",
```

```
    "Principal": {
          "AWS": ["arn:aws:iam::[source-account-id]:role/service-role/
AWSBackupDefaultServiceRole",      //[Source copy role]
                  "arn:aws:iam::[destination-account-id]:role/aws-service-role/
backup.amazonaws.com/AWSServiceRoleForBackup"], //[Destination SLR]
    },
    "Action": [
        "kms:CreateGrant"
    ],
    "Resource": "*",
    "Condition": {
        "Bool": {
            "kms:GrantIsForAWSResource": "true"
        }
    }
}
```

- When copying or restoring from a recovery account using RAM/MPA shared logically air-gapped vault

```
{
    "Sid": "Allow use of the key for copy/restore from a recovery account",
    "Effect": "Allow",
    "Principal": {
        "AWS": ["arn:aws:iam::[recovery-account-id]:role/service-role/
AWSBackupDefaultServiceRole", //[Recovery account copy/restore role]
                "arn:aws:iam::[destination-account-id]:role/aws-service-role/
backup.amazonaws.com/AWSServiceRoleForBackup"] //[Destination SLR]
    },
    "Action": [
        "kms:Encrypt",
        "kms:Decrypt",
        "kms:ReEncrypt*",
        "kms:GenerateDataKey*",
        "kms:DescribeKey"
    ],
    "Resource": "*"
},
{
    "Sid": "Allow AWS Backup to create grant on the key for copy",
    "Effect": "Allow",
    "Principal": {
```

```
        "AWS": ["arn:aws:iam::[recovery-account-id]:role/service-role/
AWSBackupDefaultServiceRole" //[Recovery account copy/restore role]
                "arn:aws:iam::[destination-account-id]:role/aws-service-role/
backup.amazonaws.com/AWSServiceRoleForBackup"], //[Destination SLR]

    },
    "Action": [
        "kms:CreateGrant"
    ],
    "Resource": "*",
    "Condition": {
        "Bool": {
            "kms:GrantIsForAWSResource": "true"
        }
    }
}
```

- When transitioning Amazon EBS Snapshot to cold storage in CMK encrypted logically air-gapped vault, the vault-owning account's SLR must have create grant permission on the CMK key policy.

```
{
    "Sid": "Allow AWS Backup to create grant on the key for snapshot archival",
    "Effect": "Allow",
    "Principal": {
        "AWS": "arn:aws:iam::[vault-owning-account-id]:role/aws-service-role/
backup.amazonaws.com/AWSServiceRoleForBackup" //[Vault owning account SLR]
    },
    "Action": [
        "kms:CreateGrant"
    ],
    "Resource": "*",
    "Condition": {
        "Bool": {
          "kms:GrantIsForAWSResource": "true"
        }
    }
}
```

IAM Role

When performing logically air-gapped vault copy operations, customers can utilize the `AWSBackupDefaultServiceRole` which includes the AWS-managed policy

AWSBackupServiceRolePolicyForBackup. However, if customers prefer to implement a least-privilege policy approach, their IAM policy must include a specific requirement:

- The source account's copy role must have access permissions to both the source and destination CMKs.

```
{
    "Version": "2012-10-17"       ,
    "Statement": [
         {
            "Sid": "KMSPermissions",
            "Effect": "Allow",
            "Action": "kms:DescribeKey",
            "Resource": [
                "arn:aws:kms:*:[source-account-id]:key/*",  //[Source logically air-
gapped vault CMK]
                "arn:aws:kms:*:[destination-account-id]:key/*"  //[Destination
 logically air-gapped vault CMK]
            ]
        },
        {
            "Sid": "KMSCreateGrantPermissions",
            "Effect": "Allow",
            "Action": "kms:CreateGrant",
            "Resource": [
                "arn:aws:kms:*:[source-account-id]:key/*",  //[Source logically air-
gapped vault CMK]
                "arn:aws:kms:*:[destination-account-id]:key/*"  //[Destination
 logically air-gapped vault CMK]
            ],
            "Condition": {
                "Bool": {
                    "kms:GrantIsForAWSResource": "true"
                }
            }
        }
    ]
}
```

Consequently, one of the most common customer errors occurs during copy when customers fail to provide sufficient permissions on their CMKs and copy roles.

Viewing encryption key types

You can view encryption key type information through both the AWS Backup console and programmatically using the AWS CLI or SDKs.

Console: When viewing logically air-gapped vaults in the AWS Backup console, the encryption key type is displayed in the vault details page under the security information section.

AWS CLI/API: The encryption key type is returned in the response of the following operations when querying logically air-gapped vaults:

- `list-backup-vaults` (including `--by-shared` for shared vaults)
- `describe-backup-vault`
- `describe-recovery-point`
- `list-recovery-points-by-backup-vault`
- `list-recovery-points-by-resource`

Considerations for vault encryption

When working with logically air-gapped vaults and encryption key types, consider the following:

- **Key selection during creation:** You can optionally specify a customer-managed KMS key when creating a logically air-gapped vault. If not specified, an AWS-owned key will be used.
- **Shared vault visibility:** Accounts with which a vault is shared can view the encryption key type but cannot modify the encryption configuration.
- **Recovery point information:** The encryption key type is also available when viewing recovery points within logically air-gapped vaults.
- **Restore operations:** Understanding the encryption key type helps you plan restore operations and understand any potential access requirements.
- **Compliance:** The encryption key type information supports compliance reporting and audit requirements by providing transparency into the encryption methods used for backup data.

Usage of service-owned key

AWS Backup creates and manages encryption keys that are used to encrypt all the backup data stored in logically air-gapped vaults, to protect and prevent loss of access of the encryption key during a data loss event.

- These keys are free of charge and do not count against the AWS KMS quotas for your account.
- A single key is only used for a specific vault and is not shared with any other account or other purpose.
- These keys are deleted once the assigned (empty) vault is also deleted.
- These keys are created using the SYMMETRIC_DEFAULT key spec.
- The default rotation policy is 90 days. You can request rotation (once every 6 months) of service-owned encryption keys for your logically air-gapped vaults through a support ticket.

Visit the AWS KMS documentation to learn more.

Considerations for security auto-remediation

When AWS Backup copies an EC2 (AMI) backup to a logically air-gapped vault, it temporarily grants `launchPermission` (on the AMI) and `createVolumePermission` (on associated EBS snapshots) to a service-owned account. These permissions are automatically revoked after the copy completes.

These operations generate `ModifyImageAttribute` and `ModifySnapshotAttribute` events in your AWS CloudTrail logs, with `userIdentity.invokedBy` set to `backup.amazonaws.com`.

If you have security auto-remediation logic (e.g., Amazon EventBridge rules with AWS Lambda) that monitors these events and revokes cross-account sharing, you must exclude events where `userIdentity.invokedBy` is `backup.amazonaws.com`. Otherwise, copy jobs to logically air-gapped vaults will fail with: "You do not have permission to access the storage of this ami."

This exclusion is safe because the copy is authorized by your vault access policies (`backup:CopyFromBackupVault` on the source vault and `backup:CopyIntoBackupVault` on the destination vault), which are evaluated before any EC2 attribute modifications occur. The temporary permissions are granted only to a fixed AWS service-owned account and are automatically revoked after the copy completes.

Example EventBridge rule event pattern that excludes AWS Backup operations:

```
{
  "source": ["aws.ec2"],
  "detail-type": ["AWS API Call via CloudTrail"],
  "detail": {
    "eventSource": ["ec2.amazonaws.com"],
    "eventName": ["ModifySnapshotAttribute", "ModifyImageAttribute"],
    "userIdentity": {
      "invokedBy": [{"anything-but": "backup.amazonaws.com"}]
    }
  }
}
```

Troubleshoot a logically air-gapped vault issue

If you encounter errors during your workflow, consult the following example errors and suggested resolutions:

`EC2 AMI copy job to logically air-gapped vault fails with permission error`

Error: `Copy job fails with "You do not have permission to access the storage of this ami."`

Possible cause: During an EC2 AMI copy job to a logically air-gapped vault, AWS Backup temporarily grants launch permission (AMI) and create volume permission (EBS snapshot) to a service-owned account, generating `ModifyImageAttribute` and `ModifySnapshotAttribute` events in your AWS CloudTrail logs. If you have security auto-remediation logic (such as EventBridge rules with Lambda) that monitors these events and automatically revokes cross-account sharing permissions, it can remove the temporary access before the copy completes.

> **Note**
>
> This can similarly happen for copy jobs for other resources like Amazon FSx.

Resolution: Update your EventBridge rule event pattern to exclude operations performed by AWS Backup. Specifically, exclude events where `userIdentity.invokedBy` is `backup.amazonaws.com` to ensure your auto-remediation logic does not revoke the temporary cross-account permissions that AWS Backup grants during the copy process.

AccessDeniedException

Error: `An error occured (AccessDeniedException) when calling the [command] operation: Insufficient privileges to perform this action."`

Possible cause: The parameter `--backup-vault-account-id` was not included when one of the following requests was run on a vault shared by RAM:

- `describe-backup-vault`
- `describe-recovery-point`
- `get-recovery-point-restore-metadata`
- `list-protected-resources-by-backup-vault`
- `list-recovery-points-by-backup-vault`

Resolution: Retry the command that returned the error, but include the parameter `--backup-vault-account-id` that specifies the account that owns the vault.

OperationNotPermittedException

Error: `OperationNotPermittedException` is returned after a `CreateResourceShare` call.

Possible cause: If you attempted to share a resource, such as a logically air-gapped vault, with another organization, you may get this exception. A vault can be shared with an account in another organization, but it cannot be shared with the other organization itself.

Resolution: Retry the operation, but specify an account as the value for `principals` instead of an organization or OU.

Encryption key type not displayed

Issue: The encryption key type is not visible when viewing a logically air-gapped vault or its recovery points.

Possible causes:

- You are viewing an older vault that was created before encryption key type support was added
- You are using an older version of the AWS CLI or SDK
- The API response does not include the encryption key type field

Resolution:

- Update your AWS CLI to the latest version
- For older vaults, the encryption key type will be populated automatically and should appear in subsequent API calls
- Verify you are using the correct API operations that return encryption key type information
- For shared vaults, verify that the vault is properly shared through AWS Resource Access Manager

"FAILED" VaultState with AccessDeniedException in CloudTrail logs

Error in CloudTrail: `"User: <assumed role> is not authorized to perform: kms:CreateGrant on this resource because the resource does not exist in this Region, no resource-based policies allow access, or a resource-based policy explicitly denies access"`

Possible causes:

- The vault was created using a customer managed key, but the assumed role does not have CreateGrant permission on the key policy required to use the key for vault creation

Resolution:

- Grant the permissions specified in the the section called “Key policy for CMK encrypted logically air-gapped vault creation” section, then retry the vault creation workflow.

Restore access backup vault in "FAILED" state cannot be removed

Error: A restore access backup vault is stuck in FAILED state. `RevokeRestoreAccessBackupVault` returns an error requiring the vault to be in AVAILABLE state, and `CreateRestoreAccessBackupVault` returns `LimitExceededException`.

Possible causes:

- The KMS key policy in the key vault account did not authorize `backup.amazonaws.com` from the recovery account at the time the vault was created.
- The IAM role used to call `CreateRestoreAccessBackupVault` did not have the `mpa:StartSession` permission required for the Multi-party approval workflow.

Resolution: Call `DeleteBackupVault` from the recovery account to remove the FAILED restore access backup vault. `RevokeRestoreAccessBackupVault` applies only to vaults in AVAILABLE state. Vault Lock does not prevent this deletion, because a restore access backup vault is a view of the underlying logically air-gapped vault and contains no recovery points of its own. After deletion, correct the KMS key policy and add the `mpa:StartSession` permission before recreating the vault.

Primary backups to logically air-gapped vaults

Overview

The logically air-gapped vault primary backup feature offers you the option to specify a logically air-gapped vault as a primary backup destination within the same account, for both scheduled and on-demand backup jobs. This eliminates the need to maintain separate copies in both a standard backup vault and a logically air-gapped vault, reducing costs and simplifying workflows while preserving the security benefits of logical air-gapping.

You can assign a logically air-gapped vault as the primary target in backup plans, organization-wide policies, or on-demand backups. Previously, to back up to a logically air-gapped vault, you first had to create a backup in a backup vault and then copy it to a logically air-gapped vault. With this feature, depending on the resource type, AWS Backup can create backups directly in your logically air-gapped vault or automatically manage temporary backups that are copied to your logically air-gapped vault and then deleted.

The behavior depends on two factors:

- Whether the resource type is supported by logically air-gapped vault.
- Whether the resource type supports full AWS Backup management.

> **Warning**
>
> We recommend integrating your logically air-gapped vaults with Multi-party approval (MPA) if you adopt this capability. This enables recovery of backups in the vault even if the vault-owning account is inaccessible.

There is no new pricing for this feature. You are only charged for backups stored in logically air-gapped vaults and for temporary snapshots (during their retention period in the system) for applicable resources at prevailing rates. See AWS Backup pricing for details.

Topics

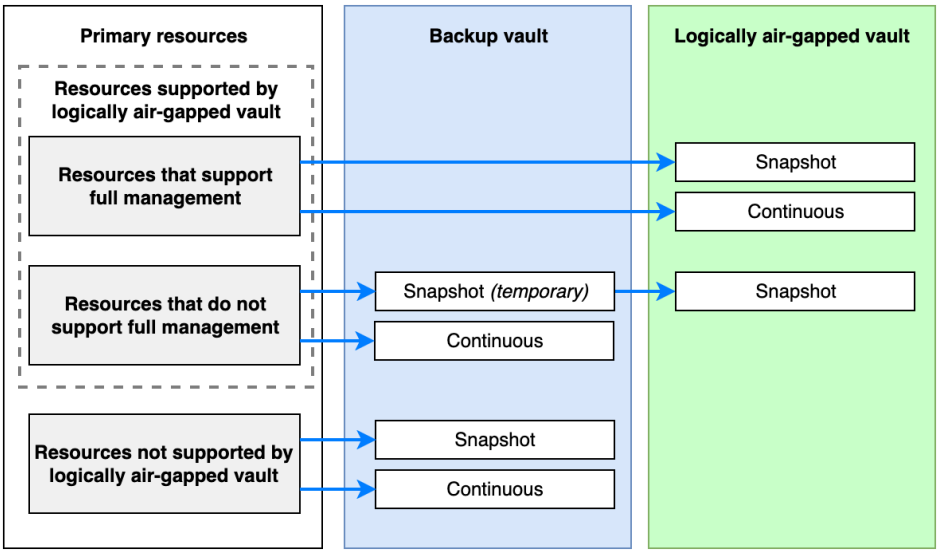
- How it works
- Cost considerations
- Configure logically air-gapped vault primary backup
- Monitor logically air-gapped vault primary backup
- Onboarding and migration
- Troubleshooting

How it works

You can onboard to this feature by updating your existing backup plan or creating a new one and adding a logically air-gapped vault ARN (field name: `TargetLogicallyAirGappedBackupVaultArn`) as a primary backup target. You can perform this operation either through the AWS Backup console or through AWS Backup CLI commands.

```
"TargetLogicallyAirGappedBackupVaultArn": "arn:aws:backup:us-
east-1:123456789012:backup-vault:AirGappedVault",
```

When you specify both a backup vault and a logically air-gapped vault as targets for your backup jobs, AWS Backup determines the appropriate workflow based on the resource type and encryption configuration.



Supported resources for primary backup to logically air-gapped vaults

To view the full list of supported resources for logically air-gapped vaults, see AWS Backup feature availability. All resources that support logically air-gapped vaults follow the principle of maintaining only one copy of your backup, rather than storing two separate copies, when this feature is used.

> **Warning**
>
> Resources not currently supported for this feature may have their support enabled in the future. When that occurs, your newly supported resource will automatically start backing up in the logically air-gapped vault using the workflow shown above.

Considerations and limitations:

- **Same account and Region only** – Your logically air-gapped vault must be in the same AWS account and Region as your resources to use this feature. You cannot create backups directly cross-account or cross-Region. We recommend backing up to a logically air-gapped vault in the same Region to enable faster recovery without requiring a copy. If you require a copy of your data in a second Region for Disaster Recovery (DR), we recommend either cross-Region replication of your primary resources for quick failover or cross-Region recovery point copies to a locked backup vault.

- **Constraints with using AWS managed keys** – Resources not supporting full AWS Backup management and encrypted with AWS managed keys (for example, `aws/ebs`, `aws/rds`) cannot be copied to logically air-gapped vaults. These resources must be encrypted with a customer managed KMS key or be unencrypted. Resources supporting full AWS Backup management **do not have this constraint.**
- **Backup frequency and concurrent copies** – For resources not supporting full AWS Backup management, ensure your backup frequency allows sufficient time for copies to complete. If backups are scheduled more frequently than copies can finish, copy jobs will queue and may eventually fail. For guidance on concurrent copy limits, see quotas.
- **Lifecycle compatibility** – The retention period specified in your backup plan must be compatible with the minimum and maximum retention periods configured for your logically air-gapped vault.
- **Locked backup vaults** – If your target backup vault has a vault lock enabled, temporary recovery points cannot be deleted manually and will be retained until either the copy completes or the retention period expires.
- **Restore testing, indexing, and scanning** – Restore testing, recovery point indexing, and malware scanning will ignore temporary recovery points with a `DELETE_AFTER_COPY` lifecycle. Recovery point indexing does not support recovery points in logically air-gapped vault. Malware scanning does not support scheduled scans of copied recovery points, which includes automatic copies taken as part of primary backups to logically air-gapped vault.

Resources supporting full AWS Backup management

Some resource types, such as Amazon EFS, Amazon S3, Amazon DynamoDB with AWS Backup advanced features etc., that support full AWS Backup management, can back up directly to your logically air-gapped vault. No recovery point is created in your backup vault, and no copy operation is required. Any scheduled copy actions in your backup plan use the recovery point in your logically air-gapped vault as the source.

Resources that support continuous backup, such as Amazon S3, can also perform continuous backup directly to a logically air-gapped vault.

For a list of resource types supporting full AWS Backup management and logically air-gapped vaults, see Feature availability by resource in the columns labeled "Full management" and "Logically air-gapped vault".

Resource not supporting full AWS Backup management

Resources such as Amazon EBS/EC2, Amazon Aurora, and Amazon FSx cannot back up directly to logically air-gapped vaults. For these resource types, AWS Backup creates a temporary recovery point in your backup vault and then automatically copies it to your logically air-gapped vault.

The temporary recovery point has a special lifecycle setting called `DELETE_AFTER_COPY`. Once the copy to your logically air-gapped vault completes successfully, AWS Backup automatically deletes the temporary recovery point. All other scheduled copy actions in your backup plan start in parallel with the copy to your logically air-gapped vault and do not impact your current copy experience.

If the copy to your logically air-gapped vault fails, the temporary recovery point is retained in your backup vault according to the retention period you specified. This helps you ensure that you always have a usable recovery point after a backup job completes. If the recovery point is later manually copied to the logically air-gapped vault, it is automatically cleaned up according to the `DELETE_AFTER_COPY` rule.

> **Warning**
>
> Resources encrypted with AWS managed keys (for example, `aws/ebs`) are not supported for copy to logically air-gapped vaults. These resources must be encrypted with an AWS Key Management Service customer managed key or be unencrypted. Resources supporting full AWS Backup management do not have this constraint.

Delete after copy lifecycle

Temporary recovery points have a new lifecycle attribute called `DeleteAfterEvent` with a value of `DELETE_AFTER_COPY`. This attribute indicates that the recovery point will be deleted automatically after all copy jobs are complete, or after the retention period you specified, whichever comes first.

A temporary recovery point is deleted when all the following conditions are true:

- All automatic and scheduled copy jobs have finished.
- There is a completed copy job to your target logically air-gapped vault with a retention period at least as long as the source recovery point.

If you need to prevent manual deletion of the temporary recovery point while copies are ongoing, consider using a locked backup vault as your target backup vault.

Continuous backup for Resource not supporting full AWS Backup management

For resources such as Amazon Aurora, if you enable continuous backup, AWS Backup creates a continuous recovery point in your backup vault and takes a temporary snapshot that is copied to your logically air-gapped vault. The temporary snapshot should be automatically deleted after the copy is completed, regardless of whether the copy succeeds or fails, because you retain a continuous recovery point in your backup vault.

If you do not want to create a continuous recovery point for Amazon Aurora in your backup vault, but do want a continuous recovery point for Amazon S3 in your logically air-gapped vault, then you can disable continuous backup (`EnableContinuousBackup`) setting in the current plan and enable S3 continuous from a different plan.

You can learn more about Aurora backup storage in Understanding Amazon Aurora backup storage usage.

Unsupported resources

If a resource type is not supported by logically air-gapped vaults, or if a non-fully managed resource is encrypted with an AWS managed key, AWS Backup creates the backup in your backup vault only. No copy to your logically air-gapped vault is attempted. The backup job completes successfully with a message indicating why the backup did not go to your logically air-gapped vault.

Cost considerations

- This feature does not incur new charges. You only pay for the storage in your vaults.
- For resources supporting full AWS Backup management, maintaining backups only in your logically air-gapped vault can result in significant cost savings compared to maintaining two backup copies in both a backup vault and a logically air-gapped vault.
- For resources not supporting full AWS Backup management, you are charged for both the temporary recovery point in your backup vault and the recovery points in your logically air-gapped vault.
  - You can still achieve significant cost savings from retention of just a single backup copy, but these savings can vary based on your backup frequency and change rate.

- Lower backup frequencies generally produce greater savings because temporary recovery points occupy storage for a shorter percentage of your billing period.
- Some resources have minimum billing durations, which increase costs for temporary recovery points.

- Disable Tags or ACLs metadata retrieval in your backup configuration if you do not use these S3 features. This reduces API calls and associated charges for metadata checks during copy operations.

Configure logically air-gapped vault primary backup

You can configure logically air-gapped vault primary backup through the AWS Backup console, the AWS CLI, AWS CloudFormation, or AWS Organizations backup policies.

Configure a Backup plan

Console

To configure logically air-gapped vault primary backup for a backup plan

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Backup plans**, and then choose **Create backup plan** or select an existing backup plan to edit.
3. In the **Backup rule configuration** section, specify your backup rule settings.
4. For **Backup vault**, choose the backup vault where temporary recovery points will be stored (for non-fully managed resources) or where backups will be stored if they cannot be placed in your logically air-gapped vault.
5. For **Logically air-gapped vault (Optional)**, choose the logically air-gapped vault where you want your backups to be stored.

   > **Note**
   >
   > The logically air-gapped vault must be in the same account and Region as your backup vault.

6. Configure the remaining backup rule settings, including lifecycle and copy options.
7. Choose **Create plan** or **Save changes**.

AWS CLI

Use the CLI command `create-backup-plan` to create a new plan, or `update-backup-plan` to update an existing plan, and include the `TargetLogicallyAirGappedBackupVaultArn` parameter in your backup rule.

Example CLI command to create a backup plan using a JSON document:

```
aws backup create-backup-plan --cli-input-json file://PATH-TO-FILE/test-backup-
plan.json
```

Example CLI command to create a backup plan directly in CLI:

```
aws backup create-backup-plan --backup-plan '{
  "BackupPlanName": "MyPlan",
  "Rules": [
    {
      "RuleName": "MyRule",
      "TargetBackupVaultName": "MyBackupVault",
      "TargetLogicallyAirGappedBackupVaultArn": "arn:aws:backup:us-
east-1:123456789012:backup-vault:MyLagVault",
      "ScheduleExpression": "cron(0 1 ? * * *)",
      "ScheduleExpressionTimezone": "America/Los_Angeles",
      "StartWindowMinutes": 60,
      "CompletionWindowMinutes": 120,
      "Lifecycle": {
        "DeleteAfterDays": 35
      }
    }
  ]
}'
```

Configure on-demand backup

Console

To configure logically air-gapped vault primary backup for an on-demand backup

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Protected resources**.
3. On the **Protected resources** page, choose **Create on-demand backup**.

4. Select the resource type and resource ARN you want to back up.
5. For **Backup vault**, choose the backup vault.
6. For **Logically air-gapped vault (Optional)**, choose the logically air-gapped vault where you want the backup to be stored.
7. Configure the remaining settings and choose **Create on-demand backup**.

AWS CLI

Use the start-backup-job command with the new `--logically-air-gapped-backup-vault-arn` parameter:

```
aws backup start-backup-job \
  --backup-vault-name MyBackupVault \
  --logically-air-gapped-backup-vault-arn arn:aws:backup:us-
east-1:123456789012:backup-vault:MyLagVault  \
  --resource-arn arn:aws:ec2:us-east-1:123456789012:volume/vol-abcd1234  \
  --iam-role-arn arn:aws:iam::123456789012:role/service-role/
AWSBackupDefaultServiceRole  \
  --lifecycle DeleteAfterDays=35
```

Monitor logically air-gapped vault primary backup

You can monitor the status of your backups and copy jobs using the AWS Backup console, AWS CLI, or Amazon EventBridge events.

Monitor for Backup Jobs

Monitor backup job status (`DescribeBackupJob`) to ensure your resources remain protected. A failed backup job indicates no recovery point was created.

- **Verify recovery point creation location** - When a backup job completes successfully, you have a recovery point in either your target backup vault or your target logically air-gapped vault. Check the `BackupVaultArn` field to determine where the recovery point was created.
- **Verify job status** - If a resource is not supported by logically air-gapped vaults, the backup job completes with a `MessageCategory` of `LOGICALLY_AIR_GAPPED_BACKUP_VAULT_NOT_SUPPORTED` and a status message explaining why the backup was created in your backup vault instead.

- **Verify temporary recovery point type** - To check if a recovery point is temporary, look for the `RecoveryPointLifecycle.DeleteAfterEvent` field with a value of `DELETE_AFTER_COPY`.

Monitor for Copy Jobs

Monitor copy jobs (`ListCopyJobs`) to your logically air-gapped vault for failures. A failed copy job means your recovery point remains in your standard backup vault without logically air-gapped vault protection.

- **Verify copy job status** - You can monitor copy job status using the existing `Copy Job State Change` EventBridge event. Optionally, filter on the destination vault (`destinationBackupVaultArn`) to focus on logically air-gapped vault copies.
- **Verify copies for a source recovery point** - Use the `ListCopyJobs` API with the new `BySourceRecoveryPointArn` filter to find all copy jobs associated with a specific recovery point, including both automatic copies to your logically air-gapped vault and scheduled copies to other destinations.
- **Verify deletion of temporary recovery point** - Track completion of temporary recovery point deletion. If the copy job state is `RUNNING`, the recovery point has not yet been deleted. If the copy to your logically air-gapped vault has `FAILED`, the recovery point will be retained per your specified retention period.

> **Note**
>
> Copy job records expire and are removed 30 days after they finish. After this period, you cannot use `ListCopyJobs` to determine historical copy status.

Monitor for Recovery point

Monitor for `EXPIRED` recovery points (`ListRecoveryPointsByBackupVault`), which may indicate AWS Backup could not delete them (possibly due to missing permissions). `EXPIRED` recovery points can have cost implications.

- **Verify recovery point state** - Use the existing Recovery Point State Change EventBridge event to monitor expirations.

- **Verify deletion of temporary recovery point** - If a recovery point with `DeleteAfterEvent: DELETE_AFTER_COPY` has not been deleted, use the `ListCopyJobs` API to determine the reason, as mentioned above.

Onboarding and migration

If you currently use copy actions to copy backups to a logically air-gapped vault, you can migrate to logically air-gapped vault primary backup to reduce costs. You can also migrate your existing Amazon S3 continuous backups from a backup vault to a logically air-gapped vault. This guide explains the prerequisites and steps required to migrate to the logically air-gapped vault primary backup feature.

Prerequisites and Best Practices

Before you can effectively use logically air-gapped vault primary backup feature, there are prerequisites and recommended best practices.

Prerequisites

Currently, a logically air-gapped vault as a primary backup target supports only backups within the same AWS account and AWS Region as your backup resources. Logically air-gapped vault backups are inherently stored in separate service accounts, providing cross-account / cross-Org isolation without requiring copies to a separate accounts. If you need cross-Region backups, continue using the logically air-gapped vault as a copy destination. Before migrating to this feature, ensure you meet the following requirements:

- **Region and account requirements**
  - Your logically air-gapped vault must be in the same AWS account and Region as your resources
- **Resource compatibility**
  - Verify that your resources are supported by logically air-gapped vaults in Feature availability by resource.
  - Check if your resources support full AWS Backup management or not. The experience of creating air-gapped backups differs between these two types of resources, even though the outcome is similar for both.
- **Encryption requirements**

- Resources that do not support full AWS Backup management must be either unencrypted or encrypted with Customer Managed Keys (CMKs). AWS Managed Key (AMK) encrypted resources are not supported by logically air-gapped vault.

Best Practices

- Start with a pilot migration of non-critical resources.
- Review and adjust backup frequencies based on copy job performance.
- Implement comprehensive monitoring before full migration.
- Regularly verify recovery point creation in intended vaults.

Planning Your Migration

- Review existing backup plans and policies.
- Identify which resources support full AWS Backup management and which do not.
  - Resources supporting full AWS Backup management (e.g., EFS, S3) - can backup directly to logically air-gapped vault
  - Resources not supporting full AWS Backup management (e.g., EC2, EBS, FSx) - require temporary backup in backup vault before being copied to Logically air-gapped vault
- Review your current backup volume and frequency and ensure your configuration aligns with concurrent copy limits for all resources not supporting full AWS Backup management.
  - Skip this step If you already copy to a logically air-gapped vault at the same frequency as your standard vault backups.
  - Consider adjusting backup frequency, if needed, to prevent copy job queuing.
  - If copy job queuing occurs, you will still have a usable recovery point in your standard backup vault while awaiting completion of the copy to your logically air-gapped vault. However, that recovery point will not provide the protection level that Logically air-gapped vault offers.

Resources supporting full AWS Backup management migration path

For fully managed resources, you can back up directly to your logically air-gapped vault without requiring copy operations.

For snapshot-based backups

This process applies to all snapshot scenarios, regardless of whether your backup vault is locked. When migrating an existing backup plan or using an existing backup vault (primary) and logically air-gapped vault (copy) in a new backup plan:

1. Maintain your existing backup vault or add it as a backup target (`TargetBackupVaultName`). This vault will not store any backups but must be provided for backwards compatibility.
2. Update your backup plan to include the logically air-gapped vault (`TargetLogicallyAirGappedBackupVaultArn`), existing in the same account, as a primary target.
3. Review any existing copy action to another logically air-gapped vault to determine if it is still needed. You can also move this vault as a primary target in Step 2 if it is in the same account.

For Amazon S3 continuous backups

Logically air-gapped vault as a primary backup target supports continuous backup for Amazon S3. However, you can maintain only one active continuous recovery point per resource in a single vault. If you have an existing active Amazon S3 continuous recovery point, you must disassociate or delete it before creating a new one in a different vault. Adding a logically air-gapped vault target (from the same account) to your backup plan, with an existing active Amazon S3 continuous recovery point, will cause the continuous backup job to fail.

To migrate your Amazon S3 continuous recovery point to your logically air-gapped vault from an **unlocked backup vault**:

1. Update your backup plan to add a copy action to your logically air-gapped vault. This will reduce your cost to generate the initial backup in your logically air-gapped vault. Skip this step if you are already copying to your local logically air-gapped vault.
2. Verify that at least one snapshot copy completes successfully in your logically air-gapped vault before proceeding.
3. Disassociate existing Amazon S3 continuous recovery point. Call the DisassociateRecoveryPoint API to change the recovery point status from AVAILABLE to STOPPED. This action preserves your existing backup data while stopping new data from being added.
4. Update backup plan to add logically air-gapped vault (`TargetLogicallyAirGappedBackupVaultArn`) as a backup target.
5. Remove any previous copy actions in the plan.

6. On the next backup plan execution, a new continuous recovery point will be created in your logically air-gapped vault. This recovery point will be incremental, based on the copied snapshot from Step 1.

To migrate your Amazon S3 continuous recovery point to your logically air-gapped vault from a **locked backup vault**:

1. Update your backup plan to add a copy action to your logically air-gapped vault. This will reduce your cost to generate the initial backup in your logically air-gapped vault. Skip this step if you are already copying to your local logically air-gapped vault.
2. Verify that at least one snapshot copy completes successfully in your logically air-gapped vault before proceeding. Ensure the copied snapshot has a retention period long enough to remain available until you complete all the steps.
3. Add the logically air-gapped vault as a primary target and remove any copy action.
   a. This step is required because locked vaults do not support disassociation of continuous recovery points.
   b. Your existing continuous recovery point in the locked backup vault will continue accumulating data until it expires according to its lifecycle.
   c. New continuous backup jobs will fail because only one active continuous recovery point can exist per resource. Since these jobs are failing, no copy actions will execute.
4. Wait for the existing continuous recovery point to expire. After expiration, a new continuous recovery point will be created in the logically air-gapped vault. This recovery point will be incremental based on the copied snapshot from step 1, provided the snapshot still exists in your logically air-gapped vault.
5. Any data accumulated in your standard vault will be lost upon expiration. Only data backed up after the new recovery point creation in the Logically air-gapped vault will be retained.

Resources not supporting full AWS Backup management migration path

Non-fully managed resources require a copy operation to the logically air-gapped vault. The process creates a temporary recovery point (billable) in your standard backup vault, which is automatically copied to your logically air-gapped vault and then deleted after the copy is completed. When you update your backup plan to include a logically air-gapped vault target:

- Backup jobs will create a recovery point in your backup vault. This has a lifecycle dictated by `DELETE_AFTER_COPY` event.

- A copy job automatically initiates transferring the recovery point to your logically air-gapped vault.
- After the copy is completed successfully, the temporary recovery point in your vault is deleted.
- If the copy fails, the temporary recovery point is retained for a maximum duration according to your specified retention period, ensuring you have a usable recovery point.

Troubleshooting

Backup or copy job fails with lifecycle incompatibility error

Error for backup jobs: `Backup job failed because the lifecycle is outside the valid range for backup vault.`

Error for copy jobs: `Copy job failed. The retention specified in the job is not within the range specified for the target Backup Vault.`

Possible cause: Backup jobs or copy jobs fail because the retention period is incompatible with the logically air-gapped vault's minimum or maximum retention settings.

Resolution: Update your backup plan to specify a retention period that falls within the minimum and maximum retention periods configured for your logically air-gapped vault.

Continuous backup jobs fail with "already has continuous backup enabled"

Error: `Bucket {bucket_name} already has continuous backup enabled for another vault Backup job failed.`

Possible cause: Continuous backup jobs fail because there is already a continuous recovery point for the resource in another vault.

Resolution: Each resource can have only one continuous recovery point. If your backup vault is unlocked, disassociate the existing continuous recovery point using the disassociate-recovery-point command. If your backup vault is locked, wait until the existing continuous recovery point expires according to its lifecycle.

Backup job completes with "Completed with issues" - Unsupported resource type

Message: `Backup job completed successfully to the target backup vault. This resource type is not supported by logically air-gapped backup vault.`

Possible cause: Backup jobs for unsupported non-fully managed resources show "Completed with issues" with a message indicating that the resource is not supported.

Resolution: Resource types that are not supported will only backup to backup vault. If you want to retain these backups in your backup vault, no action is required. If you prefer not to mix unsupported resources with your logically air-gapped vault, you can:

- Remove the resource or logically air-gapped vault target from your backup plan for these resources and continue backing up to your backup vault only. You can subsequently add the resource as part of a different plan.

Backup job completes with "Completed with issues" - Unsupported encryption key

Message: `Backup job completed successfully to the target backup vault. This resource is encrypted with an AWS managed key and cannot be copied to a logically air-gapped backup vault.`

Possible cause: Backup jobs for unsupported non-fully managed resources show "Completed with issues" with a message indicating that the resource is encrypted with an AWS managed key.

Resolution: Non-fully managed resources encrypted with AWS managed keys cannot be copied to logically air-gapped vaults. If you want to retain these backups in your backup vault, no action is required. If you prefer not to mix unsupported resources with your logically air-gapped vault, you can:

- Re-encrypt your resources with a customer managed KMS key, or
- Remove the resource or logically air-gapped vault target from your backup plan for these resources and continue backing up to your backup vault only. You can subsequently add the resource as part of a different plan.

Recovery points remain in `EXPIRED` state

Possible cause: Temporary recovery points transition to `EXPIRED` state but are not deleted.

Resolution: AWS Backup may lack permissions to delete the recovery points. Verify that your backup role has the necessary IAM permissions. You may need to manually delete `expired` recovery points.

Copy jobs are queued or failing due to high backup frequency

Possible cause: Copy jobs to logically air-gapped vaults are queuing or failing because backups are scheduled more frequently than copies can complete.

Resolution: Reduce your backup frequency or adjust your backup schedule to allow more time between backups. See the AWS Backup quotas documentation for information about concurrent copy limits.

Temporary recovery point for Amazon FSx for NetApp ONTAP remains in EXPIRED state

Error: `Amazon FSx does not support deleting the most recent AVAILABLE backup of an ONTAP volume unless all other backups of the volume have been deleted`

Possible cause: For resources not supporting full AWS Backup management, AWS Backup creates a temporary recovery point in your backup vault and deletes it after the copy to your logically air-gapped vault completes. Amazon FSx for NetApp ONTAP does not support deleting the most recent `AVAILABLE` backup of a volume unless all other backups of that volume have been deleted. When the temporary recovery point is the only backup of the volume, the deletion request fails and the recovery point remains in `EXPIRED` state in your backup vault. This restriction applies only to FSx for ONTAP volumes; other Amazon FSx file system types are not affected.

Resolution: The temporary recovery point is the volume's most recent backup when AWS Backup attempts to delete it, so Amazon FSx refuses the deletion while any other backup of that volume exists, and the recovery point remains in `EXPIRED` state. After a later backup of the volume is taken, the temporary recovery point is no longer the most recent backup and can be deleted. Manually delete any recovery point that remains in `EXPIRED` state from your backup vault; you are charged for it until you do. For more information, see Deleting backups.

Multi-party approval for logically air-gapped vaults

Overview of Multi-party approval in a logically air-gapped vault

AWS Backup offers you the option to add Multi-party approval, a capability from AWS Organizations, to your logically air-gapped vaults. Multi-party approval provides an additional option to help to protect critical operations through a distributed approval process.

Multi-party approval is designed to help protect critical resources and to minimize the time to return to full operation, such as a disruption caused by malicious actors or malware events.

This setup can help you restore the contents of a logically air-gapped vault that may have been compromised.

There is no additional cost for integrating and using Multi-party approval teams with AWS Backup logically air-gapped vaults (storage and cross-region transfers charges apply, as shown on the pricing page).

As an AWS Backup customer, you can use Multi-party approval to grant approval capabilities of some operations to a group of trusted individuals who can collaboratively approve access to a logically air-gapped vault from a separately-created recovery account in the case of suspected malicious activity that may compromise use of the primary account.

Prerequisites and best practices for using Multi-party approval with a logically air-gapped vault

Before you can effectively and securely use Multi-party approval with your logically air-gapped vaults, there are prerequisites and recommended best practices.

Best practices:

- Two (or more) AWS organizations through Organizations. One should be your primary organization where you have one or more accounts that have at least one logically air-gapped vault. The secondary organization should be your recovery organization. It is in this org where your multi-party approval team will be managed.

Prerequisites

1. You have Set up Multi-party approval and have at least one approval team.
2. At least one account in your primary organization must have a logically air-gapped vault (and the original backup vault).
3. The management account in the primary organization is opted-in to Multi-party approval.

> **Tip**
>
> AWS Backup recommends you apply a Service Control Policy (SCP) to your primary organization and configure it with the appropriate permissions to the organization and to each approval team. See Multi-party approval terms section for a sample policy.

4. Your Multi-party approval team from the secondary (recovery) organization is shared through AWS RAM with both your accounts that own the logically air-gapped vault(s) and your recovery accounts.

Required access policies

Before setting up Multi-party approval, configure the following access policies:

- **Service Control Policy (SCP):** Apply an SCP to your primary organization to control which actions are allowed across accounts. The SCP should grant appropriate permissions to the organization and to each approval team. For more information about SCPs, see Service control policies in the *Organizations User Guide*.
- **Multi-party approval team policy:** When you create an approval team, attach the appropriate resource policy. For AWS Backup logically air-gapped vaults, use the policy `arn:aws:mpa::aws:policy/backup.amazonaws.com/CreateRestoreAccessVault`. To list available policies, run `aws mpa list-policies --region us-east-1`. Multi-party approval team resources can only be created and stored in US East (N. Virginia), so MPA commands must target this Region.
- **AWS RAM sharing permissions:** When sharing your approval team across accounts, use the RAM permission `arn:aws:ram::aws:permission/AWSRAMMPAApprovalTeamAccess`. For more information, see Share an approval team using AWS RAM.
- **IAM permissions:** Users who interact with Multi-party approval need the appropriate IAM permissions, including `mpa:GetApprovalTeam`, `mpa:ListApprovalTeams`, `mpa:StartSession`, `mpa:CancelSession`, `mpa:GetSession`, and `mpa:ListSessions`. These permissions are included in the AWSBackupFullAccess managed policy.

Set up Multi-party approval

The following steps outline the recommended flow for setting up a recovery AWS organization, setting up Multi-party approval, and then using Multi-party approval with your logically air-gapped vaults:

1. An administrator creates a new organization through Organizations to be used for recovery operations.
2. In the management account of this new organization, the administrator creates and configures an IAM Identity Center (IDC) instance (to enable an organization instance, see Enable IAM

Identity Center in the *IAM Identity Center User Guide*. See also the sequence to Create a Multi-party approval identity source in the *Multi-party approval User Guide*.

3. The administrator then will create an approval team, the core group of trusted individuals who will be the primary users of Multi-party approval.
4. The administrator uses AWS RAM to share an approval team with each account that owns a logically air-gapped vault and the recovery account that needs to request access on that vault.
5. An administrator of the logically air-gapped vault owning account associates the vault with an approval team.
6. A recovery account requests access to an account that has a logically air-gapped vault with an associated Multi-party approval team ("team"). The team associated with the account approves or denies the request.
7. The admin of the account of that owns the logically air-gapped vault can request that the approval team be disassociated from the vault. The request requires current team approval.
8. An administrator can update approval team membership as necessary in accordance with their security practices or when people join or leave your organization.

Cross-Region considerations and dependencies when using Multi-party approval

When you enable Multi-party approval and your IAM Identity Center instance in different Regions, Multi-party approval makes calls across Regions to IAM Identity Center. This means that user and group information moves across Regions. Multi-party approval Team resources can only be created and stored in AWS Region US East (N. Virginia).

Other AWS Regions that reference Multi-party approval team resources will depend on AWS Region US East (N. Virginia). Accordingly, Multi-party approval will make cross-Regions calls if your Identity Center instance and/or logically air-gapped vault is not in US East (N. Virginia).

Administrator tasks

Several tasks involving AWS Backup and Multi-party overview required a user with admin permissions and access to the management account.

Create an approval team

A user at your organization with admin permissions for an AWS account needs to set up Multi-party approval (step 3 in the Overview).

Before doing this step, it is recommended as a best practice you have both a primary organization and a secondary organization (for recovery purposes) set up through AWS Organizations (step 1 in Overview.

See Create an approval team in the *Multi-party approval user guide* to create your team.

During the `aws mpa create-approval-team` operation, one of the parameters is `policies`. This is a list of ARNs (Amazon Resource Names) for Multi-party approval resource policies that define permissions that protect the team.

The policy shown in the example in the *Multi-party approval User Guide* in the procedure Create an approval team contains the policy `["arn:aws:mpa::aws:policy/backup.amazonaws.com/CreateRestoreAccessVault"]` with several necessary permissions.

Follow these steps to return a list of available policies by using `mpa list-policies`:

1. List Policies:

```
aws mpa list-policies --region us-east-1
```

2. List all policy versions:

```
aws mpa list-policy-versions --policy-arn arn:aws:mpa:::aws:policy/
backup.amazonaws.com/CreateRestoreAccessVault --region us-east-1
```

3. Get details on a policy:

```
aws mpa get-policy-version --policy-version-arn arn:aws:mpa:::aws:policy/
backup.amazonaws.com/CreateRestoreAccessVault/1 --region us-east-1
```

Expand below to see the policy that will created then attached to your approval team by this operation:

Restore access vault policy

JSON

```
{
  "Version":"2012-10-17",
```

```
  "Statement": [
    {
      "Sid": "VaultOwnerPermissions",
      "Effect": "Allow",
      "Principal": {
        "AWS": "*"
      },
      "Resource": "*",
      "Action": [
        "mpa:StartSession",
        "mpa:CancelSession"
      ],
      "Condition": {
        "StringEquals": {
          "mpa:RequestedOperation": "backup:RevokeRestoreAccessBackupVault",
          "mpa:ProtectedResourceAccount": "${aws:PrincipalAccount}"
        },
        "Bool": {
          "aws:ViaAWSService": "true"
        }
      }
    }
  ]
}
```

Share a Multi-party approval team using AWS RAM

You can share a Multi-party approval team with other AWS accounts using AWS Resource Access Manager (RAM), step 4 in the overview.

Console

Share a Multi-party approval team using AWS RAM

1. Sign in to the AWS RAM console.
2. In the navigation pane, choose **Resource shares**.
3. Choose **Create resource share**.
4. In the **Name** field, enter a descriptive name for your resource share.
5. Under **Resource type**, select **Multi-party approval Team** from the dropdown menu.
6. Under **Resources**, select the approval team you want to share.

7. Under **Principals**, specify the AWS accounts with whom you want to share the approval team.
8. To share with specific AWS accounts, select **AWS accounts** and enter the 12-digit account IDs.
9. To share with an organization or organizational unit, select **Organization** or **Organizational unit** and enter the appropriate ID.
10. (*Optional*) Under **Tags**, add any tags you want to associate with this resource share.
11. Choose **Create resource share**.

The resource share status will initially show as `PENDING`. Once the recipient accounts accept the invitation, the status will change to `ACTIVE`.

CLI

To share a Multi-party approval team using AWS RAM through the CLI, use the following commands:

First, identify the ARN of the approval team you want to share:

```
aws mpa list-approval-teams --region us-east-1
```

Create a resource share using the create-resource-share command:

```
aws ram create-resource-share \
--name "MPA-Team-Share" \
--resource-arns "arn:aws:mpa:us-east-1:ACCOUNT_ID:approval-team/TEAM_ID" \
--principals "ACCOUNT_ID_TO_SHARE_WITH" \
--permission-arns "arn:aws:ram::aws:permission/AWSRAMMPAApprovalTeamAccess" \
--region us-east-1
```

To share with an organization instead of specific accounts:

```
aws ram create-resource-share \
--name "MPA-Team-Share" \
--resource-arns "arn:aws:mpa:us-east-1:ACCOUNT_ID:approval-team/TEAM_ID" \
--permission-arns "arn:aws:ram::aws:permission/AWSRAMMPAApprovalTeamAccess" \
--allow-external-principals \
--region us-east-1
```

Check the status of your resource share:

```
aws ram get-resource-shares \
--resource-owner SELF \
--region us-east-1
```

The recipient account(s) will need to accept the resource share invitation:

```
aws ram get-resource-share-invitations --region us-east-1
```

Run in recipient account to accept an invitation:

```
aws ram accept-resource-share-invitation \
--resource-share-invitation-arn "arn:aws:ram:REGION:ACCOUNT_ID:resource-share-
invitation/INVITATION_ID" \
--region us-east-1
```

Once the invitation is accepted, the Multi-party approval team will be available for use in the recipient account.

AWS offers tools to share account access, including through AWS Resource Access Manager and Multi-party access. When you choose to share a logically air-gapped vault with another account, consider the following details:

Feature	AWS RAM based sharing	Multi-party approval based access
Access to logically air-gapped vaults	Once RAM share is complete, the vaults can be accessed.	Any attempt by a different account must be approved by a threshold number of Multi-party approval team members. The approval session automatically expires 24 hours after the request is initiated.
Access removal	The account which owns the logically air-gapped vault can	Access to a vault can only be removed by a request to the Multi-party approval team.

Feature	AWS RAM based sharing	Multi-party approval based access
	end RAM based sharing at any time.	
Copy across accounts and/or Regions	Not currently supported.	Backups can be copied within the same account or with other accounts in the same organization as the recovery account.
Cross-Region transfer billing		Cross-Region transfers are billed to the same account that owns the restore access backup vault.
Recommended use	Primary use is for data loss recovery and for restore testing.	Primary use is for situation s where account access or security is suspected to be compromised.
Regions	Available in all AWS Regions where logically air-gapped vaults are supported.	Available in all AWS Regions where logically air-gapped vaults are supported.
Restores	All supported resource types can be restored from a shared account.	All supported resource types can be restored from a shared account.
Setup	Sharing can occur as soon as the AWS Backup account sets up RAM sharing and the receiving account accepts the share.	Sharing requires the management account to first create a team, then set up RAM sharing. Then, the management account opts in to Multi-party approval and assigns that team to a logically air-gapped vault.

Feature	AWS RAM based sharing	Multi-party approval based access
Sharing	Sharing is done through RAM within same AWS organizat ion or across AWS organizat ions. Access is granted according to the 'push' model, in which the account that owns the logically air-gapped vault first grants access. Then, the other account accepts access.	Access to a logically air-gapped vault is through Organizations supported approval teams within the same AWS organization or across organizations. Access is granted according to the 'pull' model, where the receiving account first requests access, then the approval team grants or denies the request.

Requester tasks

Associate a Multi-party approval team with a logically air-gapped vault

Requester: **User with access to account that owns the logically air-gapped vault**.

You can associate a Multi-party approval team with a logically air-gapped vault to enable collaborative approval for access to the vault (step 5 in the Overview).

Console

Associate a Multi-party approval team with a logically air-gapped vault

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Navigate to the **Backup vaults** section in the left navigation pane.
3. Select the logically air-gapped backup vault you want to associate with an MPA team.
4. On the **vault details** page, select **Assign approval team**.
5. From the dropdown menu, select the approval team you want to associate with the vault
6. *Optional* Enter a comment explaining the reason for the association.
7. Select **Send request** to submit the association request.

If this is the first approval team to be associated with the vault, the team will be associated with the vault. If the vault already has an associated team, see Update Multi-party approval team for steps.

CLI

Use the CLI command `associate-backup-vault-mpa-approval-team`, modified with the following parameters:

```
aws backup associate-backup-vault-mpa-approval-team \
--backup-vault-name VAULT_NAME \
--mpa-approval-team-arn MPA_TEAM_ARN \
--requester-comment "OPTIONAL_COMMENT" \
--region REGION
```

If this is the first approval team to be associated with the vault, the team will be associated with the vault. If the vault already has an associated team, see Update Multi-party approval team for steps.

Request access to a logically air-gapped vault

Requester: **User with access to recovery account**.

You can request access to a logically air-gapped vault in another account (step 6 in the Overview).

After an approval team has granted the request, AWS Backup creates a restore access backup vault in your designated recovery account so that account will have access to recovery points in the connected logically air-gapped vault.

Required IAM permission: The IAM role used to call `CreateRestoreAccessBackupVault` (CLI: `create-restore-access-backup-vault`) must include the `mpa:StartSession` permission in addition to the required backup permissions. `mpa:StartSession` is required to start the approval session that allows the approval team to approve or decline the request. If this permission is missing, the request fails and `CreateRestoreAccessBackupVaultFailed` events appear in your AWS CloudTrail logs.

Console

Request access to a logically air-gapped vault

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.

2. Navigate to the **Backup vaults** section in the left navigation pane
3. Select the **Vaults accessible through MPA** tab
4. Select **Request vault access**.
5. Enter the source backup vault ARN of the logically air-gapped vault you want to access.
6. Enter an optional name for the restore access backup vault. If you do not input a name, AWS Backup will assign a name based on the name of the logically air gapped vault.
7. Enter an optional requester comment explaining the reason for the access request.
8. Select **Send request** to submit the access request.

The approval team members associated with the source vault will receive an email notification to approve the request.

Once the request is approved by the required number ("threshold") of team members, the restore access backup vault will be created in the recovery account.

CLI

Use the CLI command `create-restore-access-backup-vault`:

```
aws backup create-restore-access-backup-vault \
--source-backup-vault-arn SOURCE_VAULT_ARN \
--backup-vault-name OPTIONAL_VAULT_NAME \
--requester-comment "OPTIONAL_COMMENT" \
--region REGION
```

The MPA approval team members associated with the source vault will receive a notification to approve the request. Once the request is approved by the required number ("threshold") of team members, the restore access backup vault will be created in the recovery account.

You can check the status of the vault using:

```
aws backup describe-backup-vault \
--backup-vault-name VAULT_NAME \
--region REGION
```

Disassociate Multi-party approval team from logically air gapped vault

Requester: **Administrator of account that owns the logically air-gapped vault**.

You can disassociate a Multi-party approval team from a logically air-gapped vault (step 7 in the Overview).

Console

Disassociate approval team from logically air-gapped vault

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Navigate to the **Backup vaults** section in the left navigation pane.
3. Select the logically air-gapped backup vault from which you want to disassociate the approval team.
4. On the **Vault details** page, select **Disassociate approval team**.
5. Enter an optional requester comment explaining the reason for the disassociation.
6. Select **Send request** to submit the disassociation request.

The current approval team members will receive a notification to approve the request.

Once approved by the required number of team members, the team will be disassociated from the vault.

CLI

Use the CLI command `disassociate-backup-vault-mpa-approval-team`:

```
aws backup disassociate-backup-vault-mpa-approval-team \
--backup-vault-name VAULT_NAME \
--requester-comment "OPTIONAL_COMMENT" \
--region REGION
```

The current MPA approval team members will receive a notification to approve the request. Once approved by the required number of team members, the team will be disassociated from the vault.

Revoke restore access backup vault

Requester: **Administrator of account that owns the logically air-gapped vault**.

You can revoke access to a restore access backup vault from the source vault account.

Console

Revoke restore access backup vault

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Navigate to the **Backup vaults** section in the left navigation pane.
3. Select the logically air-gapped backup vault for which you want to revoke access.
4. On the **Vault details** page, scroll down to the **Access through Multi-party approval** section.
5. Find the restore access backup vault you want to revoke, then select **Request to remove vault access**.
6. Enter an optional requester comment explaining the reason for the revocation.
7. Select **Send request** to submit the revocation request.

The approval team members will receive a notification to approve the request.

Once approved by the required number of team members, the restore access backup vault will be deleted from the recovery account

CLI

First, list the restore access backup vaults associated with your source vault:

```
aws backup list-restore-access-backup-vaults \
--backup-vault-name SOURCE_VAULT_NAME \
--region REGION
```

Then, use the CLI command `revoke-restore-access-backup-vault`:

```
aws backup revoke-restore-access-backup-vault \
--backup-vault-name SOURCE_VAULT_NAME \
--restore-access-backup-vault-arn RESTORE_ACCESS_VAULT_ARN \
--requester-comment "OPTIONAL_COMMENT" \
--region REGION
```

The approval team members will receive a notification to approve the request. Once approved by the required number of team members, the restore access backup vault will be deleted from the recovery account.

Update the Multi-party approval team associated with a logically air-gapped vault

Requester: **Administrator of account that owns the logically air-gapped vault**.

You can update the Multi-party approval team associated with a logically air-gapped vault (step 8 in the Overview).

Console

Update the approval team associated with a logically air-gapped vault

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Navigate to the **Backup vaults** section in the left navigation pane.
3. Select the logically air-gapped backup vault for which you want to update the approval team.
4. On the vault details page, select **Request approval team change**.
5. From the dropdown menu, select the new approval team you want to associate with the vault.
6. Enter an optional requester comment explaining the reason for the change.
7. Select **Send request** to submit the change request.

The current approval team members will receive an email notification to approve the request.

Once approved by the required number of team members (threshold) from the current MPA team, the new team will be associated with the vault.

CLI

Use the CLI command `associate-backup-vault-mpa-approval-team` with the new team ARN:

```
aws backup associate-backup-vault-mpa-approval-team \
--backup-vault-name VAULT_NAME \
--mpa-approval-team-arn NEW_MPA_TEAM_ARN \
--requester-comment "OPTIONAL_COMMENT" \
--region REGION
```

The current approval team members will receive a notification to approve the request. Once approved by the required number of team members (threshold) from the current team, the new MPA team will be associated with the vault.

Approver tasks

A user who is a member of a Multi-party approval team can approve or deny requests that are part of a session. Other tasks include:

- Respond to requested operations
- View an approval team
- View operation history

Multi-party approval terms, concepts, and user personas

Multi-party approval in your logically air-gapped vault is an integration of AWS Organizations, AWS Account Management, and AWS Backup, along with AWS Identity and Access Management (IAM) and AWS RAM (RAM) features. Through the CLI, you can interact with each service to send the appropriate commands. You can also use the console, but you will need to navigate to the appropriate service's console to complete specific tasks.

How you interact with Multi-party approval depends on your roles and responsibilities at your organizations, as well as the permissions you have in your AWS Backup accounts.

As shown in the Multi-party approval User Guide, members of your organization who use multi-party approval will either be a ***requester***, an ***administrator***, or an ***approver***. Specific permissions apply to each job function. In accordance with security best practices, an user should only fulfill one job function.

Consoles, portals, and sessions

AWS Backup accounts with one or more logically air-gapped vaults can use multi-party approval.

Prior to the multi-party approval process, an administrator utilizes AWS Organizations to create a secondary organization for recovery purposes (a **recovery organization**) if one has not previously been set up.

Then, the administrator utilizes AWS Resource Access Manager (RAM) to set up cross-organization sharing between the primary organization and the recovery organization.

The **primary organization** is home to accounts that own and use a logically air-gapped vault, which stores the protected data.

The recovery organization is home to at least one **recovery account**. This account houses an access point that can serve as a critical 'back door' to the shared logically air-gapped vault. This access point is called a **restore access backup vault**. This access vault does not store data; instead, it serves as an access or mount point that mirrors the contents of the source logically air-gapped vault but contains no data that can be changed or deleted. For example, if a customer goes through the restore process for a recovery point in a restore access backup vault, it is the recovery point in the logically air-gapped vault that is restored through cross-account restore by way of the recovery account.

To ensure extra security, customers use this recovery account to carry out protected operations in the primary account, but only after those operations have been given approval by the associated approval team in an approval session. A session is created by AWS once an approval request has been sent, and that session ends when a threshold of approval team members approves or denies the request or when the allowed session time has passed.

A team consists of **approvers** (effectively, the *parties* portion of Multi-party approval) who receive email notifications of protected operation requests. These emails confirm that an approval session has begun for the request. Approval is granted once the required minimum threshold of approval is reached. This threshold can be set as the **multi-party approval team** ("Team") is created.

Multi-party approval teams are managed through the Organizations **multi-party approval portal** ("portal"), an AWS managed application that provides identities a centralized location where approval team members can receive and respond to approval team invitations and operation requests.

Vault access policies

With AWS Backup, you can assign policies to backup vaults and the resources they contain. Assigning policies allows you to do things like grant access to users to create backup plans and on-demand backups, but limit their ability to delete recovery points after they're created.

For information about using policies to grant or restrict access to resources, see Identity-Based Policies and Resource-Based Policies in the *IAM User Guide*. You can also control access using tags.

You can use the following example policies as a guide to limit access to resources when you are working with AWS Backup vaults. Unlike other IAM-based policies, AWS Backup access policies don't support a wildcard in the `Action` key.

For a list of Amazon Resource Names (ARNs) that you can use to identify recovery points for different resource types, see AWS Backup resource ARNs for resource-specific recovery point ARNs.

Vault access policies only control user access to AWS Backup APIs. Some backup types, such as Amazon Elastic Block Store (Amazon EBS) and Amazon Relational Database Service (Amazon RDS) snapshots, can also be accessed using the APIs of those services. You can create separate access policies in IAM that control access to those APIs to fully control the access to those backup types.

Regardless of the AWS Backup vault's access policy, cross-account access for any action other than `backup:CopyIntoBackupVault` will be rejected; that is, AWS Backup will reject any other request from an account that is different from the account of the resource that is being referenced. This restriction applies to vault access policies on both standard and logically air-gapped backup vaults. Logically air-gapped vaults additionally support cross-account sharing through AWS Resource Access Manager (RAM), which enables operations such as restore through a separate mechanism outside of vault access policies. For more information, see Logically air-gapped vault.

Topics

- Deny access to a resource type in a backup vault
- Deny access to a backup vault
- Deny access to delete recovery points in a backup vault

Deny access to a resource type in a backup vault

This policy denies access to the specified API operations for all Amazon EBS snapshots in a backup vault.

JSON

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Sid": "DenyRecoveryPointOperations",
      "Effect": "Deny",
      "Principal": {
        "AWS": "arn:aws:iam::111122223333:role/MyRole"
      },
```

```
      "Action": [
        "backup:UpdateRecoveryPointLifecycle",
        "backup:DescribeRecoveryPoint",
        "backup:DeleteRecoveryPoint",
        "backup:GetRecoveryPointRestoreMetadata",
        "backup:StartRestoreJob"
      ],
      "Resource": [
        "arn:aws:ec2:*:*:snapshot/*"
      ]
    }
  ]
}
```

Deny access to a backup vault

This policy denies access to the specified API operations targeting a backup vault.

JSON

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Sid": "DenyBackupVaultOperations",
      "Effect": "Deny",
      "Principal": {
        "AWS": "arn:aws:iam::123456789012:role/MyRole"
      },
      "Action": [
        "backup:DescribeBackupVault",
        "backup:DeleteBackupVault",
        "backup:PutBackupVaultAccessPolicy",
        "backup:DeleteBackupVaultAccessPolicy",
        "backup:GetBackupVaultAccessPolicy",
        "backup:StartBackupJob",
        "backup:GetBackupVaultNotifications",
        "backup:PutBackupVaultNotifications",
        "backup:DeleteBackupVaultNotifications",
```

```
        "backup:ListRecoveryPointsByBackupVault"
      ],
      "Resource": "arn:aws:backup:us-east-1:123456789012:backup-vault:backup
 vault name"
    }
  ]
}
```

Deny access to delete recovery points in a backup vault

Access to vaults and the ability to delete recovery points stored in them is determined by the access that you grant your users.

Follow these steps to create a resource-based access policy on a backup vault that prevents the deletion of any backups in the backup vault.

To create a resource-based access policy on a backup vault

1. Sign in to the AWS Management Console, and open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane on the left, choose **Backup vaults**.
3. Choose a backup vault in the list.
4. In the **Access policy** section, paste the following JSON example. This policy prevents anyone who is not the principal from deleting a recovery point in the target backup vault.

   JSON

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Effect": "Deny",
            "Principal": "*",
            "Action": "backup:DeleteRecoveryPoint",
            "Resource": "*",
            "Condition": {
                "StringNotEquals": {
                    "aws:userId": [
```

```
                    "AIDA1234567890EXAMPLE",
                    "AROA1234567890EXAMPLE:my-role-session",
                    "AROA1234567890EXAMPLE:*",
                    "112233445566"
                ]
            }
        }
    }
  ]
}
```

To allow list IAM identities using their ARN, use the `aws:PrincipalArn` global condition key in the following example.

JSON

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Sid": "DenyDeleteRecoveryPoint",
      "Effect": "Deny",
      "Principal": "*",
      "Action": "backup:DeleteRecoveryPoint",
      "Resource": "*",
      "Condition": {
        "ArnNotEquals": {
          "aws:PrincipalArn": [
            "arn:aws:iam::112233445566:role/mys3role",
            "arn:aws:iam::112233445566:user/shaheer",
            "arn:aws:iam::112233445566:root"
          ]
        }
      }
    }
  ]
}
```

For information about getting a unique ID for an IAM entity, see Getting the unique identifier in the *IAM User Guide*.

If you want to limit this to specific resource types, instead of `"Resource": "*"`, you can explicitly include the recovery point types to deny. For example, for Amazon EBS snapshots, change the resource type to the following.

```
"Resource": ["arn:aws:ec2:*:*:snapshot/*"]
```

5. Choose **Attach policy**.

AWS Backup Vault Lock

> **Note**
>
> AWS Backup Vault Lock has been assessed by Cohasset Associates for use in environments that are subject to SEC 17a-4, CFTC, and FINRA regulations. For more information about how AWS Backup Vault Lock relates to these regulations, see the Cohasset Associates Compliance Assessment.

AWS Backup Vault Lock is available at no additional charge. Standard AWS Backup storage charges apply to backups stored in a locked vault. Because backups in a locked vault cannot be deleted until their lifecycle completes, ensure your retention settings are configured appropriately to avoid unintended storage costs. For more information about storage pricing, see AWS Backup pricing.

AWS Backup Vault Lock is an optional feature of a backup vault, which can be helpful in giving you additional security and control over your backup vaults. When a lock is active in Compliance mode and the grace time is over, the vault configuration cannot be altered or deleted by a customer, account/data owner, or AWS as long as it contains recovery points. Each vault can have one vault lock in place.

AWS Backup ensures that your backups are available for you until they reach the expiration of their backup lifecycle retention periods. These retention periods are defined when each backup is created and are independent of the vault lock minimum and maximum retention settings. If any user (including the root user) attempts to delete a backup or change the lifecycle properties in a locked vault, AWS Backup will deny the operation.

- Vaults locked in **governance mode** can have the lock removed by users with sufficient IAM permissions.
- Vaults locked in **compliance mode** *cannot be deleted* once the cooling-off period ("**grace time**") expires if any recovery points are in the vault. During grace time, you can still remove the vault lock and change the lock configuration.

> **Warning**
>
> Once the grace time expires, the vault and its lock are immutable and cannot be changed or deleted by any user or by AWS. Backups within a locked vault cannot be deleted until their lifecycle completes, resulting in persistent costs if you are not careful. For example, ensure there are no recovery points with a retention period set to "Always"—once the grace time expires, these recovery points will be retained forever and cannot be altered or deleted.

Vault lock modes

When you create a vault lock, you have a choice of two modes: **Governance mode** or **Compliance mode**. Governance mode is intended to allow a vault to be managed only by users with sufficient IAM privileges. Governance mode helps an organization meet governance requirements, ensuring only designated personnel can make changes to a backup vault. Compliance mode is intended for backup vaults in which the vault (and by extension, its contents) is expected to never be deleted or altered until the data retention period is complete. Once a vault in compliance mode is locked, it is **immutable**, meaning the lock *cannot be removed* (the vault itself can be deleted if it is empty and does not contain any recovery points).

A vault locked in Governance mode can be managed or deleted by users who have the appropriate IAM permissions.

A vault lock in Compliance mode cannot be altered or deleted by any user or by AWS. A vault lock in compliance mode has a grace time period you set before it locks and the contents and vault lock become immutable.

Vault lock benefits

AWS Backup Vault Lock provides several benefits, including:

- WORM (*write-once, read-many*) configuration for all the backups you store and create in a backup vault.
- An additional layer of defense that protects backups (recovery points) in your backup vaults from inadvertent or malicious deletions.
- Enforcement of backup lifecycle retention periods, which prevent early deletions by privileged users (including the AWS account root user), and help meet your organization's data protection policies and procedures.

Lock a backup vault using the console

You can add a vault lock to your AWS Backup Vault using the Backup console.

To add a vault lock to your backup vault:

1. Sign in to the AWS Management Console, and open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, find **Backup vaults**. Click the link nested under Backup vaults called **Vault locks**.
3. Under **How vault locks work** or **Vault locks**, click **+ Create vault lock**.
4. In the pane **Vault lock details**, choose which vault to which you want your lock applied.
5. Under **Vault lock mode** choose in which mode you want your vault locked. For more information on choosing your modes, see Vault lock modes earlier on this page.
6. For the **Retention period**, choose the minimum and maximum retention periods (retention periods are *optional*). New backup and copy jobs created in the vault will fail if they do not conform to the retention periods you set; these periods will not apply to recovery points that are already in the vault. Recovery points that were present in the vault prior to the vault lock being enabled will follow their previous lifecycle settings.
7. If you chose compliance mode, a section called **Vault lock start date** is shown. If you chose Governance mode, this will not be displayed, and this step can be skipped.

   In compliance mode, a vault lock has a cooling-off period from the creation of the vault lock until the vault and its lock becomes immutable and unchangeable. You choose the duration of this period (called **grace time**), though it must be *at least* 3 days (72 hours).

> **Important**
>
> Once the grace time expires, the vault and its lock are immutable and cannot be changed or deleted by any user or by AWS.

8. When you are satisfied with the configuration choices, click **Create vault lock**.
9. To confirm you wish to create this lock in the chosen mode, type `confirm` in the text box, then check the box acknowledging the configuration is as intended.

If the steps have been completed successfully, a "Success" banner will appear at the top of the console.

Lock a backup vault programmatically

To configure AWS Backup Vault Lock, use the API `PutBackupVaultLockConfiguration`. The parameters to include will depend on which vault lock mode you intend. If you wish to create a vault lock in governance mode, **do not include** `ChangeableForDays`. If this parameter is included, the vault lock will be created in compliance mode.

Here is a CLI example of a compliance mode vault lock creation:

```
aws backup put-backup-vault-lock-configuration \
        --backup-vault-name my_vault_to_lock \
        --changeable-for-days 3 \
        --min-retention-days 7 \
        --max-retention-days 30
```

Here is a CLI example of a governance mode vault lock creation:

```
aws backup put-backup-vault-lock-configuration \
        --backup-vault-name my_vault_to_lock \
        --min-retention-days 7 \
        --max-retention-days 30
```

You can configure four options.

1. **`BackupVaultName`**

   The name of the vault to lock.

2. **`ChangeableForDays`** (include *only* for compliance mode)

   This parameter instructs AWS Backup to create the vault lock in **compliance mode**. Omit this parameter if you intend to create the lock in **governance mode**.

   This value is expressed in days. It must be a number no less than 3 and no greater than 36,500; otherwise, an error will return.

   From the creation of this vault lock until the expiration of the date specified, the vault lock can be removed from the vault using `DeleteBackupVaultLockConfiguration`. Alternatively, during this time, you can change the configuration using `PutBackupVaultLockConfiguration`.

   On and after the specified date determined by this parameter, the backup vault will be immutable and cannot be changed or deleted.

3. **`MaxRetentionDays`** *(optional)*

   This is a numerical value expressed in days. This is the maximum retention period that the vault retains its recovery points.

   The maximum retention time frame you choose should be in alignment with your organization's policies for retaining data. If your organization instructs data to be retained for a period, this value can be set to that period (in days). For example, financial or banking data may be required to be kept for 7 years (approximately 2,557 days, depending on leap years).

   If not specified, AWS Backup Vault Lock will not enforce a maximum retention period. If specified, backup and copy jobs to this vault with lifecycle retention periods longer than the maximum retention period will fail. Recovery points already saved in the vault prior to the vault lock's creation are not affected. The longest maximum retention period you can specify is 36500 days (approximately 100 years).

4. **`MinRetentionDays`** (*optional*; required for CloudFormation)

   This is a numerical value expressed in days. This is the minimum retention period that the vault retains its recovery points. This setting should be set to the amount of time your organization is *required* to maintain data. For example, if regulations or law requires data to be retained for at least seven years, the value in days would be approximately 2,557, depending on leap years.

   If not specified, AWS Backup Vault Lock will not enforce a minimum retention period. If specified, backup and copy jobs to this vault with lifecycle retention periods shorter than the

minimum retention period will fail. Recovery points already saved in the vault prior to AWS Backup Vault Lock are not affected. The shortest minimum retention period you can specify is 1 day.

Review a backup vault for its AWS Backup Vault Lock configuration

You can review AWS Backup Vault Lock details on a vault anytime by calling `DescribeBackupVault` or `ListBackupVaults` APIs.

To determine whether you applied a vault lock to a backup vault, call `DescribeBackupVault` and check the `Locked` property. If `"Locked": true`, like the following example, you have applied AWS Backup Vault Lock to your backup vault.

```
{
    "BackupVaultName": "my_vault_to_lock",
    "BackupVaultArn": "arn:aws:backup:us-east-1:555500000000:backup-
vault:my_vault_to_lock",
    "EncryptionKeyArn": "arn:aws:kms:us-
east-1:555500000000:key/00000000-1111-2222-3333-000000000000",
    "CreationDate": "2021-09-24T12:25:43.030000-07:00",
    "CreatorRequestId": "ac6ce255-0456-4f84-bbc4-eec919f50709",
    "NumberOfRecoveryPoints": 1,
    "Locked": true,
    "MinRetentionDays": 7,
    "MaxRetentionDays": 30,
    "LockDate": "2021-09-30T10:12:38.089000-07:00"
}
```

The preceding output confirms the following options:

1. `Locked` is a Boolean that indicates whether you have applied AWS Backup Vault Lock to this backup vault. `True` means that AWS Backup Vault Lock causes delete or update operations to the recovery points stored in the vault to fail (regardless of whether you are still in the cooling-off grace time period).
2. `LockDate` is the UTC date and time when your cooling-off grace time period ends. After this time, you cannot delete or change your lock on this vault. Use any publicly-available time converters to convert this string to your local time.

If `"Locked":false`, like the following example, you have not applied a vault lock (or a previous one has been deleted).

```
{
    "BackupVaultName": "my_vault_to_lock",
    "BackupVaultArn": "arn:aws:backup:us-east-1:555500000000:backup-
vault:my_vault_to_lock",
    "EncryptionKeyArn": "arn:aws:kms:us-
east-1:555500000000:key/00000000-1111-2222-3333-000000000000",
    "CreationDate": "2021-09-24T12:25:43.030000-07:00",
    "CreatorRequestId": "ac6ce255-0456-4f84-bbc4-eec919f50709",
    "NumberOfRecoveryPoints": 3,
    "Locked": false
}
```

Vault lock removal during grace time (Compliance mode)

To delete your vault lock during grace time (the time after locking the vault but before your `LockDate`) using the AWS Backup console,

1. Sign in to the AWS Management Console, and open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the left navigation under **My account**, click Backup vaults, then click Backup Vault Lock.
3. Click the vault lock you wish to remove, then click **Manage vault lock**.
4. Click **Delete vault lock**.
5. A warning box will appear, asking you to confirm your intent to delete the vault lock. Type `confirm` into the text box, then click **confirm**.

After the steps have all been completed successfully, a Success banner will appear at the top of the console screen.

To delete your vault lock during grace time using a CLI command, use `DeleteBackupVaultLockConfiguration` like this CLI example:

```
aws backup delete-backup-vault-lock-configuration \
          --backup-vault-name my_vault_to_lock
```

AWS account closure with a locked vault

When you close an AWS account that contains a backup vault, AWS and AWS Backup suspend your account for 90 days with your backups intact. If you do not reopen your account during those 90 days, AWS deletes the contents of your backup vault, even if AWS Backup Vault Lock was in place.

Additional security considerations

AWS Backup Vault Lock adds an additional layer of security to your data protection defense in depth. Vault lock can be combined with these other security features:

- Encryption for your recovery points
- AWS Backup vault and recovery point access policies, which allow you to grant or deny permissions at the vault level,
- AWS Backup security best practices, including its library of customer managed policies that allow you to grant or deny backup and restore permissions by AWS supported service, and
- AWS Backup Audit Manager, which allows you to automate compliance checks for your backups against a list of controls you define.

  You can work through Creating frameworks using the AWS Backup API for the control Resources are in a backup plan with an AWS Backup Vault Lock with AWS Backup Audit Manager to help ensure that your intended resources are protected with a vault lock.
- Mechanisms that render resources inactive can impact the ability to restore them. While they still cannot be deleted in a locked vault, they can be in a state other than active. For instance, the Amazon Elastic Compute Cloud setting that allows you to disable an AMI can temporarily block the ability to restore backups of EC2 instances. This affects all EC2 recovery points, even backups affected by a vault lock or a legal hold.

  If an EC2 backup is disabled, you can re-enable a disabled AMI. Once it is re-enabled, it is eligible to be restored. To block the AMI disable feature, you can use IAM policies to not allow `ec2:DisableImage`.
- Vault Lock and vault access policies govern actions taken through AWS Backup. They don't govern sharing or storage tier changes made directly through Amazon EC2 APIs. Amazon EBS snapshot recovery points can still be shared with `ModifySnapshotAttribute`, and their storage tier can be changed with `ModifySnapshotTier`. Amazon EC2 still can't delete snapshots managed by AWS Backup. To restrict sharing and tier changes, deny `ec2:ModifySnapshotAttribute` and `ec2:ModifySnapshotTier` using IAM policies or

SCPs, and exempt the IAM roles that your backup plans use, because AWS Backup uses these actions for cross-account copy and cold storage transitions. You can also encrypt your source EBS volumes and enable block public access for Amazon EBS snapshots. For more information, see Vault access policies.

Note

AWS Backup Vault Lock is not the same feature as Amazon Glacier Vault Lock, which is compatible only with Amazon Glacier.

Backup creation, maintenance, and restore

A backup, or *recovery point*, represents the content of a resource, such as an Amazon Elastic Block Store (Amazon EBS) volume or Amazon DynamoDB table, at a specified time. Recovery point is a term that refers generally to the different backups in AWS services, such as Amazon EBS snapshots and DynamoDB backups. The terms *recovery point* and *backup* are used interchangeably.

AWS Backup saves recovery points in backup vaults, which you can organize according to your business needs. For example, you can save a set of resources that contain financial information for fiscal year 2020. When you need to recover a resource, you can use either the AWS Backup console or the AWS Command Line Interface (AWS CLI) to find and recover the resource you need.

Each recovery point has a unique ID. The unique ID is at the end of the recovery point's Amazon Resource Name (ARN). For examples of recovery point ARNs and unique IDs, see the table in Resources and operations.

> **Important**
>
> To avoid additional charges, configure your retention policy with a warm storage duration of **at least one week**. For more information, see Metering, costs, and billing for AWS Backup.

The following sections provide an overview of the basic backup management tasks in AWS Backup.

Topics

- Creating an on-demand backup using AWS Backup
- Continuous backups and point-in-time recovery (PITR)
- Backup creation by resource type
- Backup and tag copy
- Add or remove tags on existing recovery points
- Backup deletion
- Backup and tag edits
- Backup search
- Backup tiering

- Access points
- Restore a backup by resource type
- Restore testing
- Stop a backup job
- View existing backups

Creating an on-demand backup using AWS Backup

On the AWS Backup console, the **Protected resources** page lists resources that have been backed up by AWS Backup at least once. If you're using AWS Backup for the first time, there aren't any resources (such as Amazon EBS volumes or Amazon RDS databases) listed on this page. This is true even if a resource was assigned to a backup plan and that backup plan has not run a scheduled backup job at least once.

Note: An on-demand backup begins to back up your resource immediately. You can choose an on-demand backup if you wish to create a backup at a time other than the scheduled time defined in a backup plan. An on-demand backup can be used, for example, to test backup and functionality at any time.

On-demand backups cannot be used with point-in-time recovery (PITR), because an on-demand backup preserves resources in the state they are in when the backup is taken, but PITR uses continuous backups, which record changes over a period of time.

Considerations

- If the AWS Backup default role is not present in your account, one is created for you with the correct permissions.
- When backups expire and are marked for deletion as part of your lifecycle policy, AWS Backup deletes the backups at a randomly chosen point over the following 8 hours. This window helps ensure consistent performance.
- For Amazon EC2 resources, AWS Backup automatically copies existing group and individual resource tags, in addition to any tags that you add in this step.
- AWS Backup takes EC2 backups with "no reboot" as the default behavior. AWS Backup currently supports resources running on Amazon EC2, and certain instance types are not supported. For more information, see Create Windows VSS backups.

To create an on-demand backup

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. On the dashboard, choose **Create an on-demand backup**. Or, in the navigation pane, choose **Protected resources** and then choose **Create an on-demand backup**.
3. For **Resource type** page, choose the resource type that you want to back up. For example, choose **DynamoDB** for Amazon DynamoDB tables.
4. Choose the name or ID of the resource to protect. For example, choose the name of the DynamoDB table for Amazon DynamoDB.
5. Ensure that **Create backup now** is selected.
6. If the resource type supports transition to cold storage, **Cold storage** is present. For more information, see the **Lifecycle to cold storage** column in table Feature availability by resource.

   To specify when this backup goes to cold storage, choose **Move backups from warm to cold storage** and then specify the time in warm storage.
7. For **Total retention period**, specify the number of days. If you specified time in cold storage, the retention period is divided between warm and cold storage.
8. Choose an existing **Backup vault** or create a new one. Choosing **Create new Backup vault** opens a new page to create a vault and then returns you to the **Create on-demand backup** page when you are finished.
9. For **IAM role**, choose the default role or a role that you created.
10. To assign a tag to your on-demand backup, expand **Tags added to recovery points**, choose **Add new tag**, and enter a tag key and tag value.
11. **Advanced backup settings** options vary by resource type:

    - For **EC2** resources: To take application-consistent snapshots using Windows Volume Shadow Copy Service (VSS), choose **Windows VSS**.
    - For **Amazon S3** resources: You can choose to exclude Access Control Lists (ACLs) from your backup by leaving **Backup Access Control Lists (ACLs)** unselected.
12. Choose **Create on-demand backup**. This opens the **Jobs** page, where you can see a list of jobs and view job status.

Continuous backups and point-in-time recovery (PITR)

For some resources, AWS Backup supports continuous backups and point-in-time recovery (PITR) in addition to snapshot backups.

With **continuous backups**, you can restore your AWS Backup-supported resource by rewinding it back to a specific time that you choose, within 1 second of precision (going back a maximum of 35 days). Continuous backup works by first creating a full backup of your resource, and then constantly backing up your resource's transaction logs. PITR works by accessing your full backup and replaying the transaction log to the time that you tell AWS Backup to recover.

Alternatively, **snapshot backups** can be taken as frequently as every hour. Snapshot backups can be stored for up to a maximum of 100 years. Snapshots can be copied for full or incremental backups.

Because continuous and snapshot backups offer different advantages, we recommend that you protect your resources with both continuous and snapshot backup rules.

An on-demand backup begins to back up your resource immediately. You can choose an on-demand backup if you wish to create a backup at a time other than the scheduled time defined in a backup plan. An on-demand backup can be used, for example, to test backup and functionality at any time.

You can't use on-demand backups with PITR, because an on-demand backup preserves resources in the state they are in when the backup is taken, while PITR uses continuous backups, which record changes over a period of time.

You can opt in to continuous backups for supported resources when you create a backup plan in AWS Backup using the AWS Backup console or the API. The continuous backup plan creates one continuous recovery point and updates that recovery point whenever the job runs.

Contents

- Point-in-time recovery considerations
- Supported services for continuous backup and PITR
- Finding a continuous backup
- Restoring a continuous backup
- Stopping or deleting continuous backups

- Copying continuous backups
- Changing your retention period
- Removing the only continuous backup rule from a backup plan

Point-in-time recovery considerations

Be aware of the following considerations for point-in-time recovery:

- **Automatic fallback to snapshots** — If AWS Backup is unable to perform a continuous backup, it tries to perform a *snapshot* backup instead.
- **No support for on-demand continuous backups** — AWS Backup doesn't support on-demand continuous backup because on-demand backup records a point in time, whereas continuous backup records changes over a period of time.
- **No support for transition to cold storage** — Continuous backups don't support transition to cold storage because transition to cold requires a minimum transition period of 90 days, whereas continuous backups have a maximum retention period of 35 days.
- **Restoring recent activity** — Amazon RDS activity allows restores up until the most recent 5 minutes of activity; Aurora allows restores up until the most recent activity as indicated by `LatestRestorableTime` (typically less than 5 minutes); Amazon S3 allows restores up until the most recent 15 minutes of activity.

> **Important**
>
> A single resource can only have one continuous backup. Expand below for additional details and best practices.

Overlapping continuous backups on the same resource

Each resource (such as an Amazon S3 bucket or an Amazon RDS database) can only have one continuous backup (recovery point); additional continuous backups are redundant. When multiple backup policies, plans, or rules instruct AWS Backup to create multiple continuous backups for the same resource, the following process applies:

- If multiple rules specify that more than one continuous backup should be in a single vault, AWS Backup follows the rule with the longest retention period (lifecycle) and ignores additional rules.

- If multiple rules specify that more than one continuous backup should be in more than one vault, AWS Backup creates one continuous backup according to the first rule processed. Each subsequent rule specifying a continuous backup for a resource that already has a continuous backup will result in a snapshot (periodic) backup instead.

When duplicate continuous backup plans occur, the snapshot backups created after the continuous recovery point can show a status of `Completed with issues`. The detailed information of this recovery point will show an error similar to "`Enabling continuous backup failed, because of the following error: PITR already configured in backup plan: [ARN]`". This error indicates that there is already at least one continuous backup configured (for a different recovery point than the one containing the error). That first continuous backup (recovery point) is able to be used for point in time restore (PITR) as long as it is has a status of `COMPLETED`.

To prevent the creation of unintended snapshots with issues (and error message), review your organization backup strategy. If necessary, adjust backup plans and policies that create multiple continuous backups of the same resource.

After you have made adjustments that result in only one continuous backup for a resource, the snapshot backups will be retained according to the specified lifecycle of the plan that created them, then they will transition to `EXPIRED` and be deleted. The continuous backup and its point-in-time recovery ability will be maintained according to the rule that created it.

Supported services for continuous backup and PITR

AWS Backup supports continuous backups and point-in-time recovery for the following services and applications:

Amazon S3

To turn on PITR for S3 backups, continuous backups need to part of the backup plan.

While this original backup of the source bucket can have PITR active, cross-Region or cross-account destination copies will not have PITR, and restoring from these copies will restore to the time they were created (the copies will be snapshot copies) instead of restoring to a specified point in time.

AWS Backup for S3 relies on receiving S3 events through Amazon EventBridge. If this setting is disabled in S3 bucket notification settings, continuous backups will stop for those buckets with the setting turned off. For more information, see Amazon EventBridge dependency for S3 continuous backups.

Disabling AWS Backup's Amazon EventBridge rule will also result in your continuous backup stopping. If you have an active backup plan with a continuous backup rule, when that rule re-triggers, AWS Backup will recreate the Amazon EventBridge rule and a new continuous backup will be created.

RDS

AWS Backup supports continuous backups and point-in-time recovery for all Amazon RDS instances and Aurora that are supported by the native Amazon RDS service. AWS Backup does not support continuous backups or point-in-time recovery for Amazon RDS Multi-AZ clusters.

Backup schedules: When you enable continuous backups for an Amazon RDS instance through AWS Backup, AWS Backup takes over the Amazon RDS automated backup window (the native daily snapshot that anchors point-in-time recovery). AWS Backup positions this automated backup window near the Amazon RDS maintenance window to prevent conflicts. You cannot directly configure the automated backup window while AWS Backup manages continuous backups, but you can influence its placement by adjusting your Amazon RDS maintenance window. The automated backup window repositions itself on the next backup cycle. RDS takes snapshots once per day, even if a backup plan has a frequency for snapshot backups other than once per day.

> **Note**
>
> AWS Backup does not modify or manage the Amazon RDS maintenance window. The maintenance window remains under your control and can be adjusted through Amazon RDS settings. Backup jobs initiated by a snapshot rule in your backup plan run on the schedule you define and can still fail if they overlap with the maintenance window. If this occurs, you receive an error similar to "Backup job could not start because it is either inside or too close to the weekly maintenance window configured in RDS instance." To avoid this error, schedule your snapshot backup rules outside of your configured Amazon RDS maintenance window.

Settings: After you apply an AWS Backup continuous backup rule to an Amazon RDS instance, you can't create or modify continuous backup settings in Amazon RDS. You must make modifications through the AWS Backup console or the AWS Backup CLI. When you turn on automated backups for the first time, an outage occurs if you change the backup retention period of the DB instance from 0 to a nonzero value. Plan this change during a maintenance window to minimize impact.

For more information about enabling automated backups, see Enabling automated backups in the *Amazon RDS User Guide*.

Transition control of continuous backup for an Amazon RDS instance back to Amazon RDS:

Console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Backup plans**.
3. Delete all the Amazon RDS backup plans with continuous backup protecting that resource.
4. Choose **Backup vaults**. Delete the continuous backup recovery point from your backup vault. Or, wait for their retention period to elapse, causing AWS Backup to automatically delete the recovery point.

After you complete these steps, AWS Backup will transition continuous backup control of your resource back to Amazon RDS.

AWS CLI

Call the `DisassociateRecoveryPoint` API operation.

To learn more, see DisassociateRecoveryPoint.

IAM permissions required for Amazon RDS continuous backups

- To use AWS Backup to configure continuous backups for your Amazon RDS database, verify that the API permission `rds:ModifyDBInstance` exists in the IAM role defined by your backup plan configuration. To restore Amazon RDS continuous backups, you must add the permission `rds:RestoreDBInstanceToPointInTime` to the IAM role that you submitted for the restore job. You can use the `AWS Backup default service role` to perform backups and restores.
- To describe the range of times available for point-in-time recovery, AWS Backup calls `rds:DescribeDBInstanceAutomatedBackups`. In the AWS Backup console, you must have the `rds:DescribeDBInstanceAutomatedBackups` API permission in your AWS Identity and Access Management (IAM) managed policy. You can use the `AWSBackupFullAccess` or `AWSBackupOperatorAccess` managed policies. Both policies have all required permissions. For more information, see Managed Policies.

Retention periods: When you change your PITR retention period, AWS Backup calls `ModifyDBInstance`to apply that change.

When AWS Backup enables PITR for the first time on an Amazon RDS instance (changing retention from 0 to a non-zero value), the operation is scheduled to occur during your database's next maintenance window to prevent unexpected downtime.

Scenarios:

- **First-time PITR enablement:** When PITR is enabled on an Amazon RDS instance for the first time (regardless of whether it's managed by AWS Backup or configured directly), the change is queued for the next maintenance window. AWS Backup automatically creates snapshot backups to maintain coverage until PITR becomes active.
- **PITR retention changes:** Non-zero to non-zero retention changes apply immediately without restart.
- **PITR disabling:** Changes from non-zero to zero retention are scheduled for the next maintenance window.

Backup coverage during transition:

- Snapshot backups provide protection while waiting for maintenance window
- Continuous recovery points become available when the backup job runs after PITR is enabled
- No gap in backup protection occurs during the transition period
- Recovery granularity may be limited to snapshot intervals until PITR is fully active

Note: Stopping the RDS instance will remove pending changes. PITR configuration changes will be requeued by the next backup job and applied during a subsequent maintenance window.

Copies of Amazon RDS continuous backups:

- **Creating copies of Amazon RDS continuous backups** — You can't create copies of Amazon RDS continuous backups because AWS Backup for Amazon RDS does not allow copying transaction logs. Instead, AWS Backup creates a snapshot and copies it with the frequency specified in the backup plan.

Restores: You can perform a point-in-time restore using either AWS Backup or Amazon RDS. For AWS Backup console instructions, see Restoring an Amazon RDS Database. For Amazon RDS instructions, see Restoring a DB Instance to a specified time in the *Amazon RDS User Guide*.

> **Tip**
>
> A multi AZ (availability zone) database instance set to `Always On` should not have a backup retention set to zero. If errors occur, use AWS CLI command `disassociate-recovery-point` instead of `delete-recovery-point`, then change the retention setting to 1 in your Amazon RDS settings.

For general information about working with Amazon RDS, see the Amazon RDS User Guide.

CLI examples for RDS and Aurora PITR restore

The following examples demonstrate how to restore RDS and Aurora databases to a point in time using the AWS Backup CLI with metadata parameters.

Example: Restore RDS database to a point in time with metadata

```
aws backup start-restore-job \
    --recovery-point-arn arn:aws:backup:us-east-1:123456789012:recovery-
point:1EB3B5E7-9EB0-435A-A80B-108B488B0D45 \
    --metadata '{"DBInstanceIdentifier":"restored-db-
instance","Engine":"mysql","UseLatestRestorableTime":"false","RestoreTime":"2024-01-15T10:30:00
 \
    --iam-role-arn arn:aws:iam::123456789012:role/service-role/
AWSBackupDefaultServiceRole \
    --resource-type RDS \
    --copy-source-tags-to-restored-resource
```

Example: Restore Aurora cluster to a point in time

```
aws backup start-restore-job \
    --recovery-point-arn arn:aws:backup:us-east-1:123456789012:recovery-
point:2FC4C6F8-0FC1-546B-B91C-209C599C1D56 \
    --metadata '{"DBClusterIdentifier":"restored-aurora-cluster","Engine":"aurora-
mysql","UseLatestRestorableTime":"true"}' \
```

```
    --iam-role-arn arn:aws:iam::123456789012:role/service-role/
AWSBackupDefaultServiceRole \
    --resource-type Aurora \
    --copy-source-tags-to-restored-resource
```

Metadata parameters for RDS PITR restore

The following metadata parameters are supported for RDS and Aurora PITR restores:

- **DBInstanceIdentifier** (RDS) or **DBClusterIdentifier** (Aurora) - Required. The name for the restored database.
- **Engine** - Required. The database engine (e.g., mysql, postgres, aurora-mysql, aurora-postgresql).
- **UseLatestRestorableTime** - Optional. Set to "true" to restore to the latest restorable time, or "false" to specify a RestoreTime.
- **RestoreTime** - Optional. The date and time to restore to (ISO 8601 format). Required if UseLatestRestorableTime is "false".

Copy tags to restored resource

Use the `--copy-source-tags-to-restored-resource` flag to copy tags from the source database to the restored database. This ensures tag-based access controls and cost allocation tags are preserved.

For complete details on RDS PITR restore parameters, see:

- RestoreDBInstanceToPointInTime in the Amazon RDS API Reference
- RestoreDBClusterToPointInTime in the Amazon RDS API Reference

Aurora

To enable continuous backup of your Aurora resources, see the steps in the first section of this page.

The procedure to restore an Aurora cluster to a point in time is a variation of the steps to restore a snapshot of an aurora cluster.

When you conduct a point in time restore, the console displays a **restore time** section. See *Restoring a continuous backup* further down on this page in Working with Continuous backups.

Important

Aurora continuous backups are supported in vaults protected by AWS Backup Vault Lock, and the vault's minimum and maximum retention settings are enforced on the recovery point. However, Aurora continuous backups do not support the logically air-gapped vault feature. To use a logically air-gapped vault with Aurora, use periodic snapshot backups instead.
The recovery point objective (RPO) for Aurora continuous backups is typically less than 5 minutes, as Aurora copies data to Amazon S3 continuously in the background. Use the `LatestRestorableTime` value to determine the most recent point to which you can restore.

Retention periods and backup windows: When you enable or change continuous backup settings for an Aurora cluster, AWS Backup calls `ModifyDBCluster` to apply those changes. This can modify the cluster's `PreferredBackupWindow`. If you have other configuration updates pending the next maintenance window, enabling continuous backups may also apply those pending changes immediately.

Note

To use AWS Backup to configure continuous backups for your Aurora cluster, verify that the API permission `rds:ModifyDBCluster` exists in the IAM role defined by your backup plan configuration.

SAP HANA on Amazon EC2 instances

You can make continuous backups , which can be used with point-in-time restore (PITR) (note that on-demand backups preserve resources in the state in which they are taken; whereas PITR uses continuous backups which record changes over a period of time).

With continuous backups, you can restore your SAP HANA database on an EC2 instance by rewinding it back to a specific time that you choose, within 1 second of precision (going back a maximum of 35 days). Continuous backup works by first creating a full backup of your resource, and then constantly backing up your resource’s transaction logs. PITR restore works by accessing your full backup and replaying the transaction log to the time that you tell AWS Backup to recover.

You can opt in to continuous backups when you create a backup plan in AWS Backup using the AWS Backup console or the API.

To enable continuous backups using the console

1. Sign in to the AWS Management Console, and open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Backup plans**, and then choose **Create Backup plan**.
3. Under **Backup rules**, choose **Add Backup rule**.
4. In the **Backup rule configuration** section, select **Enable continuous backups for supported resources**.

After you disable PITR (point-in-time restore) for SAP HANA database backups, logs will continue to be sent to AWS Backup until the recovery point expires (status equals `EXPIRED`). You can change to an alternative log backup location in SAP HANA to stop the transmission of logs to AWS Backup.

A continuous recovery point with a status of `STOPPED` indicates that a continuous recovery point has been interrupted; that is, the logs transmitted from SAP HANA to AWS Backup that show the incremental changes to a database have a gap. The recovery points that occur within this timeframe gap have a status of `STOPPED`..

For issues you may encounter during restore jobs of continuous backups (recovery points), see the SAP HANA Restore troubleshooting section of this guide.

Finding a continuous backup

You can use the AWS Backup console to find your continuous backup.

To find a continuous backup using the AWS Backup console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Backup vaults**, and then choose your backup vault in the list.
3. In the **Backups** section, in the **Backup type** column, sort for **Continuous** recovery points. You can also sort by **Recovery point ID** for the prefix *continuous*.

Restoring a continuous backup

To restore a continuous backup using the AWS Backup console

- During the PITR restore process, the AWS Backup console displays a **Restore time** section. In this section, do one of the following:
  - Choose to restore to the **Latest restorable time**.
  - Choose **Specify date and time** to enter your own date and time within your retention period.

To restore a continuous backup using the AWS Backup API

1. For Amazon S3 see Use the AWS Backup API, CLI, or SDK to restore S3 recovery points.
2. For Amazon RDS see Use the AWS Backup API, CLI, or SDK to restore Amazon RDS recovery points.

Stopping or deleting continuous backups

You can stop the creation of continuous backups or you can delete specific backups (point-in-time-recovery or PITR points).

If you want to stop continuous backups, you must delete the continuous backup rule from your backup plan. If you wish to stop continuous backups for one or more resources but not for all resources, create a new backup plan with the continuous backup rule for those resources you still want to be continuously backed up. If instead you only delete a continuous backup recovery point from your backup vault, your backup plan will still continue to execute the continuous backup rule, creating a new recovery point.

However, even after you delete your continuous backup rule, AWS Backup remembers the retention period from your now-deleted backup rule. It will automatically delete your continuous backup recovery point from your backup vault based on your specified retention period.

When deleting Amazon RDS recovery points, consider:

- A multi AZ (availability zone) database instance set to `Always On` should not have a backup retention set to zero. If errors occur, use AWS CLI command `disassociate-recovery-point` instead of `delete-recovery-point`, then change the retention setting to 1 in your Amazon RDS settings.

- When point-in-time recovery (PITR) is disabled, the change is scheduled for your maintenance window. You may continue to incur backup storage costs until the maintenance window applies the change. This process may take up to 7 days depending on your maintenance window schedule.

When deleting Aurora recovery points, consider:

If this is selected for an Amazon Aurora recovery point, AWS Backup sets the retention period to 1 day. Aurora backups cannot be completely deleted until the source cluster has also been deleted.

Transitioning continuous backups between backup plans

If you transition a resource's continuous backup protection from one backup plan to another, be aware of the following behavior:

- **Amazon RDS:** After the existing continuous recovery point's retention period elapses and the recovery point is removed from the vault, the next execution of the new backup plan configures continuous backups. There is no gap in point-in-time restore coverage during this transition. For more information, see How do I stop continuous backups for Amazon RDS in AWS Backup?
- **Amazon S3:** A disassociated Amazon S3 continuous recovery point transitions to a `STOPPED` state and remains in its backup vault. For more information, see Backup deletion.

> **Note**
>
> Calling `DisassociateRecoveryPoint` on an Amazon S3 continuous recovery point that is already in `EXPIRED` status has no effect. The recovery point state and configuration remain unchanged.

In all cases, when you remove a continuous backup rule from a backup plan, AWS Backup remembers the retention period from the deleted rule and automatically deletes the continuous backup recovery point when the retention period elapses. For more information, see Removing the only continuous backup rule from a backup plan.

Copying continuous backups

If a continuous backup rule also specifies a cross-account or cross-Region copy and AWS Backup supports the operation for the resource type, AWS Backup takes a snapshot of the resource and

copies the snapshot to the destination vault. To learn more about copying your recovery points across accounts and Regions, see Copying a backup .

Continuous backups create a periodic backups in accordance with the frequency set in the backup plan rule in the destination account and/or Region.

AWS Backup does not support on-demand copies of continuous backups.

> **Note**
>
> For database resources, ensure that your IAM role has the `rds:DeleteDBSnapshot` permission. AWS Backup temporarily creates a source snapshot during a point-in-time recovery (PITR) copy. After the copy completes, the service removes the snapshot. If the role lacks this permission, AWS Backup cannot remove the temporary snapshot. Instead, it creates an expired recovery point to prevent a snapshot leak. You must manually delete any expired recovery points that result from insufficient permissions.

Changing your retention period

You must use your backup plan to increase or decrease the retention period for your existing continuous backup rule. The minimum retention period is 1 day. The maximum retention period is 35 days. The change in retention period will take effect when the next backup is completed following this change. You cannot use the `UpdateRecoveryPointLifecycle` API or CLI to update the retention period of any continuous backup.

Retention period by service

The retention behavior can be specific to the resource type that is backed up in the recovery point.

Amazon S3

When a retention period of an Amazon S3 continuous recovery point has changed (increased or decreased), that recovery point status will become `STOPPED`. A new continuous recovery point with the altered retention settings will be created.

Amazon Aurora and Amazon RDS

For recovery points of Aurora and Amazon RDS resources, only one recovery point is possible at a time. No new recovery points are created when a retention period is changed; instead, AWS Backup updates the existing recovery point with the retention specifications within the backup plan.

SAP HANA on Amazon EC2

For recovery points of SAP HANA on EC2 resources, only one recovery point is possible at a time. No new recovery points are created when a retention period is changed; instead, AWS Backup updates the existing recovery point with the retention specifications within the backup plan.

Ensure that you set the retention period for these backups to be greater than the backup frequency to avoid a scenario where the recovery point transitions to `EXPIRED` state.

> **Tip**
>
> A backup frequency rule for a continuous backup is not the same as a periodic backup snapshot. Each backup plan, even one that doesn't create a snapshot, has a frequency you set (hourly, daily, weekly, or monthly as examples) for maintenance and syncing purposes.

Removing the only continuous backup rule from a backup plan

When you create a backup plan with a continuous backup rule and then you remove that rule, AWS Backup remembers the retention period from your now-deleted rule. It will delete the continuous backup from your backup vault when the retention period elapses.

Backup creation by resource type

With AWS Backup, you can create backups automatically using backup plans or manually by initiating an on-demand backup.

Creating automatic backups

When backups are created automatically by backup plans, they are configured with the lifecycle settings that are defined in the backup plan. They are organized in the backup vault that is specified in the backup plan. They are also assigned the tags that are listed in the backup plan. For more information about backup plans, see Backup plans.

Creating on-demand backups

When you create an on-demand backup, you can configure these settings for the backup that is being created. When a backup is created either automatically or manually, a backup *job* is initiated. For how to create an on-demand backup, see Creating an on-demand backup using AWS Backup.

Note: An on-demand backup creates a backup job; the backup job will transition in state of `Running` within an hour (or when specified). You can choose an on-demand backup if you wish to create a backup at a time other than the scheduled time defined in a backup plan. An on-demand backup can be used, for example, to test backup and functionality at any time.

On-demand backups cannot be used with point-in-time restore (PITR) since an on-demand backup preserves resources in the state they are in when the backup is taken, whereas PITR uses continuous backups which record changes over a period of time.

Backup job statuses

Each backup job has a unique ID. For example, `D48D8717-0C9D-72DF-1F56-14E703BF2345`.

You can view the status of a backup job on the **Jobs** page of the AWS Backup console. Backup job statuses include `CREATED`, `PENDING`, `RUNNING`, `ABORTING`, `ABORTED`, `COMPLETED`, `FAILED`, `EXPIRED`, and `PARTIAL`.

Incremental backups

Many resources support incremental backup with AWS Backup. A full list is available in the incremental backup section of the Feature availability by resource table.

Although each backup after the first (full) one is incremental (meaning it only captures changes from the previous backup), all backups made with AWS Backup retain the necessary reference data to allow a full restore. This is true even if the original (full) backup has reached the end of its lifecycle and been deleted.

For example, if your day 1 (full) backup was deleted due to a 3-day lifecycle policy, you would still be able to perform a full restore with the backups from days 2 and 3. AWS Backup maintains the necessary reference data from day 1 to do so.

Incremental backups and Regions

Backups of resources which are fully managed by AWS Backup can only be incremental if the vault in which the backup is created also contains an earlier backup (incremental or full); other resource types (not fully managed by AWS Backup) can have incremental backups as long as there is a previous backup of the resource within the same *Region*.

> **Note**
>
> Not all resource types support incremental backups. Some resources, such as Amazon Aurora, offer incremental backup only through continuous backups and point-in-time restore (PITR), not through snapshot-based backups. For a full list of which resources support incremental backups, see the Feature availability by resource table.

Access to source resources

AWS Backup needs access to your source resources to back them up. For example:

- To back up an Amazon EC2 instance, the instance can be in the `running` or `stopped` state, but not the `terminated` state. This is because a `running` or `stopped` instance can communicate with AWS Backup, but a `terminated` instance cannot.
- To back up a virtual machine, its hypervisor must have the Backup gateway status `ONLINE`. For more information, see Understanding hypervisor status.
- To back up an Amazon RDS database, Amazon Aurora, or Amazon DocumentDB cluster, those resources must have the status `AVAILABLE`.
- To back up an Amazon Elastic File System (Amazon EFS), it must have the status `AVAILABLE`.
- To back up an Amazon FSx file system, it must have the status `AVAILABLE`. If the status is `UPDATING`, the backup request is queued until the file system becomes `AVAILABLE`.

  FSx for ONTAP doesn't support backing up certain volume types, including DP (data-protection) volumes, LS (load-sharing) volumes, full volumes, or volumes on file systems that are full. For more information, please see FSx for ONTAP Working with backups.

AWS Backup retains previously-created backups consistent with your lifecycle policy, regardless of the health of your source resource.

Topics

- CloudFormation stack backups
- Amazon Aurora DSQL backups
- Advanced DynamoDB backup
- Amazon EBS and AWS Backup

- Amazon Relational Database Service backups
- Amazon Redshift backups
- Amazon Redshift Serverless backups
- Amazon EKS backups
- SAP HANA backup on Amazon EC2
- Amazon S3 backups
- Amazon Timestream backups
- Virtual machine backups
- Create Windows VSS backups

CloudFormation stack backups

A CloudFormation stack consists of multiple stateful and stateless resources that you can back up as a single unit. In other words, you can backup and restore an application containing multiple resources by backing up a stack and restoring the resources within it. All the resources in a stack are defined by the stack's CloudFormation template.

When a CloudFormation stack is backed up, recovery points are created for the CloudFormation template and for each additional resource supported by AWS Backup in the stack. These recovery points are grouped together within a overarching recovery point called a **composite**.

This composite recovery point cannot be restored, but nested recovery points can be restored. You can restore anywhere from one to all nested backups within a composite backup using the console or the AWS CLI.

CloudFormation application stack terminology

- **Composite recovery point**: A recovery point used to group nested recovery points together, as well other metadata.
- **Nested recovery point**: A recovery point of a resource that is part of a CloudFormation stack and is backed up as part of the composite recovery point. Each nested recovery point belongs in the stack of one composite recovery point.
- **Composite job**: A backup, copy, or restore job for a CloudFormation stack which can trigger other backup jobs for individual resources within the stack.
- **Nested job**: A backup, copy, or restore job for a resource within a CloudFormation stack.

CloudFormation stack backup jobs

The process of a backup creation is called a backup job. A CloudFormation stack backup job has a status. When a backup job has finished, it has the status of `Completed`. This signifies a CloudFormation recovery point (a backup) has been created.

CloudFormation stacks can be backed up using the console or backed up programatically. To backup any resource, including a CloudFormation stack, see Creating a backup elsewhere in this *AWS Backup Developer Guide*.

CloudFormation stacks can be backed up using the API command `StartBackupJob`. Note that the documentation and console refer to composite and nested recovery points; the API language uses the terminology "parent and child recovery points" in the same contextual relationship.

CloudFormation stacks contain all AWS resources are indicated by your CloudFormation template. Note that your template may contain resources not yet supported by AWS Backup. If your template contains a combination of AWS supported resources and unsupported resources, AWS Backup will still back up the template into a composite stack, but Backup will only create recovery points of the Backup-supported services. All resource types contained within the CloudFormation template will be included within a backup, even if you have not opted into to a particular service (toggling a service to “Enabled” in console Settings).

CloudFormation recovery point

Recovery point status

When the backup job of a stack is finished (the job status is `Completed`), a backup of the stack has been created. This backup is also known as a composite recovery point. A composite recovery point can have one of the following statuses: `Completed`, `Failed`, or `Partial`. Note that a backup job has a status, and a recovery point (also called a backup) also has a separate status.

A completed backup job means your entire stack and the resources within in are protected by AWS Backup. A failed status indicates that the backup job was unsuccessful; you should create the backup again once the issue that caused the failure is corrected.

A `Partial` status means that not all the resources in the stack were backed up. This may happen if the CloudFormation template contains resources that are not currently supported by AWS Backup, or it may happen if one or more of the backup jobs belonging to resources within the stack (nested resources) have statuses other than `Completed`. You can manually create an on-demand backup to rerun any resources that resulted in a status other than `Completed`. If you expected the stack

to have the status of `Completed` but it is marked as `Partial` instead, check to see which of the conditions above might be true about your stack.

Each nested resource within the composite recovery point has its own individual recovery point, each with its own status (either `Completed` or `Failed`). Nested recovery points with a status of `Completed` can be restored.

Manage recovery points

Composite recovery points (backups) can be copied; nested recovery points can be copied, deleted, disassociated, or restored. A composite recovery point which contains nested backups cannot be deleted. After the nested recovery points within a composite recovery point have been deleted or disassociated, you can manually delete the composite recovery point manually or let it remain until the backup plan lifecycle deletes it.

Delete a recovery point

You can delete a recovery point using the AWS Backup console or using the AWS CLI.

To delete recovery points using the AWS Backup console,

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Click on **Protected Resources** in the left-hand navigation. In the text box, type `CloudFormation` to display only your CloudFormation stacks.
3. Composite recovery points will be displayed in the Recovery points pane. The plus sign (+) to the left of each recovery point ID can be clicked to expand each composite recovery point, showing all nested recovery points contained in the composite. You can check the box to the left of any recovery point to include it in your selection of recovery points you wish to delete.
4. Click the **Delete** button.

When you use the console to delete one or more composite recovery points, a warning box will pop up. This warning box requires you to confirm your intention to delete the composite recovery points, including nested recovery points within composite stacks.

To delete recovery points using API, use the `DeleteRecoveryPoint` command.

When you use API with the AWS Command Line Interface you must delete all nested recovery points prior to deleting a composite point. If you send an API request to delete a composite stack backup (recovery point) that still contains nested recovery points within it, the request will return an error.

Disassociate a nested recovery point from composite recovery point

You can disassociate a nested recovery point from a composite recovery point (for example, you wish to keep the nested recovery point but delete the composite recovery point). Both recovery points will remain, but they will no longer be connected; that is, actions that occur on the composite recovery point will no longer apply to the nested recovery point once it has been disassociated.

You can disassociate the recovery point using the console, or you can call the API `DisassociateRecoveryPointFromParent`. [Note that the API calls use the term “parent” to refer to composite recovery points.]

Copy a recovery point

You can copy a composite recovery point, or you can copy a nested recovery point if the resource supports cross-account and cross-Region copy.

To copy recovery points using the AWS Backup console:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Click on **Protected Resources** in the left-hand navigation. In the text box, type `CloudFormation` to display only your CloudFormation stacks.
3. Composite recovery points will be displayed in the Recovery points pane. The plus sign (+) to the left of each recovery point ID can be clicked to expand each composite recovery point, showing all nested recovery points contained in the composite. You can click the radial circle button to the left of any recovery point to copy it.
4. Once it is selected, click the **Copy** button in the top-right corner of the pane.

When you copy a composite recovery point, nested recovery points that don’t support copy functionality won’t end up in the copied stack. The composite recovery point will have a status of `Partial`.

Frequently Asked Questions

1. *"What is included as part of the application backup?"*

   As part of each backup of an application defined using CloudFormation, the template, the processed value of each parameter in the template, and the nested resources supported by AWS

Backup are backed up. A nested resource is backed up in the same way as an individual resource not part of a CloudFormation stack is backed up. Note that values of parameters marked as `no-echo` will not be backed up.

2. *"Can I back up my CloudFormation stack that has nested stacks?"*

   Yes. Your CloudFormation stacks which contain nested stacks can be in your backup.

3. *"Does a `Partial` status mean the creation of my backup failed?"*

   No. A partial status indicates that some of the recovery points were backed up, while some were not. There are three conditions to check if you were expecting a `Completed` backup result:

   a. Does your CloudFormation stack contain resources currently unsupported by AWS Backup? For a list of supported resources, see Supported AWS resources and third-party applications in our Developer Guide.

   b. One or more of the backup jobs belonging to resources within the stack were not successful and the job has to be rerun.

   c. A nested recovery point was deleted or disassociated from the composite recovery point.

4. *"How do I exclude resources in my CloudFormation stack backup?"*

   When you back up your CloudFormation stack, you can exclude resources from being part of the backup. In the console, during the create a backup plan and update a backup plan processes, there is an assign resources step. In this step, there is a **Resource selection** section. If you choose **include specific resource types** and have included CloudFormation as a resource to backup, you can **exclude specific resource IDs from the selected resource types**. You can also use tags to exclude resources within the stack.

   > **Important**
   >
   > Resources that belong to a CloudFormation stack must be in the exclusion list of *every* resource selection in the backup plan that targets that resource type. If a resource is excluded only from the CloudFormation backup selection but is matched by another resource selection (for example, an S3 selection using `arn:aws:s3:::*`), the resource will still be backed up by that other selection.

   Using CLI, you can use

- `NotResources` in your backup plan to exclude a specific resource from your CloudFormation stacks.
- `StringNotLike` to exclude items through tags.

5. *"What types of backups are supported for nested resources?"*

   Backups of nested resources may be either full or incremental backups, depending on which kind of backup is supported by AWS Backup for these resources. For more information, see How incremental backups work. However, note that PITR (point-in-time restore) is not supported for Amazon S3 and Amazon RDS nested resources.

6. *"Are change sets that are part of the CloudFormation stack backed up?"*

   No. Change sets are not backed up as part of CloudFormation stack backup.

7. *"How does the status of the CloudFormation stack impact the backup?"*

   The status of the CloudFormation stack may impact the backup. A stack with a status that includes `COMPLETE` can be backed up, such as statuses `CREATE_COMPLETE`, `ROLLBACK_COMPLETE`, `UPDATE_COMPLETE`, `UPDATE_ROLLBACK_COMPLETE`, `IMPORT_COMPLETE`, or `IMPORT_ROLLBACK_COMPLETE`.

   In the case where an upload of a new template fails and the stack move to the status of `ROLLBACK_COMPLETE`, the new template will be backed up but backups of the nested resources will be based on the rolled-back resources.

8. *"How do application stack lifecycles differ from other recovery point lifecycles?"*

   Nested recovery point lifecycles are determined by the backup plan to which they belong. The composite recovery point is determined by the longest lifecycle of all nested recovery points. When the last remaining nested recovery point within a composite recovery point is deleted or disassociated, the composite recovery point will also be deleted.

9. *“How are tags of a CloudFormation copied to recovery points?”*

   Yes. Those tags will be copied to each respective nested recovery point.

10. *“Is there an order for deleting composite and nested recovery points (backups)?”*

Yes. Some backups must be deleted before others can be deleted. Composite backups which contain nested recovery points cannot be deleted until all recovery points within the composite have been deleted. Once a composite recovery point is no longer contains nested recovery points, you can delete it manually. Otherwise, it will be deleted in accordance with its backup plan lifecycle.

Restore applications within a stack

See How to restore application stack backups for information on restoring nested recovery points.

Amazon Aurora DSQL backups

You can use AWS Backup to create backups of your Amazon Aurora DSQL single-Region and multi-Region clusters. Amazon Aurora DSQL cluster backups are always full backups.

Backup creation for Amazon Aurora DSQL clusters use the standard creation backup process. For more information, see the following:

- Creating an on-demand backup using AWS Backup
- Create a backup plan

To use AWS Backup to create backups of your Amazon Aurora DSQL clusters, you must enable protection for Aurora DSQL. For more information, see Service Opt-in.

When you backup a multi-Region cluster, consider the following items:

- A multi-Region cluster backup requires a separate backup for each Region within the cluster; a backup in one Region doesn't create a recovery point for all Regions in a multi-Region cluster.
- As a best practice, AWS Backup recommends you create a recovery point in one Region and copy it to another related Region. For multi-Region restore, you need a recovery point in one supported Region, and a copy of that recovery point in another Region within the same continent group.

  The following supported continent groups are available:

  - **Americas**: US East (N. Virginia), US East (Ohio), US West (Oregon), Canada West (Calgary), Canada (Central)

- **Europe**: Europe (Frankfurt), Europe (Stockholm), Europe (Spain), Europe (Ireland), Europe (London), Europe (Paris)
- **Asia-Pacific**: Asia Pacific (Hong Kong), Asia Pacific (Tokyo), Asia Pacific (Seoul), Asia Pacific (Osaka), Asia Pacific (Mumbai), Asia Pacific (Singapore)

AWS Backup recommends that you add the backup copy rule to the backup plan. If you do not add the copy rule to the backup plan, you must manually copy the backup to the required Region in which to perform the restore, which will increase your Recovery Time Objective (RTO) times.

For information about restoring an Aurora DSQL recovery point (backup), see Amazon Aurora DSQL restore.

Advanced DynamoDB backup

AWS Backup supports additional, advanced features for your Amazon DynamoDB data protection needs.

Customers who started using AWS Backup after November 2021 have advanced DynamoDB backup features enabled by default. Specifically, advanced DynamoDB backup features are enabled by default to customers who have not created a backup vault prior to November 21, 2021.

It's best practice for existing AWS Backup customers to enable advanced features for DynamoDB. There is no difference in warm backup storage pricing after you enable advanced features. You can potentially save money by moving backups to cold storage and optimize your costs by using cost allocation tags. You can also start taking advantage of AWS Backup's cross-Region and cross-account copy and security features.

Topics

- Benefits of advanced DDB backup
- Considerations for Advanced DynamoDB backup
- Understanding backup overlap and costs
- Enabling advanced DynamoDB backup using the console
- Enabling advanced DynamoDB backup programmatically
- Editing an advanced DynamoDB backup
- Restoring an advanced DynamoDB backup

- Deleting an advanced DynamoDB backup
- Other benefits of full AWS Backup management when you enable advanced DynamoDB backup

Benefits of advanced DDB backup

After you enable AWS Backup's advanced features in your AWS Region, you unlock the following features for all new for DynamoDB table backups you create:

- Cost savings and optimization:
  - Tiering backups to cold storage to reduce storage costs
  - Cost allocation tagging for use with Cost Explorer
- Additional copy options:
  - Cross-Region copy
  - Cross-account copy
- Security:
  - Backups inherit tags from their source DynamoDB tables, allowing you to use those tags to set permissions and service control policies (SCPs).

Considerations for Advanced DynamoDB backup

Opting in

Backups, including those of Advanced DDB resources, can be created by a backup plan, an on-demand backup, or through a backup policy. Backups created by a plan or on-demand will automatically opt-in your account to allow backups of Advanced DDB resources.

If your backup job is created by a backup policy, you need to manually opt-in to Advanced DynamoDB backups, either through the Backup console or through CLI.

Custom policies and roles

If you use a custom role or policy instead of AWS Backup's default service role, you must add or use the following permissions policies (or add their equivalent permissions) to your custom role:

- `AWSBackupServiceRolePolicyForBackup` to perform advanced DynamoDB backup.
- `AWSBackupServiceRolePolicyForRestores` to restore advanced DynamoDB backups.

To learn more about AWS-managed policies and view examples of customer-managed policies, see Managed policies for AWS Backup.

Understanding backup overlap and costs

AWS Backup advanced DynamoDB features provide periodic snapshot-based backups with additional capabilities. These do not replace or duplicate native DynamoDB point-in-time recovery (PITR), which provides continuous backup with second-level granularity. If your recovery strategy requires both periodic snapshots (for long-term retention, cross-Region DR, or compliance) and continuous recovery (for operational recovery within a short window), you may choose to enable both AWS Backup advanced features and native DynamoDB PITR. In this case, you will incur charges for both — this is intentional complementary protection, not unintended duplication. If your recovery requirements are met by periodic snapshots alone, native DynamoDB PITR can be disabled to reduce costs.

Enabling advanced DynamoDB backup using the console

You can enable AWS Backup advanced features for DynamoDB backups using either the AWS Backup or DynamoDB console.

To enable advanced DynamoDB backup features from the AWS Backup console:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the left navigation menu, choose **Settings**.
3. Under the **Supported services** section, verify that **DynamoDB** is **Enabled**.

   If it is not, choose **Opt-in** and enable DynamoDB as an AWS Backup supported service.
4. Under the **Advanced features for DynamoDB backups** section, choose **Enable**.
5. Choose **Enable features**.

For how to enable AWS Backup advanced features using the DynamoDB console, see Enabling AWS Backup features in the *Amazon DynamoDB User Guide*.

Enabling advanced DynamoDB backup programmatically

You can also enable AWS Backup advanced features for DynamoDB backups using the AWS Command Line Interface (CLI). You enable advanced DynamoDB backups when you set both of the following values to `true`:

To programmatically enable AWS Backup advanced features for DynamoDB backups:

1. Check if you already enabled AWS Backup advanced features for DynamoDB using the following command:

```
$ aws backup describe-region-settings
```

If `"DynamoDB":true` under both `"ResourceTypeManagementPreference"` and `"ResourceTypeOptInPreference"`, you have already enabled advanced DynamoDB backup.

If, like the following output, you have at least one instance of `"DynamoDB":false`, you have not yet enabled advanced DynamoDB backup, proceed to the next step.

```
{
  "ResourceTypeManagementPreference":{
    "DynamoDB":false,
    "EFS":true
  }
  "ResourceTypeOptInPreference":{
    "Aurora":true,
    "DocumentDB":false,
    "DynamoDB":false,
    "EBS":true,
    "EC2":true,
    "EFS":true,
    "FSx":true,
    "Neptune":false,
    "RDS":true,
    "Storage Gateway":true
  }
}
```

2. Use the following UpdateRegionSettings operation to set both `"ResourceTypeManagementPreference"` and `"ResourceTypeOptInPreference"` to `"DynamoDB":true`:

```
aws backup update-region-settings \
              --resource-type-opt-in-preference DynamoDB=true \
              --resource-type-management-preference DynamoDB=true
```

Editing an advanced DynamoDB backup

When you create a DynamoDB backup after you enable AWS Backup advanced features, you can use AWS Backup to:

- Copy a backup across Regions
- Copy a backup across accounts
- Change when AWS Backup tiers a backup to cold storage
- Tag the backup

To use those advanced features on an existing backup, see Editing a backup.

If you later disable AWS Backup advanced features for DynamoDB, you can continue to perform those operations to DynamoDB backups that you created during the period of time when you enabled advanced features.

Restoring an advanced DynamoDB backup

You can restore DynamoDB backups taken with AWS Backup advanced features enabled in the same way you restore DynamoDB backups taken prior to enabling AWS Backup advanced features. You can perform a restore using either AWS Backup or DynamoDB.

You can specify how to encrypt your newly-restored table with the following options:

- When you restore in the same Region as your original table, you can optionally specify an encryption key for your restored table. If you do not specify an encryption key, AWS Backup will automatically encrypt your restored table using the same key that encrypted your original table.
- When you restore in a different Region than your original table, you must specify an encryption key.

To restore using AWS Backup, see Restore a Amazon DynamoDB table.

To restore using DynamoDB, see Restoring a DynamoDB table from a backup in the *Amazon DynamoDB User Guide*.

Deleting an advanced DynamoDB backup

You cannot delete backups created using these advanced features in DynamoDB. You must use AWS Backup to delete backups to maintain global consistency throughout your AWS environment.

To delete a DynamoDB backup, see Backup deletion.

Other benefits of full AWS Backup management when you enable advanced DynamoDB backup

When you enable AWS Backup advanced features for DynamoDB, you give full management of your DynamoDB backups to AWS Backup. Doing so gives you the following, additional benefits:

Encryption

AWS Backup automatically encrypts the backups with the KMS key of your destination AWS Backup vault. Previously, they were encrypted using the same encryption method of your source DynamoDB table. This increases the number of defenses you can use to safeguard your data. See Encryption for backups in AWS Backup for more information.

Amazon Resource Name (ARN)

Each backup ARN's service namespace is `awsbackup`. Previously, the service namespace was `dynamodb`. Put another way, the beginning of each ARN will change from `arn:aws:dynamodb` to `arn:aws:backup`. See ARNs for AWS Backup in the *Service Authorization Reference*.

With this change, you or your backup administrator can create access policies for backups using the `awsbackup` service namespace that now apply to DynamoDB backups created after you enable advanced features. By using the `awsbackup` service namespace, you can also apply policies to other backups taken by AWS Backup. See Access control for more information.

Location of charges on billing statement

Charges for backups (including storage, data transfers, restores, and early deletion) appear under "Backup" in your AWS bill. Previously, charges appeared under "DynamoDB" in your bill.

This change ensures that you can use AWS Backup billing to centrally monitor your backup costs. See Metering, costs, and billing for AWS Backup for more information.

Amazon EBS and AWS Backup

The backup process for Amazon EBS resources is similar to the steps used to back up other resources types:

- Create an on-demand backup
- Create a scheduled backup

Resource-specific information is noted in the following sections.

Amazon EBS Archive Tier for cold storage

EBS is one of the resource that supports a transition of backups to cold storage. For more information, see Lifecycle and storage tiers.

Amazon EBS multi-volume, crash-consistent backups

By default, AWS Backup creates crash-consistent backups of Amazon EBS volumes that are attached to an Amazon EC2 instance. Crash consistency means that the snapshots for every Amazon EBS volume attached to the same Amazon EC2 instance are taken at the exact same moment. You no longer have to stop your instances or coordinate between multiple Amazon EBS volumes to ensure crash-consistency of your application state.

Since multi-volume, crash-consistent snapshots are a default AWS Backup functionality, you don't need to do anything different to use this feature.

The role used to create an EBS snapshot recovery point is associated with that snapshot. This same role must be used to delete recovery points created by it or to transition recovery points of it to an archive tier.

Amazon EBS Snapshot Lock and AWS Backup

AWS Backup managed Amazon EBS snapshots and snapshots associated with a AWS Backup managed Amazon EC2 AMI which have Amazon EBS Snapshot Lock applied may not be deleted as part of the recovery point lifecycle if the snapshot lock duration exceeds the backup lifecycle. Instead, these recovery points will have the status of `EXPIRED`. These recovery points can be deleted manually if you choose to first remove the Amazon EBS snapshot lock.

Restoring Amazon EBS resources

To restore your Amazon EBS volumes, follow the steps in Restoring an Amazon EBS volume.

Amazon Relational Database Service backups

Amazon RDS and AWS Backup

When you consider the options to back up your Amazon RDS instances and clusters, it's important to clarify which kind of backup you want to create and use. Several AWS resources, including Amazon RDS, offer their own native backup solutions.

Amazon RDS gives the option of making automated backups and manual backups. Recovery points created by AWS Backup are classified differently depending on the backup type:

- **Periodic snapshots** created by AWS Backup are considered manual backups in Amazon RDS. These are snapshot-based backups taken according to your backup plan schedule.
- **Continuous backups** created by AWS Backup are considered automated backups in Amazon RDS. These enable point-in-time restore (PITR) by maintaining transaction logs alongside automated snapshots.

This distinction is important because manual and automated backups have different retention behaviors and lifecycle management in Amazon RDS.

When you use AWS Backup to create a backup (recovery point) of an Amazon RDS instance, AWS Backup creates a manual DB snapshot of that instance by using the `create-db-snapshot` operation in Amazon RDS. All snapshots of a DB instance, whether they come from Amazon RDS automated backups, manual snapshots, or AWS Backup, belong to the same incremental snapshot chain for that instance.

Only the first snapshot ever taken of a DB instance stores the full contents of the database. Every later snapshot of the same instance is incremental, which means it stores only the data that changed since the most recent snapshot. As a result, if the instance already has automated backups or earlier snapshots, the first recovery point that AWS Backup creates is also incremental, not a new full snapshot. In this context, "full" means a snapshot that you can restore the entire database from. It does not mean a full storage-level copy of the data.

When using cross account or cross Region copies, incremental snapshot copy jobs process faster than full snapshot copy jobs. Keeping a previous snapshot copy until the new copy job is complete may reduce the copy job duration. If you choose to copy snapshots from RDS database instances, it is important to note that deleting previous copies first will cause full snapshot copies to be made (instead of incremental). For more information on optimizing copying, see Incremental snapshot copying in the *Amazon RDS User Guide*.

> **Important**
>
> When a AWS Backup backup plan is scheduled to create multiple daily snapshots of an Amazon RDS instance, and when one of those scheduled AWS Backup Start Backup window coincides with the Amazon RDS Backup window, the data lineage of the backups can branch off into non-identical backups, creating unplanned and conflicting backups. To

prevent this, ensure your AWS Backup backup plan or Amazon RDS window do not coincide in their times.

Considerations

AWS Backup supports creating on-demand backups of RDS Custom for SQL Server instances. However, restoring RDS Custom for SQL Server through AWS Backup is not natively supported. To restore, use the `restore-db-instance-from-db-snapshot` operation in Amazon RDS with the AWS Backup-created snapshot. For more information, see Restore an Amazon RDS Custom for SQL Server instance using a backup from AWS Backup.

RDS Custom for Oracle is not currently supported by AWS Backup.

AWS Backup does not support backup and restore of RDS on Outposts or in Local Zones, including Dedicated Local Zones. AWS Backup requires RDS instances to have `BackupTarget` set to `region` (the default).

Understanding backup overlap and costs

AWS Backup periodic snapshots are classified as manual backups in Amazon RDS. While they share the same incremental snapshot chain as automated backups, they count toward your total backup storage alongside automated backups. Amazon RDS provides a free backup storage allocation equal to your provisioned DB instance storage — this covers both automated backups and manual snapshots combined. Storage beyond that allocation is billed. If you run both scheduled AWS Backup snapshots and Amazon RDS automated backups, both contribute to this total, and you should factor this into your cost planning.

When you use AWS Backup to back up an Amazon RDS DB instance, an automated snapshot might appear with an unexpected timestamp; it matches the AWS Backup snapshot creation time rather than the Amazon RDS automated backup window. If Amazon RDS does not find a recent automated backup within the past day, it consolidates the backup. Amazon RDS attaches automated backup metadata to the manual snapshot that AWS Backup created. This ensures the DB instance maintains a regular automated backup. This most commonly occurs when an AWS Backup operation runs long enough to overlap with the Amazon RDS automated backup window.

When a backup is consolidated, both the manual and automated snapshot entries reference the same underlying snapshot data. You are charged only once for that single underlying snapshot and

your data remains fully protected. Also, backup retention and point-in-time restore functionality are unaffected regardless of whether entries are consolidated.

Amazon RDS continuous backups and point in time restore

Continuous backups involve using AWS Backup to create a full backup of your Amazon RDS resource, then capturing all changes through a transaction log. You can achieve a greater granularity by rewinding to the point in time you desire to restore to instead of choosing a previous snapshot taken at fixed time intervals.

See continuous backups and PITR supported services and managing continuous backup settings for more information.

> **Important**
>
> Enabling continuous backups for Amazon RDS using AWS Backup when they were previously disabled (or disabling continuous backups when they were previously enabled) takes the Amazon RDS instance offline to make the changes. Plan this change during a maintenance window to minimize impact. If automated backups were enabled from Amazon RDS and that backup was simply moved to AWS Backup, then no downtime is required.

Amazon RDS Multi-Availability Zone backups

AWS Backup backs up and supports Amazon RDS for MySQL and for PostgreSQL Multi-AZ (Availability Zone) deployment options with one primary and two readable standby database instances.

For a list of Regions where Multi-Availability Zone backups are available, see the Amazon RDS Multi-AZ column in Supported services by AWS Region.

The Multi-AZ deployment option optimizes write transactions and is ideal when your workloads require additional read capacity, lower write transaction latency, more resilience from network jitter (which impacts the consistency of write transaction latency), and high availability and durability.

To create a Multi-AZ cluster, you can choose either MySQL or PostgreSQL as the engine type.

In the AWS Backup console, there are three deployment options:

- **Multi-AZ DB cluster:** Creates a DB cluster with a primary DB instances and two readable standby DB instances, which each DB instance in a different Availability Zone. Provides high availability, data redundancy, and increases capacity to server-ready workloads.
- **Multi-AZ DB instance:** Creates a primary DB instance and a standby DB instance in a different Availability Zone. This provides high availability and data redundancy, but the standby DB instance doesn't support connections for read workloads.
- **Single DB instance:** Creates a single DB instance with no standby DB instances.

Backup behavior with instances and clusters

- Point-in-Time Recovery (PITR) can support instances, but not clusters.
- Copying a Multi-AZ DB cluster snapshot is not supported.
- The Amazon Resource Name (ARN) for an RDS recovery point depends on whether an instance or cluster is used:

  An RDS instance ARN: `arn:aws:rds:`*`region`*`:` *`account`*`:db:`*`name`*

  An RDS Multi-Availability Cluster: `arn:aws:rds:`*`region`*`:`*`account`*`:cluster:`*`name`*

For more information, consult Multi-AZ DB cluster deployments in the *Amazon RDS User Guide*.

For more information on Creating a Multi-AZ DB cluster snapshot, see the Amazon RDS User Guide.

Amazon Aurora Global Databases

AWS recommends maintaining backups in every Region where your global database is deployed.

Common questions for Amazon RDS and Aurora continuous backups

How are Amazon RDS and Aurora continuous backups linked to automated backups?

AWS Backup point-in-time recovery (continuous backups) centralizes the management of automated backups in Amazon RDS and Aurora. When you enable continuous backup in AWS Backup, AWS Backup takes over control of automated backups from Amazon RDS or Aurora. You can no longer change the backup start time (preferred backup window) or retention using the Amazon RDS or Aurora console or APIs. You can still manage the maintenance window from the Amazon RDS or Aurora console or API. AWS Backup intelligently schedules automated backups to

avoid overlapping with other scheduled backups (snapshots) and the maintenance window where possible.

How do continuous backups appear in the AWS Backup console?

In the AWS Backup console, navigate to the vault specified in your backup plan and look for *Continuous* in the **Backup type** column. The backup shows as *Available* after the first continuous backup completes, indicating that PITR is available for that resource.

How can control of automated backups be returned to Amazon RDS or Aurora?

To return control of automated backups to Amazon RDS or Aurora, first delete the backup plan or rule with continuous backups enabled for Amazon RDS or Aurora, or modify the rule to remove continuous backup. Then choose one of the following options:

- **Preserve backup data:** Run the `DisassociateRecoveryPoint` API. This disassociates the continuous backup and releases control to Amazon RDS, preserving the backup data in Amazon RDS that you can use to perform a restore.
- **Remove backup data:** Run the `DeleteRecoveryPoint` API. This deletes the continuous recovery point, removes backup data, and gives back control to Amazon RDS.

> **Important**
>
> When you run `DisassociateRecoveryPoint` or `DeleteRecoveryPoint`, AWS Backup calls the `ModifyDBInstance` API for Amazon RDS instances or the `ModifyDBCluster` API for Aurora clusters, and applies the change immediately. If there are any pending configuration changes on Amazon RDS or Aurora, those changes are also applied immediately. There may be a brief downtime during this operation.

For more information, see How can I stop an Amazon RDS continuous backup in AWS Backup?

Amazon Redshift backups

Amazon Redshift is a fully managed, scalable cloud data warehouse that accelerates your time to insights with fast, easy, and secure analytics. You can use AWS Backup to protect your data warehouses with immutable backups, separate access policies, and centralized organizational governance of backup and restore jobs.

An Amazon Redshift data warehouse is a collection of computing resources called nodes, which are organized into a group called a cluster. AWS Backup can backup these clusters.

For information on Amazon Redshift , see the Amazon Redshift Getting Started Guide, the Amazon Redshift Database Developer Guide, and the Amazon Redshift Cluster Management Guide.

Back up Amazon Redshift provisioned clusters

You can protect your Amazon Redshift clusters using the AWS Backup console or programmatically using API or CLI. These clusters can be backed up on a regular schedule as part of a backup plan, or they can be backed up as needed via on-demand backup.

You can restore a single table (also known as item-level restore) or an entire cluster. Note that tables cannot be backed up by themselves; tables are backed up as part of a cluster when the cluster is backed up.

Using AWS Backup allows you to view your resources in a centralized way; however, if Amazon Redshift is the only resource you use, you can continue to use the automated snapshot scheduler in Amazon Redshift. Note that you cannot continue to manage manual snapshot settings using Amazon Redshift if you choose to manage these via AWS Backup.

You can backup Amazon Redshift clusters either through the AWS Backup console or using the AWS CLI.

There are two ways to use the AWS Backup console to backup a Amazon Redshift cluster: on demand or as part of a backup plan.

Create on-demand Amazon Redshift backups

See Creating an on-demand backup type page for more information.

To create a manual snapshot, leave the continuous backup checkbox unchecked when you create a backup plan that includes Amazon Redshift resources.

Create scheduled Amazon Redshift backups in a backup plan

Your scheduled backups can include Amazon Redshift clusters if they are a protected resource. To opt into protecting Amazon Redshift clusters:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.

2. Using the navigation pane, choose **Protected resources**.
3. Toggle Amazon Redshift to **On**.
4. See Assigning resources to the console to include Amazon Redshift clusters in an existing or new plan.

Under **Manage Backup plans**, you can choose to create a backup plan and include Amazon Redshift clusters, or you can update an existing one to include Amazon Redshift clusters. When adding the resource type *Amazon Redshift*, you can choose to add **All Amazon Redshift clusters**, or check the boxes next to the clusters you wish to include in your backup plan.

Back up programmatically

You can also define your backup plan in a JSON document and provide it using the AWS Backup console or AWS CLI. See Creating backup plans using a JSON document and the AWS Backup CLI for information on how to create a backup plan programatically.

You can do the following operations using API:

- Start a backup job
- Describe a backup job
- Get recovery point metadata

> **Note**
>
> `BackupSizeInBytes` metadata is supported for the following resource types: Amazon EBS volumes, Amazon EFS file systems, Amazon RDS databases, DynamoDB tables, Amazon EC2 instances, Amazon FSx file systems, and Amazon S3 buckets. This field provides the size of the backup in bytes and is available through the `DescribeRecoveryPoint` API and AWS Backup console. For unsupported resource types, this field will not be populated.

- List recovery points by resources
- List tags for the recovery point

View Amazon Redshift cluster backups

To view and modify your Amazon Redshift table backups within the console:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Choose **Backup vaults**. Then, click on the backup vault name that contains your Amazon Redshift clusters.
3. The backup vault will display a summary and a list of backups. You can click on the link in the column **Recovery point ID**.
4. To delete one or more recovery points, check the box(es) you wish to delete. Under the button **Actions**, you can select **Delete**.

Restore a Amazon Redshift cluster

See how to Restore a Amazon Redshift cluster for more information.

Amazon Redshift Serverless backups

Overview

AWS Backup offers full backup management of your Amazon Redshift Serverless namespaces. Through AWS Backup, you can schedule and restore Redshift Serverless manual snapshots through the console or through CLI.

Redshift Serverless data protection through AWS Backup provides several options for backing up and restoring your data warehouses. You can create a scheduled or on-demand snapshot of your namespace. Then, you can choose to restore all the databases in that snapshot to a Amazon Redshift provisioned cluster or a Serverless namespace. Alternatively, you can restore a single table.

Redshift Serverless offers both automated and manual snapshots. Currently, AWS Backup can be used to manage manual snapshots but not automated ones.

Backup options for Redshift Serverless

You can use the AWS Backup console or CLI to create backups on demand or as part of a backup plan.

Create on-demand backup

You can create on-demand backups of Redshift Serverless namespaces through the following steps:

Console

1. Open the AWS Backup console.
2. On the dashboard, choose **Create an on-demand backup**.
3. Choose **Redshift Serverless** in the resource type dropdown menu.
4. Select the namespace you plan to back up.
5. Ensure **Create backup now** is selected.
6. Specify the retention period for the backup.
7. Choose an existing backup vault or create a new one.
8. Select the IAM role to use for the backup.
9. Optionally, add tags to the backup. To assign a tag to your on-demand backup, expand **Tags added to recovery points**, choose **Add new tag**, and enter a tag key and tag value.
10. Select **Create on-demand backup** to begin the backup job.
11. Once the job is initiated, the console will show the Jobs screen where you can see a list of your backup jobs and their statuses.

AWS CLI

Use the **start-backup-job** command.

Required parameters

- `BackupVaultName`
- `IamRoleArn`
- `ResourceArn`

Optional parameters

- `CompleteWindowMinutes`
- `IdempotencyToken`
- `Lifecyle`
- `StartWindowMinutes`

Example

The following example creates an on-demand backup of a Redshift Serverless namespace.

```
aws backup start-backup-job \
    --backup-vault-name sample-vault \
    --iam-role-arn arn:aws:iam::account:role/service-role/
AWSBackupDefaultServiceRole \
    --resource-arn arn:aws:redshift-serverless:region:account:namespace/namespace-
name-UUID
```

Create scheduled Redshift Serverless backups in a backup plan

You can create a new backup plan for their Redshift Serverless namespaces through the AWS Backup console or through CLI, or you can add Redshift Serverless to an existing backup plan.

Your scheduled backups can include Redshift Serverless namespaces if they are a protected resource. To opt into protecting Redshift Serverless in the AWS Backup console, complete the following steps:

Console

To opt into protecting Redshift Serverless in the AWS Backup console, complete the following steps:

1. Open the AWS Backup console.
2. Using the navigation pane, choose **Protected resources**.
3. Toggle **Amazon Redshift Serverless** to **On**.
4. See Select AWS services to backup to include Redshift Serverless namespaces in an existing or new plan. When you add the resource type *Redshift Serverless*, you can choose to add **All Amazon Redshift namespaces**, or check the boxes next to the namespaces you wish to back up.

Under **Manage Backup plans**, you can:

- Create a backup plan and include Redshift Serverless;
- Update an existing backup plan to include Redshift Serverless.

AWS CLI

See Create backup plans using the AWS CLI for guidance to use **create-backup-plan**.

If you want to alter an existing plan to include your Serverless resources, use the command **update-backup-plan**.

The ARN (Amazon Resource Name) for Serverless resources to include in "BackupSelection": { "Resources" has the following format:

```
arn:aws:redshift-serverless:Region:account:snapshot/a12bc34d-567e-890f-123g-
h4ijk56l78m9
```

See Amazon Redshift Serverless restore for information to restore data from a snapshot to a Serverless namespace.

Amazon EKS backups

An Amazon Elastic Kubernetes Service (Amazon EKS) cluster consists of multiple resources that you can back up as a single unit. When you back up an Amazon EKS cluster, AWS Backup creates a composite recovery point that includes both EKS cluster state and persistent volume backups.

When an Amazon EKS cluster is backed up, recovery points are created for the Amazon EKS cluster state and persistent volumes supported by AWS Backup. These recovery points are grouped together within an overarching recovery point called a **composite**.

There are two distinct components of an Amazon EKS backup:

- *Amazon EKS Cluster State:* This is a backup of the Amazon EKS cluster state. See Amazon EKS backup terminology below for what is included.
- *Persistent Storage:* This is a backup of persistent storage (Amazon EBS, Amazon S3, Amazon Elastic File System) attached to the Amazon EKS cluster via Persistent Volume Claims and supported by EKS Add Ons CSI Driver.

Amazon EKS backup terminology

The following terms are used throughout the Amazon EKS backup documentation. For Amazon EKS Specific terminology, please refer to Amazon EKS Documentation.

EKS Backup Terminology

- **Composite recovery point** – A recovery point used to group nested recovery points together for an Amazon EKS cluster backup.
- **Nested recovery point** – A recovery point of a resource that is part of an Amazon EKS cluster and is backed up as part of the composite recovery point.
- **EKS Cluster State** – The Kubernetes manifests (YAML or JSON files) that define the desired state of Kubernetes resources in your cluster. This includes Kubernetes resources and deployments such as: secrets, config maps, stateful sets, DaemonSets, storage classes, storage maps, replica sets, persistent volume claims, custom resource definitions, roles, and role bindings.
- **Amazon EKS Cluster Configuration Child Recovery Point** – Contains Amazon EKS cluster state.
- **Persistent Volume Child Recovery Points** – Contains persistent volume backups for supported storage types (EBS, S3, EFS) supported by EKS Add Ons CSI Driver.

Amazon EKS backup structure

Amazon EKS backups include the following components:

- Amazon EKS Cluster State
- Persistent Storage: Backups of supported storage types including Amazon EBS, Amazon EFS, and Amazon S3

Amazon EKS Backups will not include the following components:

- Container images from external repositories (ECR, Docker)
- EKS cluster infrastructure components (e.g. VPCs, Subnets)
- Auto-generated EKS resources like nodes, auto-generated pods, events, leases, and jobs.

EKS backup setup and prerequisites ("Before you backup")

- **EKS Cluster Settings:**
  - EKS Cluster authorization mode set to API or API_AND_CONFIG_MAP for AWS Backup to create Access Entries to access the EKS cluster.
- **Permissions:**

- AWS Backup's managed policy AWSBackupServiceRolePolicyForBackup contains the required permissions to backup your Amazon EKS cluster and EBS and EFS persistent storage
- If your EKS Cluster contains an S3 bucket you will need to ensure the following policies and prerequisites for your S3 bucket are added and enabled as documented:
  - AWSBackupServiceRolePolicyForS3Backup
  - Prerequisites for S3 Backups
- **Encryption:**
  - Amazon EKS child recovery points will be encrypted with the Amazon KMS key set of the target Backup Vault
  - Persistent Storage recovery points will be encrypted as per the current support for each storage class: EBS Snapshots, S3 Backups, EFS Backups. See Encryption for backups in AWS Backup

Create an Amazon EKS backup

The process of a backup creation is called a backup job. An Amazon EKS cluster backup job has a status. When a backup job has finished, it has the status of Completed. This signifies a recovery point (a backup) has been created.

Creating an on-demand Amazon EKS backup

Console

To create an on-demand backup of your Amazon EKS cluster:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Protected resources**.
3. Under **Resource type**, select **Amazon EKS**.
4. Select the checkbox next to the Amazon EKS cluster you want to back up.
5. Choose **Create on-demand backup**.
6. Configure your backup settings, including backup window, transition to cold storage, and retention period.
7. Choose **Create on-demand backup**.

AWS CLI

To create an on-demand backup of your Amazon EKS cluster using the AWS CLI:

Use the **start-backup-job** command:

```
aws backup start-backup-job \
    --backup-vault-name my-backup-vault \
    --resource-arn arn:aws:eks:us-west-2:123456789012:cluster/my-cluster \
    --iam-role-arn arn:aws:iam::123456789012:role/AWSBackupDefaultServiceRole \
    --region us-west-2
```

Optionally, specify additional parameters such as lifecycle settings:

```
aws backup start-backup-job \
    --backup-vault-name my-backup-vault \
    --resource-arn arn:aws:eks:us-west-2:123456789012:cluster/my-cluster \
    --iam-role-arn arn:aws:iam::123456789012:role/AWSBackupDefaultServiceRole \
    --lifecycle MoveToColdStorageAfterDays=30,DeleteAfterDays=365 \
    --region us-west-2
```

Monitor the backup job status:

```
aws backup describe-backup-job \
    --backup-job-id backup-job-id \
    --region us-west-2
```

Amazon EKS backup ARN format

Composite Recovery Point arn:*partition*:backup:*region*:*accountId*:recovery-point:composite:eks/*cluster-name*-*timestamp*

Child Recovery Point arn:*partition*:backup:*region*:*accountId*:recovery-point:eks/*cluster-name*-*timestamp*

Amazon EKS recovery points

Recovery point status

When the backup job of an Amazon EKS cluster is finished (the job status is `Completed`), a backup of the cluster has been created. This backup is also known as a composite recovery point.

A composite recovery point can have one of the following statuses: `Completed`, `Failed`, or `Partial`.

Each Amazon EKS backup creates a parent backup job for the composite recovery point and child backup jobs for each child recovery point (cluster configuration and persistent volumes).

- A completed backup job means your entire Amazon EKS cluster and the resources within it are protected by AWS Backup.
- A failed status indicates that the backup job was unsuccessful; you should create the backup again once the issue that caused the failure is corrected.
- A `Partial` status means that not all the resources in the cluster were backed up. This may happen if one or more of the backup jobs belonging to resources within the cluster (nested resources) have statuses other than `Completed`. You can manually create an on-demand backup to rerun any resources that resulted in a status other than `Completed`.
- A `Completed with issues` status means that not all the resources in the cluster were backed up. This can happen when we fail to backup some Kubernetes objects in the cluster. You can subscribe to **Notification Events** for failed objects for backup. For more information, see Notification options with AWS Backup.

Each nested resource within the composite recovery point has its own individual recovery point, each with its own status (either `Completed` or `Failed`). Nested recovery points with a status of `Completed` can be restored.

AWS Backup supports lifecycle transitions to cold storage for persistent volume recovery points. You can subscribe to notifications to receive alerts on backup job status.

Manage recovery points

Composite recovery points (backups) can be copied; persistent volume child recovery points can be copied, deleted, disassociated, or restored. The Amazon EKS cluster state child recovery point cannot be copied, deleted, or disassociated as it maintains a 1:1 relationship with its parent composite recovery point.

A composite recovery point which contains nested backups cannot be deleted. After the nested recovery points within a composite recovery point have been deleted or disassociated, you can manually delete the composite recovery point manually or let it remain until the backup plan lifecycle deletes it.

Delete a recovery point

You can delete a recovery point using the console or using the AWS CLI.

To delete recovery points using the console:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Click on Protected Resources in the left-hand navigation. In the text box, type EKS to display only your Amazon EKS clusters.
3. Composite recovery points will be displayed in the Recovery points pane. The plus sign (+) to the left of each recovery point ID can be clicked to expand each composite recovery point, showing all nested recovery points contained in the composite. You can check the box to the left of any recovery point to include it in your selection of recovery points you wish to delete.
4. Click the Delete button.

When you use the console to delete one or more composite recovery points, a warning box will pop up. This warning box requires you to confirm your intention to delete the composite recovery points, including nested recovery points within composite stacks.

To delete recovery points using API, use the DeleteRecoveryPoint command.

When you use API with the AWS Command Line Interface you must delete all nested recovery points prior to deleting a composite point.

Disassociate a nested recovery point from composite recovery point

You can disassociate a nested recovery point from a composite recovery point (for example, you wish to keep the nested recovery point but delete the composite recovery point). Both recovery points will remain, but they will no longer be connected; that is, actions that occur on the composite recovery point will no longer apply to the nested recovery point once it has been disassociated. The Amazon EKS cluster state child recovery point cannot be disassociated as it maintains a 1:1 relationship with its parent composite recovery point.

You can disassociate the recovery point using the console, or you can call the API DisassociateRecoveryPointFromParent.

Copy a recovery point

You can copy a composite recovery point, or you can copy a nested recovery point if the resource supports cross-account and cross-Region copy.

To copy recovery points using the AWS Backup console:

Open the AWS Backup console at https://console.aws.amazon.com/backup.

1. Click on **Vaults** in the left-hand navigation, and go to the vault that contains the recovery point you want to copy. In the text box, type EKS to display only your recovery points for Amazon EKS clusters.
2. Both composite and nested recovery points will be displayed under the Recovery point ID pane. Note you cannot select and copy a nested EKS recovery point.
3. The arrow sign to the left of each composite recovery point ID can be clicked to expand, showing all nested recovery points contained in the composite. You can click the square checkbox to the left of any recovery point to copy it.
4. Once it is selected, click the **Actions** dropdown in the top-right corner of the pane and click **Copy**.

Amazon EKS backups support all copy types:

- Same region/account
- Cross account
- Cross region
- Opt-in regions

Limitations

- Persistent volumes using a CSI Driver via CSI migration, in-tree storage plugins or ACK controllers are not supported. Note that the annotation `volume.kubernetes.io/storage-provisioner: ebs.csi.aws.com` is metadata indicating which provisioner could manage the volume, not that the volume uses CSI. The actual provisioner is determined by the storageClass.
- Amazon S3 buckets with specific prefixes attached to CSI Driver MountPoints cannot be backed up. Only Amazon S3 buckets as targets are supported, not specific prefixes.
- Amazon S3 bucket backups as part of an EKS cluster backup will only support snapshot backups.
- Backups of EFS file systems in a cross-account are not supported via EKS Backups.
- Amazon EFS persistent volumes with non-root subpath mounts are not supported. Only root file system handles or access point handles are supported for backup. Unsupported configurations will produce a failed child backup job with the corresponding error message.

- Amazon FSx via CSI driver is not supported via EKS Backups.
- AWS Backup does not support Amazon EKS on AWS Outposts.
- Subject to backup and restore quotas.

Backup Jobs Completed with Issues

When backing up an Amazon EKS cluster, some Kubernetes objects may fail to be retrieved. In this case, the backup job will complete with a `Completed with issues` status rather than failing entirely, with the following status message:

- Some Kubernetes Objects failed to be backed up. To get notified of these failures, enable SNS event notifications.

The following Kubernetes object types may be skipped during a backup job due to Amazon EKS Metrics Server Add On unavailability issues resulting in a 503 Service Unavailable error. See here for troubleshooting guidance.

- `metrics.k8s.io`
- `custom.metrics.k8s.io`
- `external.metrics.k8s.io`
- `metrics.eks.amazonaws.com`

Frequently Asked Questions

1. *"What is included as part of the Amazon EKS backup?"*

   As part of each backup of an Amazon EKS cluster, the Amazon EKS cluster state and persistent volumes supported by AWS Backup are backed up. The Amazon EKS cluster state includes details like cluster name, IAM role, Amazon VPC configuration, network settings, logging, encryption, add-ons, access entries, managed node groups, Fargate profiles, pod identity associations, and Kubernetes manifest files.

2. *"Does a `Partial` status mean the creation of my backup failed?"*

   No. A partial status indicates that some of the recovery points were backed up, while some were not. There are two conditions to check if you were expecting a `Completed` backup result:

a. One or more of the backup jobs belonging to resources within the cluster were not successful and the job has to be rerun.

b. A nested recovery point was deleted or disassociated from the composite recovery point.

3. *"Do I need to have an agent or Amazon EKS Add-on installed on my Amazon EKS cluster before backup?"*

   No. AWS Backup does not require any agents or add-ons to be installed on your Amazon EKS cluster. The only pre-requisite is to have your EKS Cluster's authorization mode set to API or API_AND_CONFIG_MAP for AWS Backup to create Access Entries to access the EKS cluster.

4. *"Does Amazon EKS Backups include Amazon EKS infrastructure components or Amazon ECR images?"*

   No. Amazon EKS backups focus on the EKS cluster state and application workloads, not the underlying infrastructure components or container images.

5. *"Can I lifecycle my EKS Composite Recovery Point to cold storage?"*

   You can transition to cold storage for underlying child recovery points that support cold storage tiers. See the AWS Backup feature availability matrix for full list of support.

6. *"Are my EKS backups incremental?"*

   AWS Backup will take incremental backups of each child recovery point where supported today, this includes EBS volumes, EFS Filesystems and S3 buckets. The EKS cluster state child recovery point will be a full backup. See the AWS Backup feature availability matrix.

7. *"Can I create an index and search my EKS backups?"*

   No, however you can create on-demand indexes and search persistent volumes where the underlying storage type supports this capability through AWS Backup. See the AWS Backup feature availability matrix.

SAP HANA backup on Amazon EC2

> **Note**
>
> Supported services by AWS Region contains the currently supported Regions where SAP HANA database backups on Amazon EC2 instances are available.

AWS Backup supports backups and restores of SAP HANA databases on Amazon EC2 instances.

Topics

- Overview of SAP HANA databases with AWS Backup
- Prerequisites for backing up SAP HANA databases through AWS Backup
- SAP HANA backup operations in the AWS Backup console
- View SAP HANA database backups
- Use AWS CLI for SAP HANA databases with AWS Backup
- Troubleshooting backups of SAP HANA databases
- Glossary of SAP HANA terms when using AWS Backup
- AWS Backup support of SAP HANA databases on EC2 instances release notes

Overview of SAP HANA databases with AWS Backup

In addition to the ability to create backups and to restore databases, AWS Backup integration with Amazon EC2 Systems Manager for SAP allows customers to identify and tag SAP HANA databases.

AWS Backup is integrated with AWS Backint Agent to perform SAP HANA backups and restores. For more information, see AWS Backint.

When you take backups of SAP HANA, your snapshots and on-demand backups are full backups. However, you can achieve incremental backups by enabling continuous backups for point-in-time recovery (PITR).

Prerequisites for backing up SAP HANA databases through AWS Backup

Several prerequisites must be completed before backup and restore activities can be performed. Note you will need administrative access to your SAP HANA database and permissions to create new IAM roles and policies in your AWS account to perform these steps.

Complete these prerequisites at Amazon EC2 Systems Manager.

1. Set up required permissions for Amazon EC2 instance running SAP HANA database
2. Register credentials in AWS Secrets Manager
3. Install AWS Backint and AWS Systems Manager for SAP Agents

4. Verify SSM Agent
5. Verify parameters
6. Register SAP HANA database

It is best practice to register each HANA instance only once. Multiple registrations can result in multiple ARNs for the same database. Maintaining a single ARN and registration simplifies backup plan creation and maintenance and can also help reduce unplanned duplication of backups.

SAP HANA backup operations in the AWS Backup console

Once the prerequisites and SSM for SAP setups are complete, you can back up and restore your SAP HANA on EC2 databases.

Opt in to protect SAP HANA resources

To use AWS Backup to protect your SAP HANA databases, SAP HANA must be toggled on as one of the protected resources. To opt in:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the left navigation pane, choose **Settings**.
3. Under **Service opt-in**, select **Configure resources**.
4. Opt in to **SAP HANA on Amazon EC2.**.
5. Click **Confirm**.

Service opt-in for SAP HANA on Amazon EC2 will now be enabled.

Create a scheduled backup of SAP HANA databases

You can edit an existing backup plan and add SAP HANA resources to it, or you can create a new backup plan just for SAP HANA resources.

If you choose to create a new backup plan, you will have three options:

1. **Option 1: Start with a template**
    1. Choose a backup plan template.
    2. Specify a backup plan name.

3. Click **Create plan**.

2. **Option 2: Build a new plan**

   1. Specify a backup plan name.
   2. Optionally specify tags to add to backup plan.
   3. Specify the backup rule configuration.

      a. Specify a backup rule name.

      b. Select an existing vault or create a new backup vault. This is where your backups are stored.

      c. Specify a backup frequency.

      d. Specify a backup window.

      e. Specify the retention period.

      f. (*Optional*) Configure transition to cold storage and copy to destination settings.

      > **Note**
      >
      > Continuous backups (transaction logs) cannot be copied to other Regions or accounts. Snapshot backups can be copied from full backups.

      g. (*Optional*) Specify tags to add to recovery points.

   4. Click **Create plan**.

3. **Option 3: Define a plan using JSON**

   1. Specify the JSON for your backup plan by either modifying the JSON expression of an existing backup plan or creating a new expression.
   2. Specify a backup plan name.
   3. Click **Validate JSON**.

   Once the backup plan is created successfully, you can assign resources to the backup plan in the next step.

Whichever plan you use, ensure you assign resources. You can choose which SAP HANA databases to assign, including system and tenant databases. You also have the option to exclude specific resource IDs.

Create an on-demand backup of SAP HANA databases

You can create a full on-demand backup that runs immediately after creation. Note that on-demand backups of SAP HANA databases on Amazon EC2 instances are full backups; incremental backups are not supported.

Your on-demand backup is now created. It will begin backing up your specified resources. The console will transition you to the **Backup jobs** page where you can view the job progress. Take note of the backup job ID from the blue banner at the top of your screen, as you will need it to easily find the status of your backup job. When the backup is completed, the status will progress to `Completed`. Backups can take up to several hours.

Refresh the **Backup jobs list** to see the status change. You can also search for and click on your **backup job ID** to view detailed job status.

Continuous backups of SAP HANA databases

You can make continuous backups , which can be used with point-in-time restore (PITR) (note that on-demand backups preserve resources in the state in which they are taken; whereas PITR uses continuous backups which record changes over a period of time).

With continuous backups, you can restore your SAP HANA database on an EC2 instance by rewinding it back to a specific time that you choose, within 1 second of precision (going back a maximum of 35 days). Continuous backup works by first creating a full backup of your resource, and then constantly backing up your resource's transaction logs. PITR restore works by accessing your full backup and replaying the transaction log to the time that you tell AWS Backup to recover.

You can opt in to continuous backups when you create a backup plan in AWS Backup using the AWS Backup console or the API.

To enable continuous backups using the console

1. Sign in to the AWS Management Console, and open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Backup plans**, and then choose **Create Backup plan**.
3. Under **Backup rules**, choose **Add Backup rule**.
4. In the **Backup rule configuration** section, select **Enable continuous backups for supported resources**.

After you disable PITR (point-in-time restore) for SAP HANA database backups, logs will continue to be sent to AWS Backup until the recovery point expires (status equals `EXPIRED`). You can change to an alternative log backup location in SAP HANA to stop the transmission of logs to AWS Backup.

A continuous recovery point with a status of `STOPPED` indicates that a continuous recovery point has been interrupted; that is, the logs transmitted from SAP HANA to AWS Backup that show the incremental changes to a database have a gap. The recovery points that occur within this timeframe gap have a status of `STOPPED`..

For issues you may encounter during restore jobs of continuous backups (recovery points), see the SAP HANA Restore troubleshooting section of this guide.

View SAP HANA database backups

View the status of backup and restore jobs:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Jobs**.
3. Choose backup jobs, restore jobs or copy jobs to see the list of your jobs.
4. Search for and click on your job ID to view detailed job statuses.

View all recovery points in a vault:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Backup vaults**.
3. Search for and click on a backup vault to view all the recovery points within the vault.

View details of protected resources:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Protected resources**.
3. You may also filter by resource type to view all backups of that resource type.

Use AWS CLI for SAP HANA databases with AWS Backup

Each action within the Backup console has a corresponding API call.

To programmatically configure and manage AWS Backup and its resources, use the API call `StartBackupJob` to backup an SAP HANA database on an EC2 instance.

Use `start-backup-job` as the CLI command.

Troubleshooting backups of SAP HANA databases

If you encounter errors during your workflow, consult the following example errors and suggested resolutions:

Python prerequisites

- **Error: Zypper error related to Python version** since SSM for SAP and AWS Backup require Python 3.6 but SUSE 12 SP5 by default supports Python 3.4.

  Resolution: Install multiple versions of Python on SUSE12 SP5 by doing the following steps:

  1. Run an update-alternatives command to create a symlink for Python 3 in '/usr/local/bin/' instead of directly using '/usr/bin/python3'. This commands will set Python 3.4 as the default version. The command is: `# sudo update-alternatives --install /usr/local/bin/python3 python3 /usr/bin/python3.4 5`
  2. Add Python 3.6 to alternatives configuration by running the following command: `# sudo update-alternatives --install /usr/local/bin/python3 python3 /usr/bin/python3.6 2`
  3. Change the alternative configuration to Python 3.6 by running the following command: `# sudo update-alternatives --config python3`

     The following output should be displayed:

     ```
     There are 2 choices for the alternative python3 (providing /usr/local/bin/python3).
      Selection Path Priority Status
     * 0 /usr/bin/python3.4 5 auto mode
      1 /usr/bin/python3.4 5 manual mode
      2 /usr/bin/python3.6 2 manual mode
     Press enter to keep the current choice[*], or type selection number:
     ```

  4. Enter the number corresponding to Python 3.6.
  5. Check the Python version and confirm Python 3.6 is being used.
  6. (*Optional, but recommended*) Verify Zypper commands work as expected.

Amazon EC2 Systems Manager for SAP discovery and registration

- **Error: SSM for SAP failed to discover workload** due to blocked access to public endpoint for AWS Secrets Manager and SSM.

  Resolution: Test if endpoints are reachable from your SAP HANA database. If they cannot be reached, you can create Amazon VPC endpoints for AWS Secrets Manager and SSM for SAP.

  1. Test access to Secrets Manager from Amazon EC2 host for HANA DB by running the following the command: `aws secretsmanager get-secret-value --secret-id hanaeccsbx_hbx_database_awsbkp` . If the command fails to return a value, the firewall is blocking access to Secrets Manager service endpoint. The log will stop at the step “Retrieving secrets from Secrets Manager”.
  2. Test connectivity to SSM for SAP endpoint by running the command `aws ssm-sap list-registration` . If the command fails to return a value, the firewall is blocking access to the SSM for SAP endpoint.

     Example error: `Connection was closed before we received a valid response from endpoint URL: “https://ssm-sap.us-west-2.amazonaws.com/register-application"`.

  There are two options to proceed if the endpoints are not reachable.

  - Open firewall ports to allow access to public service endpoint for Secrets Manager and SSM for SAP; or,
  - Create VPC endpoints for Secrets Manager and SSM for SAP, then:
    - Ensure Amazon VPC is enabled for DNSSupport and DNSHostname.
    - Ensure your VPC endpoint has enabled Allow Private DNS Name.
    - If the SSM for SAP discovery completed successfully, the log will show the host is discovered.
- **Error: AWS Backup and Backint connection fails due to blocked access to AWS Backup service public endpoints.** `aws-backint-agent.log` can show errors similar to this: `time="2024-01-03T11:39:15-08:00" level=error msg="Storage configuration validation failed: missing backup data plane Id"` or `level=fatal msg="Error performing backup missing backup data plane Id`. Also, the AWS Backup console can show `Fatal Error: An internal error occured.`

  Resolution: Open firewall ports to allow access to public service endpoints (HTTPS). After this option is used, DNS will resolve requests to AWS services through public IP addresses.

- **Error: SSM for SAP registration fails due to HANA password containing special characters.** Example errors can include `Error connecting to database HBX/HBX when validating its credentials.` or `Discovery failed because credentials for HBX/SYSTEMDB either not provided or cannot be validated.` after testing a connection using hdbsql for `systemdb` and `tenantdb` that was tested from HANA database Amazon EC2 instance.

  In the AWS Backupconsole on the Jobs page, the backup job details can show a status of `FAILED` with the error `Miscellaneous: b'* 10: authentication failed SQLSTATE: 28000\n'`.

  Resolution: Ensure your password does not have special characters, such as $.

- **Error: `b'* 447: backup could not be completed: [110507] Backint exited with exit code 1 instead of 0. console output: time...`**

  Resolution: The AWS BackInt Agent for SAP HANA installation might not have completed successfully. Retry the process to deploy the AWS Backint Agent and Amazon EC2 Systems Manager Agent on your SAP application server.

- **Error: Console does not match log files after registration.**

  The discovery log shows failed registration when trying to connect to HANA DB due to the password containing special characters, though the SSM for SAP Application Manager for SAP console displays successful registration. it does not confirm that registration was successful. If the console shows successful registration but the logs do not, backups will fail.

  Confirm the registration status:

  1. Log into the SSM console
  2. Select **Run Command** from the left side navigation.
  3. Under text field **Command history**, input `Instance ID:Equal:`, with the value equal to the instance you used for registration. This will filter command history.
  4. Use the command id column to find commands with status `Failed`. Then, find the document name of **AWSSystemsManagerSAP-Discovery**.
  5. In the AWS CLI, run the command `aws ssm-sap register-application status`. If returned value shows `Error`, the registration was unsuccessful.

  Resolution: Ensure your HANA password does not have special characters (such as '$').

Creating a backup of an SAP HANA database

- **Error: AWS Backup console displays message “Fatal Error” when an on-demand backup for SystemDB or TenantDB is created.** This occurs because the public endpoint cannot be accessed. This is caused by a client side firewall that blocks access to this endpoint.

  `aws-backint-agent.log` can show errors such as `level=error msg="Storage configuration validation failed: missing backup data plane Id"` or `level=fatal msg="Error performing backup missing backup data plane Id."`

  Resolution: Open firewall access to public endpoint .

- **Error:** `Database cannot be backed up while it is stopped.`

  Resolution: Ensure the database to be backed up is active. Database data and logs can be backed up only while the database is online.

- **Error:** `Getting backup metadata failed. Check the SSM document execution for more details.`

  Resolution: Ensure the database to be backed up is active. Database data and logs can be backed up only while the database is online.

Monitoring backup logs

- **Error:** `Encountered an issue with log backups, please check SAP HANA for details.`

  Resolution: Check SAP HANA to ensure log backups are being sent to AWS Backup from SAP HANA.

- **Error:** `One or more log backup attempts failed for recovery point.`

  Resolution: Check SAP HANA for details. Ensure log backups are being sent to AWS Backup from SAP HANA.

- **Error:** `Unable to determine the status of log backups for recovery point.`

  Resolution: Check SAP HANA for details. Ensure log backups are being sent to AWS Backup from SAP HANA.

- **Error:** `Log backups for recovery point %s were interrupted due to a restore operation on the database.`

Resolution: Wait for the restore job to complete. The log backups should resume.

Glossary of SAP HANA terms when using AWS Backup

Data Backup Types: SAP HANA supports two types of data backups: Full and INC (incremental). AWS Backup optimizes which type is used during each backup operation.

Catalog Backups: SAP HANA maintains its own manifest called a *catalog*. AWS Backup interacts with this catalog. Each new backup will create an entry in the catalog.

Continuous Log Backup (Transaction Logs): For Point in Time Recovery (PITR) functions, SAP HANA tracks all transactions since the most recent backup.

System Copy: A restore job in which the restore target database is different from the source database from which the recovery point was created.

Destructive Restore: A destructive restore is a type of restore job during which a restored database deletes or overwrites the source or existing database.

FULL: A full backup is a backup of a complete database.

INC: An incremental backup is a backup of all changes to an SAP HANA database since the previous backup.

AWS Backup support of SAP HANA databases on EC2 instances release notes

Certain functionalities are not supported at this time:

- Continuous backups (which use transaction logs) cannot be copied to other Regions or accounts. Snapshot backups can be copied to supported Regions and accounts from full backups.
- Backup Audit Manager and reporting are not currently supported.
- Supported services by AWS Region contains the currently supported Regions for SAP HANA database backups on Amazon EC2 instances.

Amazon S3 backups

Overview

AWS Backup supports centralized backup and restore of applications storing data in S3 alone or alongside other AWS services for database, storage, and compute. Many features are available for S3 backups, including Backup Audit Manager.

You can use a single backup policy in AWS Backup to centrally automate the creation of backups of your application data. AWS Backup automatically organizes backups across different AWS services and third-party applications in one centralized, encrypted location (known as a backup vault) so that you can manage backups of your entire application through a centralized experience. For S3, you can create continuous backups and restore your application data stored in S3 and restore the backups to a point-in-time with a single click.

Backup access points

AWS Backup lets you access S3 backup data directly through S3 access points, without initiating a restore. For more information, see Access points.

Backup tiering

Amazon S3 is the only resource that supports backup tiering to a lower cost warm storage tier. For more information, see Backup tiering.

Prerequisites for S3 backups

Permissions and policies for Amazon S3 backup and restore

To backup, copy, and restore S3 resources, you must have the correct policies in your role. To add these policies, go to AWS managed policies. Add the AWSBackupServiceRolePolicyForS3Backup and AWSBackupServiceRolePolicyForS3Restore to the roles that you intend to use to backup and restore S3 buckets.

If you do not have sufficient permission, please request the manager of your organization's administrative (admin) account to add the policies to the intended roles.

For more information, please see Managed policies and inline policies in the *IAM User Guide*.

Backups and versioning

You must enable S3 Versioning on your S3 bucket to use AWS Backup for Amazon S3.

We recommend that you set a lifecycle expiration period for your S3 versions.

All objects (including all versions) in the bucket when the backup begins will be stored in the recovery point (completed backup). These can include the current version of each object, older versions, delete markers, and objects pending lifecycle actions.

The storage cost will be calculated for all objects in the backup, including objects scheduled for deletion (objects that will expire). You can use CLI or scripts to remove the inclusion of objects scheduled for expiration.

To learn more about setting up S3 lifecycle policies, follow the instructions on this page.

Considerations for Amazon S3 backups

The following points should be considered when you backup S3 resources:

- **Focused object metadata support** – AWS Backup supports the following metadata: tags, access control lists (ACLs), user-defined metadata, original creation date, and version ID. Object annotations are not supported. You may also restore all backed-up data and metadata except original creation date, version ID, storage class, and e-tags.
- When you restore an S3 object, AWS Backup applies a checksum value, even if the original object did not use the checksum feature.
- An S3 object key name can be made up of most UTF-8 encodable strings. The following Unicode characters are allowed: `#x9` | `#xA` | `#xD` | `#x20 to #xD7FF` | `#xE000 to #xFFFD` | `#x10000 to #x10FFFF`.

  Object key names that include characters not in this list might be excluded from backups.
- **Cold storage transition** – Use AWS Backup lifecycle management policy to define the timeline for backup expiration. Cold storage transition of S3 backups is not supported.
- For periodic backups, AWS Backup makes a best effort to track all changes to your object metadata. However, if you update a tag or ACL multiple times within 1 minute, AWS Backup might not capture all intermediate states.
- AWS Backup does not offer support for backups of SSE-C-encrypted objects. AWS Backup also does not support backups of bucket configurations, including bucket policies, settings, names, or access points.
- AWS Backup does not support backups of S3 on AWS Outposts.
- **Object annotations** – AWS Backup does not support backup or restore of S3 object annotations.

- **CloudTrail logging** – If you log data read events, you must have CloudTrail logs delivered to a different target bucket. If you save CloudTrail logs in the bucket that they log, there is an infinite loop, which can cause unexpected charges.

  For more information, see Data events in the *CloudTrail User Guide*.
- **Server access logging** – If you enable server access logging, you must have the logs delivered to a different target bucket. If you save these logs in the bucket that they log, there is an infinite loop. For more information, see Enabling Amazon S3 server access logging.
- **KMS encryption behavior during backups** – When backing up S3 objects encrypted with AWS KMS keys (SSE-KMS), AWS Backup calls `kms:Decrypt` against the source KMS key. This is because AWS Backup uses an independent encryption model for S3 backups. To achieve this, AWS Backup must first decrypt the data using the source KMS key, then re-encrypt it using the backup vault's KMS key. As a result, you will see `kms:Decrypt` calls against your source KMS key in CloudTrail logs during backup operations. For more information, see Encryption for backups in AWS Backup.
- **Continuous backup integrity checks** – AWS Backup periodically runs integrity checks on S3 continuous backups to verify data consistency with your source bucket. If you use a custom IAM policy, ensure `s3:PutInventoryConfiguration` and `s3:GetInventoryConfiguration` permissions are included. See S3 continuous backup integrity checks.

Supported bucket types, quantities, and object sizes

AWS Backup supports backup and restore operations for S3 objects of any size, up to the maximum object size supported by Amazon S3.

AWS Backup supports backup and restore of general purpose S3 buckets. Directory buckets are not supported at this time.

AWS Backup supports all general purpose S3 buckets in your account, matching the Amazon S3 bucket quota you have configured. For information about S3 bucket quotas, see Amazon S3 quotas.

Supported S3 Storage Classes

AWS Backup allows you to backup your S3 data stored in the following S3 Storage Classes:

- S3 Standard
- S3 Standard - Infrequent Access (IA)

- S3 One Zone-IA
- S3 Glacier Instant Retrieval
- S3 Intelligent-Tiering (S3 INT)

Backups of an object in the storage class S3 Intelligent-Tiering (INT) access those objects. This access triggers S3 Intelligent-Tiering to automatically move those objects to Frequent Access.

Backups that access Infrequent Access tiers, including S3 Standard - Infrequent Access (IA) and S3 One Zone-IA classes, move under the S3 storage charge of Frequent Access (applies to Infrequent Access or Archive Instant Access tiers).

The archived storage classes S3 Glacier Flexible Retrieval and S3 Glacier Deep Archive are not supported.

For more information about storage pricing for Amazon S3, see Amazon S3 Pricing.

S3 backup types

With AWS Backup, you can create the following types of backups of your S3 buckets, including object data, tags, Access Control Lists (ACLs), and user-defined metadata:

- **Continuous backups** allow you to restore to any point in time within the last 35 days. Continuous backups for an S3 bucket should only be configured in one backup plan.

  See Point-in-Time Recovery for a list of supported services and instructions on how to use AWS Backup to take continuous backups.
- **Periodic backups** use snapshots of your data to allow you to retain data for your specified duration up to 99 years. You can schedule periodic backups in frequencies such as 1 hour, 12 hours, 1 day, 1 week, or 1 month. AWS Backup takes periodic backups during the backup window you define in your backup plan.

  See Creating a backup plan to understand how AWS Backup applies your backup plan to your resources.

Cross-account and cross-Region copies are available for S3 backups, but copies of continuous backups do not have point-in-time restore capabilities.

Continuous and periodic backups of S3 buckets must both reside in the same backup vault.

AWS Backup for S3 relies on receiving S3 events through Amazon EventBridge. If this setting is disabled in S3 bucket notification settings, continuous backups will stop for those buckets with the setting turned off. For more information, see Amazon EventBridge dependency for S3 continuous backups.

Disabling AWS Backup's Amazon EventBridge rule will also result in your continuous backup stopping. If you have an active backup plan with a continuous backup rule, when that rule re-triggers, AWS Backup will recreate the Amazon EventBridge rule and a new continuous backup will be created.

For both backup types, the first backup is a full backup, while subsequent backups are incremental at object-level.

Amazon EventBridge dependency for S3 continuous backups

AWS Backup uses Amazon EventBridge to receive S3 object events that drive backup operations. AWS Backup supports two methods of Amazon EventBridge event configuration:

- **System tag-based event configuration (recommended)**: AWS Backup tags your bucket with the `aws:backup:enabled` system tag and creates a single managed Amazon EventBridge rule (named `AwsBackupManagedRule-BucketSystemTag`) that filters on this tag. S3 includes the tag in every event notification sent to Amazon EventBridge. This method supports all general purpose buckets in your account, matching your Amazon S3 bucket quota.
- **Bucket-ARN-based event configuration**: If the backup role does not have `s3:ListTagsForResource` permission, AWS Backup creates managed Amazon EventBridge rules (named `AwsBackupManagedRule-`*N*) that list individual bucket ARNs. This method supports up to 1,000 buckets per account.

In both methods, the managed rule subscribes to the following S3 events and forwards them to the AWS Backup service:

- Object Created
- Object ACL Updated
- Object Tags Added
- Object Tags Deleted
- Object Deleted

These events allow AWS Backup to maintain the continuous (point-in-time restore) recovery window.

For AWS Backup to receive these events, S3 must be configured to publish them to the default Amazon EventBridge event bus. This is controlled by the bucket's Amazon EventBridge notification setting, which is independent of the AWS Backup-managed rule.

> **Important**
>
> If the bucket-level Amazon EventBridge notification setting is disabled, S3 stops publishing object-level events. The AWS Backup-managed Amazon EventBridge rule remains enabled but receives no events for that bucket. This failure is silent – no notification, alert, or backup job failure is generated by AWS Backup.

Verify Amazon EventBridge is enabled on your bucket

Console: Open the S3 console, select your bucket, choose the **Properties** tab, scroll to **Event notifications**, and verify that **Amazon EventBridge** shows **On**.

CLI:

```
aws s3api get-bucket-notification-configuration --bucket amzn-s3-demo-bucket
```

Interpret the response as follows:

- If the response contains `"EventBridgeConfiguration": {}`, Amazon EventBridge is **enabled**.
- If the `EventBridgeConfiguration` field is absent from the response, Amazon EventBridge is **disabled**.

Enable Amazon EventBridge notifications

Console: Open the S3 console, select your bucket, choose **Properties**, scroll to **Event notifications**, choose **Amazon EventBridge**, choose **Edit**, turn **On**, and choose **Save**.

CLI:

```
aws s3api put-bucket-notification-configuration \
```

```
    --bucket amzn-s3-demo-bucket \
    --notification-configuration '{"EventBridgeConfiguration": {}}'
```

> **Important**
>
> `put-bucket-notification-configuration` is a replace operation, not a merge. If your bucket has other notification targets (SNS, SQS, Lambda), you must include them in the same call. Otherwise, they will be removed.

What happens if Amazon EventBridge is disabled on a bucket under continuous backup

- S3 stops publishing object-level events for the bucket.
- The AWS Backup-managed Amazon EventBridge rule remains enabled but receives no events for that bucket.
- New object changes made after the disable are not reflected in the continuous recovery window.
- Existing recovery points and prior continuous backup history are retained (not deleted).
- Scheduled snapshot (periodic) backup jobs defined in the plan continue to run on schedule – only the continuous (PITR) tracking is affected.
- No notification or alert is generated. The failure is silent from the AWS Backup side.

Detect changes to Amazon EventBridge notification settings

Because AWS Backup cannot detect when bucket-level Amazon EventBridge is disabled, consider setting up proactive monitoring to alert you when notification settings change. You can use either or both of the following approaches:

- **CloudTrail + Amazon EventBridge** – Create an Amazon EventBridge rule that matches `PutBucketNotificationConfiguration` API calls via CloudTrail to alert in real time when someone changes a bucket's notification configuration. The following event pattern can be used:

```
{
  "source": ["aws.s3"],
  "detail-type": ["AWS API Call via CloudTrail"],
  "detail": {
    "eventSource": ["s3.amazonaws.com"],
    "eventName": ["PutBucketNotificationConfiguration"]
```

```
    }
  }
```

This requires CloudTrail to be enabled and logging S3 management events (enabled by default). The target (SNS, Lambda, etc.) should inspect the event detail to determine whether Amazon EventBridge notifications were specifically removed.

- **AWS Config** – Use the AWS Config managed rule `s3-event-notifications-enabled` to flag non-compliant buckets on configuration changes.

Required Amazon EventBridge permissions for the backup role

The IAM role used for S3 continuous backup (PITR) jobs must have the following Amazon EventBridge permissions to allow AWS Backup to manage the `AwsBackupManagedRule`:

```
{
  "Sid": "EventBridgePermissionsForAwsBackupManagedRule",
  "Effect": "Allow",
  "Action": [
    "events:DeleteRule",
    "events:PutTargets",
    "events:DescribeRule",
    "events:EnableRule",
    "events:PutRule",
    "events:RemoveTargets",
    "events:ListTargetsByRule",
    "events:DisableRule"
  ],
  "Resource": [
    "arn:aws:events:*:*:rule/AwsBackupManagedRule*"
  ]
},
{
  "Sid": "EventBridgeListRulesPermissions",
  "Effect": "Allow",
  "Action": "events:ListRules",
  "Resource": "*"
}
```

The AWS managed policy `AWSBackupServiceRolePolicyForS3Backup` already includes these permissions. If you use a custom IAM role, you must add them explicitly.

Required permissions for system tag-based event configuration

To use system tag-based event configuration and protect more than 1,000 S3 buckets, the IAM role used for S3 backup jobs must have the following permission:

```
{
  "Sid": "S3BucketTagReadPermissions",
  "Effect": "Allow",
  "Action": [
    "s3:ListTagsForResource"
  ],
  "Resource": "arn:aws:s3:::*"
}
```

The AWS managed policy `AWSBackupServiceRolePolicyForS3Backup` includes this permission. If you use a custom IAM role, you must add it explicitly. Without this permission, AWS Backup falls back to bucket-ARN-based event configuration, which supports up to 1,000 buckets.

AWS Backup applies and removes the `aws:backup:enabled` system tag on your S3 buckets automatically. Tagging and untagging of system tags does not require additional IAM permissions in the backup role.

> **Note**
>
> The `aws:backup:enabled` system tag does not count toward the 50 user-tag limit per S3 bucket.

Re-enable continuous backup coverage

Re-enable bucket-level Amazon EventBridge using the console or CLI steps above. S3 will resume publishing events, and the managed rule will resume forwarding them to AWS Backup. Changes made while Amazon EventBridge was disabled will not be retroactively captured. The existing continuous recovery point and prior backup history remain valid.

S3 continuous backup integrity checks

With S3 continuous backups, AWS Backup tracks changes to your bucket through Amazon EventBridge events. To complement this event-driven approach, AWS Backup periodically runs an integrity check that compares your source bucket with your continuous backup data using Amazon

S3 Inventory. If the integrity check identifies any discrepancies, AWS Backup reconciles them. No configuration or opt-in is required.

If an object cannot be backed up during reconciliation – for example, due to a bucket policy or KMS key policy – AWS Backup follows the same notification path as regular backup operations. The object is identified again in the next integrity check run. To prevent repeated failures, resolve the access restriction. Objects in the S3 Glacier Flexible Retrieval or S3 Glacier Deep Archive storage classes are not processed. When a continuous recovery point is reconciled, the next periodic snapshot taken from that recovery point also reflects the updated state.

AWS Backup emits CloudTrail events when integrity checks start, complete, or fail. For the list of event names, see Logging AWS Backup API calls with CloudTrail.

If you use the AWS managed policy AWSBackupServiceRolePolicyForS3Backup, no action is required. If you use a custom IAM policy for your backup role, add the `s3:PutInventoryConfiguration` and `s3:GetInventoryConfiguration` permissions.

```
{
  "Sid": "S3InventoryPermissionsForIntegrityChecks",
  "Effect": "Allow",
  "Action": [
    "s3:PutInventoryConfiguration",
    "s3:GetInventoryConfiguration"
  ],
  "Resource": "arn:aws:s3:::*"
}
```

There is no additional charge for the integrity check. Standard Amazon S3 charges apply for Amazon S3 Inventory reports generated during the process.

Compare S3 backup types

Your backup strategy for S3 resources can involve just continuous backups, just periodic (snapshot) backups, or a combination of both. The information below can help you choose what works best for your organization:

Continuous backups only:

- After the first full backup of your existing data is complete, changes in your S3 bucket data are tracked as they occur.

- The tracked changes allow you to use PITR (point-in-time restore) for the retention period of the continuous backup. To perform a restore job, you choose the point in time to which you wish to restore.
- The retention period of each continuous backup has a maximum of 35 days.
- For backup plans you create through CLI, advanced backup settings for Amazon S3 (which include the option to include tags and ACLs in the backup) are turned on by default. You may exclude these in the backup options. See Advanced Amazon S3 backup settings for an example of the syntax.
- Because continuous backups track changes through events rather than scanning the bucket, AWS Backup periodically runs integrity checks to verify data consistency with the source bucket. Periodic (snapshot) backups do not require this check because each snapshot performs a full bucket walk. See S3 continuous backup integrity checks.

Periodic (snapshot) backups only, scheduled or on-demand:

- AWS Backup scans the entire S3 bucket, retrieves each object's ACL and tags, and initiates a Head request for every object that was in the prior snapshot but was not found in the snapshot being created.
- The backup is point-in-time consistent.
- The backup date and time recorded is the time at which AWS Backup completes the traversal of the bucket, not at the time which a backup job was created.
- The first backup of a bucket is a full backup. Each subsequent backup is incremental, representing the change in data since the last snapshot.
- The snapshot made by the periodic backup can have a retention period of up to 99 years.

Continuous backups combined with periodic/snapshot backups:

- After the first full backup of your existing data (each bucket) is complete, changes in your bucket are tracked as they occur.
- You can perform a point-in-time restore from a continuous recovery point.
- Snapshots are point-in-time consistent.
- Snapshots are taken directly from the continuous recovery point, eliminating the need to rescan a bucket to allow for faster processes.

- Snapshots and continuous recovery points share data lineage; storage of data between snapshot and continuous recovery points is not duplicated.
- When advanced Amazon S3 backup settings, such as including tags and ACLs in a backup, are changed for a `continuous` recovery point, AWS Backup stops that recovery point and creates a new one with the updated setting(s).
- AWS Backup periodically runs integrity checks on continuous recovery points to verify data consistency with the source bucket. When a continuous recovery point is reconciled, the next periodic snapshot taken from that recovery point also reflects the updated state.

When a continuous backup job is running for an S3 bucket, you can still initiate periodic (snapshot) backup jobs. However, the following behavior applies:

- Snapshot backup jobs will use the same backup options (ACLs and object tags settings) as the existing continuous backup.
- If you specify different backup options for a snapshot job than what the continuous backup uses, the snapshot job will still use the continuous backup's settings and complete with a "Completed with issues" status.

  When this occurs, you'll see the following status message: `"Periodic/snapshot backup for bucket <bucket name> has different backup options than the continuous backup. When using continuous backups along with snapshot backups for the same bucket, the snapshot will use the same settings for backing up ACLs and Object tags as the continuous backup."`

The following table shows when a full scan is required when changing BackupOptions for existing continuous recovery points:

Full scan behavior when BackupOptions is modified

Previous BackupOptions	New BackupOptions	Full scan
backupACLs and backupObjectTags enabled	backupACLs and backupObjectTags disabled	No
backupACLs and backupObjectTags enabled	backupACLs enabled; backupObjectTags disabled	No

Previous BackupOptions	New BackupOptions	Full scan
backupACLs and backupObjectTags enabled	backupACLs disabled; backupObjectTags enabled	No
backupACLs and backupObjectTags disabled	backupACLs and backupObjectTags enabled	Yes
backupACLs enabled; backupObjectTags disabled	backupACLs and backupObjectTags enabled	Yes
backupACLs disabled; backupObjectTags enabled	backupACLs and backupObjectTags enabled	Yes

S3 backup completion windows

The table below shows sample buckets of various sizes to help you guide estimates of the completion time of the initial full backup of an S3 bucket. Backup times will vary with the size, content, configuration, and settings of each bucket.

Bucket size	Number of objects	Estimated time to complete initial backup
425 GB (gigabytes)	135 million	31 hours
800 TB (terabytes)	670 million	38 hours
6 PB (petabytes)	5 billion	100 hours
370 TB (terabytes)	7.5 billion	180 hours

Best practices and cost considerations for S3 backups

Large bucket best practices

For buckets with more than 300 million objects:

- For buckets with greater than 300 million objects, the backup rate can reach up to 17,000 objects per second during the initial full backup of the bucket (incremental backups will have a different speed); buckets containing fewer than 300 million objects back up at a rate close to 1,000 objects per second.
- Continuous backups are recommended.
- If backup lifecycle is planned for more than 35 days, you can also enable snapshot backups for the bucket in the same vault in which your continuous backups are stored.

Backup strategy optimization

- For accounts which make backups at least daily or more frequently, cost benefits can be realized by using continuous backups if the data within the backups has minimal changes between backups.
- Larger buckets that do not change frequently can benefit from continuous backups, since this can result in lower costs when scans of the whole bucket along with multiple requests per objects don't need to be performed on pre-existing objects (objects that are unchanged from the previous backup).
- Buckets that contain more than 100 million objects and that have a small delete rate compared to the overall backup size might realize cost benefits with a backup plan that contains both a continuous backup with a retention period of 2 days along with snapshots of a longer retention.
- Periodic (snapshot) backup time aligns with the start of the backup process when a bucket scan is not needed. Scans are not needed in a bucket that contains both continuous backup and snapshots since in these cases snapshots are taken from a continuous recovery point.

Object lifecycle and delete markers

- S3 lifecycle policies have an optional feature called **Delete expired object delete markers**. When this feature is left off, delete markers, sometimes in the millions, expire with no cleanup plan. When buckets without this feature are backed up, two issues impact time and cost:
  - Delete markers are backed up, just like objects. Backup time and restore time can be impacted depending on the ratio of objects to delete markers.
  - Each object and marker that is backed up has a minimum charge. Each delete marker is charged the same as a 128KiB object.

Storage class cost considerations

- For each object in a single S3-GIR (Amazon S3 Glacier Instant Retrieval), AWS Backup performs multiple calls, which will result in retrieval charges when a backup is conducted.

  Similar retrieval costs apply to buckets with objects in S3-IA and S3 One Zone-IA storage classes.

AWS service cost optimization

- Using features of AWS KMS, CloudTrail, Amazon CloudWatch, and Amazon GuardDuty as part of your backup strategy can result in additional costs beyond S3 bucket data storage. See the following for information on adjusting these features:
  - Reducing the cost of SSE-KMS with Amazon S3 Bucket keys in the *Amazon S3 User Guide*.
  - You can reduce CloudTrail costs by excluding AWS KMS events and by disabling S3 data events:
    - **Exclude AWS KMS events:** In the *CloudTrail User Guide*, Creating a trail in the console (basic event selectors) allows the option to exclude AWS KMS events to filter these events out of your trail (default setting includes all KMS events):
      - The option to log or exclude KMS events is available only if you log management events on your trail. If you choose not to log management events, KMS events are not logged, and you cannot change KMS event logging settings.
      - AWS KMS actions such as `Encrypt`, `Decrypt`, and `GenerateDataKey` typically generate a large volume (more than 99%) of events. These actions are now logged as **Read** events. Low-volume, relevant KMS actions such as `Disable`, `Delete`, and `ScheduleKey` (which typically account for less than 0.5% of KMS event volume) are logged as **Write** events.
      - To exclude high-volume events like `Encrypt`, `Decrypt`, and `GenerateDataKey`, but still log relevant events such as `Disable`, `Delete`, and `ScheduleKey`, choose to log **Write** management events, and clear the check box for **Exclude AWS KMS events**.
    - **Disable S3 data events:** By default, trails and event data stores do not log data events. Disable S3 data events before your initial backup to reduce costs.
  - To reduce CloudWatch costs, you can stop sending CloudTrail events to CloudWatch Logs when you update a trail to disable CloudWatch Logs settings.
  - Estimating GuardDuty usage cost in the *Amazon GuardDuty User Guide*.

S3 backup messages

When a backup job completes or fails, you may see the following message. The following table can help you determine the possible cause of the status message.

Scenario	Job Status	Message	Example
All objects failed to be backed up for a snapshot or initial continuous backup	`FAILED`	"No objects were backed up from the source bucket *`BucketName`* . To get notified of these failures, enable SNS event notifications."	Backup role does not have the permission to get object version ACL. Consequently, none of the objects are backed up.
All objects failed to be backed up for a subsequent continuous backup.	`COMPLETED`	"No objects were backed up from the source bucket *`BucketName`* . To get notified of these failures, enable SNS event notifications."	

Advanced Amazon S3 backup settings

AWS Backup provides advanced settings to control what metadata is included in your Amazon S3 backups. You can optionally exclude Access Control Lists (ACLs) and object tags, which can be helpful if your objects are set up without ACLs and object tags. In other words, if you do not use ACLs or objects tags for your S3 resources, you may find it beneficial to exclude them from your backups.

Configuring backup of ACLs and object tags

You can configure ACL and object tag backup options either through the AWS Backup console or through the AWS CLI.

Console

Configure ACL and tag options using the console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup/.
2. In the navigation pane, choose **Backup plans**, then choose **Create backup plan**.
3. In your backup plan settings, expand **Advanced backup settings**.
4. For Amazon S3 resources, configure the following options:

   - **Back up ACLs**: Select the check box to include ACLs or leave it unselected to exclude them.

     Backup up object tags: Select the check box to include object tags in your backup.
5. Complete the backup plan configuration and choose **Create plan**.

AWS CLI

You can selectively include or exclude Access Control Lists (ACLs) and object tags from your Amazon S3 backups using the following backup options:

BackupACLs

Controls whether object ACLs are included in the backup. Set to `disabled` to exclude ACLs. Default: `enabled`

BackupObjectTags

Controls whether object tags are included in the backup. Set to `disabled` to exclude tags. Default: `enabled`

Configure ACL and tag options using the AWS CLI

To configure ACL and object tag backup options using the AWS CLI, use the `update-backup-plan` command with advanced backup settings:

```
aws backup update-backup-plan \
    --backup-plan-id "your-backup-plan-id" \
    --backup-plan '{
        "BackupPlanName": "MyS3BackupPlan",
```

```
        "Rules": [{
            "RuleName": "MyS3BackupRule",
            "TargetBackupVaultName": "MyBackupVault",
            "ScheduleExpression": "cron(0 2 ? * * *)",
            "Lifecycle": {
                "DeleteAfterDays": 30
            },
            "RecoveryPointTags": {},
            "CopyActions": [],
            "EnableContinuousBackup": false
        }],
        "AdvancedBackupSettings": [{
            "ResourceType": "S3",
            "BackupOptions": {
                "BackupACLs": "disabled",
                "BackupObjectTags": "disabled"
            }
        }]
    }'
```

The `BackupOptions` parameters control metadata inclusion:

- `"BackupACLs": "disabled"` - Excludes ACLs from backups
- `"BackupObjectTags": "disabled"` - Excludes object tags from backups
- `"BackupACLs": "enabled"` - Includes ACLs in backups (default)
- `"BackupObjectTags": "enabled"` - Includes object tags in backups (default)

Amazon Timestream backups

Amazon Timestream is a scalable time series database that allows storage and analysis of up to trillions of time series data points daily. Timestream is optimized for cost and time savings by keeping recent data in memory and by storing historical data in a cost-optimized storage tier in accordance with your policies.

A Timestream database has tables. These tables contain records, and each record is a single data point in a time series. A time series is a sequence of records recorded over a time interval, such as a stock price, usage level of memory of an Amazon EC2 instance, or a temperature reading. AWS Backup can centrally backup and restore Timestream tables. You can copy these table backups to other accounts and several other AWS Regions within the same organization.

Timestream does not currently offer native backup and restore services, so using AWS Backup to create secure copies of your Timestream tables can add an extra layer of security and resilience to your resources.

> **Important**
>
> AWS Backup supports Amazon Timestream for LiveAnalytics only. AWS Backup does not support Amazon Timestream for InfluxDB.
>
> - For information about backing up Amazon Timestream for LiveAnalytics tables, see Backups in the *Amazon Timestream for LiveAnalytics Developer Guide*.
> - Amazon Timestream for InfluxDB uses its own native backup solution and is not supported by AWS Backup. For more information, see Resilience in Amazon Timestream for InfluxDB in the *Amazon Timestream for InfluxDB Developer Guide*.

Back up Timestream tables

You can backup Timestream tables either through the AWS Backup console or using the AWS CLI.

There are two ways to use the AWS Backup console to backup a Timestream table: on demand or as part of a backup plan.

Create on-demand Timestream backups

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Using the navigation pane, choose **Protected resources**, and then **Create on-demand backup**.
3. On the **Create on-demand backup** page, choose Amazon Timestream.
4. Choose **Resource type** Timestream, and then choose the table name you want to back up.
5. In Backup window, ensure that **Create backup now** is selected. This initiates a backup immediately and enables you to see your cluster sooner on the **Protected resources** page.
6. In the drop down menu **Transition to cold storage**, you can set your transition settings.
7. In **Retention Period**, you can choose how long to retain your backup.
8. Choose an existing backup vault or create a new backup vault. Choosing **Create new backup vault** opens a new page to create a vault and then returns you to the **Create on-demand backup page** when you are finished.

9. Under **IAM role**, choose **Default role** (if the AWS Backup default role is not present in your account, it will be created for you with the correct permissions).
10. *Optionally,* tags can be added to your recovery point. If you want to assign one or more tags to your on-demand backup, enter a **key** and optional **value**, and choose **Add tag**.
11. Choose **Create on-demand backup**. This takes you to the **Jobs** page, where you will see a list of jobs.
12. Choose the **Backup job ID** for the cluster to see the details of that job. It will display a status of `Completed`, `In Progress`, or `Failed`. You can click the refresh button to update the displayed status.

Create scheduled Timestream backups in a backup plan

Your scheduled backups can include Timestream tables if they are a protected resource. To opt into protecting Amazon Timestream tables:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Using the navigation pane, choose **Protected resources**.
3. Toggle Amazon Timestream to **On**.
4. See Assigning resources to the console to include Timestream tables in an existing or new plan.

Under **Manage Backup plans**, you can choose to create a backup plan and include Timestream tables, or you can update an existing one to include Timestream tables. When adding the resource type *Timestream*, you can choose to add **All Timestream tables**, or check the boxes next to the tables you wish to add under **Select specific resource types**.

The first backup made of Timestream tables will be a full backup. Subsequent backups will be incremental backups.

After you've created or modified your backup plan, navigate to Backup plans in the left navigation. The backup plan you specified should display your clusters under **Resource Assignments**.

Backing up programmatically

You can use the operation name `start-backup-job`. Include the following parameters:

```
aws backup start-backup-job \
--backup-vault-name backup-vault-name \
```

```
--resource-arn arn:aws:timestream:region:account:database/database-name/table/table-
name \
--iam-role-arn arn:aws:iam::account:role/role-name \
--region AWS Region \
--endpoint-url URL
```

View Timestream table backups

To view and modify your Timestream table backups within the console:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Choose **Backup vaults**. Then, click on the backup vault name that contains your Timestream tables.
3. The backup vault will display a summary and a list of backups.
   a. You can click on the link in the column **Recovery point ID**, or
   b. You can check the box to the left of the recovery point ID and click **Actions** to delete the recovery points that are no longer needed.

Restore a Timestream table

See how to restore a Timestream table

Virtual machine backups

AWS Backup supports centralized and automated data protection for on-premises VMware virtual machines (VMs) along with VMs in the VMware Cloud™ (VMC) on AWS and VMware Cloud™ (VMC) on AWS Outposts. You can back up from your on-premises and VMC virtual machines to AWS Backup. Then, you can restore from AWS Backup to on-premises VMs, VMs in the VMC, or the VMC on AWS Outposts.

AWS Backup also provides you with fully-managed, AWS-native VM backup management capabilities, such as VM discovery, backup scheduling, retention management, a low-cost storage tier, cross-Region and cross-account copy, support for AWS Backup Vault Lock and AWS Backup Audit Manager, encryption that is independent from source data, and backup access policies. For a full list of capabilities and details, see the Feature availability by resource table.

You can use AWS Backup to protect your virtual machines on VMware Cloud™ on AWS Outposts. AWS Backup stores your VM backups in the AWS Region to which your VMware Cloud™ on AWS

Outposts is connected. You can use AWS Backup to protect your VMware Cloud™ on AWS Backup VMs when you're using VMware Cloud™ on AWS Outposts to meet your low-latency and local data-processing needs for your application data. Based on your data residency requirements, you may choose AWS Backup to store backups of your application data in the parent AWS Region to which your AWS Outposts is connected.

Supported VMs

AWS Backup can back up and restore virtual machines managed by a VMware vCenter.

Currently supported:

- vSphere 8, 7.0, and 6.7
- Virtual disk sizes that are multiples of 1 KiB
- NFS, VMFS, and VSAN datastores on premises and in VMC on AWS
- SCSI Hot-Add and Network Block Device Secure Sockets Layer (NBDSSL) transport modes for copying data from source VMs to AWS for on-premises VMware
- Hot-Add mode to protect VMs on VMware Cloud on AWS

Not currently supported:

- RDM (raw disk mapping) disks or NVMe controllers and their disks
- Independent-persistent and independent-non persistent disk modes

VM discovery requirements

Backup gateway persists each virtual machine's vCenter folder path as an AWS tag on the discovered VM record. For a VM to be included in the hypervisor's VM list, every segment of its vCenter folder path must contain only characters valid in AWS tag values: letters, digits, spaces, and the following special characters: `_ . : / = + - @`.

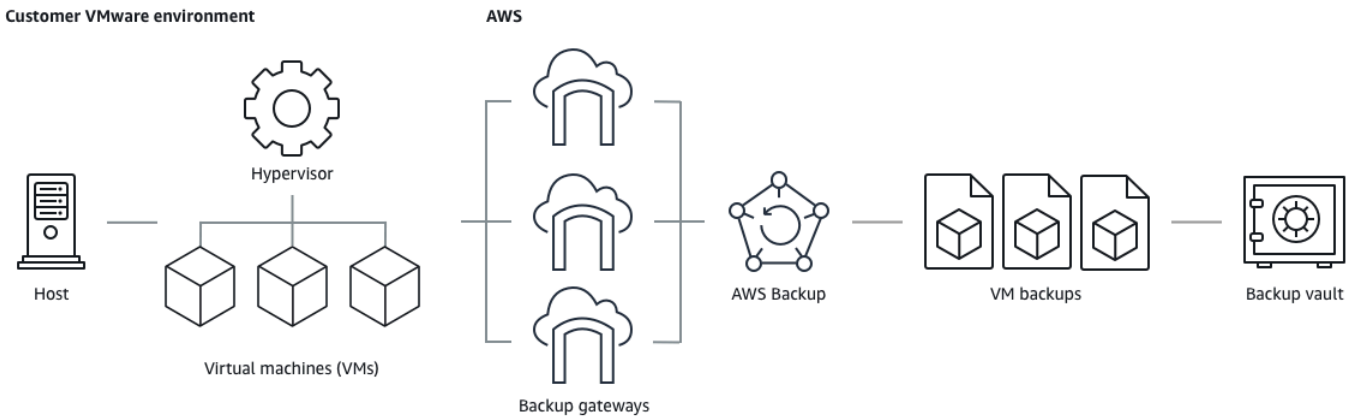
VMs in folders with other characters (such as parentheses ( ), commas, or question marks) are excluded from discovery. The exclusion does not appear in the metadata-sync status or in error metrics. To include these VMs, rename the folder in vCenter to remove the unsupported characters, then wait for the next metadata sync.

Backup consistency

AWS Backup, by default, captures application-consistent backups of VMs using the VMware Tools quiescence setting on the VM. Your backups are application consistent if your applications are compatible with VMware Tools. If the quiescence capability is not available, AWS Backup captures crash-consistent backups. Validate that your backups meet your organization's needs by testing your restores.

Backup gateway

Backup gateway is downloadable AWS Backup software that you deploy to your VMware infrastructure to connect your VMware VMs to AWS Backup. The gateway connects to your VM management server to discover VMs, discovers your VMs, encrypts data, and efficiently transfers data to AWS Backup. The following diagram illustrates how Backup gateway connects to your VMs:



To download the Backup gateway software, follow the procedure for Working with gateways.

Download VM software

Backup gateway is distributed as an OVF (Open Virtualization Format) template that you deploy to your VMware infrastructure. The gateway software connects your VMware VMs to AWS Backup by discovering VMs, encrypting data, and efficiently transferring data to AWS Backup.

To obtain the OVF template, use the AWS Backup console:

1. Sign in to the AWS Management Console and open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the left navigation pane, under **External resources**, choose **Gateways**.

3. Choose **Create gateway**.
4. In the **Set up gateway** section, download the OVF template and deploy it to your VMware environment.

For information on VPC (Virtual Private Cloud) endpoints, see AWS Backup and AWS PrivateLink connectivity.

Backup gateway comes with its own API which is separately maintained from the AWS Backup API. To view a list of Backup gateway API actions, see Backup gateway actions. To view a list of Backup gateway API data types, see Backup gateway data types.

Endpoints

Existing users who currently use a public endpoint and who wish to switch to a VPC (Virtual Private Cloud) endpoint can create a new gateway with a VPC endpoint using AWS PrivateLink, associate the existing hypervisor to the gateway, and then delete the gateway containing the public endpoint.

Configure your infrastructure to use Backup gateway

Backup gateway requires the following network, firewall, and hardware configurations to back up and restore your virtual machines.

Network configuration

Backup gateway requires certain ports to be allowed for its operation. Allow the following ports:

1. **TCP 443 Outbound**
   - Source: Backup gateway
   - Destination: AWS
   - Use: Allows Backup gateway to communicate with AWS.
2. **TCP 80 Inbound**
   - Source: The host you use to connect to the AWS Management Console
   - Destination: Backup gateway
   - Use: By local systems to obtain the Backup gateway activation key. Port 80 is only used during activation of Backup gateway. AWS Backup does not require port 80 to be publicly accessible. The required level of access to port 80 depends on your network configuration. If you activate

your gateway from the AWS Management Console, the host from which you connect to the console must have access to your gateway's port 80.

3. **UDP 53 Outbound**
   - Source: Backup gateway
   - Destination: Domain Name Service (DNS) server
   - Use: Allows Backup gateway to communicate with the DNS.
4. **TCP 22 Outbound**
   - Source: Backup gateway
   - Destination: Support
   - Use: Allows Support to access your gateway to help you with issues. You don't need to open this port for the normal operation of your gateway, but you must open it for troubleshooting.
5. **UDP 123 Outbound**
   - Source: NTP client
   - Destination: NTP server
   - Use: Used by local systems to synchronize virtual machine time to the host time.
6. **TCP 443 Outbound**
   - Source: Backup gateway
   - Destination: VMware vCenter
   - Use: Allows Backup gateway to communicate with VMware vCenter.
7. **TCP 443 Outbound**
   - Source: Backup gateway
   - Destination: ESXi hosts
   - Use: Allows Backup gateway to communicate with ESXi hosts.
8. **TCP 902 Outbound**
   - Source: Backup gateway
   - Destination: VMware ESXi hosts
   - Use: Used for data transfer via Backup gateway.

The above ports are necessary for Backup gateway. See Create a VPC endpoint for more information on how to configure Amazon VPC endpoints for AWS Backup.

Firewall configuration

Backup gateway requires access to the following service endpoints to communicate with Amazon Web Services. If you use a firewall or router to filter or limit network traffic, you must configure your firewall and router to allow these service endpoints for outbound communication to AWS. Use of an HTTP proxy in between Backup gateway and service points is not supported.

Endpoint types

Standard endpoints: Support IPv4 traffic between your gateway appliance and AWS.

The following service endpoints are required by all gateways for control path (`anon-cp`, `client-cp`, `proxy-app`) and data path (`dp-1`) operations.

```
anon-cp.backup-gateway.region.amazonaws.com:443
client-cp.backup-gateway.region.amazonaws.com:443
proxy-app.backup-gateway.region.amazonaws.com:443
dp-1.backup-gateway.region.amazonaws.com:443
```

Dual-stack endpoints: Support both IPv4 and IPv6 traffic between your gateway appliance and AWS.

The following dual-stack service endpoints are required by all gateways for control path (activation, controlplane, proxy) and data path (dataplane) operations.

```
activation-backup-gateway.region.api.aws:443
controlplane-backup-gateway.region.api.aws:443
proxy-backup-gateway.region.api.aws:443
dataplane-backup-gateway.region.api.aws:443
```

Configure your gateway for multiple NICs in VMware

You can maintain separate networks for your internal and external traffic by attaching multiple virtual network interface connections (NICs) to your gateway and then directing internal traffic (gateway to hypervisor) and external traffic (gateway to AWS) separately.

By default, virtual machines connected to AWS Backup gateway have one network adapter (`eth0`). This network includes the hypervisor, the virtual machines, and network gateway (Backup gateway) which communicates with the broader Internet.

Here is an example of a setup with multiple virtual network interfaces:

```
eth0:
- IP: 10.0.3.83
- routes: 10.0.3.0/24

eth1:
- IP: 10.0.0.241
- routes: 10.0.0.0/24
- default gateway: 10.0.0.1
```

- In this example, the connection is to a hypervisor with IP `10.0.3.123`, the gateway will use `eth0` as the hypervisor IP is part of the `10.0.3.0/24` block
- To connect to a hypervisor with IP `10.0.0.234`, the gateway will use `eth1`
- To connect to an IP outside of the local networks (ex. `34.193.121.211`), the gateway will fall back to the default gateway, `10.0.0.1`, which is in the `10.0.0.0/24` block and thus go through `eth1`

The first sequence to add an additional network adapter occurs in the vSphere client:

1. In the VMware vSphere client, open the context menu (with a right-click) for your gateway virtual machine, and choose **Edit Settings**.
2. On the **Virtual Hardware** tab of the **Virtual Machine Properties** dialog box, open the **Add New Device** menu, and select **Network Adapter** to add a new network adapter.
3.
   a. Expand the **New Networ**k details to configure the new adapter.
   b. Ensure that **Connect At Power On** is selected.
   c. For **Adapter Type**, see Network Adapter Types in the ESXi and vCenter Server Documentation.
4. Click **Okay** to save the new network adapter settings.

The next sequence of steps to configure an additional adapter occurs in the AWS Backup gateway console (note this is not the same interface as the AWS management console where backups and other services are managed).

Once the new NIC is added to the gateway VM, you need to

- Go to `Command Prompt` and turn on the new adapters

- Configure static IPs for each new NIC
- Set the preferred NIC as the default

To do these:

1. In the VMware vSphere client, select your gateway virtual machine and **Launch Web Console** to access the Backup gateway local console.
   - For more information on accessing a local console, see Accessing the Gateway Local Console with VMware ESXi
2. Exit Command Prompt and go to Network Configuration > Configure Static IP and follow the setup instructions to update the routing table.
   a. Assign a static IP within the network adapter's subnet.
   b. Set up a network mask.
   c. Enter the IP address of the default gateway. This is the network gateway that connects to all traffic outside of the local network.
3. Select **Set Default Adapter** to designate the adapter that will be connected to the cloud as the default device.
4. All IP addresses for the gateway can be displayed in both the local console and on the VM summary page in VMware vSphere.

VMware permissions

This section lists the minimum VMware permissions required to use AWS Backup gateway. These permissions are necessary for Backup gateway to discover, backup, and restore virtual machines.

To use Backup gateway with VMware Cloud™ on AWS or VMware Cloud™ on AWS Outposts, you must use the default admin user `cloudadmin@vmc.local` or assign the CloudAdmin role to your dedicated user.

To use Backup gateway with VMware on-premises virtual machines, create a dedicated user with the permissions listed below.

Global

- Disable methods
- Enable methods

- Licenses
- Log event
- Manage custom attributes
- Set custom attributes

vSphere Tagging

- Assign or Unassign vSphere Tag

DataStore

- Allocate space
- Browse datastore
- Configure datastore (for vSAN datastore)
- Low level file operations
- Update virtual machine files

Host

- Configuration
  - Advanced settings
  - Storage partition configuration

Folder

- Create folder

Network

- Assign network

dvPort Group

- Create

- Delete

Resource

- Assign virtual machine to resource pool

Virtual Machine

- Change Configuration
  - Acquire disk lease
  - Add existing disk
  - Add new disk
  - Advanced configuration
  - Change settings
  - Configure raw device
  - Modify device settings
  - Remove disk
  - Set annotation
  - Toggle disk change tracking
- Edit Inventory
  - Create from existing
  - Create new
  - Register
  - Remove
  - Unregister
- Interaction
  - Power Off
  - Power On
- Provisioning
  - Allow disk access
  - Allow read-only disk access
  - Allow virtual machine download

- Snapshot Management
  - Create snapshot
  - Remove Snapshot
  - Revert to snapshot

Working with gateways

To back up and restore your virtual machines (VMs) using AWS Backup, you must first install a Backup gateway. A gateway is software in the form of an OVF (Open Virtualization Format) template that connects Amazon Web Services Backup to your hypervisor, allowing it to automatically detect your virtual machines, and enables you to back up and restore them.

A single gateway can run up to 4 backup or restore jobs at once. To run more than 4 jobs at once, create more gateways and associate them with your hypervisor.

Creating a gateway

You can create a backup gateway using two approaches:

- **Console method (standard)**: Creates gateways through the AWS Backup console with automatic activation
- **Manual method**: Creates gateways using gateway VM's local console by obtaining activation keys and using AWS CLI commands

Both methods require downloading and deploying the OVF template first (see Download VM software).

Both methods allow gateway to communicate over IPv6, which requires gateway appliance version 2.x+ and additional firewall configuration on dual-stack endpoints.

> **Important**
>
> **IPv6 hypervisor requirement:** If your gateway is activated through IPv6, you **must** create a hypervisor with an IPv6 address. For example, use `2607:fda8:1001:210::252` instead of `10.0.0.252`. If you associate an IPv6 gateway with an IPv4 hypervisor, backup and restore jobs will likely fail.

Console method

To create a gateway:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the left navigation pane, under the **External resources** section, choose **Gateways**.
3. Choose **Create gateway**.
4. In the **Set up gateway** section, follow these instructions to download and deploy the OVF template.

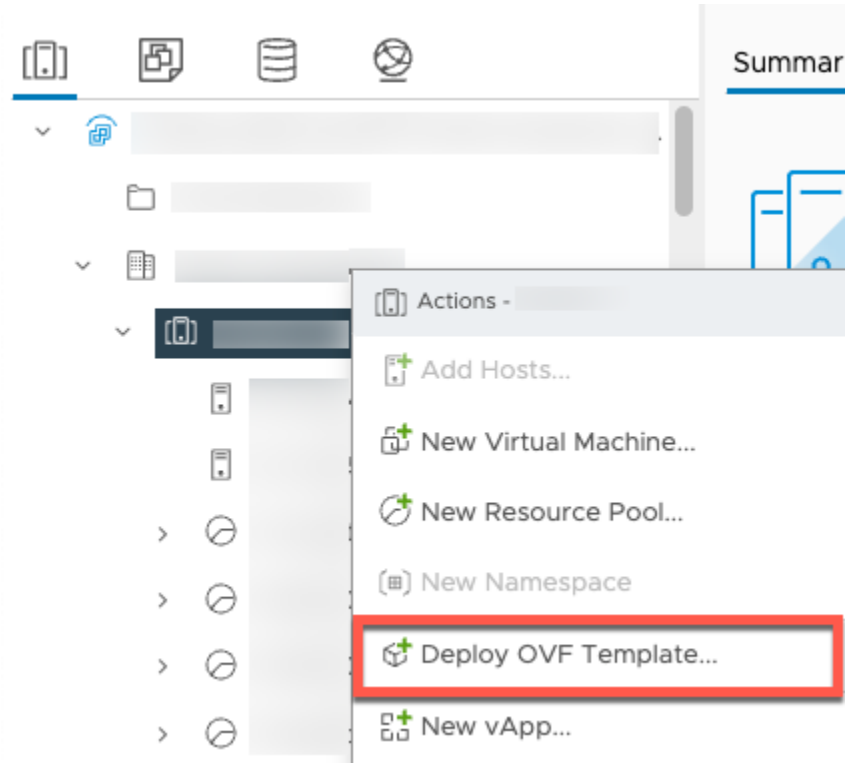
Downloading VMware software

Connecting the hypervisor

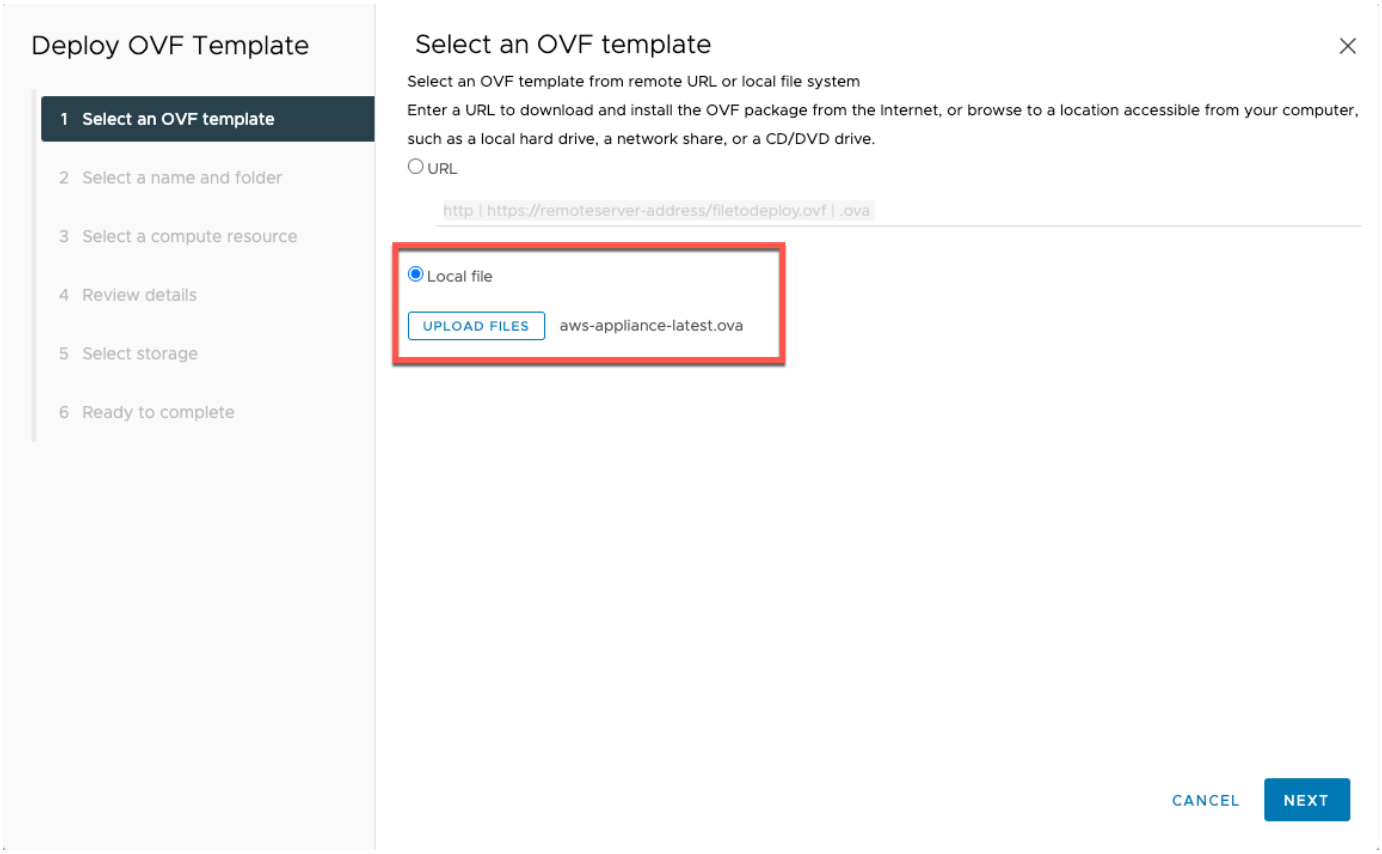
Gateways connect AWS Backup to your hypervisor so you can create and store backups of your virtual machines. To set up your gateway on VMware ESXi, download the OVF template. The download may take about 10 minutes.

After it is complete, proceed with the following steps:

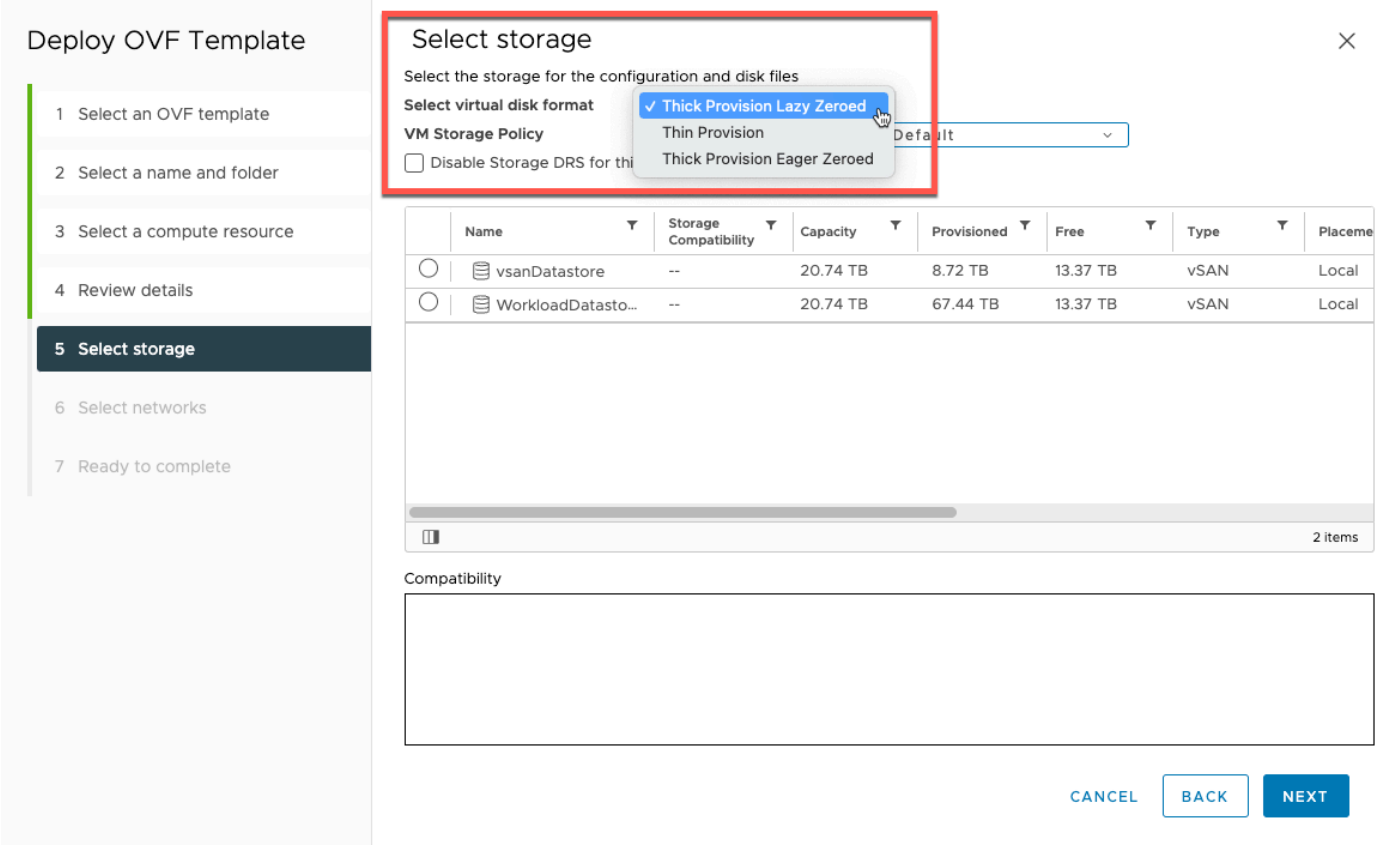
1. Connect to your virtual machine hypervisor using VMware vSphere.
2. Right-click a parent object of a virtual machine and select *Deploy OVF Template.*



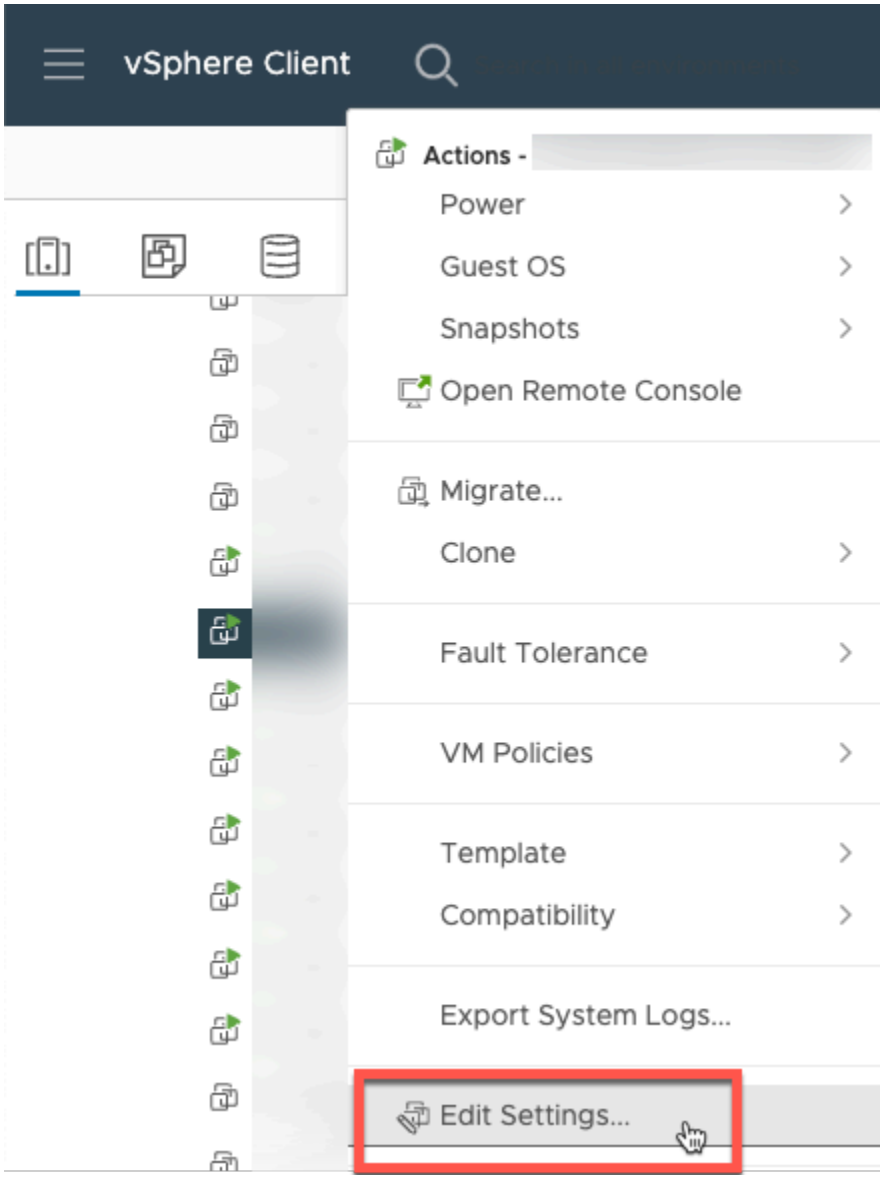
3. Choose **Local file**, and upload the **aws-appliance-latest.ova** file you downloaded.



4. Follow the deployment wizard steps to deploy it. On the **Select storage** page, select virtual disk format **Thick Provision Lazy Zeroed**.

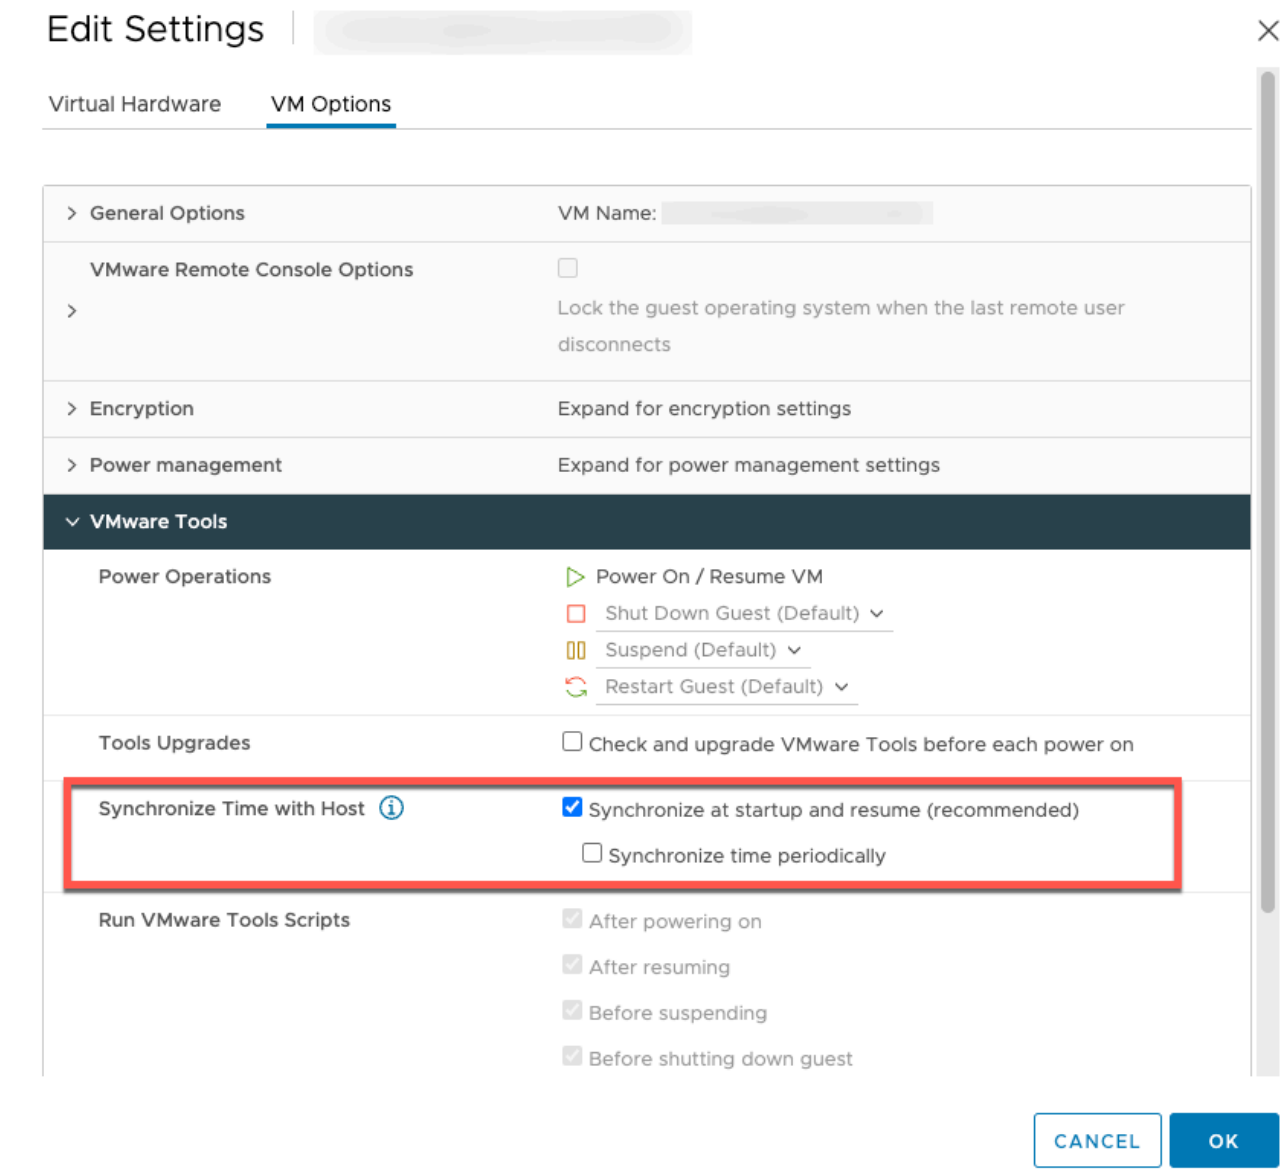


5. After deploying the OVF, right-click the gateway and choose **Edit Settings**.

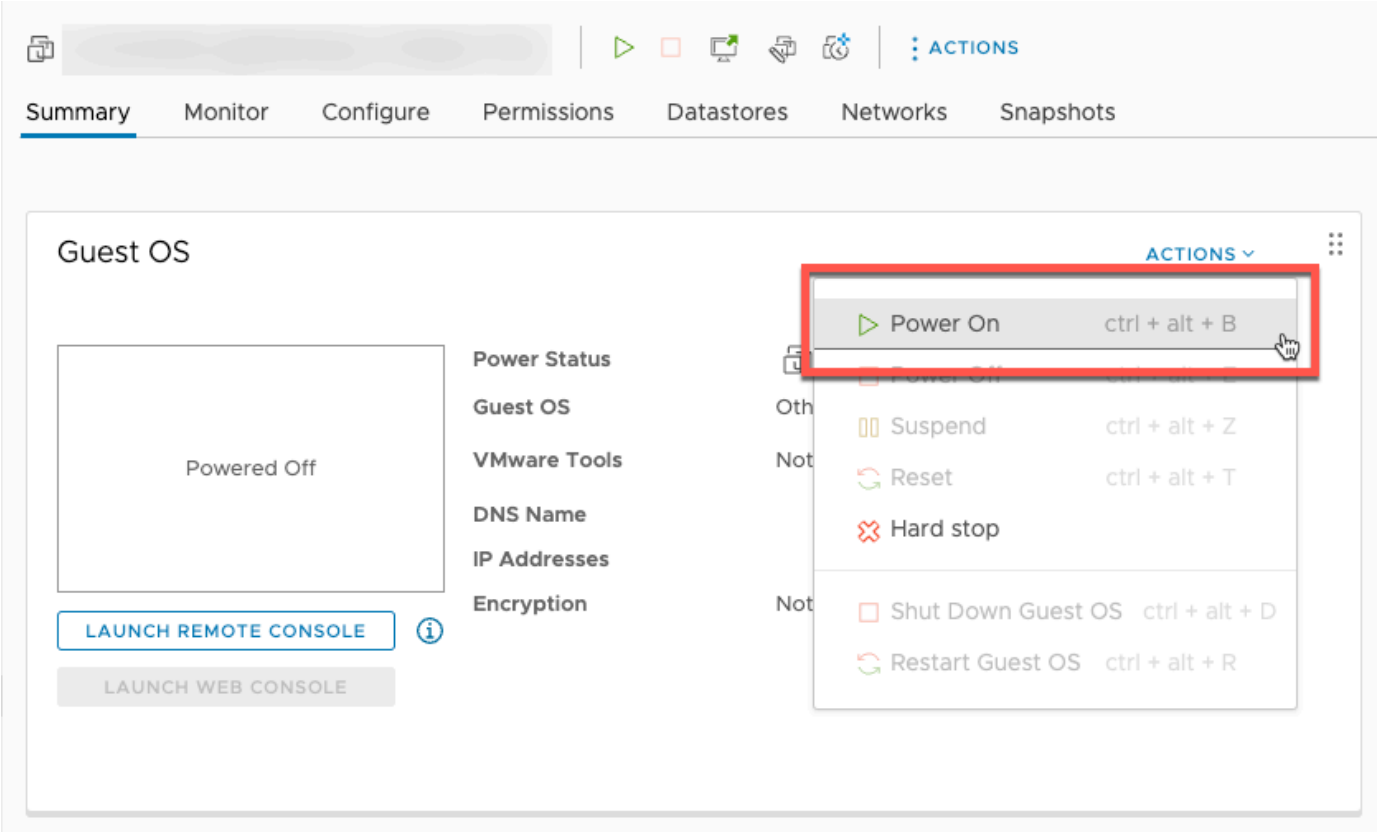


a. Under **VM Options**, go to **VM Tools**.

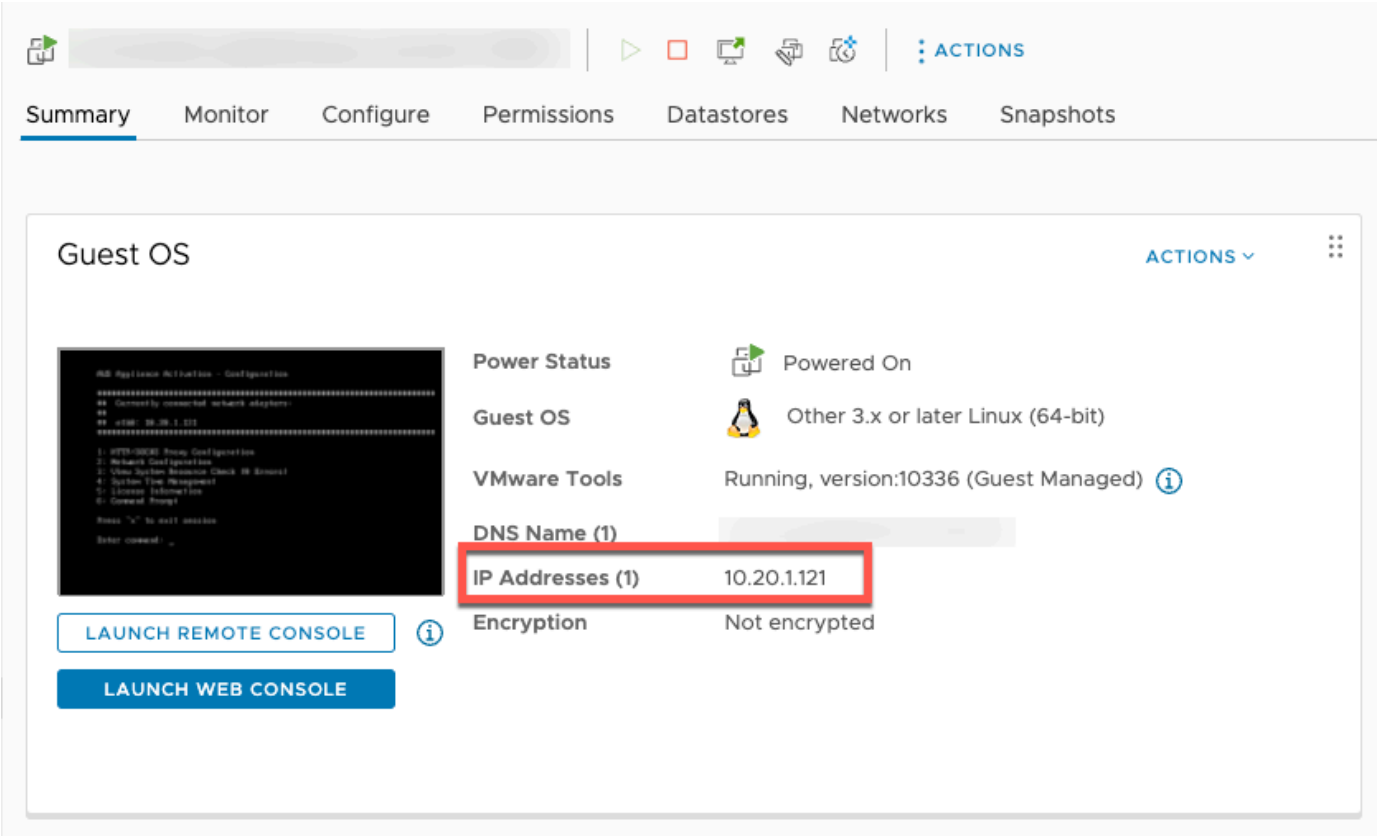
b. Ensure that for **Synchronize Time with Host**, **Synchronize at start up and resume** is selected.



6. Turn on the virtual machine by selecting “Power On” from the **Actions** menu.



7. Copy the IP address from the VM summary and enter it below.



Once the VMWare software is downloaded, complete the following steps:

1. In the **Gateway connection** section, type in the **IP address** of the gateway.
   a. To find this IP address, go to the vSphere Client.
   b. Select your gateway under the **Summary** tab.
   c. Copy the **IP address** and paste it in the AWS Backup console text bar.
2. In the **Gateway settings** section,
   a. Type in a **Gateway name**.
   b. Verify the AWS Region.
   c. Choose whether the endpoint is publicly accessible or hosted with your virtual private cloud (VPC).
      - If **publicly accessible** is selected, choose the IP version (IPv4 or IPv6) for gateway connectivity.
      - If **VPC** is selected, enter the VPC endpoint DNS Name. For more information, see Create a VPC endpoint.
3. *[Optional]* In the **Gateway tags** section, you can assign tags by inputting the **key** and *optional* **value**. To add more than one tag, click **Add another tag**.
4. To complete the process, click **Create gateway**, which takes you to the gateway detail page.

Manual gateway creation

Getting an activation key

To receive an activation key for your gateway, make a web request to the gateway virtual machine (VM) or use the gateway local console. The gateway VM returns a response that contains the activation key, which is then passed as one of the parameters for the `CreateGateway` API to specify the configuration of your gateway.

> **Tip**
>
> Gateway activation keys expire in 30 minutes if unused.

Getting an activation key using web request

The following examples show you how to get an activation key using HTTP request. You can either use a web browser or Linux curl or equivalent command using the following URLs.

> **Note**
>
> Replace the highlighted variables with actual values for your gateway. Acceptable values are as follows:
>
> - *gateway_ip_address* - The IPv4 address of your gateway, for example 172.31.29.201
> - *region_code* - The Region where you want to activate your gateway. See Regional endpoints in the *AWS General Reference Guide*. If this parameter is not specified, or if the value provided is misspelled or doesn't match a valid region, the command will default to the us-east-1 region.

Dual-stack endpoints (IPv6 support)

IPv4:

```
curl "http://gateway_ip_address/?
activationRegion=region_code&gatewayType=BACKUP_VM&endpointType=DUALSTACK&ipVersion=ipv4&no_red
```

IPv6:

```
curl "http://gateway_ip_address/?
activationRegion=region_code&gatewayType=BACKUP_VM&endpointType=DUALSTACK&ipVersion=ipv6&no_red
```

Getting an activation key using local console

The following examples show you how to get an activation key using gateway host's local console

1. Log in to your virtual machine console.
2. From the **AWS Appliance Activation - Configuration** main menu, select 0 to choose **Get activation key**
3. Select 2 **Backup Gateway** for gateway family option
4. Enter the AWS Region where you want to activate your gateway
5. For network type, enter 1 for Public or 2 for VPC endpoint

6. For endpoint type, enter 1 for standard endpoint or 2 for dual-stack endpoint
   - For dual-stack endpoint, select 1 for IPv4 or 2 for IPv6
7. Activation key will be populated automatically

Creating the gateway

Use the AWS CLI to create the gateway after obtaining an activation key:

1. Obtain activation key using curl commands or local console method
2. Create gateway using AWS CLI, for more information, see CreateGateway in the *Backup gateway API Reference*.

```
aws backup-gateway create-gateway \
                    --region region_code \
                    --activation-key activation_key \
                    --gateway-display-name gateway_name \
                    --gateway-type BACKUP_VM
```

3. Verify gateway appears in AWS Backup console under **External Resources** → **Gateways**

Editing or deleting a gateway

To edit or delete a gateway:

1. In the left navigation pane, under the **External resources** section, choose **Gateways**.
2. In the **Gateways** section, choose a gateway by its **Gateway name**.
3. To edit the gateway name, choose **Edit**.
4. To delete the gateway, choose **Delete**, then choose **Delete gateway**.

   You cannot reactivate a deleted gateway. If you want to connect to the hypervisor again, follow the procedure in Creating a gateway .
5. To connect to a hypervisor, in the **Connected hypervisor** section, choose **Connect**.

   Each gateway connects to a single hypervisor. However, you can connect multiple gateways to the same hypervisor to increase the bandwidth between them beyond that of the first gateway.
6. To assign, edit, or manage tags, in the **Tags** section, choose **Manage tags**.

Backup gateway Bandwidth Throttling

> **Note**
>
> This feature will be available on new gateways deployed after December 15, 2022. For existing gateways, this new capability will be available through an automatic software update on or before January 30, 2023. To update the gateway to the latest version manually, use AWS CLI command `UpdateGatewaySoftwareNow`.

You can limit the upload throughput from your gateway to AWS Backup to control the amount of network bandwidth the gateway uses. By default, an activated gateway has no rate limits.

You can configure a bandwidth rate-limit schedule using the AWS Backup Console or using API through the AWS CLI (`PutBandwidthRateLimitSchedule`). When you use a bandwidth rate limit schedule, you can configure limits to change automatically throughout the day or week.

Bandwidth rate limiting works by balancing the throughput of all data being uploaded, averaged over each second. While it is possible for uploads to cross the bandwidth rate limit briefly for any given micro- or millisecond, this does not typically result in large spikes over longer periods of time.

You can add up to a maximum of 20 intervals. The maximum value for the upload rate is 8,000,000 Mbps.

View and edit the bandwidth rate-limit schedule for your gateway using the AWS Backup console.

This section describes how to view and edit the bandwidth rate limit schedule for your gateway.

To view and edit the bandwidth rate limit schedule

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the left navigation pane, choose **Gateways**. In the Gateways pane, gateways are displayed by name. Click the radio button adjacent to the gateway name you want to manage.
3. Once you select a radio button, the drop-down menu **Action** is available to click. Click **Actions**, then click **Edit bandwidth rate limit schedule**. The current schedule is displayed. By default, a new or unedited gateway has no defined bandwidth rate limits.

> **Note**
>
> You can also click **Manage schedule** in the gateway details page to navigate to the Edit bandwidth page.

4. *(Optional)* Choose **Add interval** to add a new configurable interval to the schedule. For each interval, input the following information:

   a. **Days of week** — Select the recurring day or days on which you want the interval to apply. When chosen, the days will display below the drop-down menu. You can remove them by clicking the **X** next to the day.

   b. **Start time** — Enter the start time for the bandwidth interval, using the *HH:MM* 24-hour format. Time is rendered in Universal Coordinated Time (UTC).

      Note: Your bandwidth-rate-limit interval begins at the start of the specified minute.

   c. **End time** — Enter the end time for the bandwidth interval, using the *HH:MM* 24-hour format. Time is rendered in Universal Coordinated Time (UTC).

      > **Important**
      >
      > The bandwidth-rate-limit interval ends at the end of the minute specified. To schedule an interval that ends at the end of an hour, enter 59. To schedule consecutive continuous intervals, transitioning at the start of the hour, with no interruption between the intervals, enter 59 for the end minute of the first interval. Enter `00` for the start minute of the succeeding interval.

   d. **Upload rate** — Enter the upload rate limit, in megabits per second (Mbps). The minimum value is 102 megabytes per second (Mbps).

5. *(Optional)* Repeat the previous step as desired until your bandwidth rate-limit schedule is complete. If you need to delete an interval from your schedule, choose **Remove**.

> **Important**
>
> Bandwidth rate-limit intervals cannot overlap. The start time of an interval must occur after the end time of a preceding interval and before the start time of a following interval; its end time must occur before the start time of the following interval.

6. When you are finished, click the **Save changes** button.

View and edit the bandwidth rate-limit schedule for your gateway using AWS CLI.

The GetBandwidthRateLimitSchedule action can be used to view the bandwidth throttle schedule for a specified gateway. If there is no schedule set, the schedule will be an empty list of intervals. Here is an example using the AWS CLI to fetch the bandwidth schedule of a gateway:

```
aws backup-gateway get-bandwidth-rate-limit-schedule --gateway-arn "arn:aws:backup-gateway:region:account-id:gateway/bgw-gw id"
```

To edit a gateway's bandwidth throttle schedule, you can use the PutBandwidthRateLimitSchedule action. Note that you can only update a gateway's schedule as a whole, rather than modifying, adding, or removing individual intervals. Calling this action will overwrite the gateway's previous bandwidth throttle schedule.

```
aws backup-gateway put-bandwidth-rate-limit-schedule --gateway-arn "arn:aws:backup-gateway:region:account-id:gateway/gw-id" --bandwidth-rate-limit-intervals ...
```

Working with hypervisors

After you finish Creating a gateway, you can connect it to a hypervisor to enable AWS Backup to work with the virtual machines managed by that hypervisor. For example, the hypervisor for VMware VMs is VMware vCenter Server. Ensure your hypervisor is configured with the necessary permissions for AWS Backup.

Adding a hypervisor

To add a hypervisor:

1. In the left navigation pane, under the **External resources** section, choose **Hypervisors**.
2. Choose **Add hypervisor**.
3. In the **Hypervisor settings** section, type in a **Hypervisor name**.
4. For **vCenter server host**, use the dropdown menu to select either **IP address** or **FQDN** (fully-qualified domain name). Type in the corresponding value.
5. To allow AWS Backup to discover the virtual machines on the hypervisor, enter the hypervisor's **Username** and **Password**.

6. Encrypt your password. You can specify this encryption by selecting a specific service-managed KMS key or a customer-managed KMS key using the dropdown menu or choose **Create KMS key**. If you do not select a specific key, AWS Backup will encrypt your password using a service-owned key.
7. In the **Connecting gateway** section, use the dropdown list to specify which Gateway to connect to your hypervisor.
8. Choose **Test gateway connection** to verify your previous inputs.
9. *Optionally*, in the **Hypervisor tags** section, you can assign tags to the hypervisor by choosing **Add new tag**.
10. *Optional* **VMware tag mapping**: You can add up to 10 VMware tags you currently use on your virtual machines to generate AWS tags.
11. In the **Log group setting** panel, you may choose to integrate with Amazon CloudWatch Logs to maintain logs of your hypervisor (standard CloudWatch Logs pricing will apply based on usage). Each hypervisor can belong to one log group.
    a. If you have not yet created a log group, select the **Create a new log group** radio button. The hypervisor you are editing will be associated with this log group.
    b. If you have previously created a log group for a different hypervisor, you can use that log group for this hypervisor. Select **Use an existing log group**.
    c. If you do not want CloudWatch logging, select **Deactivate logging**.
12. Choose **Add hypervisor**, which takes you to its detail page.

> **Tip**
>
> You can use Amazon CloudWatch Logs (see step 11 above) to obtain information about your hypervisor, including error monitoring, network connection between the gateway and the hypervisor, and network configuration information. For information about CloudWatch log groups, see Working with Log Groups and Log Streams in the *Amazon CloudWatch User Guide*.

Viewing virtual machines managed by a hypervisor

To view virtual machines on a hypervisor:

1. In the left navigation pane, under the **External resources** section, choose **Hypervisors**.

2. In the **Hypervisors** section, choose a hypervisor by its **Hypervisor name** to go to its detail page.
3. In the section under **Hypervisor summary**, choose the **Virtual machines** tab.
4. In the **Connected virtual machines** section, a list of virtual machines populates automatically.

Viewing gateways connected to a hypervisor

To view gateways connected to the hypervisor:

1. Choose the **Gateways** tab.
2. In the **Connected gateways** section, a list of gateways populates automatically.

Connecting a hypervisor to additional gateways

Your backup and restore speeds might be limited by the bandwidth of the connection between your gateway and hypervisor. You can increase these speeds by connecting one or more additional gateways to your hypervisor. You can do this in the **Connected gateways** section as follows:

1. Choose **Connect**.
2. Select another gateway using the dropdown menu. Alternatively, choose **Create gateway** to create a new gateway.
3. Choose **Connect**.

Editing a hypervisor configuration

If you do not use the **Test gateway connection** feature, you might add a hypervisor with an incorrect username or password. In that case, the hypervisor's connection status is always `Pending`. Alternatively, you might rotate the username or password to access your hypervisor. Update this information using the following procedure:

To edit an already-added hypervisor:

1. In the left navigation pane, under the **External resources** section, choose **Hypervisors**.
2. In the **Hypervisors** section, choose a hypervisor by its **Hypervisor name** to go to its detail page.
3. Choose **Edit**.

4. The top panel is named **Hypervisor settings**.

   a. Under **vCenter server host**, you can also edit the FQDN (Fully-Qualified Domain Name) or the IP address.

   b. *Optionally,* enter the hypervisor's **Username** and **Password**.

5. In the **Log group setting** panel, you may choose to integrate with Amazon CloudWatch to maintain logs of your hypervisor (standard CloudWatch pricing will apply based on usage). Each hypervisor can belong to one log group.

   a. If you have not yet created a log group, select the **Create a new log group** radio button. The hypervisor you are editing will be associated with this log group.

   b. If you have previously created a log group for a different hypervisor, you can use that log group for this hypervisor. Select **Use an existing log group**.

   c. If you do not want CloudWatch logging, select **Deactivate logging**.

> **Tip**
>
> You can use Amazon CloudWatch Logs (see step 5 above) to obtain information about your hypervisor, including error monitoring, network connection between the gateway and the hypervisor, and network configuration information. For information about CloudWatch log groups, see Working with Log Groups and Log Streams in the *Amazon CloudWatch User Guide*.

To update a hypervisor programmatically, use the CLI command update-hypervisor and UpdateHypervisor API call.

Deleting a hypervisor configuration

If you need to remove an already-added hypervisor, remove the hypervisor configuration and add another. This remove operation applies to the configuration to connect to the hypervisor. It does not delete the hypervisor.

To delete the configuration to connect to an already-added hypervisor:

1. In the left navigation pane, under the **External resources** section, choose **Hypervisors**.
2. In the **Hypervisors** section, choose a hypervisor by its **Hypervisor name** to go to its detail page.

3. Choose **Remove**, then choose **Remove hypervisor**.
4. Optional: replace the removed hypervisor configuration using the procedure for Adding a hypervisor.

Understanding hypervisor status

The following describes each of the possible hypevisor statuses and, if applicable, remediation steps. The `ONLINE` status is the normal status of the hypervisor. A hypervisor should have this status all or most of the time it's in use for backup and recovery of VMs managed by the hypervisor.

Hypervisor statuses

Status	Meaning and remediation
`ONLINE`	You added a hypervisor to AWS Backup, associated with it a gateway, and can connect with that gateway over your network to perform backup and recovery of virtual machines managed by the hypervisor. You can perform on-demand and scheduled backups of those virtual machines at any time.
`PENDING`	You added a hypervisor to AWS Backup but: • It is not associated with any gateway, or • It is associated with one or more gateways, but all those gateways were deleted or are otherwise not active. To change a hypervisor status from `PENDING` to `ONLINE`, create a gateway and connect your hypervisor to that gateway.
`OFFLINE`	You added a hypervisor to AWS Backup and associated it with a gateway, but the gateway

Status	Meaning and remediation
	cannot connect to the hypervisor over your network. To change a hypervisor status from `OFFLINE` to `ONLINE`, verify the correctness of your network configuration. If the issue persists, verify that your hypervisor's IP address or fully-qualified domain name is correct. If they are incorrect, add your hypervisor again using the correct information and test your gateway connection.
`ERROR`	You added a hypervisor to AWS Backup and associated it with a gateway, but the gateway cannot communicate with the hypervisor. To change a hypervisor status from `ERROR` to `ONLINE`, verify that hypervisor's username and password are correct. If they are incorrect, edit your hypervisor configuration.

Next steps

To back up virtual machines on your hypervisor, see Backing up virtual machines.

Backing up virtual machines

After Adding a hypervisor, Backup gateway automatically lists your virtual machines. You can view your virtual machines by choosing either **Hypervisors** or **Virtual machines** in the left navigation pane.

- Choose **Hypervisors** to view only the virtual machines managed by a specific hypervisor. With this view, you can work with one virtual machine at a time.
- Choose **Virtual machines** to view all the virtual machines across all the hypervisors you added to your AWS account. With this view, you can work with some or all your virtual machines across multiple hypervisors.

Regardless of which view you choose, to perform a backup operation on a specific virtual machine, choose its **VM name** to open its detail page. The VM detail page is the starting point for the following procedures.

Creating an on-demand backup of a virtual machine

An on-demand backup is a one-time, full backup you manually initiate. You can use on-demand backups to test AWS Backup's backup and restore capabilities.

To create an on-demand backup of a virtual machine:

1. Choose **Create on-demand backup**.
2. Configure your on-demand backup.
3. Choose **Create on-demand backup**.
4. Check when your backup job has the status `Completed`. In the left navigation menu, choose **Jobs**.
5. Choose the **Backup Job ID** to view backup job information such as the **Backup size** and time elapsed between the **Creation date** and **Completion date**.

Incremental VM backups

Newer VMware versions contain a feature called Changed Block Tracking, which keeps track of the storage blocks of virtual machines as they change over time. When you use AWS Backup to back up a virtual machine, AWS Backup attempts to use the CBT data if it is available. AWS Backup uses CBT data to speed up the backup process; without CBT data, backup jobs are often slower and use more hypervisor resources. The backup can still be successfully completed even when the CBT data is not valid or available. For example, the CBT data might not be valid or might be unavailable if the virtual machine or ESXi host experiences a hard shutdown.

On the occasions CBT data is invalid or unavailable, the backup status will read `Successful` with a message. In these cases, the message will indicate that, in the absence of CBT data, AWS Backup used its own proprietary change detection mechanism to complete the backup instead of VMware's CBT data. Subsequent backups will reattempt to use CBT data, and in most cases the CBT data will be successfully valid and available. If the issue persists, see VMware Troubleshooting for steps to remedy.

For CBT to function correctly, the following must be true:

- Host needs to be ESXi 4.0 or later
- The VM owning the disks must have hardware version 7 or later
- CBT must be enabled for the virtual machine (it is enabled by default)

> **Note**
>
> The ESXi and hardware version requirements listed above are VMware's minimum requirements for the CBT feature to function. They do not represent the vSphere versions supported by AWS Backup. For the list of vSphere versions supported by AWS Backup, see Supported VMs.

To verify if a virtual disk has CBT enabled:

1. Open the vSphere Client and select a powered-off virtual machine.
2. Right-click the virtual machine and navigate to **Edit Settings** > **Options** > **Advanced/General** > **Configuration Parameters**.
3. The option `ctkEnabled` needs to equal `True`.

Automating virtual machine backup by assigning resources to a backup plan

A backup plan is a user-defined data protection policy that automates data protection across many AWS services and third-party applications. You first create your backup plan by specifying its backup frequency, retention period, lifecycle policy, and many other options. To create a backup plan, see Getting started tutorial.

After you create your backup plan, you assign AWS Backup-supported resources, including virtual machines, to that backup plan. AWS Backup offers many ways to assign resources, including assigning all the resources in your account, including or excluding single specific resources, or adding resources with certain tags.

In addition to its existing resource assignment features, AWS Backup support for virtual machines introduces several new features to help you quickly assign virtual machines to backup plans. From the **Virtual machines** page, you can assign tags to multiple virtual machines or use the new **Assign resources to plan** feature. Use these features to assign your virtual machines already discovered by AWS Backup gateway.

If you anticipate discovering and assigning additional virtual machines in the future, and would like to automate the resource assignment step to include those future virtual machines, use the new **Create group assignment** feature.

VMware Tags

Tags are key-value pairs you can use to manage, to filter, and to search for your resources.

A VMware tag is composed of a **category** and a **tag name**. VMware tags are used to group virtual machines. A tag name is a label assigned to a virtual machine. A category is a collection of tag names.

In AWS tags, you can use characters among UTF-8 letters, numbers, spaces, and special characters `+ - = . _ : / .`.

If you use tags on your virtual machines, you can add up to 10 matching tags in AWS Backup to help with organization. You can map up to 10 VMware tags to AWS tags. In the AWS Backup console, these can be found in **External resources > Virtual Machines > AWS tags** or **VMware tags**.

VMware tag mapping

If you use tags on your virtual machines, you can add up to 10 matching tags in AWS Backup for additional clarity and organization. Mappings apply to any virtual machine on the hypervisor.

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the console, go to **edit Hypervisor** (Click **External resources**, then **Hypervisors**, then click the Hypervisor name, then click **Manage mappings**).
3. The last pane, **VMware tag mapping**, contains four textbox fields into which you can enter your VMware tag information into corresponding AWS tags. The four fields are **Vmware tag category**, **VMware tag name**, **AWS tag key**, and **AWS tag value** (*example: Category = OS; Tag name = Windows; AWS tag key = OS-Windows, and AWS tag value = Windows)*.
4. After you have entered your preferred values, click **Add mapping**. If you make an error, you can click **Remove** to delete entered information.
5. After adding mapping(s), specify the IAM role you intend to use to apply these AWS tags to the VMware virtual machines.

   The policy `AWSBackupGatewayServiceRolePolicyForVirtualMachineMetadataSync` contains needed permissions. You can attach this policy to the role you are using (or have an administrator attached it) or you can create a custom policy for the role being used.

6. Lastly, click **Add hypervisor** or **Save**.

The IAM role trust relationship should be modified to add the backup-gateway.amazonaws.com and backup.amazonaws.com services. Without this service, you will likely experience an error when you map tags. To edit the trust relationship for an existing role,

1. Log into the IAM console.
2. In the navigation pane of the console, choose **Roles**.
3. Choose the name of the role you wish to modify, then select the **Trust relationships** tab on the details page.
4. Under **Policy Document, paste the following:**

   JSON

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "backup.amazonaws.com",
          "backup-gateway.amazonaws.com"
        ]
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

5. Choose **Update Trust Policy**.

See Editing the trust relationship for an existing role in the *AWS Directory Service Administration Guide* for more detail.

View VMware tag mappings

In the AWS Backup console, click on **External Resources**, then click on **Hypervisors**, then click on the Hypervisor name link to view properties for the selected hypervisor. Under the summary

pane, there are four tabs, the last of which is **VMware tag mappings**. Note if you do not yet have mappings, "No VMware tag mappings." will be displayed.

From here, you can sync the metadata of virtual machines discovered by the hypervisor, you can copy mappings to your hypervisor(s), you can add AWS tags mapped to teh VMware tags to the backup selection of a backup plan, or you can manage mappings.

In the console, to see which tags are applied to a selected virtual machine, click **Virtual machines**, then the virtual machine name, then **AWS tags** or **VMware tags**. You can view the tags associated with this virtual machine, and additionally you can manage the tags.

Assign virtual machines to plan using VMware tag mappings

To assign virtual machines to a backup plan using mapped tags, do the following:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the console go to VMware tag mappings on the hypervisor details page (click **External resources**, then click **Hypervisors** then click on the hypervisor name).
3. Select the checkbox next to multiple mapped tags to assign those tags to the same backup plan.
4. Click **Add to resource assignment**.
5. Choose an existing **Backup plan** from the dropdown list. Alternatively, you can choose **Create backup plan** to create a new backup plan.
6. Click **Confirm**. This opens the **Assign resources** page with **Refine selection using tags** fields with values pre-populated.

VMware tags using the AWS CLI

AWS Backup uses the API call `PutHypervisorPropertyMappings` to map hypervisor entity properties in on-premise to properties in AWS.

In the AWS CLI, use the operation `put-hypervisor-property-mappings`:

```
aws backup-gateway put-hypervisor-property-mappings \
--hypervisor-arn arn:aws:backup-gateway:region:account:hypervisor/hypervisorId \
--vmware-to-aws-tag-mappings list of VMware to AWS tag mappings \
--iam-role-arn arn:aws:iam::account:role/roleName \
--region AWSRegion
--endpoint-url URL
```

Here is an example:

```
aws backup-gateway put-hypervisor-property-mappings \
--hypervisor-arn arn:aws:backup-gateway:us-east-1:123456789012:hypervisor/hype-12345 \
--vmware-to-aws-tag-mappings VmwareCategory=OS,VmwareTagName=Windows,AwsTagKey=OS-
Windows,AwsTagValue=Windows \
--iam-role-arn arn:aws:iam::123456789012:role/SyncRole \
--region us-east-1
```

You can also use GetHypervisorPropertyMappings to assist with property mappings information. In the AWS CLI, use the operation `get-hypervisor-property-mappings`. Here is an example template:

```
aws backup-gateway get-hypervisor-property-mappings --hypervisor-arn HypervisorARN
--region AWSRegion
```

Here is an example:

```
aws backup-gateway get-hypervisor-property-mappings \
--hypervisor-arn arn:aws:backup-gateway:us-east-1:123456789012:hypervisor/hype-12345 \
--region us-east-1
```

Sync metadata of virtual machines discovered by the hypervisor in AWS using API, CLI, or SDK

You can sync the metadata of virtual machines. When you do, the VMware tags present on the virtual machine that are part of the mappings will be synched. Also, AWS tags mapped to the VMware tags present on the virtual machine will be applied to the AWS Virtual Machine resource.

AWS Backup uses the API call StartVirtualMachinesMetadataSync to sync the metadata of the virtual machines discovered by the hypervisor. To sync metadata of virtual machines discovered by the hypervisor using AWS CLI, use the operation `start-virtual-machines-metadata-sync`.

Example template:

```
aws backup-gateway start-virtual-machines-metadata-sync \
--hypervisor-arn Hypervisor ARN
--region AWSRegion
```

Example:

```
aws backup-gateway start-virtual-machines-metadata-sync \
--hypervisor-arn arn:aws:backup-gateway:us-east-1:123456789012:hypervisor/hype-12345 \
--region us-east-1
```

You can also use GetHypervisor to assist with hypervisor information, such as host, state, status of latest metadata sync, and also to retrieve the last successful metadata sync time. In the AWS CLI, use the operation `get-hypervisor`.

Example template:

```
aws backup-gateway get-hypervisor \
--hypervisor-arn Hypervisor ARN
--region AWSRegion
```

Example:

```
aws backup-gateway get-hypervisor \
--hypervisor-arn arn:aws:backup-gateway:us-east-1:123456789012:hypervisor/hype-12345 \
--region us-east-1
```

For more information, see API documentation VmwareTag and VmwareToAwsTagMapping.

This feature will be available on new gateways deployed after December 15, 2022. For existing gateways, this new capability will be available through an automatic software update on or before January 30, 2023. To update the gateway to the latest version manually, use AWS CLI command UpdateGatewaySoftwareNow.

Example:

```
aws backup-gateway update-gateway-software-now \
--gateway-arn arn:aws:backup-gateway:us-east-1:123456789012:gateway/bgw-12345 \
--region us-east-1
```

Assigning virtual machines using tags

You can assign your virtual machines currently discovered by AWS Backup, along with other AWS Backup resources, by assigning them a tag that you have already assigned to one of your existing backup plans. You can also create a new backup plan and a new tag-based resource assignment. Backup plans check for newly-assigned resources each time they run a backup job.

To tag multiple virtual machines with the same tag:

1. In the left navigation pane, choose **Virtual machines**.
2. Select the checkbox next to **VM name** to choose all your virtual machines. Alternatively, select the checkbox next to the VM names you want to tag.
3. Choose **Add tags**.
4. Type in a tag **Key**.
5. Recommended: type in a tag **Value**.
6. Choose **Confirm**.

Assigning virtual machines using the Assign resources to plan feature

You can assign virtual machines currently discovered by AWS Backup to an existing or new backup plan using the **Assign resources to plan** feature.

To assign virtual machines using the Assign resources to plan feature:

1. In the left navigation pane, choose **Virtual machines**.
2. Select the checkbox next to **VM name** to choose all your virtual machines. Alternatively, select the checkbox next to multiple VM names to assign them to the same backup plan.
3. Choose **Assignments**, then choose **Assign resources to plan**.
4. Type in a **Resource assignment name**.
5. Choose a resource assignment **IAM role** to create backups and manage recovery points. If you do not have a specific IAM role to use, we recommend the **Default role** which has the correct permissions.

   > **Note**
   >
   > If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

6. In the **Backup plan** section, choose an existing **Backup plan** from the dropdown list. Alternatively, choose **Create backup plan** to create a new backup plan.
7. Choose **Assign resources**.

8. Optional: Verify your virtual machines are assigned to a backup plan by choosing **View Backup plan**. Then, in the **Resource assignments** section, choose the resource assignment **Name**.

Assigning virtual machines using the Create group assignment feature

Unlike the preceding two resource assignment features for virtual machines, the **Create group assignment** feature not only assigns virtual machines currently discovered by AWS Backup, but also virtual machines discovered in the future in a folder or hypervisor you define.

Also, you do not need to select any checkboxes to use the **Create group assignment** feature.

To assign virtual machines using the Assign resources to plan feature:

1. In the left navigation pane, choose **Virtual machines**.
2. Choose **Assignments**, then choose **Create group assignment**.
3. Type in a **Resource assignment name**.
4. Choose a resource assignment **IAM role** to create backups and manage recovery points. If you do not have a specific IAM role to use, we recommend the **Default role** which has the correct permissions.

   > **Note**
   >
   > If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

5. In the **Resource group** section, select the **Group type** dropdown menu. Your options are **Folder** or **Hypervisor**.

   a. Choose **Folder** to assign all the virtual machines in a folder on a hypervisor. Select a folder **Group name**, such as `datacenter/vm`, using the dropdown menu. You can also choose to include **Subfolders**.

      > **Note**
      >
      > To make Folder-based assignments, during the discovery process, AWS Backup tags virtual machines with the folder it finds them in during the discovery process. If you later move a virtual machine to a different folder, AWS Backup cannot

update the tag for you due to AWS tagging best practices. This assignment method might result in continuing to take backups of virtual machines you moved out of your assigned folder.

   b. Choose **Hypervisor** to assign all the virtual machines managed by a hypervisor. Select a hypervisor ID **Group name** using the dropdown menu.

6. In the **Backup plan** section, choose an existing **Backup plan** from the dropdown list. Alternatively, choose **Create backup plan** to create a new backup plan.
7. Choose **Create group assignment**.
8. Optional: verify your virtual machines are assigned to a backup plan by choosing **View Backup plan**. In the **Resource assignments** section, choose the resource assignment **Name**.

Next steps

To restore a virtual machine, see Restore a virtual machine using AWS Backup.

Information about third-party source components for Backup gateway

In this section, you can find information about third party tools and licenses that we depend on to deliver Backup gateway functionality.

The source code for certain third-party source software components that are included with the Backup gateway software is available for download at the following locations:

- For gateways deployed on VMware ESXi, download sources.tgz.

This product includes software developed by the OpenSSL project for use in the OpenSSL Toolkit (https://www.openssl.org/).

This product includes software developed by VMware® vSphere Software Development Kit (https://www.vmware.com).

For the relevant licenses for all dependent third-party tools, see Third-Party Licenses.

Open-source components for AWS Appliance

Several third-party tools and licenses are used to deliver functionality for Backup gateway.

Use the following links to download source code for certain open-source software components that are included with AWS Appliance software:

- For gateways deployed on VMware ESXi, download sources.tar

This product includes software developed by the OpenSSL project for use in the OpenSSL Toolkit (https://www.openssl.org/). For the relevant licenses for all dependent third-party tools, see Third-Party Licenses.

Troubleshoot VM issues

Incremental Backups / CBT issues and messages

Failure message: `"The VMware Change Block Tracking (CBT) data was invalid during this backup, but the incremental backup was successfully completed with our proprietary change detection mechanism."`

If this message continues, reset CBT as directed by VMware.

Message notes CBT was not turned on or was unavailable: *"VMware Change Block Tracking (CBT) was not available for this virtual machine, but the incremental backup was successfully completed with our proprietary change mechanism."*

Check to make sure CBT is turned on. To verify if a virtual disk has CBT enabled:

1. Open the vSphere Client and select a powered-off virtual machine.
2. Right-click the virtual machine and navigate to **Edit Settings** > **Options** > **Advanced/General** > **Configuration Parameters**.
3. The option `ctkEnabled` needs to equal `True`.

If it is turned on, ensure you are using up-to-date VMware features. The host must be ESXi 4.0 or later and the virtual machine owning the disks to be tracked must be hardware version 7 or later.

> **Note**
>
> The ESXi and hardware version requirements above are VMware's minimum requirements for the CBT feature to function. They do not represent the vSphere versions supported by AWS Backup. For the list of vSphere versions supported by AWS Backup, see Supported VMs.

If CBT is turned on (enabled) and the software and hardware are up to date, turn off the virtual machine and then turn it back on again. Ensure that CBT is turned on. Then, perform the backup again.

VMware backup failure

When a VMware backup fails, it may be related to one of the following:

Failure message: `"Failed to process backup data. Aborted backup job."` or `"Error opening disk on the virtual machine"`.

Possible causes: This error may occur because of a configuration issue; or, the VMware version or disk isn't supported.

Remedy 1: Ensure your infrastructure is configured to use a gateway and ensure all required ports are open.

1. Access the backup gateway console. Note this is different from the AWS Backup console.
2. On the **Backup gateway configuration** page enter option **3** to test the network connectivity.
3. If the network test is successful, enter **X**.
4. Return to the Backup gateway configuration page.
5. Enter **7** to access the command prompt.
6. Run the following commands to verify network connectivity:

```
ncport -d ESXi Host-p 902

ncport -d ESXi Host-p 443
```

Remedy 2: Use Supported VMs versions.

Remedy 3: If a gateway appliance is configured with incorrect DNS servers, then the backup fails. To verify the DNS configuration, complete the following steps:

1. Access the backup gateway console.
2. On the **Backup gateway configuration** page enter option **2** to navigate to the network configuration.
3. In **Network configuration**, enter **7** to view the DNS configuration.
4. Review the DNS server IP addresses. If the DNS server IP address are incorrect then exist the prompt to return to **Network Configuration**.

5. In **Network Configuration**, enter **6** to edit the DNS configuration.
6. Enter the correct DNS server IP addresses. Then, enter **X** to complete your network configuration.

To obtain more information about your hypervisor, such as errors and network configuration and connection, see Editing a hypervisor configuration to configure the hypervisor to integrate with Amazon CloudWatch Logs.

Backup failures from network connection issues

Failure message: `"Failed to upload backup during data ingestion. Aborted backup job."` or `"Cloud network request timed out during data ingestion"`.

Possible causes: This error can occur if the network connection is insufficient to handle data uploads. If network bandwidth is low, the link between the VM and AWS Backup can become congested and cause backups to fail.

Required network bandwidth depends on several factors, including the size of the VM, the incremental data generated for each VM backup, the backup window, and restore requirements.

Remedy: Best practices and recommendations include having a minimum bandwidth of 100 Mbps upload bandwidth for on-premises VMs connected to AWS Backup. This is the minimum required for the gateway to function. Bandwidth below 100 Mbps prevents the gateway from completing any backup or restore job. With a large number of VMs, backup jobs might still fail due to timeout even with bandwidth above 100 Mbps. Consider increasing bandwidth based on the number of VMs and the size of data being backed up. Once the bandwidth is confirmed, retry the backup job.

Aborted backup job

Failure message: `"Failed to create backup during snapshot creation. Aborted backup job."`

Possible cause: The VMware host where the gateway appliance resides may have an issue.

Remedy: Check the configuration of your VMware host and review the it for issues. For additional information, see Editing a hypervisor configuration.

No available gateways

Failure message: `"No gateways available to work on job."`

Possible cause: all connected gateways are busy with other jobs. Each gateway has a limit of four concurrent jobs (backup or restore).

For **remedies**, see the next section for steps on increasing number of gateways and steps to increase backup plan window time.

VMware backup job failure

Failure message: `"Abort signal detected"`

Possible causes:

- **Low Network Bandwidth**: Insufficient network bandwidth can impede the completion of backups within the completion window. When the backup job requires more bandwidth than available, it can result in failure and trigger the "Abort Signal Detected" error.
- **Inadequate Number of Backup Gateways**: If the number of backup gateways is not sufficient to handle the backup rotation for all the configured VMs, the backup job may fail. This can occur when the backup plan's window for completing backups is too short or the number of backup gateways are not enough.
- Backup Plan completion window is too small.

Remedies:

Increase bandwidth: Consider increasing the network bandwidth between AWS and the on-premises environment. This step will provide more bandwidth for the backup process, allowing data to transfer smoothly without triggering the error. You must have at least 100 Mbps bandwidth to AWS to back up on-premises VMware VMs using AWS Backup. This is the minimum required for the gateway to complete any backup or restore job. If you have a large number of VMs, consider increasing bandwidth beyond 100 Mbps, as jobs might still time out with insufficient bandwidth.

If a bandwidth rate limit is configured for the backup gateway, it can restrict the flow of data and lead to backup failures. Increasing the bandwidth rate limit to ensure sufficient data transfer capacity may help reduce failures. This adjustment can mitigate the occurrence of the "Abort Signal Detected" error. For more information, see Backup gateway Bandwidth Throttling.

Increase the number of Backup gateways: A single backup gateway can process up to 4 backup and restore jobs at a time. Additional jobs will queue and wait for the gateway to free up until the backup start window passed. If the backup window passes and the queued jobs have not started,

those backup jobs will fail with "abort signal detected". You can increase the number of backup gateways to alleviate the number of failed jobs. See Working with gateways for more detail.

Increase backup plan window time: You can increase the **complete within duration** of the backup window in your backup plan. See Backup plan options and configuration for more detail.

Restore to EC2 fails due to missing ENA drivers

Failure message: `"Restore failed because the instance type selected for restore requires that the VirtualMachine image have Elastic Network Adapter (ENA) drivers installed, but no ENA drivers were present in the backed up VM."`

Possible cause: Your restore job failed because you are attempting to restore a VMware backup that does not have ENA drivers installed to an Amazon EC2 instance type that requires Elastic Network Adapters (ENA).

Remedies:

Restore to an instance type that does not require ENA: Restore this recovery point to an instance type that does not require ENA. For a list of instance types and their requirements, see Amazon EC2 instance types.

Install ENA drivers and take a new backup: To restore to instance types that require ENA, install ENA drivers on the virtual machine first, take a new backup, and then restore from that backup. For more information about driver requirements, see VM Import/Export requirements.

For help resolving these issues, see AWS Knowledge Center.

Create Windows VSS backups

With AWS Backup, you can back up and restore VSS (Volume Shadow Copy Service)-enabled Windows applications running on Amazon EC2 instances. If the application has VSS writer registered with Windows VSS, then AWS Backup creates a snapshot that will be consistent for that application.

You can perform consistent restores, while using the same managed backup service that is used to protect other AWS resources. With application-consistent Windows backups on EC2, you get the same consistency settings and application awareness as traditional backup tools.

> **Note**
>
> AWS Backup only supports application-consistent backups of resources running on Amazon EC2, specifically backup scenarios where application data can be restored by replacing an existing instance with a new instance created from the backup. Not all instance types or applications are supported for Windows VSS backups.

For more information, see Create VSS based snapshots in the *Amazon EC2 User Guide*.

To back up and restore VSS-enabled Windows resources running Amazon EC2, follow these steps to complete the required prerequisite tasks. For instructions, see Prerequisites to create Windows VSS based EBS snapshots in the *Amazon EC2 User Guide*.

1. Download, install, and configure the SSM agent in AWS Systems Manager. This step is required. For instructions, see Working with SSM agent on EC2 instances for Windows Server in the *AWS Systems Manager User Guide*.
2. Add an IAM policy to the IAM role and attach the role to the Amazon EC2 instance before you take the Windows VSS (Volume Shadow Copy Service) backup. For instructions, see Use an IAM managed policy to grant permissions for VSS based snapshots in the *Amazon EC2 User Guide*. For an example of the IAM policy, see Managed policies for AWS Backup.
3. Download and install VSS components to the Windows on Amazon EC2 instance
4. Enable VSS in AWS Backup:
    1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
    2. On the dashboard, choose the type of backup you want to create, either **Create an on-demand backup** or **Manage Backup plans**. Provide the information needed for your backup type.
    3. When you're assigning resources, choose **EC2**. Windows VSS backup is currently supported for EC2 instances only.
    4. In the **Advanced settings** section, choose **Windows VSS**. This enables you to take application-consistent Windows VSS backups.
    5. Create your backup.

A backup job with a status of `Completed` does not guarantee that the VSS portion is successful; VSS inclusion is made on a best-effort basis. Proceed with the following steps to determine if a backup is application-consistent, crash-consistent, or failed:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Under **My account** in the left navigation, click **Jobs**.
3. A status of `Completed` indicates a successful job that is application-consistent (VSS).

   A status of `Completed with issues` indicates that the VSS operation has failed, so only a crash-consistent backup has been successful. This status will also have a popover message `"Windows VSS Backup Job Error encountered, trying for regular backup"`.

   If the backup was unsuccessful, the status will be `Failed`.
4. To view additional details of the backup job, click on the individual job. For example, the details may read `Windows VSS Backup attempt failed because of timeout on VSS enabled snapshot creation`.

VSS-enabled backups with a target that is non-Windows or non-VSS component Windows that is successful job will be crash-consistent without VSS.

Unsupported Amazon EC2 instances

The following Amazon EC2 instance types are not supported for VSS-enabled Windows backups because they are small instances and might not take the backup successfully.

- t3.nano
- t3.micro
- t3a.nano
- t3a.micro
- t2.nano
- t2.micro

Backup and tag copy

You can copy backups to multiple AWS accounts or AWS Regions on demand or automatically as part of a scheduled backup plan for most resource types, though backups in cold storage or archive

tiers cannot be copied. See the section called "Feature availability by resource" and Encryption for a backup copy to a different account or AWS Region for details.

> **Note**
>
> If a resource type does not support cross-Region or cross-account copy in AWS Backup, then copy operations for that resource type are not supported in any scenario, including same-Region and same-account copies to a different vault. To determine which resource types support copy operations, see the Cross-Region backup and Cross-account backup columns in Feature availability by resource. If the underlying AWS service provides its own native copy or replication feature, refer to that service's documentation.

> **Note**
>
> AWS Backup does not provide any service-level agreements (SLAs) for copy job completion times. Copy times can vary based upon system load, capacity, and the size of the resources being copied, even for copy jobs containing the same resources.

Some resource types have both continuous backup capability and cross-Region and cross-account copy available. When a cross-Region or cross-account copy of a continuous backup is made, the copied recovery point (backup) becomes a snapshot (periodic) backup (not available for all resource types that support both backup types). Depending on the resource type, the snapshots may be an incremental copy or a full copy. PITR (Point-in-Time Restore) is not available for these copies.

> **Important**
>
> Copies retain their source configuration, including creation dates and retention period. The creation date refers to when the source was created, not when the copy was created. You can override the retention period.
> The configuration of the source backup being copied overrides its copy's expiration setting if the copy retention period is set to **Always** in the AWS Backup console (or `DeleteAfterDays` value is set to `-1` in the API request); that is, a copy with a retention setting set to never expire will retain its source recovery point's expiration date.
> If you want your backup copies to never expire, either set your source backups to never expire or specify your copy to expire 100 years after its creation.

Important

For resources that support incremental cross-Region copy and cold storage, a copy configuration with only 1-day warm retention can intermittently fall back to a full copy instead of an incremental one, because the previous destination recovery point may transition to cold storage before the next copy runs. To avoid these full copies, we recommend the copy action should be configured with a warm retention period of at least 3 days or more.

Copy job retry

AWS Backup implements the following retry strategy for copy jobs: If AWS Backup encounters any system errors, the copy job enters a retry phase that lasts for 2 hours. During this time, the copy job status remains in the `CREATED` state while the system periodically attempts to initiate the job. If the job successfully starts within this window, it transitions to the `RUNNING` state.

If issues persist beyond the 2-hour retry period, the AWS Backup service team is automatically alerted. The team then investigates and addresses any underlying problems. After the issue is resolved, they manually retry the copy request, ensuring that the copy jobs are completed as requested.

The copy job retry process differs from backup job retry process, which uses a defined start window with regular retry attempts until either success or expiration. The copy job mechanism provides an additional layer of reliability by incorporating direct service team intervention for persistent issues.

Copy job concurrency

Only one backup or copy job can run at a time for a given resource. Additional copy jobs for the same resource remain in `CREATED` status until the running job completes. For more information about concurrency limits, see AWS Backup quotas.

Copy jobs for large resources can take several hours to complete, which can result in additional copy jobs waiting in `CREATED` status. For resource types that support incremental copies, a short retention period can lead to situations where the only recovery point in the destination vault expires. If no recovery point exists in the destination vault, then the next copy must be a full copy instead of an incremental copy. To avoid this, set the copy retention period to at least one week. For more information, see Metering, costs, and billing. To determine which resource types support incremental copies, see Feature availability by resource.

If copy jobs continue to queue in CREATED status for the same resource, reduce the copy frequency.

Contents

- Creating backup copies across AWS Regions
- Creating backup copies across AWS accounts
- Copy tags onto backups

Creating backup copies across AWS Regions

Using AWS Backup, you can copy backups to multiple AWS Regions on demand or automatically as part of a scheduled backup plan. Cross-Region replication is particularly valuable if you have business continuity or compliance requirements to store backups a minimum distance away from your production data. For a video tutorial, see Managing cross-Region copies of backups.

When you copy a backup to a new AWS Region for the first time, AWS Backup copies the backup in full. In general, if a service supports incremental backups, subsequent copies of that backup in the same AWS Region will be incremental. AWS Backup will re-encrypt your copy using the customer managed key of your destination vault.

An exception is Amazon EBS, where copying a snapshot to a vault that uses a different AWS KMS encryption key results in a full (not incremental) copy. If you consistently copy to the same vault with the same encryption key, subsequent copies remain incremental.

Requirements

- Most AWS Backup-supported resources support cross-Region backup. For specifics, see Feature availability by resource.
- Most AWS Regions support cross-Region backup. For specifics, see Feature availability by AWS Region.
- AWS Backup does not support cross-Region copies for storage in cold tiers.

Cross-Region copy encryption

See Encryption for a backup copy to a different account or AWS Region for details on how encryption works for copy jobs.

Cross-Region copy considerations with specific resources

Amazon RDS

AWS Backup does not pass the option group when performing a cross-Region copy. Instead, AWS Backup copies the default option group, even if you have configured a custom option group.

If your custom option group uses persistent options, the cross-Region copy job fails unless the destination Region has the same option group as the source Region. In this case, AWS Backup still copies the default option group.

You can't copy an option group to a different AWS Region. You must manually create the same option group in the destination AWS Region before performing the cross-Region copy.

If you attempt a cross-Region copy without a matching option group in the target Region, the copy job fails with an error message such as "The snapshot requires a target option group with the following options: ...."

Performing on-demand cross-Region backup

To copy an existing backup on-demand

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Choose **Backup vaults**.
3. Choose the vault that contains the recovery point you want to copy.
4. In the **Backups** section, select a recovery point to copy.
5. Using the **Actions** dropdown button, choose **Copy**.
6. Enter the following values:

   Copy to destination

   Choose the destination AWS Region for the copy. You can add a new copy rule per copy to a new destination.

   Destination Backup vault

   Choose the destination backup vault for the copy.

Transition to cold storage

Choose when to transition the backup copy to cold storage. Backups transitioned to cold storage must be stored there for a minimum of 90 days. This value cannot be changed after a copy has transitioned to cold storage.

To see the list of resources that you can transition to cold storage, see the "Lifecycle to cold storage" section of the Feature availability by resource table. The cold storage expression is ignored for other resources.

Retention period

Choose specifies the number of days after creation that the copy is deleted. This value must be greater than 90 days beyond the **Transition to cold storage** value.

IAM role

Choose the IAM role that AWS Backup will use when creating the copy. The role must also have AWS Backup listed as a trusted entity, which enables AWS Backup to assume the role. If you choose **Default** and the AWS Backup default role is not present in your account, one will be created for you with the correct permissions.

7. Choose **Copy**.

Scheduling cross-Region backup

You can use a scheduled backup plan to copy backups across AWS Regions.

To copy a backup using a scheduled backup plan

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In **My account**, choose **Backup plans**, and then choose **Create Backup plan**.
3. On the **Create Backup plan** page, choose **Build a new plan**.
4. For **Backup plan name**, enter a name for your backup plan.
5. In the **Backup rule configuration** section, add a backup rule that defines a backup schedule, backup window, and lifecycle rules. You can add more backup rules later.
   a. For **Backup rule name**, enter a name for your rule.
   b. For **Backup vault**, choose a vault from the list. Recovery points for this backup will be saved in this vault. You can create a new backup vault.

c. For **Backup frequency**, choose how often you want to take backups.

d. For services that support PITR, if you want this feature, choose **Enable continuous backups for point-in-time recovery (PITR)**. For a list a services that support PITR, see that section of the Feature availability by resource table.

e. For **Backup window**, choose **Use backup window defaults - *recommended***. You can customize the backup window.

f. For **Copy to destination**, Choose the destination AWS Region for your backup copy. Your backup will be copied to this Region. You can add a new copy rule per copy to a new destination. Then enter the following values:

Copy to another account's vault

Do not toggle this option. To learn more about cross-account copy, see Creating backup copies across AWS accounts

Destination Backup vault

Choose the backup vault in the destination Region where AWS Backup will copy your backup.

If you would like to create a new backup vault for cross-Region copy, choose **Create new Backup vault**. Enter the information in the wizard. Then choose **Create Backup vault**.

6. Choose **Create plan**.

Creating backup copies across AWS accounts

Using AWS Backup, you can back up to multiple AWS accounts on demand or automatically as part of a scheduled backup plan. Use a cross-account backup if you want to securely copy your backups to one or more AWS accounts in your organization for operational or security reasons. If your original backup is inadvertently deleted, you can copy the backup from its destination account to its source account, and then start the restore. Before you can do this, you must have two accounts that belong to the same organization in the AWS Organizations service. For more information, see Tutorial: Creating and configuring an organization in the *Organizations User Guide*.

In your destination account, you must create a backup vault. Then, you assign an encryption key to encrypt backups in the destination account, and a resource-based access policy to allow AWS Backup to access the resources you would like to copy. In the source account, if your resources are

encrypted with a customer managed key, you must share this customer managed key with the destination account. You can then create a backup plan and choose a destination account that is part of your organizational unit in AWS Organizations.

When you copy a backup to cross-account for the first time, AWS Backup copies the backup in full. In general, if a service supports incremental backups, subsequent copies of that backup to the same destination account and vault are incremental. AWS Backup re-encrypts your copy using the encryption key of your destination vault.

> **Note**
>
> For resources fully managed by AWS Backup, you can use either a customer managed key or the AWS Backup managed key (`aws/backup`) for the destination vault. For all other resource types, a customer managed key is required because AWS managed key policies are immutable and cannot be shared cross-account. For more information, see Encryption for a backup copy to a different account or AWS Region.

Requirements

- Before you manage resources across multiple AWS accounts in AWS Backup, your accounts must belong to the same organization in the AWS Organizations service.
- Most resources supported by AWS Backup support cross-account backup. For specifics, see Feature availability by resource.
- Most AWS Regions support cross-account backup. For specifics, see Feature availability by AWS Region.
- AWS Backup does not support cross-account copies for storage in cold tiers.

Setting up cross-account backup

What do you need to create cross-account backups?

- **A source account**

  The source account is the account where your production AWS resources and primary backups reside.

The source account user initiates the cross-account backup operation. The source account user or role must have appropriate API permissions to initiate the operation. Appropriate permissions might be the AWS managed policy `AWSBackupFullAccess`, which enables full access to AWS Backup operations, or a customer managed policy that allows actions such as `ec2:ModifySnapshotAttribute`. For more information about policy types, see AWS Backup Managed Policies.

For cross-account copy jobs, the IAM role in the source account must also have the following identity-based permissions:

- `backup:CopyFromBackupVault` – Authorizes copying the source recovery point.
- `backup:CopyIntoBackupVault` – Authorizes writing into the destination vault. Because this is a cross-account call, this permission must be present in both the identity-based policy (attached to the IAM role) and the resource-based policy (on the destination backup vault).

The AWS managed policy `AWSBackupServiceRolePolicyForBackup` includes both of these permissions. If you use a custom IAM role, you must add these permissions explicitly. The following is an example identity-based policy:

```
{
  "Statement": [
    {
      "Sid": "BackupVaultCopyPermissions",
      "Effect": "Allow",
      "Action": "backup:CopyFromBackupVault",
      "Resource": "*"
    },
    {
      "Sid": "BackupVaultPermissions",
      "Effect": "Allow",
      "Action": "backup:CopyIntoBackupVault",
      "Resource": "arn:aws:backup:*:*:backup-vault:*"
    }
  ]
}
```

- **A destination account**

The destination account is the account where you would like to keep a copy of your backup. You can choose more than one destination account. The destination account must be in the same organization as the source account in AWS Organizations.

You must “Allow” the access policy `backup:CopyIntoBackupVault` for your destination backup vault. The absence of this policy will deny attempts to copy into the destination account.

- **A management account in AWS Organizations**

  The management account is the primary account in your organization, as defined by AWS Organizations, that you use to opt-in to cross-account backup across your AWS accounts. Before your organization can start with cross-account backups, you must enable cross-account backup in the AWS Backup console or through the UpdateGlobalSettings API.

For information about security, see Security considerations for cross-account backup.

To use cross-account backup, you must enable the cross-account backup feature. Then, you must "Allow" the access policy `backup:CopyIntoBackupVault` into your destination backup vault.

Amazon EC2 offers EC2 Allowed AMIs. If this setting is enabled in your account, add your source account ID to your allowlist. Otherwise, the copy operation will fail with an error message, such as "Source AMI not found in Region".

Enable cross-account backup

1. Log in using your AWS Organizations management account credentials. Cross-account backup can only be enabled or disabled using these credentials.
2. Open the AWS Backup console at https://console.aws.amazon.com/backup.
3. In **My account**, choose **Settings**.
4. For **Cross-account backup**, choose **Enable**.
5. In **Backup vaults**, choose your destination vault.

   For cross-account copy, the source vault and the destination vault are in different accounts. Switch to the account which owns the destination account, as necessary.
6. In the **Access policy** section, "Allow" `backup:CopyIntoBackupVault`. For an example, choose **Add permissions** and then **Allow access to a Backup vault from organization**. Any cross-account action other than `backup:CopyIntoBackupVault` will be rejected.

7. Now, any account in your organization can share the contents of their backup vault with any other account in your organization. For more information, see Configuring backup vault access for cross-account copies. To limit which accounts can receive the contents of other accounts' backup vaults, see Configuring your account as a destination account.

Scheduling cross-account backup

You can use a scheduled backup plan to copy backups across AWS accounts.

To copy a backup using a scheduled backup plan

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In **My account**, choose **Backup plans**, and then choose **Create Backup plan**.
3. On the **Create Backup plan** page, choose **Build a new plan**.
4. For **Backup plan name**, enter a name for your backup plan.
5. In the **Backup rule configuration** section, add a backup rule that defines a backup schedule, backup window, and lifecycle rules. You can add more backup rules later.

   For **Rule name**, enter a name for your rule.
6. In the **Schedule** section under **Frequency**, choose how often you want the backup to be taken.
7. For **Backup window**, choose **Use backup window defaults** (recommended). You can customize the backup window.
8. For **Backup vault**, choose a vault from the list. Recovery points for this backup will be saved in this vault. You can create a new backup vault.
9. In the **Generate copy - optional** section, enter the following values:

   Destination Region

   Choose the destination AWS Region for your backup copy. Your backup will be copied to this Region. You can add a new copy rule per copy to a new destination.

   Copy to another account's vault

   Toggle to choose this option. The option turns blue when selected. The **External vault ARN** option will appear.

External vault ARN

Enter the Amazon Resource Name (ARN) of the destination account. The ARN is a string that contains the account ID and its AWS Region. AWS Backup will copy the backup to the destination account's vault. The **Destination region** list automatically updates to the Region in the external vault ARN.

For **Allow Backup vault access**, choose **Allow**. Then choose **Allow** in the wizard that opens.

AWS Backup needs permissions to access the external account to copy backup to the specified value. The wizard shows the following example policy that provides this access.

JSON

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Sid": "AllowAccountCopyIntoBackupVault",
      "Effect": "Allow",
      "Action": "backup:CopyIntoBackupVault",
      "Resource": "*",
      "Principal": {
        "AWS": "arn:aws:iam::123456789012:root"
      }
    }
  ]
}
```

Transition to cold storage

Choose when to transition the backup copy to cold storage and when to expire (delete) the copy. Backups transitioned to cold storage must be stored in cold storage for a minimum of 90 days. This value cannot be changed after a copy has transitioned to cold storage.

To see the list of resources that you can transition to cold storage, see the "Lifecycle to cold storage" section of the Feature availability by resource table. The cold storage expression is ignored for other resources.

Expire specifies the number of days after creation that the copy is deleted. This value must be greater than 90 days beyond the **Transition to cold storage** value.

> **Note**
>
> When backups expire and are marked for deletion as part of your lifecycle policy, AWS Backup deletes the backups at a randomly chosen point over the following 8 hours. This window helps ensure consistent performance.

10. Choose **Tags added to recovery points** to add tags to your recovery points.
11. For **Advanced backup settings**, choose **Windows VSS** to enable application-aware snapshots for the selected third-party software running on EC2.
12. Choose **Create plan**.

Performing on-demand cross-account backup

You can copy a backup to a different AWS account on demand.

To copy a backup on-demand

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. For **My account**, choose **Backup vault** to see all your backup vaults listed. You can filter by the backup vault name or tag.
3. Choose the **Recovery point ID** of the backup you want to copy.
4. Choose **Copy**.
5. Expand **Backup details** to see information about the recovery point you are copying.
6. In the **Copy configuration** section, choose an option from the **Destination region** list.
7. Choose **Copy to another account's vault**. The option turns blue when selected.
8. Enter the Amazon Resource Name (ARN) of the destination account. The ARN is a string that contains the account ID and its AWS Region. AWS Backup will copy the backup to the destination account's vault. The **Destination region** list automatically updates to the Region in the external vault ARN.
9. For **Allow Backup vault access**, choose **Allow**. Then choose **Allow** in the wizard that opens.

To create the copy, AWS Backup needs permissions to access the source account. The wizard shows an example policy that provides this access. This policy is shown following.

JSON

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Sid": "AllowAccountCopyIntoBackupVault",
      "Effect": "Allow",
      "Action": "backup:CopyIntoBackupVault",
      "Resource": "*",
      "Principal": {
        "AWS": "arn:aws:iam::123456789012:root"
      }
    }
  ]
}
```

10. For **Transition to cold storage**, choose when to transition the backup copy to cold storage and when to expire (delete) the copy. Backups transitioned to cold storage must be stored in cold storage for a minimum of 90 days. This value cannot be changed after a copy has transitioned to cold storage.

    To see the list of resources that you can transition to cold storage, see the "Lifecycle to cold storage" section of the Feature availability by resource table. The cold storage expression is ignored for other resources.

    Expire specifies the number of days after creation that the copy is deleted. This value must be greater than 90 days beyond the **Transition to cold storage** value.

11. For **IAM role**, specify the IAM role (such as the default role) that has the permissions to make your backup available for copying. The act of copying is performed by your destination account's service linked role.

12. Choose **Copy**. Depending on the size of the resource you are copying, this process could take several hours to complete. When the copy job completes, you will see the copy in the **Copy jobs** tab in the **Jobs** menu.

Encryption keys and cross-account copies

See Encryption for a backup copy to a different account or AWS Region for details on how encryption works for copy job.

For additional help troubleshooting cross-account copy failures, please see the AWS Knowledge Center.

Restoring a backup from one AWS account to another

AWS Backup does not support recovering resources from one AWS account to another. However, you can copy a backup from one account to a different account and then restore it in that account. For example, you can't restore a backup from account A to account B, but you can copy a backup from account A to account B, and then restore it in account B.

Before restoring a backup from one account to another, ensure that the destination account has the service-linked role (SLR) for the resource type you are restoring. If the destination account has never used that AWS service before, the SLR may not exist. You can create the SLR by using the service in the destination account, which automatically creates it. Once the SLR requirement is addressed, restoring a backup from one account to another is a two-step process:

To restore a backup from one account to another

1. Copy the backup from the source AWS account to the account you want to restore to. For instructions, see Setting up cross-account backup.
2. Use the appropriate instructions for your resource to restore the backup.

Configuring backup vault access for cross-account copies

AWS Backup allows you to configure your backup vault to grant access to other AWS accounts, allowing them to copy recovery points to your vault for cross-account backup. This access configuration uses resource-based policies to permit specific accounts to perform backup operations.

> **Note**
>
> This is different from AWS Backup vault sharing for Logically Air Gapped (LAG) vaults, which uses AWS Resource Access Manager (RAM) to share vault resources directly.

You can grant vault access to one or multiple accounts, or your entire organization in AWS Organizations. You can configure a destination backup vault with access for a source AWS Account, user, or IAM role.

To configure vault access for a destination Backup vault

1. Choose **AWS Backup**, and then choose **Backup vaults**.
2. Choose the name of the backup vault that you want to configure access for.
3. In the **Access policy** pane, choose the **Add permissions** dropdown.
4. Choose **Allow account level access to a Backup vault**. Or, you can choose to allow organization-level or role-level access.
5. Enter the **AccountID** of the account you'd like to grant access to this destination backup vault.
6. Choose **Save policy**.

You can use IAM policies to configure vault access.

Configure destination backup vault access for an AWS account or IAM role

The following policy configures vault access for account number 4444555566666 and the IAM role `SomeRole` in account number 111122223333.

JSON

```
{
  "Version":"2012-10-17",
  "Statement":[
    {
      "Effect":"Allow",
      "Principal":{
        "AWS":[
          "arn:aws:iam::444455556666:root",
          "arn:aws:iam::111122223333:role/SomeRole"
        ]
      },
      "Action":"backup:CopyIntoBackupVault",
      "Resource":"*"
    }
  ]
}
```

Share a destination backup vault an organizational unit in AWS Organizations

The following policy shares a backup vault with organizational units using their `PrincipalOrgPaths`.

JSON

```
{
  "Version":"2012-10-17",
  "Statement":[
    {
      "Effect":"Allow",
      "Principal":"*",
      "Action":"backup:CopyIntoBackupVault",
      "Resource":"*",
      "Condition":{
        "ForAnyValue:StringLike":{
          "aws:PrincipalOrgPaths":[
            "o-a1b2c3d4e5/r-f6g7h8i9j0example/ou-def0-awsbbbbb/",
            "o-a1b2c3d4e5/r-f6g7h8i9j0example/ou-def0-awsbbbbb/ou-jkl0-awsddddd/
*"
          ]
        }
      }
    }
  ]
}
```

Share a destination backup vault with an organization in AWS Organizations

The following policy shares a backup vault with the organization with `PrincipalOrgID` "o-a1b2c3d4e5".

JSON

```
{
  "Version":"2012-10-17",
  "Statement":[
    {
      "Effect":"Allow",
```

```
      "Principal":"*",
      "Action":"backup:CopyIntoBackupVault",
      "Resource":"*",
      "Condition":{
        "StringEquals":{
          "aws:PrincipalOrgID":[
            "o-a1b2c3d4e5"
          ]
        }
      }
    }
  ]
}
```

Configuring your account as a destination account

When you first enable cross-account backups using your AWS Organizations management account, any user of a member account can configure their account to be a destination account. We recommend setting one or more of the following service control policies (SCPs) in AWS Organizations to limit your destination accounts. To learn more about attaching service control policies to AWS Organizations nodes, see Attaching and detaching service control policies.

Limit destination accounts using tags

When attached to an AWS Organizations root, OU, or individual account, this policy limits copies destinations from that root, OU, or account to only those accounts with backup vaults you've tagged `DestinationBackupVault`. The permission `"backup:CopyIntoBackupVault"` controls how a backup vault behaves and, in this case, which destination backup vaults are valid. Use this policy, along with the corresponding tag applied to approved destination vaults, to control the destinations of cross-account copies to only approved accounts and backup vaults.

JSON

```
{
  "Version":"2012-10-17",
  "Statement":[
    {
      "Effect":"Deny",
      "Action":"backup:CopyIntoBackupVault",
      "Resource":"*",
```

```
      "Condition":{
        "Null":{
          "aws:ResourceTag/DestinationBackupVault":"true"
        }
      }
    }
  ]
}
```

Limit destination accounts using account numbers and vault names

When attached to an AWS Organizations root, OU, or individual account, this policy limits copies originating from that root, OU, or account to only two destination accounts. The permission `"backup:CopyFromBackupVault"` controls how a recovery point in the backup vault behaves, and, in this case, the destinations where you can copy that recovery point to. The source vault will only permit copies to the first destination account (112233445566) if one or more destination backup vault names begin with `cab-`. The source vault will only permit copies to the second destination account (123456789012) if the destination is the single backup vault named `fort-knox`.

JSON

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Sid": "DenyCopyFromBackupVault",
      "Effect": "Deny",
      "Action": "backup:CopyFromBackupVault",
      "Resource": "arn:aws:ec2:*:*:snapshot/*",
      "Condition": {
        "ForAllValues:ArnNotLike": {
          "backup:CopyTargets": [
            "arn:aws:backup:*:112233445566:backup-vault:cab-*",
            "arn:aws:backup:us-east-1:123456789012:backup-vault:fort-knox"
          ]
        }
      }
    }
```

```
  ]
}
```

Limit destination accounts using organizational units in AWS Organizations

When attached to an AWS Organizations root or OU that contains your source account, or when attached to your source account, the following policy limits the destination accounts to those accounts within the two specified OUs.

JSON

```
{
  "Version":"2012-10-17",
  "Statement":[
    {
      "Effect":"Deny",
      "Action":"backup:CopyFromBackupVault",
      "Resource":"*",
      "Condition":{
        "ForAllValues:StringNotLike":{
          "backup:CopyTargetOrgPaths":[
            "o-a1b2c3d4e5/r-f6g7h8i9j0example/ou-def0-awsbbbbb/",
            "o-a1b2c3d4e5/r-f6g7h8i9j0example/ou-def0-awsbbbbb/ou-jkl0-awsddddd/
*"
          ]
        }
      }
    }
  ]
}
```

Security considerations for cross-account backup

Be aware of the following when using performing cross-account backups in AWS Backup:

- The destination vault cannot be the default vault. This is because the default vault is encrypted with a key that cannot be shared with other accounts.

- Cross-account backups might still run for up to 15 minutes after you disable cross-account backup. This is due to eventual consistency, and might result in some cross-account jobs starting or completing even after you disable cross-account backup.
- If the destination account leaves the organization at a later date, that account will retain the backups. To avoid potential data leakage, place a deny permission on the `organizations:LeaveOrganization` permission in a service control policy (SCP) attached to the destination account. For detailed information about SCPs, see Removing a member account from your organization in the *Organizations User Guide*.
- If you delete a copy job role during a cross-account copy, AWS Backup can't unshare snapshots from the source account when the copy job completes. In this case, the backup job finishes, but the copy job status shows as Failed to unshare snapshot.

Copy tags onto backups

In general, AWS Backup copies tags from the resources it protects to your *recovery points*. For more information on how to copy tags during a restore, see Copy tags during a restore.

For example, when you back up an Amazon EC2 volume, AWS Backup copies its group and individual resource tags to the resulting snapshot, subject to the following:

- For a list of resource-specific permissions that are required to save metadata tags on backups, see Permissions required to assign tags to backups.
- Tags that are originally associated with a resource and tags that are assigned during backup are assigned to recovery points stored in a backup vault, up to a maximum of 50 (this is an AWS limitation). Tags that are assigned during backup have priority, and both sets of tags are copied in alphabetical order.

  For continuous backup, tags added to the backups from the primary resource won't be removed if the tags are removed from the primary resource. You will need to remove the tags from backups manually. Make sure that the number of tags on the backup is up to a maximum of 50.
- DynamoDB does not support assigning tags to backups unless you first enable Advanced DynamoDB backup.
- Amazon EBS volumes that are attached to Amazon EC2 instances are nested resources. Tags on the Amazon EBS volumes that are attached to Amazon EC2 instances are nested tags. If AWS Backup can't copy nested tags, the backup job fails.

- When an Amazon EC2 backup creates an image recovery point and a set of snapshots, AWS Backup copies tags to the resulting AMI. If AWS Backup can't copy the tags from the volumes associated with the Amazon EC2 instance to the resulting snapshots, the backup job fails.

If you copy your backup to another AWS Region, AWS Backup copies all tags of the original backup to the destination AWS Region.

Add or remove tags on existing recovery points

In addition to copying tags from source resources at backup creation time, you can add, modify, or remove tags on existing recovery points.

> **Important**
>
> You can only manage tags via the AWS Backup console and API for recovery points of resources that are fully managed by AWS Backup. For non-fully managed resources (such as Amazon EC2, Amazon EBS, Amazon RDS, and Aurora), the **Manage tags** option is not available in the console. To manage tags on these recovery points, use the underlying service's tagging capabilities instead.

Required permissions

The IAM role or user performing the tagging operation must have the following permissions:

- `backup:TagResource` – Required to add or modify tags on a recovery point.
- `backup:UntagResource` – Required to remove tags from a recovery point.
- `backup:ListTags` – Required to view existing tags on a recovery point.

Add or remove tags using the console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup/.
2. In the navigation pane, choose **Backup vaults**.
3. Choose the backup vault that contains the recovery point you want to tag.
4. Under **Recovery points**, choose the **Recovery point ID** of the backup you want to tag.

5. In the **Tags** section, choose **Manage tags**.
6. To add a tag, choose **Add tag**, then enter a **Key** and an optional **Value**.

   To remove a tag, choose **Remove** next to the tag you want to delete.
7. Choose **Save**.

Add or remove tags using the AWS CLI

Add tags to a recovery point:

```
aws backup tag-resource \
    --resource-arn arn:aws:backup:us-east-1:123456789012:recovery-
point:abcdef12-3456-7890-abcd-ef1234567890 \
    --tags Key1=Value1,Key2=Value2
```

Remove tags from a recovery point:

```
aws backup untag-resource \
    --resource-arn arn:aws:backup:us-east-1:123456789012:recovery-
point:abcdef12-3456-7890-abcd-ef1234567890 \
    --tag-key-list Key1 Key2
```

List tags on a recovery point:

```
aws backup list-tags \
    --resource-arn arn:aws:backup:us-east-1:123456789012:recovery-
point:abcdef12-3456-7890-abcd-ef1234567890
```

For more information, see TagResource, UntagResource, and ListTags in the *AWS Backup API Reference*.

Backup deletion

We recommend you use AWS Backup to automatically delete the backups that you no longer need by configuring your lifecycle when you created your backup plan. For example, if you set your backup plan's lifecycle to retain your recovery points for one year, AWS Backup will automatically delete on January 1, 2022 the recovery points it created on or within several hours of January 1, 2021. (AWS Backup randomizes its deletions within 8 hours following recovery point expiration

to maintain performance.) To learn more about configuring your lifecycle retention policy, see Creating a backup plan.

However, you might want to manually delete one or more recovery points. For example:

- You have `EXPIRED` recovery points. These are recovery points AWS Backup was unable to delete automatically because you deleted or modified the original IAM policy you used to create your backup plan. When AWS Backup attempted to delete them, it lacked permission to do so.

  Expired recovery points might also be created if an AWS managed Amazon EBS or Amazon EC2 recovery point has an Amazon EBS Snapshot Lock applied and AWS Backup is unable to complete the lifecycle process that would normally result in the recovery point being deleted. Note these expired recovery points can be restored from the Amazon EC2 console and API or Amazon EBS console and API.

  > **Warning**
  >
  > You will continue to store expired recovery points in your account. This might increase your storage costs.

  > **Note**
  >
  > AWS Backup periodically identifies orphaned AMIs - AMIs created by AWS Backup that are no longer associated with a recovery point - and places them as expired recovery points in the account's Default vault. You can then delete these expired recovery points, which also deletes the associated AMI and snapshots. If the Default vault does not exist in the Region, create one so that orphaned AMIs can be placed there as expired recovery points. You must then delete these recovery points manually.

  After August 6, 2021, AWS Backup will show the target recovery point as **Expired** in its backup vault. You can hover your mouse over the red **Expired** status for a popover status message that explains why it was unable to delete the backup. You can also choose **Refresh** to receive the most recent information.
- You no longer want a backup plan to operate the way you configured it. Updating the backup plan affects the future recovery points it will create, but does not affect the recovery point it already created. To learn more, see Updating a backup plan.

- You need to clean up after finishing a test or tutorial.

Deleting backups manually

To manually delete recovery points

1. In the AWS Backup console, in the navigation pane, choose **Backup vaults**.
2. On the **Backup vaults** page, choose the backup vault where you stored the backups.
3. Choose recovery points, then choose **Actions**, **Delete**.
4. 1. If your list contains a continuous backup, choose one of following options. Each continuous backup has a single recovery point.
      - **Permanently delete my backup data** or **Delete recovery point**. By selecting one of these options, you stop future continuous backups and also delete your existing continuous backup data.

        > **Note**
        >
        > See Continuous backups and point-in-time recovery (PITR) for Amazon S3, Amazon RDS, and Aurora continuous backup considerations.

      - **Keep my continuous backup data** or **Disassociate recovery point**. By selecting one of these options, you stop future continuous backups but retain your existing continuous backup data until it expires as defined by your retention period.

        A disassociated Amazon S3 continuous recovery point (backup) will remain in its backup vault, but its state will transition to `STOPPED`.

        > **Note**
        >
        > Calling `DisassociateRecoveryPoint` on an Amazon S3 continuous recovery point that is already in `EXPIRED` status has no effect. The recovery point state and configuration remain unchanged.

   2. To delete all the recovery points selected, type delete, and then choose **Delete recovery points**.

3. AWS Backup begins to submit your recovery points for deletion and displays a progress bar. Keep your browser tab open and do not navigate away from this page during the submission process.
4. At the end of the submission process, AWS Backup presents you a status in the banner. The status can be:
   - **Successfully submitted**. You can choose to **View progress** about each recovery point's deletion status.
   - Failed to submit. You can choose to **View progress** about each recovery point's deletion status or **Try again** with your submission.
   - A mixed result where some recovery points were successfully submitted while other recovery points failed to submit.
5. If you choose **View progress**, you can review the **Deletion status** of each backup. If a deletion status is **Failed** or **Expired**, you can click that status to see the reason. You can also choose to **Retry failed deletions**.

Troubleshooting manual deletions

In rare situations, AWS Backup might not complete your delete request. AWS Backup uses the service-linked role `AWSServiceRoleForBackup` to perform deletions.

If your delete request fails, verify that your IAM role has the permission to create service-linked roles. Specifically, verify your IAM role has the `iam:CreateServiceLinkedRole` action. If it does not, add this permission to the role used to create a backup. Adding this permission allows AWS Backup to perform manual deletions.

If, after you confirm that your IAM role has the `iam:CreateServiceLinkedRole` action, your recovery points are still stuck in the `DELETING` status, we are likely investigating your issue. Complete your manual deletion with the following steps:

1. Set up a reminder to come back in 2-3 days.
2. After 2-3 days, check for recently `EXPIRED` deletion points that are the result of your first manual deletion operation.
3. Manually delete those `EXPIRED` recovery points.

For more information on roles, see Using service-linked roles and Adding and removing IAM identity permissions.

Backup and tag edits

After you create a backup using AWS Backup, you can change the lifecycle or tags of the backup. The lifecycle defines when a backup is transitioned to cold storage and when it expires. AWS Backup transitions and expires backups automatically according to the lifecycle that you define.

To see the list of resources that you can transition to cold storage, see the "Lifecycle to cold storage" section of the Feature availability by resource table. The cold storage expression is ignored for other resources.

> **Note**
>
> Editing the tags of a backup using the AWS Backup console is only supported for backups of Amazon Elastic File System (Amazon EFS) file systems and Advanced Amazon DynamoDB.
> Tags that were added to the recovery point on creation for other resources will still appear, but will be greyed out and uneditable. Even though these tags are not editable in the AWS Backup console, you can edit the tags of these other services' backups using the service's console or API.

Backups that are transitioned to cold storage must be stored in cold storage for a minimum of 90 days. Therefore, the "retention" setting must be 90 days greater than the "transition to cold after days" setting. When you update the "transition to cold after days" setting, the value must be a minimum of the backup's age plus one day. The "transition to cold after days" setting cannot be changed after a backup has been transitioned to cold.

The following is an example of how to update the lifecycle of a backup.

To edit the lifecycle of a backup

1. Sign in to the AWS Management Console, and open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Backup vaults**.
3. In the **Backups** section, choose a backup.
4. On the backup details page, choose **Edit**.
5. Configure the lifecycle settings, and then choose **Save**.

Backup search

Overview

With AWS Backup, you can create backups, also known as recovery points, of AWS resources. You can search for backups of certain resource types such as Amazon S3 and Amazon EBS, as well as items and files within those backups using the AWS Backup console or the command line.

AWS Backup offers the ability for you to search the metadata of your backups of supported resource types at a granular level for files or objects that match the properties you define in your search, such as size, creation date, and resource type. You can dive even deeper by defining the properties of the items you want to locate.

First, create a backup index you want to be able to include in a future search. Backup index creation can be automated through a backup plan or you can manually create one for any existing recovery point. When you're ready to search, set the backup and item properties you want to see in the search results. Optionally, you can restore the backup or item you sought in the search.

This document outlines the steps to create a backup index, search indexed backups, restore from your search results, and troubleshoot any issues with the index and search functions in AWS Backup.

Use cases for backup indexes and search

You may be an administrator who wants to recover a specific file or object. Instead of manually identifying or guessing which backups contain the data, you can search the metadata of your recovery points and restore the exact backup, files, or objects you need.

Restoring a full backup just to find the specific item that might be in it can take hours or days. Instead, with a backup search, you can find and restore just the specific file or object you require.

Backup searches are useful for backup administrators, backup operators, data owners, and other IT professionals who interact with data backup, restore, and compliance.

Access

Before you create an index and a search, your account must have required permissions for the operations.

Index permissions

For index operations, AWS Backup authenticates based on the IAM role, not user credentials (for IAM user and IAM role specifics, see Authentication).

The following permissions are required to create an index of an EBS backup. These permissions are contained in the managed policy AWSBackupServiceRolePolicyForIndexing:

- `ec2:DescribeSnapshots`
- `ebs:ListSnapshotBlocks`
- `ebs:GetSnapshotBlock`
- `kms:Decrypt`

No index permissions are required to create an S3 index.

Search permissions

The following permissions are required to create a search. These permissions are contained in the managed policy AWSBackupSearchOperatorAccess:

- `backup:ListIndexedRecoveryPointsForSearch`
- `backup:SearchRecoveryPoint`

If you choose to encrypt the search results with a customer managed AWS KMS key, ensure the following permissions are in the key:

- `kms:GenerateDataKey`
- `kms:Decrypt`

Process Flow

A backup search involves three steps, plus an optional fourth restore step for when you want to restore the items returned in your search.

Index your backups: Enable indexing in your backup plan(s) or manually create a backup index through the console or CLI for each existing backup (recovery point) you want to be eligible for searches.

Search backup metadata for a recovery point, file or, object: Specify the properties of the backups and items you want to find in your search, such as your searching your S3 buckets created

between April 2 and 6. with tags of `Administration` and for objects greater than 100 MB with the key name containing `Admin`.

Review search results: If you find the recovery point or item you were seeking, you have the option to restore it. If you haven't found the recovery point or item, you can refine the backup properties and item properties, then initiate a new search.

Restore specific items *(optional)*: Specify file paths or items to restore, as well as the restore conditions.

Backup indexes

To be searchable, a backup (recovery point) must first have a corresponding index.

Backup index creation can be enabled in a backup plan so that each future backup will also have an associated backup index. You can also create an index as you create an on-demand backup.

Alternatively, you can retroactively create an index for an existing recovery point, either from the Vault recovery point detail screen in the AWS Backup console or through AWS CLI.

Recovery points of supported resource types can have a backup index if they are stored in a standard backup vault (recovery points in a logical air-gapped vault do not currently support backup indexes).

S3 backup indexes

An S3 backup can be periodic, where it is scheduled at a fixed interval according to your backup plan. Each time a periodic backup is created, a backup index is created for it. An S3 backup can also be continuous, where each change in the backup is logged. Since there can be numerous changes daily, only one backup index is created daily for a continuous backup.

The first backup index that is created for a continuous S3 recovery point is full; subsequent indexes for the same recovery point may be incremental.

EBS backup indexes

Each backup index created for an EBS recovery point is full (not incremental).

AWS Backup attempts to automatically repair snapshot issues during the creation of a backup index. If a file system was in a dirty state when the recovery point was created, AWS Backup will automatically attempt to recovery the file system. If this recovery fails, the index creation job will also fail.

The nature of the snapshot determines if it can be indexed:

Can be indexed:

- File systems: ext2, ext3, ext4, vfat, xfs, and ntfs

Cannot be indexed:

- Snapshots in archive tier (cold storage)
- RAID and other multi disk storage options
- Symbolic links
- Hard links

Backup index creation steps

Console

Add backup index creation to your backup plan.

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Select **Backup plans** under **My account** in the left navigation bar.
3. Select the link in the **Backup plans** pane with the name of the plan where you want to add index creation.
4. In the second pane **Backup rules**, select **Add backup rule**.
5. Scroll down to the pane **Backup indexes**. Check the box next to the resource type(s) for which you want to create an index.

   With each new backup this plan creates, a corresponding index for that recovery point will also be concurrently created.

Create an index for an existing recovery point

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Select **Vault** in the left navigation bar.
3. Select the link of under the **Vault** name column in which the backup where you want to make a backup index is stored.

4. Place a checkmark next to the recovery point for which you want to create a backup index.
5. Select the **Action** button, and then select **Create index**.

While the index is being created, it will have the index status of `In progress`. Once the status has transitioned to `Available`, the recovery point can be included in a search.

Create an index as you create an on demand backup.

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. See the steps for Creating an on-demand backup using AWS Backup Creating an on demand backup using AWS Backup.
3. In **Settings**, if you have chosen resource type that supports index and search, the line item **Backup search index** will be display. Toggle on **Create backup search** index to have an index be created concurrently with this on-demand backup.

AWS CLI

Create a backup index through CLI

Use the AWS CLI command `create-backup-plan` to make a new backup plan. Or, use `update-backup-plan` to modify an existing plan.

For either operation, within the parameter `--backup-plan -rules`, include `IndexActions`.

See `IndexActions` in `BackupRuleInput` in the *AWS Backup API Reference Guide* for more information.

Once a recovery point has an index, you can update its settings.

Example:

```
aws backup update-recovery-point-index-settings
--recovery-point-arn arn:aws:ec2:us-west-2::snapshot/snap-012345678901234567

--backup-vault-name [vaultname] //
--index ENABLED
--endpoint-url [URL]
--iam-role-arn arn:aws:iam::012345678901:role/Admin
```

Searches

Once you have one or more backups with an index, you can search those indexed backups through the AWS Backup console or through AWS CLI.

As you create a search, you'll select one resource type. The results will only return recovery points containing that type, such as S3 buckets or EBS snapshots.

You then specify the properties of the backups (recovery points) you wish to include in the search. You can specify up to 9 properties. Property types included more than once will return results that match all included values.

Specify the properties of the items you wish to find within the returned recovery points, such as bucket name or file size. Narrow your results by including multiple properties.

If one value for an item property is included when you create a search through the AWS Backup console, the results will return only items that match that item property (AND logic). If you repeat the same item property, but with different values, the results will return all items that match *any* of the included values (OR logic). For example, if you include two EBS file paths, all items of recovery points that are included in the search that match *either* file path will be in the search results.

- S3 item properties include creation time, Etags, object key, object size, and version ID.
- EBS item properties you can use to help filter your search include creation time, file path, last modification time, and size.

Optionally, you can include an AWS KMS key ID to encrypt your results. If a key is not included, AWS Backup will use a service-owned key to encrypt results.

Console

Search for items in your backups

There are multiple paths to create a search of indexed backups:

You can find your preferred recovery point by navigating to **Backup vaults** and selecting the specific recovery point you wish to search. Then, select **Search**. You can also start a search from the **Recovery point details** page.

During a restore where you have specific items you wish to include, you can search your backups to help locate the URL(s) or file paths that contain the items.

To search through more than one backup, review the following steps:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Navigate to **Search** in the left navigation menu.
3. Select **Create search** in the Search history section.
4. Select a **resource type**. You must select one resource type for each search. If you change the resource type after additional fields have been entered, your entries will be lost and you will need to re-enter them.
5. Choose 1 to 9 **backup properties** to help narrow the recovery points that will be returned in your search.

   AWS Backup will scan all of your backups that have an index. It will return only recovery points that match all different backup properties. For example, `backup tag = "savings"` and `backup creation date = May 20, 2019 through May 23, 2019`, inclusive.

   You may include multiple values of the same property, such as three different tags. If the property is repeated with different values, the search will return all items that match any of the values specified (known as "OR" logic). For example, `backup tag = "savings"`; `backup tag = "checking"`; `backup creation date = May 20, 2019 through May 23, 2019`, inclusive; and `backup creation date = May 20, 2020 through May 23, 2020`, inclusive.

   A backup creation date range counts as one backup property. Only one backup creation date range can be included as a backup property.
6. Choose 1 to 9 **item properties** to help further narrow the returns in your search.

   You may include multiple values of the same property. If the property is repeated with different values, the search will return all items that match any of the values specified.
7. *Optional*: To encrypt your search results, you can specify an extant AWS KMS key by the dropdown menu or ARN, or you can create a new KMS key.
8. AWS Backup recommends creating a unique search job name.
9. Select **Start search**.

   You may see a warning saying that your search may include a large number of recovery points. The best practice is to go back to the backup properties and select additional criteria to narrow the search. Fewer backups in a search may result in lower costs.

AWS CLI

Use the AWS CLI command start-search-job.

Required parameters:

```
--search-scope // defines the backup properties you wish to include in the search
```

Optional parameters:

```
--client-token // string
--encryption-key-arn // if not included, AWS Backup uses service-owned key to
 encrypt results
--item-filters // accepted keys and values depend on which resource type is included
 in the search
--name // If not included, AWS Backup auto-assigns a name
```

Accepted S3 item filters include:

```
--object-keys // string
--sizes // long condition
--creation-times // time condition
--version-ids // string
--etags // string
```

Accepted EBS item filters include:

```
--file-paths //
--sizes //
--creation-times //
--last-modification-time //
```

Stop a search

You can stop a search job if is in status `RUNNING`.

A search job will continue until it reaches `COMPLETED` status (or `FAILED` status if there is an error). You can interrupt a `RUNNING` search job if you wish to end an in-progress search job, which may be desirable if you have found the backup or item you were seeking before the job completed.

- In the AWS Backup console, Select the **Stop search job** button.
- In the CLI, send the operation `stop-search-job` with the search job identifier you want to stop.

Search results

Once a search job has begun, it will begin aggregating results even while has an `Running` status. While a search job is running and until it completes, partial results are available:

- In the console, results will display as they are retrieved during the search. Results do not auto-refresh, but you can view the latest results by selecting the refresh button. To view results beyond the first 1,000 items, select **Export results**.
- The CLI operations get-search-job and list-search-jobs return search job statuses. If the job status is `RUNNING`, the operation returns an incomplete list.

Results from a search job are available in the console and through CLI for 7 days once a search is stopped or has completed. During this time, you can export the results to your preferred Amazon S3 bucket so you can access the results past this timeframe.

Each search job contains detailed information, available in the console or through CLI, including which recovery points were searched, the search name and status, its description, its creation and completion date and time, and information about the objects or items returned as well as the number of items and recovery points scanned.

If the results do not contain the recovery point, item, or object you were seeking, you can create a new search with different backup and item properties. Each search is charged individually.

Each resource type has unique considerations for the results the search returns:

- A search of S3 recovery points will not return delete markers as part of its search result, even if those objects match the search's specified item properties.
- Results of an EBS search may have a null value for creation time for file systems in which that field is unsupported. Those file systems may include, but are not limited to, vfat, ext2/3, and versions of XFS prior to v5.

Export search results to an S3 bucket

AWS Backup retains search results for 7 days, starting with the completion time and date. These results are viewable in the AWS Backup console or retrievable through the CLI operation list-search-job-results.

A best practice is to export your search results to an S3 bucket to retain results beyond the 7 day retention. The export job will create a folder named `Export Job ID` in your designated bucket, then export the results to that folder. Once the results are exported there, they are available for as long as you retain the bucket.

You can export the search results of any supported resource type, not just an S3 search.

Console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Navigate to **Jobs > Search jobs**.
3. Select **search job results**.
4. Place a checkmark next to the result(s) you wish to export.
5. Select **Export to S3**.
6. Choose the destination S3 bucket for the export job.
7. Once you have configured all the fields, select **Export**.

   The export action creates an export job. These are viewable in **Jobs > Export jobs**. Once an export job has reached `COMPLETED` status, the search result information is available in the S3 bucket to retrieve or to download as one or more .csv files.

AWS CLI

Use the AWS CLI command start-search-result-export-job.

Required parameters:

```
--search-job-identifier
--export-specification
```

Optional parameters:

```
--client-token
```

```
--role-arn
--tags
```

Operation will return `ExportJobArn` and `ExportJobIdentifier`.

Use list-search-result-export-jobs to retrieve the statuses of export jobs.

Cost considerations and best practices

Each backup index creation and each search job incurs a charge. Each backup index has a storage charge. Each restore from search results (as with all other all restore jobs) is charged. Learn more at AWS Backup pricing.

You can narrow the possible returned results of a search job by including multiple backup and item properties; this may result in a lower cost than a search than spans all possible recovery points.

Restore from search

Many customers choose to search through their backups - and the objects or files within them - to find a specific recovery point or items to restore. See Restore a backup by resource type for information on restores in general.

You can restore from your search results in the AWS Backup console by navigating to **Jobs > Search job results > Restore**. To restore through AWS CLI, use `start-restore-job` with metadata specific to the resource type, recovery point, and items involved in the restore.

See Restore S3 data using AWS Backup for information on how to restore a recovery point with S3 data, to restore an S3 bucket, or to restore up to five objects or folders with an S3 bucket.

See Restore an Amazon EBS volume for information about restoring an EBS snapshot to a new volume that attaches to an EC2 instance.

Backup tiering

Overview

AWS Backup offers a lower-cost warm storage tier for Amazon S3 backups that reduces long-term storage costs by up to 30% while maintaining enterprise-grade protection and recovery

capabilities. The low-cost tier provides the same performance and features as the warm storage tier. You can configure tiering to move S3 backup data to cost-optimized storage based on the age of objects within your backup vaults.

AWS Backup tiering provides the ability to optimize storage costs for S3 backup data that is retained for extended periods due to regulatory compliance, disaster readiness, and ransomware protection strategies. You can configure tiering for all S3 backups for all vaults in an account, or create targeted configurations for specific vaults and protected resources.

First, create a tiering configuration that specifies which S3 resources should be tiered and after how many days (minimum 60 days). Tiering configurations can be automated to apply to all backups or targeted to specific resources. When backup data reaches the specified age threshold, it transitions to the lower-cost storage tier while maintaining identical restore capabilities.

This document outlines the steps to create tiering configurations, manage tiered backup data, monitor cost savings, and troubleshoot any issues with the tiering function in AWS Backup.

> **Important**
>
> Cost considerations for backup tiering:
>
> - Backup tiering has three cost components: warm storage tier, low-cost warm storage tier, and transition costs. When backup data transitions to the lower-cost tier, you'll incur a one-time per-object transition fee based on the number of objects eligible for tiering.
> - For large datasets with many objects, transition costs may be significant initially but are typically offset by ongoing storage savings for data retained beyond the minimum 60-day threshold.

Tiering Configurations

S3 backup tiering involves creating a tiering configuration that specifies which resources should be tiered and the number of days before transitioning (minimum 60 days) to the lower-cost tier. To enable cost optimization, backup data must be covered by a tiering configuration.

Tiering configuration creation can be set up to apply broadly across backups of all S3 resources in your account, or targeted to specific vaults and resources. You can create multiple configurations to handle different data retention and cost optimization requirements.

Tiering configurations apply to both existing backup data in vaults and new backups created after the configuration is established.

S3 backup tiering configurations specify:

- **Resource scope:** All resources across all vaults, all resources in a specific vault, or selected individual resources. A tiering configuration that applies to all vaults and all resources is considered the default configuration.
- **Transition timing:** Minimum 60 days before data moves to the lower-cost tier
- **Vault assignment:** Which backup vaults the configuration applies to (for all vaults or specific vault name)
- **Resource selection:** Up to 5 different resource selection rules per configuration

Configuration constraints:

- **One configuration per vault:** Each vault can only have one tiering configuration apart from the default configuration
- **Maximum 5 resource selections:** Vault specific configuration supports up to 5 different resource group and corresponding tiering settings
- **Maximum 100 resources:** Up to 100 specific resources across all resource groups can be selected per configuration
- **Vault priority:** If both "all vaults" and specific vault configurations exist, the specific vault configuration takes priority

Creating tiering configurations

Creating tiering configurations

Console

To create an all vaults tiering configuration (default)

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **S3 backup tiering**.
3. Choose **Create configuration**.
4. For configuration name, enter a unique descriptive name.

5. Choose **All S3 resources in all vaults**.
6. For tiering down settings in days, enter the number of days (minimum 60) before data transitions to the lower-cost tier.
7. (Optional) Add tags.
8. Choose **Create configuration**.

To create a specific vault tiering configuration

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **S3 backup tiering**.
3. Choose **Create configuration**.
4. For configuration name, enter a unique descriptive name.
5. Choose **S3 resources in a specific vault**.
6. For vault selection, select a specific backup vault from the dropdown.
7. For resource selection, choose either:
   a. **All S3 resources in this vault** to apply to all S3 resources in the vault
   b. **Specific S3 resources in this vault** to select individual S3 buckets
8. If selecting specific resources:
   a. Choose individual S3 resources (up to 100 total across resource groups in a configuration)
   b. Set tiering down settings in days for each resource group
   c. Choose **Add tiering setting** to create additional rules (up to 5 total)
9. (Optional) Add tags.
10. Choose **Create configuration**.

AWS CLI

To create an all vaults tiering configuration (default) using the AWS CLI

```
aws backup create-tiering-configuration \
--tiering-configuration '{
  "TieringConfigurationName":"MyTieringConfig",
  "BackupVaultName":"*",
```

```
  "ResourceSelection":[{
    "Resources":["*"],
    "TieringDownSettingsInDays":60,
    "ResourceType":"S3"
  }]
}'
```

Managing tiering configurations

Viewing tiering configurations

You can view existing tiering configurations through the AWS Backup console, AWS CLI, or REST API.

Console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **S3 backup tiering**.
3. View the list of configurations with their scope, transition settings, and status.

AWS CLI

To list all tiering configurations using the AWS CLI

```
aws backup list-tiering-configurations --max-results 50
```

To get specific tiering configuration details

```
aws backup get-tiering-configuration --tiering-configuration-name "MyTieringConfig"
```

Modifying tiering configurations

You can update existing tiering configurations to change transition timing, resource selection, or vault assignments.

Console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.

2. In the navigation pane, choose **S3 backup tiering**.
3. Select the configuration to modify.
4. Choose Edit.
5. Update the desired settings.
6. For Tiering down settings in days, enter the number of days (minimum 60) before data transitions to the lower-cost tier.
7. Choose Save changes.

AWS CLI

To update a tiering configuration using the AWS CLI

```
aws backup update-tiering-configuration \
--tiering-configuration-name "MyTieringConfig" \
--tiering-configuration '{
  "BackupVaultName":"*",
  "ResourceSelection":[{
    "Resources":["*"],
    "TieringDownSettingsInDays":60,
    "ResourceType":"S3"
  }]
}'
```

Deleting tiering configurations

You can delete tiering configurations when they are no longer needed.

Console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **S3 backup tiering**.
3. Select the configuration to delete.
4. Choose Delete.
5. Enter the tiering configuration name to confirm deletion.
6. Choose Delete tiering configuration.

AWS CLI

To delete a tiering configuration using the AWS CLI

```
aws backup delete-tiering-configuration \
--tiering-configuration-name "MyTieringConfig"
```

Note: Deleting a tiering configuration keeps existing data in the lower-cost tier but prevents new data from tiering down.

How S3 backup tiering configurations apply

AWS Backup evaluates objects within backups for tiering eligibility based on their age. The service checks object age on a daily basis and transitions eligible objects to the lower-cost tier according to your configuration settings. Tiering evaluation occurs automatically in the background. When objects within a backup reach the specified age threshold (minimum 60 days), they become eligible for transition during the next evaluation cycle.

Both objects in existing backups and newly created backup objects are subject to tiering configurations. If you have multiple configurations that could apply to the same backup objects, vault-specific settings take precedence over the configurations that applies to all vaults. The tiering process is irreversible - once objects move to the lower-cost tier, they remain there until the backup is deleted according to your retention policies.

Cost structure and monitoring

Pricing model

S3 backup tiering uses a cost-optimized pricing structure:

- **Storage cost:** Lower /GB-month cost compared to standard warm tier
- **Transition fee:** One-time per-object fee when moving to lower-cost tier
- **Restore cost:** /GB charge when restoring data same as the warm tier restore
- **No retrieval fees:** There are no additional retrieval charges

Cost monitoring

Monitor your tiering cost savings through:

- **AWS Cost Explorer:** Separate usage types for each storage tier
- **AWS Cost and Usage Reports:** Detailed breakdown with cost allocation tags
- **AWS Backup console:** Configuration information

Example cost savings

For a 500TB S3 bucket with 1 billion objects where 60% are eligible for tiering:

- **Before tiering:** $25,600/month
- **After tiering:** $21,000/month
- **Monthly savings:** $4,600/month (18% reduction)
- **One-time transition fee:** $6,000

Configuration Examples

Example 1: Account-wide tiering

Apply tiering to all S3 resources across all backup vaults:

```
{
  "TieringConfigurationName":"MyTieringConfig",
  "BackupVaultName":"*",
  "ResourceSelection":[
    {
      "Resources":["*"],
      "TieringDownSettingsInDays":60,
      "ResourceType":"S3"
    }
  ]
}
```

Example 2:

Tier all resources in a MyBackupVault4 vault:

```
{
  "TieringConfigurationName":"MyTieringConfig",
  "BackupVaultName":"MyBackupVault4",
```

```
  "ResourceSelection":[
    {
      "Resources":["*"],
      "TieringDownSettingsInDays":60,
      "ResourceType":"S3"
    }
  ]
}
```

Example 3:

Tier specific buckets with different rules:

```
{
  "TieringConfigurationName":"MyTieringConfig",
  "BackupVaultName":"MyBackupVault",
  "ResourceSelection":[
    {
      "Resources": ["arn:aws:s3:::mybucket1", "arn:aws:s3:::mybucket2"],
      "TieringDownSettingsInDays": 60,
      "ResourceType": "S3"
    },
    {
      "Resources": ["arn:aws:s3:::mybucket3"],
      "TieringDownSettingsInDays": 120,
      "ResourceType": "S3"
    }
  ]
}
```

Example 4:

Set the rule to not tier a bucket (set tiering to 36500):

```
{
  "TieringConfigurationName":"MyTieringConfig",
  "BackupVaultName":"*",
  "ResourceSelection":[
    {
      "Resources":["arn:aws:s3:::mybucket7", "arn:aws:s3:::mybucket8"],
      "TieringDownSettingsInDays":36500,
      "ResourceType":"S3"
```

```
    }
  ]
}
```

Supported features and limitations

Supported features

- **Backup types:** Both continuous and periodic backups
- **Vault types:** Standard backup vaults and logically air-gapped vaults
- **Vault lock:** Full compatibility with locked backup vaults
- **Cross-region/account:** Copying tiered data (copies land in standard tier at destination)
- **Restore capabilities:** Point-in-time recovery and item-level restores
- **Search and indexing:** Full compatibility with backup search functionality
- **Compliance:** Maintains all compliance and audit capabilities

Limitations

- **Minimum transition time:** 60 days before data can be moved to lower-cost tier
- **Resource limit:** Up to 100 specific resources per configuration
- **Configuration limit:** Up to 5 different resource selection rules per configuration
- **One configuration per vault:** Each vault can only have one vault specific tiering configuration, apart from the default
- **One-way transition:** Data moved to lower-cost tier remains there until deletion

Troubleshooting

Common issues

Configuration not applying to existing backups

- Verify that the configuration is properly assigned to the correct vaults
- Check that resources are correctly selected in targeted configurations
- Ensure backup data meets the minimum age requirement (60 days)

`AlreadyExistsException` when creating configuration

- Ensure the tiering configuration name is unique within your account
- Check if the target vault already has an active tiering configuration

`LimitExceededException` errors

- Verify you have not exceeded the maximum of 5 resource selection group per configuration
- Check that you have not selected more than 100 specific resources

Higher than expected transition costs

- Review the number of objects being transitioned
- Consider the transition fee impact for frequently changing data
- Evaluate whether the minimum threshold settings is appropriate for your use case

Access points

Overview

AWS Backup now lets you access Amazon S3 backup data directly through S3 access points, without initiating a restore. Using a backup access point, you can read S3 backup data immediately using familiar S3 APIs such as `GetObject`, `HeadObject`, `ListObjectsV2`, and `ListObjectVersions`.

Backup access points provide on-demand, read-only connections to recovery points. You create a backup access point for a specific recovery point, and AWS Backup provisions an S3 access point that routes S3 API calls to your backup data. This enables use cases such as targeted file recovery, data validation, compliance auditing, and forensic investigation without waiting for a full restore to complete.

While a backup access point is active for a recovery point, AWS Backup pauses lifecycle deletion and blocks manual deletion of that recovery point. This protects your data from being removed while applications are actively reading it. Backup tiering to lower-cost warm storage, which is configured separately, is not affected. You must delete all backup access points associated with a recovery point before the recovery point can be deleted or lifecycled.

This page explains how to create, manage, and delete backup access points using the AWS Backup console and AWS CLI.

> **Important**
>
> Access points provide read-only access to backup data. Write operations through backup access points are not supported.

How backup access points work

When you create a backup access point, AWS Backup performs the following steps:

1. Validates that the recovery point exists and is in a usable state (AVAILABLE, STOPPED, or COMPLETED)
2. Creates a backup access point resource in AWS Backup with status CREATING
3. Provisions an S3 access point on your behalf (asynchronously)
4. Pauses lifecycle deletion and blocks manual deletion of the associated recovery point
5. Updates the backup access point status to AVAILABLE

Once the backup access point is AVAILABLE, you call `DescribeBackupAccessPoint` to retrieve the S3 access point ARN and alias. You then use the S3 access point alias or ARN with standard S3 APIs to read your backup data.

When you are done, call `DeleteBackupAccessPoint`. AWS Backup deletes the S3 access point and this resumes recovery point lifecycle transitions (if no other access points remain).

Access point statuses

Status	Description
CREATING	Access point creation is in progress. S3 access point is being provisioned.
AVAILABLE	Access point is ready. Use `DescribeBackupAccessPoint` to get the S3 ARN and alias.

Status	Description
DELETING	Deletion is in progress.
FAILED	Creation or deletion failed. Use `DescribeBackupAccessPoint` to get the error details.
EXPIRED	The backup access point is no longer usable.
DISASSOCIATING	Access is being revoked. The logically air-gapped vault is being unshared or the restore access backup vault is being deleted
DISASSOCIATED	Access has been revoked. The backup access point is no longer usable.

Continuous recovery points

For continuous (PITR) recovery points, you can specify an `AccessPointInTime` timestamp when creating the access point. This determines the point-in-time view of your backup data. The timestamp must be within the continuous backup retention window *at the time of access point creation. As long as the access point exists, the data for that point in time is retained and accessible, even if the timestamp later falls outside the retention window.*

Prerequisites

Before creating/managing a backup access point, ensure:

- Your IAM role has the required permissions. You can use the managed policy AWSBackupAccessPointOperatorAccess, or ensure the following permissions are applied:
  - `backup:CreateBackupAccessPoint`
  - `backup:DescribeBackupAccessPoint`
  - `backup:DeleteBackupAccessPoint`
  - `backup:ListBackupAccessPoints`
  - `backup:ListBackupAccessPointsByRecoveryPoint`
  - `s3:GetAccessPoint` (required during create, describe, and delete)

- `s3:CreateAccessPoint` (required during create)
- `s3:DeleteAccessPoint` (required during delete)
- `s3:PutAccessPointPolicy` (required only if an access point policy is included in the request)

Creating a backup access point

Console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup/.
2. In the navigation pane, choose **Protected resources** or navigate to a recovery point through **Backup vaults**.
3. Select the S3 recovery point you want to access.
4. Choose **Create access point**.
5. For **Access point name**, enter a unique name. This name is shared with the S3 access point namespace, so it must not conflict with existing S3 access points in the same Region and account. *Note: For naming rules, see Naming rules for Amazon S3 access points.*
6. (Optional) For continuous recovery points, specify the **Access point in time** to set the point-in-time view.
7. (Optional) Add access point policy.
8. (Optional) Add tags.
9. Choose **Create access point**.
10. The backup access point status shows CREATING. When status changes to AVAILABLE, choose the backup access point to view the S3 access point ARN and alias.

AWS CLI

```
aws backup create-backup-access-point \
  --recovery-point-arn "arn:aws:backup:us-east-1:123456789012:recovery-
point:rp-1234567890abcdef0" \
  --name "my-access-point" \
  --access-point-metadata '{"AccessPointInTime": "2026-07-01T12:00:00Z"}'
```

Response:

```
{
  "AccessPointArn": "arn:aws:backup:us-east-1:123456789012:accesspoint/my-access-
point",
  "Status": "CREATING"
}
```

After creation, describe the access point to get the S3 ARN and alias:

```
aws backup describe-backup-access-point \
  --access-point-arn "arn:aws:backup:us-east-1:123456789012:accesspoint/my-access-
point"
```

Response (when AVAILABLE):

```
{
  "AccessPointArn": "arn:aws:backup:us-east-1:123456789012:accesspoint/my-access-
point",
  "AccessPointMetadata": {
    "S3AccessPointAlias": "my-access-point-abc123-ext-s3alias",
    "S3AccessPointArn": "arn:aws:s3:us-east-1:123456789012:accesspoint/my-access-
point"
  },
  "BackupVaultName": "MyVault",
  "CreationTime": "2026-07-28T12:00:25.938000-07:00",
  "Name": "my-access-point",
  "RecoveryPointArn": "arn:aws:backup:us-east-1:123456789012:recovery-
point:rp-1234567890abcdef0",
  "ResourceArn": "arn:aws:s3:::my-bucket",
  "ResourceType": "S3",
  "Status": "AVAILABLE"
}
```

Accessing backup data through the access point

Once the access point is in AVAILABLE status, use the S3 access point alias or ARN with S3 APIs to read backup data.

Using the alias

```
# List objects in the backup
```

```
aws s3api list-objects-v2 \
  --bucket "my-access-point-abc123-ext-s3alias"

# Get a specific object
aws s3api get-object \
  --bucket "my-access-point-abc123-ext-s3alias" \
  --key "path/to/my-file.txt" \
  output-file.txt

# Head an object (metadata only)
aws s3api head-object \
  --bucket "my-access-point-abc123-ext-s3alias" \
  --key "path/to/my-file.txt"
```

Using the S3 access point ARN

```
aws s3api list-objects-v2 \
  --bucket "arn:aws:s3:us-east-1:123456789012:accesspoint/my-access-point"
```

Managing access points

Listing access points

List all backup access points in your account:

```
aws backup list-backup-access-points
```

List backup access points for a specific recovery point:

> **Note**
>
> If you are the recovery point owner and have shared the recovery point with other accounts, the response includes access points created by those accounts.

```
aws backup list-backup-access-points-by-recovery-point \
  --recovery-point-arn "arn:aws:backup:us-east-1:123456789012:recovery-
point:rp-1234567890abcdef0"
```

List backup access points for a specific resource:

```
aws backup list-backup-access-points-by-resource \
  --resource-arn "arn:aws:s3:::my-bucket"
```

Deleting an access point

```
aws backup delete-backup-access-point \
  --access-point-arn "arn:aws:backup:us-east-1:123456789012:accesspoint/my-access-
point"
```

When you delete a backup access point:

- The S3 access point is deleted
- Recovery point lifecycle resumes (if no other access points remain for that recovery point)
- The backup access point resource is removed from AWS Backup

> **Important**
>
> Always delete backup access points using the `DeleteBackupAccessPoint` API or the AWS Backup console. Do not delete the underlying S3 access point directly through S3 APIs. Deleting the S3 access point directly leaves the backup access point in EXPIRED status and lifecycle protection remains in place until you delete the orphaned backup access point.

Access point policies

Each access point can have its own resource-based access point policy. Access point policies control how objects in the backup data source can be accessed through the access point by resource, user, or conditions.

How access point policies work

When you create a backup access point, the S3 access point is created without an access point policy by default. Without a policy, access to backup data through the access point is governed by the caller's IAM permissions.

You can optionally specify a policy during creation or add one afterward to:

- Grant read access to additional IAM principals in your account
- Restrict access by source IP address or VPC endpoint
- Add conditions such as requiring MFA or specific tags

Restrictions that you include in an access point policy apply only to requests made through that access point. An access point policy grants permissions only if the underlying data source also allows the same access.

Editing an access point policy

You can edit an access point policy through the AWS Backup console or the S3 `PutAccessPointPolicy` API.

Console

1. Open the AWS Backup console.
2. Navigate to **Access points** and select your access point.
3. Choose the **Access point policy** tab.
4. Choose **Edit**.
5. Modify the JSON policy.
6. Choose **Save changes**.

   When you edit a policy, AWS Identity and Access Management Access Analyzer runs policy checks to validate your policy against IAM policy grammar and best practices. Resolve security warnings, errors, and suggestions before saving.

AWS CLI

```
aws s3control put-access-point-policy \
  --account-id 123456789012 \
  --name "my-access-point" \
  --policy '{
    "Version": "2012-10-17",
    "Statement": [
      {
        "Effect": "Allow",
        "Principal": {
          "AWS": "arn:aws:iam::123456789012:role/DataAnalyst"
        },
```

```
        "Action": ["s3:GetObject", "s3:ListBucket"],
        "Resource": [
          "arn:aws:s3:us-east-1:123456789012:accesspoint/my-access-point",
          "arn:aws:s3:us-east-1:123456789012:accesspoint/my-access-point/object/*"
        ]
      }
    ]
}'
```

Block Public Access

Amazon S3 Block Public Access settings apply to backup access points. If Block Public Access is turned on for your account, public access through backup access points is blocked regardless of the access point policy. To check your settings, review your account-level Block Public Access configuration.

Learn more

- Configuring IAM policies for using access points
- Amazon S3 Access Points
- IAM Access Analyzer policy validation
- Using Amazon S3 Block Public Access

Supported S3 operations

The following table shows which S3 API operations are compatible with backup access points. Backup access points are read-only, so all write, delete, and multipart upload operations are not supported.

S3 operation	Backup access point
`GetObject`	Supported
`HeadObject`	Supported
`GetObjectAttributes`	Supported
`GetObjectTagging`	Supported

S3 operation	Backup access point
GetBucketLocation	Supported
ListObjectsV2	Supported
ListObjects	Supported
ListObjectVersions	Supported
PutObject	Not supported
DeleteObject	Not supported
DeleteObjects	Not supported
CopyObject	Not supported
CreateMultipartUpload	Not supported
UploadPart	Not supported
CompleteMultipartUpload	Not supported
AbortMultipartUpload	Not supported
PutObjectTagging	Not supported
DeleteObjectTagging	Not supported
PutObjectAcl	Not supported
GetObjectAcl	Not supported
GetObjectLegalHold	Not supported
GetObjectRetention	Not supported

Recovery point lifecycle protection

While one or more active access points exist for a recovery point:

- **Scheduled lifecycle transitions are paused.** The recovery point will not be automatically deleted by lifecycle rules.
- **Manual deletion is blocked.** Calls to `DeleteRecoveryPoint` return an error until all access points are deleted.
- **Tiering is not affected.** Backup tiering to lower-cost warm storage is configured independently and continues to apply normally while access points are active.

You can still update the recovery point lifecycle (for example, changing the delete-after-days value), but the deletion will not be enforced until all backup access points for the recovery point are removed.

Once all access points for a recovery point are deleted, lifecycle processing resumes automatically.

> **Note**
>
> Each account can have up to 5 backup access points per recovery point, regardless of their status. To create a new backup access point when the limit is reached, delete an existing one (including any in FAILED or EXPIRED status).

Access points for recovery points in logically air-gapped vaults

You can create backup access points for recovery points stored in logically air-gapped vaults. Access points can be created by the vault-owning account directly, or from another account that has been granted access through AWS Resource Access Manager (RAM) sharing or a restore access backup vault (via Multi-Party Approval).

Same-account access

The account that owns the logically air-gapped vault can create access points for recovery points in that vault, the same way as for standard backup vaults.

RAM sharing

When a logically air-gapped vault is shared with an account via RAM, that account can create access points for recovery points in the shared vault. The RAM permission `AWSRAMPermissionBackupVaultReadOnly` includes `backup:CreateBackupAccessPoint`. Vault owners can restrict this by adding a vault access policy that denies `backup:CreateBackupAccessPoint`.

For more information, see Share a logically air-gapped vault.

Restore access backup vaults (Multi-Party Approval)

An account that has a restore access backup vault for a logically air-gapped vault can create backup access points for recovery points accessible through that vault.

For more information, see Multi-party approval for logically air-gapped vaults.

When access is revoked

If access is removed (logically air-gapped vault unshared via RAM, or restore access backup vault deleted or revoked):

- All backup access points created by the formerly authorized account transition to DISASSOCIATED status
- Lifecycle protection for associated recovery points is removed if the recovery point has no more associated access points.
- S3 API calls through existing S3 access points fail

Monitoring and notifications

AWS Backup emits CloudTrail events, EventBridge events, and Amazon SNS notifications for backup access point status changes (creation, deletion, failure etc.). For details on event formats and how to subscribe, see Monitoring AWS Backup.

Supported features and limitations

Supported features

- **Recovery point types**: Both snapshot and continuous (PITR) recovery points
- **Vault types**: Standard backup vaults and logically air-gapped vaults
- **Cross-account**: Access points can be created for recovery points shared through RAM

Limitations

- **Read-only**: Only read operations (GET, HEAD, LIST) are supported. Write operations are not available.

- **S3 only**: Access points are supported for Amazon S3 recovery points only.
- **Access point limit**: Up to 5 access points per recovery point in an account.
- **Name uniqueness**: Access point names share the S3 access point namespace within a Region and account. Names cannot conflict with existing S3 access points.
- **No scheduled deletion**: Access points do not have a delete-after-days setting. You must delete them manually.
- **External deletion**: If the S3 access point is deleted directly through S3 APIs, the backup access point moves to EXPIRED status and becomes unusable.
- **Access**: The access point creator's account must have access to the recovery point (either owns it, has RAM share access or has a restore access backup vault for the logically air-gapped vault via Multi-Party Approval).
- **Access point name**: After you delete a backup access point, wait for sometime before you attempt to create a new backup access point with the same name in the same account and Region. Attempting to reuse the name before it becomes available returns a `ConflictException`.

Troubleshooting

Common issues

`InvalidParameterValueException` when creating

- Verify the recovery point is in `AVAILABLE`, `STOPPED`, or `COMPLETED` state.
- For continuous recovery points, ensure `AccessPointInTime` is within the retention window.
- Verify the access point name follows S3 naming rules (lowercase, no underscores, 3-50 characters).

Cannot delete recovery point

- Check if active access points exist for the recovery point using `list-backup-access-points-by-recovery-point`.
- Delete all access points before attempting to delete the recovery point.

Access point in `EXPIRED` status

- The backup access point is no longer usable. Delete the expired backup access point using `DeleteBackupAccessPoint`, then create a new one if needed.

S3 API calls return `409 InvalidBucketState`

- The backup access point is not yet in `AVAILABLE` status. Wait for creation to complete, then retry.

`LimitExceededException`

- You have reached the maximum of 5 access points for this recovery point. Delete an existing access point (including any in `FAILED` or `EXPIRED` status) before creating a new one.

`ConflictException` when creating

- You are attempting to create a backup access point with a name that was recently deleted. Wait for the name to become available, or choose a different name.

Restore a backup by resource type

How to restore

For console restore instructions and links to documentation for each AWS Backup-supported resource type, see the links at the bottom of this page.

To restore a backup programmatically, use the StartRestoreJob API operation.

The configuration values ("restore metadata") that you need to restore your resource varies depending on the resource that you want to restore. To get the configuration metadata that your backup was created with, you can call GetRecoveryPointRestoreMetadata. Restore metadata examples are also available in the links at the bottom of this page.

Restoring from cold storage typically takes 4 hours more than restoring from warm storage.

For each restore, a restore job is created with a unique job ID—for example, `1323657E-2AA4-1D94-2C48-5D7A423E7394`.

Note

AWS Backup does not provide any service-level agreements (SLAs) for a restore time. Restore times can vary based upon system load and capacity, even for restores containing the same resources.

Non-destructive restores

When you use AWS Backup to restore a backup, it creates a new resource with the backup that you are restoring. This is to protect your existing resources from being destroyed by your restore activity.

Important

The following resource types perform destructive restores that overwrite existing data:

- **Amazon Redshift Serverless** – Namespace restores overwrite all existing data in the target namespace. For more information, see Amazon Redshift Serverless restore considerations.
- **SAP HANA on Amazon EC2** – Restores overwrite the database at the specified target restore location. For more information, see Restore an SAP HANA database on an Amazon EC2 instance.

For all other resource types, AWS Backup creates a new resource during restore.

Restore testing

You can conduct tests on your resources to simulate a restore experience. This helps determine if you meet your organizational Restore Time Objective (RTO) and helps prepare for future restore needs.

For more information, see Restore testing.

Copy tags during a restore

> **Note**
>
> Restores of Amazon DynamoDB, Amazon S3, SAP HANA on Amazon EC2 instances, virtual machines, and Amazon Timestream resources currently do not have this feature available.

Introduction

You can copy tags as you restore a resource if the tags belonged to the protected resource at the time of backup. Tags, which are labels containing a key and value pair, can help you identify and search for resources. When you start a restore job, tags that belonged to the original backed-up resources can be added to the resource being restored.

When you choose to include tags during a restore job, this step can replace the overhead and labor of manually applying tags to resources after a restore job is completed. Note this is distinct from adding new tags to restored resources.

When you restore a backup in the console flow, your source tags will be copied by default. In the console, uncheck the box if you wish to opt out of copying tags to a restored resource

In the API operation `StartRestoreJob`, the parameter `CopySourceTagsToRestoredResource` is set to `false` by default, which will exclude the original source tags from the resource you are restoring. If you wish to *include* tags from the original source, set this to `True`.

Considerations

- A resource can have up to 50 tags, including restored resources. Please see Tagging your AWS resources for more information about tag limits.
- Ensure the correct permissions are present in the role used for restores to copy tags. The default role for restores contains the necessary permissions. A custom role must include additional permissions to tag resources.
- The following resources are not currently supported for restore tag inclusion: VMware Cloud™ on AWS, VMware Cloud™ on AWS Outposts, on-premises systems, SAP HANA on Amazon EC2 instances, Timestream, DynamoDB, Advanced DynamoDB, and Amazon S3.
- For continuous backups, the tags on the original resource as of the most recent backup will be copied to the restored resource.

- Tags will not be copied for item-level restores.
- Tags that were added to a backup after the backup job was completed but were not present on the original resource prior to the backup will not be copied to the restored resource. Only Backups created after May 22, 2023 are eligible for tag copy on restore.

Tag interaction with specific resources

- **Amazon EC2**
  - Tags applied to restored **Amazon EC2** instances are also applied to the attached restored **Amazon EBS** volumes.
  - Tags applied to the EBS volumes attached to source instances are not copied to the volumes attached to restored instances. If you have IAM policies that allow or deny users access to EBS volumes based on their tags, you must manually reassign the required tags to the restored volumes to ensure your policies remain in effect.
- When you restore an **Amazon EFS** resource, it must be copied to a new file system. Restorations to an existing file system cannot have tags copied to it.
- **Amazon RDS**
  - If the RDS cluster that was backed up is still active, tags from this cluster will be copied.
  - If the original cluster is no longer active, tags from the snapshot of the cluster will be copied instead.
  - Tags which were present on the resource at the time of the backup will be copied during the restore regardless if the Boolean parameter for `CopySourceTagsToRestoredResource` is set to `True` or `False`. However, if the snapshot does not contain tags, then the above Boolean setting will be used.
- **Amazon Redshift** clusters, by default, always include tags during a restore job.

Copy tags via the console

1. Open the AWS Backup console
2. In the navigation pane, choose **Protected resources**, and select the Amazon S3 resource ID that you want to restore.
3. On the **Resource details** page, you will see a list of recovery points for the selected resource ID. To restore a resource:

a. In the **Backup** pane, choose the recovery point ID of the resource.

b. In the upper-right corner of the pane, choose **Restore** (alternatively, you can go to the backup vault, find the recovery point, and then click **Actions** then click **Restore**).

4. On the **Restore backup page**, locate the panel named Restore with tags. To include all tags from the original resource, retain the check the box (note in the console this box is checked by default).
5. Click **Restore backup** after you have selected all your preferred settings and roles.

To include tags programmatically

Use the API operation `StartRestoreJob` . Ensure the following Boolean parameter is set to `True`:

```
CopySourceTagsToRestoredResource = true
```

If the boolean parameter `CopySourceTagsToRestoredResource` = `True`, the restore job will copy the tags from the original resource(s) to the restored material.

> **Important**
>
> The restore job will fail if this parameter is included for an unsupported resource (VMware, AWS Outposts, on-premises systems, SAP HANA on EC2 instances, Timestream, DynamoDB, Advanced DynamoDB, and Amazon S3).

```
{
    "RecoveryPointArn": "arn:aws:ec2:us-east-1::image/ami-1234567890a1b234",
    "Metadata": {
        "InstanceInitiatedShutdownBehavior": "stop",
        "DisableApiTermination": "false",
        "EbsOptimized": "false",
        "InstanceType": "t1.micro",
        "SubnetId": "subnet-123ab456cd7efgh89",
        "SecurityGroupIds": "[\"sg-0a1bc2d345ef67890\"]",
        "Placement": "{\"GroupName\":null,\"Tenancy\":\"default\"}",
        "HibernationOptions": "{\"Configured\":false}",
        "IamInstanceProfileName": "UseBackedUpValue",
        "aws:backup:request-id": "1a2345b6-cd78-90e1-2345-67f890g1h2ij"
    },
```

```
    "IamRoleArn": "arn:aws:iam::123456789012:role/EC2Restore",
    "ResourceType": "EC2",
    "IdempotencyToken": "34ab5678-9012-3c4d-5678-efg9h01f23i4",
    "CopySourceTagsToRestoredResource": true
}
```

Troubleshoot tag restore issues

ERROR: Insufficient Permissions

REMEDY: Ensure you have the necessary permissions in your restore role so you can include tags on your restored resource. The default AWS managed service role policy for restores, AWSBackupServiceRolePolicyForRestores, contains the necessary permissions for this task.

If you choose to use a custom role, ensure the following permissions are present:

- `elasticfilesystem:TagResource`
- `storagegateway:AddTagsToResource`
- `rds:AddTagsToResource`
- `ec2:CreateTags`
- `cloudformation:TagResource`

For more information, see API permissions.

Restore job statuses

You can view the status of a restore job on the **Jobs** page of the AWS Backup console. Restore job statuses include **pending**, **running**, **completed**, **aborted**, and **failed**.

Topics

- Restoring an Amazon Aurora cluster
- Amazon Aurora DSQL restore
- Restore CloudFormation stacks
- Restoring a DocumentDB cluster
- Restore a Amazon DynamoDB table
- Restore an Amazon EBS volume

- Restore an Amazon EC2 instance
- Restore an Amazon EFS file system
- Restore an Amazon EKS cluster
- Restore an FSx file system
- Restore a Neptune cluster
- Restore an RDS database
- Restore an Amazon Redshift cluster
- Amazon Redshift Serverless restore
- Restore an SAP HANA database on an Amazon EC2 instance
- Restore S3 data using AWS Backup
- Restore a Storage Gateway volume
- Restore an Amazon Timestream table
- Restore a virtual machine using AWS Backup

Restoring an Amazon Aurora cluster

Use the AWS Backup console to restore Aurora recovery points

AWS Backup restores your Aurora cluster; it does not create or attach an Amazon RDS instance to your cluster. In the following steps, you will create and attach an Amazon RDS instance to your restored Aurora cluster using the CLI.

Restoring an Aurora cluster requires that you specify multiple restore options. For information about these options, see Overview of Backing Up and Restoring an Aurora DB Cluster in the *Amazon Aurora User Guide*. Specifications for the restore options can be found in the API guide for `RestoreDBClusterFromSnapshot`.

To restore an Amazon Aurora cluster

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Protected resources** and the Aurora resource ID that you want to restore.
3. On the **Resource details** page, a list of recovery points for the selected resource ID is shown. To restore a resource, in the **Backups** pane, choose the radio button next to the recovery point ID of the resource. In the upper-right corner of the pane, choose **Restore**.

4. In the **Instance specifications** pane, accept the defaults or specify the options for the **DB engine**, **DB engine version**, and **Capacity type** settings.

   > **Note**
   >
   > If **Serverless** capacity type is selected, a **Capacity settings** pane appears. Specify the options for the **Minimum Aurora capacity unit** and **Maximum Aurora capacity unit** settings, or choose different options from the **Additional scaling configuration** section.

5. In the **Settings** pane, specify a name that is unique for all DB cluster instances owned by your AWS account in the current Region.
6. In the **Network & Security** pane, accept the defaults or specify the options for the **Virtual Private Cloud (VPC)**, **Subnet group**, and **Availability zone** settings.
7. In the **Database options** pane, accept the defaults or specify the options for **Database port**, **DB cluster parameter group**, and **IAM DB Authentication Enabled** settings.
8. In the **Backup** pane, accept the default or specify the option for the **Copy tags to snapshots** setting.
9. In the **Backtrack** pane, accept the default or specify the options for the **Enable Backtrack** or **Disable Backtrack** settings.
10. In the **Encryption** pane, accept the default or specify the options for the **Enable encryption** or **Disable encryption** settings.
11. In the **Log exports** pane, choose the log types to publish to Amazon CloudWatch Logs. The **IAM role** is already defined.
12. In the **Restore role** pane, choose the IAM role that AWS Backup will assume for this restore.

    > **Note**
    >
    > If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

13. After specifying all your settings, choose **Restore backup**.

    The **Restore jobs** pane appears. A message at the top of the page provides information about the restore job.

14. After your restore finishes, attach your restored Aurora cluster to an Amazon RDS instance.

    Using the AWS CLI:

    - For Linux, macOS, or Unix:

    ```
    aws rds create-db-instance --db-instance-identifier sample-instance \
                  --db-cluster-identifier sample-cluster --engine aurora-mysql --db-
    instance-class db.r4.large
    ```

    - For Windows:

    ```
    aws rds create-db-instance --db-instance-identifier sample-instance ^
                  --db-cluster-identifier sample-cluster --engine aurora-mysql --db-
    instance-class db.r4.large
    ```

See continuous backups and point-in-time restore (PITR) for information about continuous backups and restoring to a chosen point in time.

Use the AWS Backup API, CLI, or SDK to restore Amazon Aurora recovery points

Use `StartRestoreJob`. The metadata you can include for a restore job will depend if you are restoring a continuous backup to a point in time (PITR) or if you are restoring a snapshot.

Restore a cluster from a snapshot

You can specify the following metadata for an Aurora snapshot restore job. See `RestoreDBClusterFromSnapshot` in the *Amazon Relational Database Service API Reference* for additional information and accepted values.

```
// Required metadata:
dbClusterIdentifier // string
engine // string

// Optional metadata:
availabilityZones // array of strings
backtrackWindow // long
copyTagsToSnapshot // Boolean
databaseName // string
dbClusterParameterGroupName // string
dbSubnetGroupName // string
```

```
enableCloudwatchLogsExports // array of strings
enableIAMDatabaseAuthentication // Boolean
engineMode // string
engineVersion // string
kmsKeyId // string
optionGroupName // string
port // integer
scalingConfiguration // object
vpcSecurityGroupIds // array of strings
```

Example:

```
"restoreMetadata":"{\"EngineVersion\":\"5.6.10a\",\"KmsKeyId\":\"arn:aws:kms:us-
east-1:234567890123:key/45678901-ab23-4567-8cd9-012d345e6f7\",\"EngineMode\":
\"serverless\",\"AvailabilityZones\":\"[\\\"us-east-1b\\\",\\\"us-east-1e\\\",\\
\"us-east-1c\\\"]\",\"Port\":\"3306\",\"DatabaseName\":\"\",\"DBSubnetGroupName\":
\"default-vpc-05a3b07cf6e193e1g\",\"VpcSecurityGroupIds\":\"[\\\"sg-012d52c68c6e88f00\
\\"]\",\"ScalingConfiguration\":\"{\\\"MinCapacity\\\":2,\\\"MaxCapacity\\\":64,
\\\"AutoPause\\\":true,\\\"SecondsUntilAutoPause\\\":300,\\\"TimeoutAction\\\":
\\\"RollbackCapacityChange\\\"}\",\"EnableIAMDatabaseAuthentication\":\"false\",
\"DBClusterParameterGroupName\":\"default.aurora5.6\",\"CopyTagsToSnapshot\":\"true\",
\"Engine\":\"aurora\",\"EnableCloudwatchLogsExports\":\"[]\"}"
```

Restore a cluster to a point in time (PITR)

You can specify the following metadata when you want to restore an Aurora continuous backup (recovery point) to a specific point in time (PITR). See RestoreDBClusterToPointInTime in the *Amazon Relational Database Service API Reference* for additional information and accepted values.

```
// Required metadata:
dbClusterIdentifier // string
engine // string
restoreToTime // timestamp; must be specified if UseLatestRestorableTime parameter
 isn't provided

// Optional metadata:
backtrackWindow // long
copyTagsToSnapshot // Boolean
dbClusterParameterGroupName // string
dbSubnetGroupName // string
enableCloudwatchLogsExports // array of strings
enableIAMDatabaseAuthentication // Boolean
```

```
engineMode // string
engineVersion // string
kmsKeyId // string
optionGroupName // string
port // integer
scalingConfiguration // object
vpcSecurityGroupIds // array of strings
```

Amazon Aurora DSQL restore

Topics

- Overview
- Restore Aurora DSQL single Region cluster
- Restore an Aurora DSQL multi-Region cluster
- Troubleshoot Aurora DSQL restore issues
- Aurora DSQL restore frequently asked questions

Overview

To restore a Amazon Aurora DSQL single-Region cluster, use the AWS Backup console or CLI to select the recovery point (backup) you wish to restore. To restore a Aurora DSQL multi-Region cluster, you can now use either the AWS Backup console or CLI.

For single-Region restore, include the name, cluster encryption, and deletion protection, then initiate the restore to a newly created cluster.

For multi-Region restore, you'll need to specify additional parameters including a witness Region, peer Region(s), and regional configuration settings. Multi-Region restore creates a cluster that spans multiple AWS Regions, providing enhanced availability and disaster recovery capabilities.

Restore Aurora DSQL single Region cluster

You can restore an Aurora DSQL cluster to a single Region by using the AWS Backup console or AWS CLI.

Console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.

2. Select the "Restore" button next to the recovery point you wish to restore.
3. Configure the settings for the new cluster to which your recovery point will be restored.
    a. By default, the AMK (AWS managed key) will be used to encrypt the restored data. You may alternatively specify a different key.
    b. Deletion protection for your Aurora clusters is enabled by default, but unselect the box to turn off the option.
4. Review the settings; when they are satisfactory, select the **Restore backup** button.

AWS Backup will create a new Aurora DSQL cluster.

AWS CLI

Single Region restore

1. Use the CLI command `aws backup start-restore-job` to restore an Aurora cluster from the specified recovery point.
2. Include the necessary metadata for the restore job. Example:

Example

```
aws backup start-restore-job \
    --recovery-point-arn "arn:aws:dsql:us-east-1:123456789012:cluster/example-
cluster/backup/example-backup" \
    --iam-role-arn "arn:aws:iam::123456789012:role/service-role/
AWSBackupDefaultServiceRole" \
    --metadata '{"regionalConfig":"[{\"region\":\"us-east-1\",
\"isDeletionProtectionEnabled\":true,\"kmsKeyId\":\"my_key\"}]"}' \
    --copy-source-tags-to-restored-resource
```

Restore an Aurora DSQL multi-Region cluster

Aurora DSQL multi-Region cluster restore occurs within a continent group, which is a set of AWS Regions peers. Multi-Region restore requires that the Regions you specify in the operation are contained in one continent group. For more information about multi-Region clusters, see Configuring multi-Region clusters.

The following continent groups are supported:

- **Americas**: US East (N. Virginia), US East (Ohio), US West (Oregon), Canada West (Calgary), Canada (Central)
- **Europe**: Europe (Frankfurt), Europe (Stockholm), Europe (Spain), Europe (Ireland), Europe (London), Europe (Paris)
- **Asia-Pacific**: Asia Pacific (Hong Kong), Asia Pacific (Tokyo), Asia Pacific (Seoul), Asia Pacific (Osaka), Asia Pacific (Mumbai), Asia Pacific (Singapore)

To complete multi-Region restore, ensure you have the following permissions:

- `backup:StartRestoreJob`
- `dsql:UpdateCluster`
- `dsql:AddPeerCluster`
- `dsql:RemovePeerCluster`

You can restore a backup of an Aurora DSQL cluster to multiple Regions using either the AWS Backup console or CLI commands.

> **Tip**
>
> If you have a backup plan with a rule that automatically creates a cross-Region copy to one of the indicated Regions, the created copy can be used for this multi-Region restore.

Multi-Region restore starts with your current Region. You will also need a:

- Peer Region with an identical cross-Region copy of the recovery point in your current Region
- Witness Region, a designated AWS Region that participates in multi-Region cluster configurations by supporting transaction log-only writes without consuming storage for the actual data. For more information about witness Regions, see Creating a multi-Region cluster.

The individual steps are shown below:

Console

The AWS Backup console now supports multi-Region restore for Aurora DSQL clusters, providing a streamlined process for creating clusters that span multiple Regions. For more information about multi-Region clusters, see Configuring multi-Region clusters.

1. Sign in to the AWS Management Console and open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Backup vaults**.
3. Choose the backup vault that contains the Aurora DSQL recovery point you want to restore.
4. Select the recovery point you want to restore, then choose **Restore**.
5. On the restore page, under **Restore options**, select **Add peer Regions** to enable multi-Region restore.
6. Select a **Peer cluster Region** from the dropdown menu. This Region must be within the same continent group as your current Region and also must contain a cross-Region copy from the recovery point in the current (first) Region.
7. Select a **Witness Region** from the dropdown menu. This Region must also be within the same continent group.
8. Configure the **Cluster settings** for both the primary and peer Region clusters:
   a. For the primary cluster, configure:
      - **Cluster encryption** (optional): Select a KMS key for encryption.
      - **Deletion protection**: Enable or disable deletion protection.
   b. For the peer Region cluster, configure:
      - **Peer Region cluster encryption** (optional): Select a KMS key for encryption.
      - **Peer Region cluster deletion protection**: Enable or disable deletion protection.
9. Review your settings and choose **Restore backup**.
10. The console will initiate the multi-Region restore process, which creates clusters in both Regions and automatically links them together.

AWS CLI

Multi-Region restore can now be achieved using the new orchestrated restore metadata with AWS Backup CLI commands. This approach simplifies the process by handling the cluster linking

automatically. For more information about creating multi-Region clusters programmatically, see Configuring multi-Region clusters in the Aurora DSQL User Guide.

> **Important**
>
> Both the primary cluster and peer cluster must be in Regions within the same continent group. The operation will fail if the clusters are in Regions outside the continent group. Supported continent groups include:
>
> - **Americas**: US East (N. Virginia), US East (Ohio), US West (Oregon), Canada West (Calgary), Canada (Central)
> - **Europe**: Europe (Frankfurt), Europe (Stockholm), Europe (Spain), Europe (Ireland), Europe (London), Europe (Paris)
> - **Asia-Pacific**: Asia Pacific (Hong Kong), Asia Pacific (Tokyo), Asia Pacific (Seoul), Asia Pacific (Osaka), Asia Pacific (Mumbai), Asia Pacific (Singapore)

Multi-Region restore through AWS CLI using orchestrated restore metadata

1. Create a restore job using the CLI command `aws backup start-restore-job` with the new multi-Region orchestration metadata:

   Example

```
aws backup start-restore-job \
--recovery-point-arn "arn:aws:backup:us-east-1:123456789012:recovery-
point:abcd1234" \
--iam-role-arn "arn:aws:iam::123456789012:role/service-role/
AWSBackupDefaultServiceRole" \
--metadata '{
    "witnessRegion":"us-west-2",
    "useMultiRegionOrchestration":"true",
    "peerRegions":"[\"us-east-2\"]",
    "regionalConfig":"[{\"region\":\"us-east-1\",\"isDeletionProtectionEnabled
\":true,\"kmsKeyId\":\"arn:aws:kms:us-east-1:123456789012:key/
ba4b3773-4bb8-4a7a-994c-46ede70202f5\"},{\"region\":\"us-west-2\",
\"isDeletionProtectionEnabled\":true,\"kmsKeyId\":\"arn:aws:kms:us-
west-2:123456789012:key/ba4b3773-4bb8-4a7a-994c-46ede70202f5\"}]"
}' \
--copy-source-tags-to-restored-resource
```

The metadata structure includes:

- `witnessRegion`: The Region that will serve as the witness for the multi-Region cluster. For more information, see Resilience in Amazon Aurora DSQL.
- `useMultiRegionOrchestration`: Set to `true` to enable multi-Region orchestration.
- `peerRegions`: An array containing one or more AWS Regions with peer clusters in the multi-Region cluster. For more information, see MultiRegionProperties in the Aurora DSQL API Reference.
- `regionalConfig`: An array containing configuration for each Region:
  - `region`: The AWS Region identifier.
  - `isDeletionProtectionEnabled`: Boolean flag to enable/disable deletion protection. For more information, see CreateCluster in the Aurora DSQL API Reference.
  - `kmsKeyId`: The KMS key ARN for encryption (optional).

  If `regionalConfig` properties are not specified, then default values will be applied: default encryption and `isDeletionProtectionEnabled` = TRUE.

2. Monitor the restore job status using the `aws backup describe-restore-job` command:

```
aws backup describe-restore-job --restore-job-id job-12345678
```

3. Once the restore job completes, you can verify the multi-Region cluster configuration using the Aurora DSQL CLI:

```
aws dsql describe-cluster --cluster-identifier your-cluster-id
```

For more information about multi-Region cluster operations, see UpdateCluster in the Aurora DSQL API Reference.

Troubleshoot Aurora DSQL restore issues

Error: Insufficient permissions

Possible cause: If you try to copy an Aurora DSQL recovery point into an account (cross-account copy) that has never interacted with DSQL API, you may get a permission issue error since the DSQL service-linked role isn't set up in the destination account.

Remedy: Attach the DSQL managed policy that includes the DSQL service-linked role, AuroraDsqlServiceLinkedRolePolicy, to a role in the destination account.

If you encounter any other issues with the backup or restore process, you can check the status of your backup and restore jobs in the AWS Backup console or using the AWS CLI. Additionally, you can review the AWS CloudTrail logs for any relevant error messages or events related to your AWS Backup operations.

Aurora DSQL restore frequently asked questions

1. *"Can I use AWS Backup for Aurora DSQL from the Aurora DSQL console?"*

   No, you can only perform backups and restores, as well as managing backups, from AWS Backup console, SDK, or CLI.

2. *"What backup granularity is available for Aurora DSQL? Can I backup specific tables or databases in my cluster"*

   You can only back up and restore a whole Aurora DSQL cluster.

3. *"Are backups of Aurora DSQL full backups or incremental backups?"*

   Recovery points of Aurora DSQL clusters (backups) are full backups of your clusters.

4. *"Can I create backups for my Aurora DSQL multi-Region clusters?"*

   Yes, you can create backups for each cluster in multi-Region clusters in the using the same steps as when you create a backup of a single cluster in a single Region.

   AWS Backup recommends as a best practice to create a cross-Region copy of your backup in the other Region from which you plan to restore the Multi-Region cluster, as multi-Region restore requires an identical copy of the same recovery point [*identical* in this operation means the recovery points have the same resource name and creation time].

5. *"Will my restored cluster overwrite my existing cluster?"*

   No. When you restore your Aurora DSQL data, AWS Backup creates a new cluster from your snapshots; the restored cluster won't overwrite the source cluster.

Restore CloudFormation stacks

A CloudFormation composite backup is a combination of a CloudFormation template and all associated nested recovery points. Any number of nested recovery points can be restored, but the composite recovery point (which is the top-level recovery point) cannot be restored.

When you restore a CloudFormation template recovery point, you create a new stack with a change set to represent the backup.

Restore CloudFormation with the AWS Backup console

From the CloudFormation console you can see the new stack and change set. To learn more about change sets, see Updating stacks using change sets in the *CloudFormation User Guide*.

Determine which nested recovery points you want to restore from with your CloudFormation stack, and then restore them using the AWS Backup console.

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Go to **Backup vaults**, select the backup vault containing your desired recovery point, then click on **Recovery points**.
3. Restore the CloudFormation template recovery point.
    a. Click the composite recovery point containing the nested recovery points you want to restore to bring up the Details page for the composite recovery point.
    b. Under **Nested recovery points**, the nested recovery points will be displayed. Each recovery point will have a recovery point ID, a status, a resource ID, a resource type, a backup type, and the time that recovery point was created. Click the radio button next to the CloudFormation recovery point, then click **Restore**. Ensure that you are selecting the recovery point that has **resource type: CloudFormation** and **backup type: backup.**
4. Once the restore job for the CloudFormation template is completed, your restored CloudFormation template will be visible in the CloudFormation console under **Stacks**.
5. Under **Stack names** you should find the restored template with the status of `REVIEW_IN_PROGRESS`.
6. Click on the name of the stack to see the stack's details.
7. There are tabs under the stack name. Click on **Change sets**.
8. Execute the change set.

9. After this processes, the resources in the original stack will be recreated in the new stack. The stateful resources will be recreated empty. For example, a stateful resource is an Amazon EBS volume or an Amazon EFS file system. To recover the stateful resources, go back to the list of recovery points in the AWS Backup console, select the recovery point you need, and initiate a restore.

Important

CloudFormation performs pre-deployment validation, including resource existence checks, when a change set is created. If a resource defined in the restored template already exists in your account (for example, a DynamoDB table with a fixed name), the restore operation fails at the create change set step (step 8) rather than at the execute change set step (step 9). For more information, see Considerations for validating stack deployments in the *AWS CloudFormation User Guide*.

Note

If a CloudFormation restore operation fails, the stack may remain in `REVIEW_IN_PROGRESS` status with a `FAILED` change set. Delete these stacks manually to avoid naming conflicts when you retry the restore operation.
For more information, see Deleting a stack in the *AWS CloudFormation User Guide*.

Restore CloudFormation with AWS CLI

In the command line interface, `start-restore-job` allows you to restore a CloudFormation stack.

The following list is the accepted metadata to restore an CloudFormation resource.

```
// Mandatory metadata:
ChangeSetName // This is the name of the change set which will be created
StackName // This is the name of the stack that will be created by the new change set

// Optional metadata:
ChangeSetDescription // This is the description of the new change set
StackParameters // This is the JSON of the stack parameters required by the stack
```

```
aws:backup:request-id
```

Restoring a DocumentDB cluster

Use the AWS Backup console to restore Amazon DocumentDB recovery points

Restoring a Amazon DocumentDB cluster requires that you specify multiple restore options. For information about these options, see Restoring from a Cluster Snapshot in the *Amazon DocumentDB Developer Guide*.

To restore a Amazon DocumentDB cluster

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Protected resources** and the Amazon DocumentDB resource ID that you want to restore.
3. On the **Resource details** page, a list of recovery points for the selected resource ID is shown. To restore a resource, in the **Backups** pane, choose the radio button next to the recovery point ID of the resource. In the upper-right corner of the pane, choose **Restore**.
4. Ensure you are on the console page **Restore Amazon DocumentDB cluster snapshots**.
5. For **Restore options**, you can configure the following:

   - **Engine version** - Select the DocumentDB engine version for the restored cluster.

   > **Note**
   >
   > Instance class and number of instances cannot be configured during the restore process. The restored DocumentDB cluster will use the default instance configuration. You can modify the instance class and add or remove instances after the restore completes by using the Amazon DocumentDB console or API.

6. In the **Settings** pane, input a unique name for your DB cluster identifier.

   You can use letters, numbers, and hyphens, though you cannot have two consecutive hyphens or end the name with a hyphen. The final name will be all lowercase.

7. In the **Database options** pane, select the database port.

   This is the TCP/IP port that the DB instance or cluster will use for application connections. The connection string of any application connecting to the DB instance or cluster must specify

its port number. Both the security group applied to the DB instance or cluster and your organization firewalls must allow connections to the port. All DB instances in a DB cluster use the same port.

8. Also in the **Database options** pane, select the DB cluster parameter group.

   This is the parameter group associated with this instance's DB cluster. The DB cluster parameter group acts as a container for engine configuration values that are applied to every DB instance in the cluster.

9. In the **Encryption** pane, select the key that will be used to encrypt this database volume. The default is `aws/rds`. You may alternatively use a customer managed key (CMK).

10. In the **Log exports** pane, choose the log types to publish to Amazon CloudWatch Logs. The **IAM role** is already defined.

11. In the **Restore role** pane, choose either the default IAM role for the restore job or a different IAM role.

    > **Note**
    >
    > If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

12. In the Protected resource tags pane, you may optionally choose to copy tags from the backup to the restored database cluster.

13. After specifying all your settings, choose **Restore backup**.

    The **Restore jobs** pane appears. A message at the top of the page provides information about the restore job.

14. After your restore finishes, attach your restored Amazon DocumentDB cluster to an Amazon RDS instance.

Use the AWS Backup API, CLI, or SDK to restore Amazon DocumentDB recovery points

First, restore your cluster. Use `StartRestoreJob`. You can specify the following metadata during Amazon DocumentDB restores:

```
availabilityZones
backtrackWindow
copyTagsToSnapshot // Boolean
databaseName // string
dbClusterIdentifier // string
dbClusterParameterGroupName // string
dbSubnetGroupName // string
enableCloudwatchLogsExports // string
enableIAMDatabaseAuthentication // Boolean
engine // string
engineMode // string
engineVersion // string
kmsKeyId // string
port // integer
optionGroupName // string
scalingConfiguration
vpcSecurityGroupIds // string
```

Then, attach your restored Amazon DocumentDB cluster to an Amazon RDS instance using `create-db-instance`.

- For Linux, macOS, or Unix:

```
aws docdb create-db-instance --db-instance-identifier sample-instance /
                  --db-cluster-identifier sample-cluster --engine docdb --db-
instance-class db.r5.large
```

- For Windows:

```
aws docdb create-db-instance --db-instance-identifier sample-instance ^
                  --db-cluster-identifier sample-cluster --engine docdb --db-
instance-class db.r5.large
```

Restore a Amazon DynamoDB table

Use the AWS Backup console to restore DynamoDB recovery points

To restore a DynamoDB table

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Protected resources** and the DynamoDB resource ID you want to restore.
3. On the **Resource details** page, a list of recovery points for the selected resource ID is shown. To restore a resource, in the **Backups** pane, choose the radio button next to the recovery point ID of the resource. In the upper-right corner of the pane, choose **Restore**.
4. For **Settings**, **New table name** text field, enter a new table name.
5. For **Restore role**, choose the IAM role that AWS Backup will assume for this restore.

   > **Note**
   >
   > If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

6. For **Encryption settings**:

   a. If your backup is managed by DynamoDB (its ARN begins with `arn:aws:dynamodb`), AWS Backup encrypts your restored table using an AWS-owned key.

      To choose a different key to encrypt your restored table, you can either use the AWS Backup StartRestoreJob operation or perform the restore from the DynamoDB console.

   b. If your backup supports full AWS Backup management (its ARN begins with `arn:aws:backup`), you can choose any of the following encryption options to protect your restored table:

      - (Default) DynamoDB-owned KMS key (no additional charge for encryption)
      - DynamoDB-managed KMS key (KMS charges apply)
      - Customer-managed KMS key (KMS charges apply)

"DynamoDB-owned" and "DynamoDB-managed" keys are the same as "AWS-owned" and "AWS-managed" keys, respectively. For clarification, see Encryption at Rest: How It Works in the *Amazon DynamoDB Developer Guide*.

For more information about full AWS Backup management, see Advanced DynamoDB backup.

> **Note**
>
> The following guidance applies only if you restore a copied backup AND want to encrypt the restored table with the same key you used to encrypt your original table.
> When restoring a cross-Region backup, to encrypt your restored table using the same key you used to encrypt your original table, your key must be a multi-Region key. AWS-owned and AWS-managed keys are not multi-Region keys. To learn more, see Multi-Region keys in the *AWS Key Management Service Developer Guide*.
> When restoring a cross-account backup, to encrypt your restored table using the same key you used to encrypt your original table, you must share the key in your source account with your destination account. AWS-owned and AWS-managed keys cannot be shared between accounts. To learn more, see Allowing users in other accounts to use a KMS key in the *AWS Key Management Service Developer Guide*.

7. Choose **Restore backup**.

   The **Restore jobs** pane appears. A message at the top of the page provides information about the restore job.

Use the AWS Backup API, CLI, or SDK to restore DynamoDB recovery points

Use `StartRestoreJob`. You can specify the following metadata during any DynamoDB restore. The metadata is not case-sensitive.

```
targetTableName
encryptionType
kmsMasterKeyArn
aws:backup:request-id
```

The following is an example of the `restoreMetadata` argument for a `StartRestoreJob` operation in the CLI:

```
aws backup start-restore-job \
--recovery-point-arn "arn:aws:backup:us-east-1:123456789012:recovery-point:abcdef12-
g3hi-4567-8cjk-012345678901" \
--iam-role-arn "arn:aws:iam::123456789012:role/YourIamRole" \
--metadata
 'TargetTableName=TestRestoreTestTable,EncryptionType=KMS,kmsMasterKeyArn=arn:aws:kms:us-
east-1:123456789012:key/abcdefg' \
--region us-east-1 \
--endpoint-url https://endpointurl.com
```

The preceding example encrypts the restored table using a customer-managed key.

To encrypt your restored table using an AWS-owned key, specify the following restore metadata: `"encryptionType\":\"Default\"`.

To encrypt your restored table using an AWS-managed key, omit the `kmsMasterKeyArn` parameter and specify: `"encryptionType\":\"KMS\"`.

To encrypt your restored table using a customer-managed key, specify the following restore metadata: `"encryptionType\":\"KMS\",\"kmsMasterKeyArn\":\"`*`arn:aws:kms:us-west-2:111122223333:key/1234abcd-12ab-34cd-56ef-1234567890ab`*`\"`.

Restore an Amazon EBS volume

When you restore an Amazon Elastic Block Store (EBS) snapshot, you can choose to restore it as an EBS volume, restore it to a AWS Storage Gateway volume, or restore selected items from it to an Amazon S3 bucket.

Restore to an EBS volume

When you restore a snapshot (periodic backup of EBS data) to a new volume, you will specify the volume type, size in GiB, and an availability zone. You can optionally choose to encrypt the new volume with an existing or new AWS KMS key.

Restore to a gateway volume

When you restore to a gateway volume, you will need to specify a gateway in a reachable state, choose your iSCSI target name, and choose a disk ID if your gateway is volume stored or a capacity equal or greater than your snapshot if your gateway is volume cached.

File level restore to an Amazon S3 bucket

Prior to starting a restore job of EBS resources to an Amazon S3 bucket, review EBS permissions and Amazon S3 restore permissions for access requirements. The necessary permissions are contained in the managed policy AWSBackupServiceRolePolicyForItemRestores and should be included in the IAM role used for the restore operation.

All new object uploads, including restored data, to an S3 bucket is automatically encrypted. When you choose this type of restore, specify SSE-S3 (server-side Amazon S3 managed key) or SSE-KMS (server-side AWS KMS managed key). SSE-S3 is the default.

You can input up to five paths when restoring from the AWS Backup console; you can specify multiple paths through the command line. A path must have a length less than 1024 bytes in UTF-8 encoded strings, including the user-designated and AWS Backup-designated prefixes

If your snapshot contains multiple partitions, specify the file system identifier of the partition that contains the data you plan to restore. This identifier can be found using Backup search and is the same of the UUID or file system Disk ID.

	To new EBS volume	To gateway	File level restore to S3 bucket
Encryption	**Optional**. You can choose an existing AWS KMS key or create a new KMS key.		Required. Choose from SSE-S3, SSE-KMS, or the default destination bucket encryption[1].
Permissions and roles	Choose existing role; If none exists, default role with correct permissions is created.	Choose existing role;If none exists, default role with correct permissions is created	Role choice must have sufficient EBS and Amazon S3 restore permissions.
Restore from cold storage (EBS Archive Tier)	Available	Unavailable	Unavailable

	To new EBS volume	To gateway	File level restore to S3 bucket
Settings to specify	Volume type; size (GiB); Availability zone; Throughput	Gateway (in a reachable state); iSCSI target name; Disk id (for volume stored gateways); Capacity (for volume cached gateways)	Restore type, including: Destination bucket name; Path(s) to restore; Encryption type; File level restore KMS Key Id if SSE-KMS is set as encryption type

[1]In the AWS Backup console, you select one of the three encryption options; if you use CLI to restore, omit `encryptionType` to restore to the default destination bucket encryption.

Restore an EBS snapshot with the AWS Backup console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Protected resources** and then choose the EBS resource ID you want to restore.
3. On the **Resource details** page, a list of recovery points for the selected resource ID is shown. To restore a resource, in the **Backups** pane, choose the radio button next to the recovery point ID of the resource. In the upper-right corner of the pane, choose **Restore**.
4. Specify the restore parameters for your resource. The restore parameters you enter are specific to the resource type that you selected.

   For **Resource type**, choose the AWS resource to create when restoring this backup.
5. If you choose **EBS volume**, provide the values for **Volume type**, **Size (GiB)**, and choose an **Availability zone**. After **Throughput**, there will be an optional checkbox **Encrypt this volume**. *This option will stay active if the EBS recovery point is encrypted.*. You may specify a KMS key or you may create an AWS KMS key.

   If you choose **Storage Gateway volume**, choose a **Gateway** in a reachable state. Also choose your **iSCSI target name**. For *Volume stored* gateways, choose a **Disk Id**. For *Volume cached* gateways, choose a capacity that is at least as large as your protected resource.

If you choose **file level restore**, you can include up to 5 objects or folders from the snapshot. You can search your indexed backups to find the file name or path.

- Input the file paths.
- Choose to use an existing Amazon S3 bucket or create a new bucket for the destination where the objects or folders will be restored.
- Set the encryption of the restored object(s). You can choose the default destination bucket encryption, SSE-S3, or SSE-KMS. For additional detail, see Restore S3 data using AWS Backup.

6. For **Restore role**, choose the IAM role that AWS Backup will assume for this restore. If the AWS Backup default role is not present in your account, a **Default role** is created for you with the correct permissions. You can delete this default role or make it unusable.

> **Note**
>
> If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

7. Choose **Restore backup** (**Restore items** is displayed for file level restore).

   The **Restore jobs** pane will appear. A message at the top of the page provides information about the restore job.

Restore from archived EBS snapshots

Restoring an archived EBS snapshot moves it from cold to warm storage temporarily to create a new EBS volume. This type of restore incurs a one-time retrieval charge. Storage costs for both warm and cold storage are billed during this restore period.

> **Tip**
>
> EBS volumes in cold storage can't be restored to a gateway volume or be restored at the file level.

You can restore an archived EBS snapshot in cold storage by using the AWS Backup console or the command line. A restore from cold storage can take up to 72 hours. For more information, see Archive Amazon EBS snapshots in the *Amazon EBS User Guide*.

Console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Navigate to **Backup vaults** > *Vault* > **Restore archived EBS snapshot**.
3. In the **Settings** section, input a value from 0 to 180, inclusive, that specifies the number of days to temporarily restore an archived snapshot.
4. Input other settings: volume type, size, IOPS, availability zone, throughput, and encryption.
5. Choose your **restore role**.

> **Note**
>
> If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

6. Select **Restore backup**. On the confirmation pop up, confirm the snapshots and restore type. Then, select **Restore snapshot**.

AWS CLI

1. Use `start-restore-job`
2. Include the required parameters for archived EBS snapshot restore:

```
--recovery-point-arn arn:aws:backup:region:account-id:recovery-point:recovery-
point-id
--metadata
 '{"temporaryRestoreDays":"value","volumeType":"value","volumeSize":"value","availabilit
--iam-role-arn arn:aws:iam::account-id:role/service-role/
AWSBackupDefaultServiceRole
--resource-type EBS
```

3. Specify the temporary restore duration (0-180 days) in the `temporaryRestoreDays` parameter. This determines how long the archived snapshot will be available in warm storage.

4. Configure the new EBS volume settings including `volumeType` (gp2, gp3, io1, io2, st1, sc1), `volumeSize` in GiB, and target `availabilityZone`.
5. Monitor the restore job status using `describe-restore-job` with the returned restore job ID. Archive restores can take up to 72 hours to complete.

Restore an EBS snapshot by AWS CLI

To restore Amazon EBS using the API or CLI, use `StartRestoreJob`. You can specify the following metadata during an Amazon EBS restore:

```
aws:backup:request-id
availabilityZone
encrypted // if set to true, encryption will be enabled as volume is restored
iops
kmsKeyId // if included, this key will be used to encrypt the restored volume instead
 of default KMS Key Id
restoreType // include for file level restore - see details below
throughput
temporaryRestoreDays
volumeType
volumeSize
```

Example:

```
"restoreMetadata": "{\"encrypted\":\"false\",\"volumeId\":\"vol-04cc95f3490b5ceea\",
\"availabilityZone\":null}"
```

File level restore specifications

`restoreType` is required for file level restore. For this type of restore, the following unique metadata is required:

```
destinationBucketName //
pathsToRestore //
encryptionType // You can specify SSE-S3 or SSE-KMS; do not include if you want to
 restore to default encryption
kmsKeyId //
```

Filesystem identifier is optional for single partition Snapshots. If this information is not passed, then just the absolute path without the ":" separator (such as `{"/data/process/abc.txt", "/data/department/xyz.txt"}`) will be accepted.

Restore an Amazon EC2 instance

When you restore an EC2 instance, AWS Backup creates an Amazon Machine Image (AMI), an instance, the Amazon EBS root volume, Amazon EBS data volumes (if the protected resource had data volumes), and Amazon EBS snapshots. You can customize some instance settings using the AWS Backup console, or a larger number of settings using the AWS CLI or an AWS SDK.

The following considerations apply to restoring EC2 instances:

- AWS Backup configures the restored instance to use the same key pair that the protected resource used originally. You can't specify a different key pair for the restored instance during the restore process.
- AWS Backup does not back up and restore user-data that is used while launching an Amazon EC2 instance.
- When configuring the restored instance, you can choose between using the same instance profile that the protected resource used originally or launching without an instance profile. This is to prevent the possibility of privilege escalations. You can update the instance profile for the restored instance using the Amazon EC2 console.

  If you use the original instance profile, you must grant AWS Backup the following permissions, where the resource ARN is the ARN of the IAM role associated with the instance profile.

  ```
  {
        "Effect": "Allow",
        "Action": "iam:PassRole",
        "Resource": "arn:aws:iam::account-id:role/role-name"
  },
  ```

  Replace *role-name* with the name of the EC2 instance profile role that will be attached to the restored EC2 instance. This is not the AWS Backup service role, but rather the IAM role that provides permissions to applications running on the EC2 instance.
- During a restore, all Amazon EC2 quotas and configuration restrictions apply.
- If the vault containing your Amazon EC2 recovery points has a vault lock, see Additional security considerations for more information.

- AWS Backup does not automatically join restored EC2 instances to an Active Directory domain. If the original instance was domain-joined, you must rejoin the restored instance manually or by using automation. For example, you can trigger an AWS Lambda function in response to the restore job completed event. For more information, see State: COMPLETED. To join an instance to a directory during launch, see Seamlessly join an EC2 instance to your AD Connector directory.

Use the AWS Backup console to restore Amazon EC2 recovery points

You can restore an entire Amazon EC2 instance from a single recovery point, including the root volume, data volumes, and some instance configuration settings, such as the instance type and key pair.

To restore Amazon EC2 resources using the AWS Backup console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Protected resources**, then choose the ID of the Amazon EC2 resource to open the resource details page.
3. In the **Recovery points** pane, choose the radio button next to the ID of the recovery point to restore. In the upper-right corner of the pane, choose **Restore**.
4. In the **Network settings** pane, we use the settings from the protected instance to select the default values for the instance type, VPC, subnet, security group, and instance IAM role. You can use these default values or change them as needed.
5. In the **Restore role** pane, use the **Default role** or use **Choose an IAM role** to specify an IAM role that grants AWS Backup permission to restore the backup.

   > **Note**
   >
   > If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

6. In the **Protected resource tags** pane, we select **Copy tags from the protected resource to the restored resource** by default. If you do not want to copy these tags, clear the check box.
7. In the **Advanced settings** pane, accept the default values for the instance settings or change them as needed. For information about these settings, choose **Info** for the setting to open its help pane.

8. When you are finishing configuring the instance, choose **Restore backup**.

Restore Amazon EC2 with AWS CLI

In the command line interface, `start-restore-job` allows you to restore with up to 32 parameters (including some parameters that are not customizable through the AWS Backup console).

The following list is the accepted metadata you can pass to restore an Amazon EC2 recovery point.

```
InstanceType
KeyName
SubnetId
Architecture
EnaSupport
SecurityGroupIds
IamInstanceProfileName
CpuOptions
InstanceInitiatedShutdownBehavior
HibernationOptions
DisableApiTermination
CreditSpecification
Placement
RootDeviceType
RamdiskId
KernelId
UserData
Monitoring
NetworkInterfaces
ElasticGpuSpecification
CapacityReservationSpecification
InstanceMarketOptions
LicenseSpecifications
EbsOptimized
VirtualizationType
Platform
RequireIMDSv2
BlockDeviceMappings
aws:backup:request-id
```

AWS Backup accepts the following information-only attributes. However, including them will not affect the restore:

```
vpcId
```

`BlockDeviceMappings` is an optional parameter you can include. AWS Backup supports the following `BlockDeviceMappings` attributes.

> **Note**
>
> `SnapshotId` and `OutpostArn` are not supported.

```
{
  "BlockDeviceMappings": [
    {
        "DeviceName" : string,
        "NoDevice" : string,
        "VirtualName" : string,
        "Ebs": {
            "DeleteOnTermination": boolean,
            "Iops": number,
            "VolumeSize": number,
            "VolumeType": string,
            "Throughput": number,
            "Encrypted": boolean,
            "KmsKeyId": string
        }
    }
 }
```

For example:

```
{
  "BlockDeviceMappings": [
    {
      "DeviceName": "/def/tuvw",
      "Ebs": {
        "DeleteOnTermination": true,
        "Iops": 3000,
        "VolumeSize": 16,
        "VolumeType": "gp3",
        "Throughput": 125,
        "Encrypted": true,
```

```
        "KmsKeyId": "arn:aws:kms:us-west-2:123456789012:key/ab3cde45-67f8-9g01-
hi2j-3456klmno7p8"
      }
    },
    {
      "DeviceName": "/abc/xyz",
      "Ebs": {
        "DeleteOnTermination": false,
        "Iops": 3000,
        "VolumeSize": 16,
        "VolumeType": "gp3",
        "Throughput": 125,
        "Encrypted": false
      }
    }
  ]
}
```

You can also restore an Amazon EC2 instance without including any stored parameters. This option is available on the **Protected resource** tab on the AWS Backup console.

> **Important**
>
> If you do not override the AWS KMS key in the `BlockDeviceMappings` when restoring from cross-account or cross-Region backups, your restore might fail. For more information, see Troubleshoot Amazon EC2 instance restore issues.

Troubleshoot Amazon EC2 instance restore issues

Troubleshooting issues

- Cross-account restore failures
- Cross-Region restore failures

Cross-account restore failures

Description: Amazon EC2 instance restore fails when attempting to restore from a backup that is shared with your account.

Possible issues: Your account might not have access to the AWS KMS keys used to encrypt the source volumes in the sharing account. The KMS keys might not be shared with your account.

Or, the volumes attached to the source instance are unencrypted.

Solution: To resolve this issue, set the `encrypted` attribute to `true`, and do one of the following:

- Override the KMS keys in the `BlockDeviceMappings` and specify a KMS key that you own in your account.
- Request the owning account to grant you access to the KMS keys used to encrypt the volumes by updating the KMS key policy. For more information, see Allow users in other accounts to use a KMS key.

Cross-Region restore failures

Description: Amazon EC2 instance restore fails when attempting to restore from a cross-Region backup.

Issue: The volumes in the backup might be encrypted with single-Region AWS KMS keys that are not available in the destination Region. Or, the volumes attached to the source instance are unencrypted.

Solution: To resolve this issue, set the `encrypted` attribute to `true`, and override the KMS key in the `BlockDeviceMappings` with a KMS key in the destination Region.

Restore an Amazon EFS file system

If you are restoring an Amazon Elastic File System (Amazon EFS) instance, you can perform a full restore or an item-level restore.

Full Restore

When you perform a full restore, the entire file system is restored.

AWS Backup does not support destructive restores with Amazon EFS. A destructive restore is when a restored file system deletes or overwrites the source or existing file system. Instead, AWS Backup restores your file system to a recovery directory off of the root directory.

Important

After a full restore completes, all restored files are located in a recovery subdirectory under the root of the file system. An example subdirectory is `aws-backup-restore_`*`datetime`*. Files are not restored to their original locations. You must manually move or copy the restored files from this recovery directory to their intended locations. The same behavior applies to item-level restores.
For additional troubleshooting guidance related to Amazon EFS restores, see the following AWS Knowledge Center articles:

- Troubleshooting EFS file system restore issues for AWS Backup
- Getting accurate Amazon EFS file system and object sizes after a restore in AWS Backup

Item-Level Restore

When you perform an item-level restore, AWS Backup restores a specific file or directory. You must specify the path relative to the file system root. For example, if the file system is mounted to `/user/home/myname/efs` and the file path is `user/home/myname/efs/file1`, you enter **`/file1`**. Paths are case sensitive. Wildcard characters and regex strings are not supported. Your path may be different from what is in the host if the file system is mounted using an access point.

You can select up to 10 items when you use the console to perform an EFS restore. There is no item limit when you use CLI to restore; however, there is a 200 KB limit on the length of the restore metadata that can be passed.

You can restore those items to either a new or existing file system. Either way, AWS Backup creates a new Amazon EFS directory (`aws-backup-restore_`*`datetime`*) off of the root directory to contain the items. The full hierarchy of the specified items is preserved in the recovery directory. For example, if directory A contains subdirectories B, C, and D, AWS Backup retains the hierarchical structure when A, B, C, and D are recovered. Regardless of whether you perform an Amazon EFS item-level restore to an existing file system or to a new file system, each restore attempt creates a new recovery directory off of the root directory to contain the restored files. If you attempt multiple restores for the same path, several directories containing the restored items might exist.

> **Note**
>
> As with full restores, item-level restored files are placed in a recovery subdirectory and not in their original locations. You must manually move the restored files to their intended locations.

> **Note**
>
> If you only keep one weekly backup, you can only restore to the state of the file system at the time you took that backup. You can't restore to prior incremental backups.

Use the AWS Backup console to restore an Amazon EFS recovery point

To restore an Amazon EFS file system

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Your EFS backup vault receives the access policy Deny `backup:StartRestoreJob` upon creation. If you are restoring your backup vault for the first time, you must change your access policy as follows.
   a. Choose **Backup vaults**.
   b. Choose the backup vault containing the recovery point you would like to restore.
   c. Scroll down to the vault **Access policy**
   d. If present, delete `backup:StartRestoreJob` from the `Statement`. Do this by choosing **Edit**, deleting `backup:StartRestoreJob`, then choosing **Save policy**.
3. In the navigation pane, choose **Protected resources** and the EFS file system ID you want to restore.
4. On the **Resource details** page, a list of recovery points for the selected file system ID is shown. To restore a file system, in the **Backups** pane, choose the radio button next to the recovery point ID of the file system. In the upper-right corner of the pane, choose **Restore**.
5. Specify the restore parameters for your file system. The restore parameters you enter are specific to the resource type that you selected.

You can perform a **Full restore**, which restores the entire file system. Or, you can restore specific files and directories using **Item-level restore**.

- Choose the **Full restore** option to restore the file system in its entirety including all root level folders and files.
- Choose the **Item-level restore** option to restore a specific file or directory. You can select and restore up to five items within your Amazon EFS.

  To restore a specific file or directory, you must specify the relative path related to the mount point. For example, if the file system is mounted to `/user/home/myname/efs` and the file path is `user/home/myname/efs/file1`, enter **`/file1`**. Paths are case sensitive and cannot contain special characters, wildcard characters, and regex strings.

  1. In the **Item path** text box, enter the path for your file or folder.
  2. Choose **Add item** to add additional files or directories. You can select and restore up to five items within your EFS file system.

6. For **Restore location**

   - Choose **Restore to directory in source file system** if you want to restore to the source file system.
   - Choose **Restore to a new file system** if you want to restore to a different file system.

7. For **File system type**

   - (Recommended) Choose **Regional** if you want to restore your file system across multiple AWS Availability Zones.
   - Choose **One Zone** if you want to restore your file system to a single Availability Zone. Then, in the **Availability Zone** dropdown, choose the destination for your restore.

   For more information, see Managing Amazon EFS storage classes in the *Amazon EFS User Guide.*

8. For **Performance**

   - If you chose to perform a **Regional** restore, choose either **(Recommended) General purpose** or **Max I/O**.
   - If you chose to perform a **One Zone** restore, you must choose **(Recommended) General purpose**. One Zone restores do not support **Max I/O**.

9. For **Enable encryption**

   - Choose **Enable encryption**, if you want to encrypt your file system. KMS key IDs and aliases appear in the list after they have been created using the AWS Key Management Service (AWS KMS) console.
   - In the **KMS key** text box, choose the key you want to use from the list.

10. For **Restore role**, choose the IAM role that AWS Backup will assume for this restore.

> **Note**
>
> If the AWS Backup default role is not present in your account, a **Default role** is created for you with the correct permissions. You can delete this default role or make it unusable.

> **Note**
>
> If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

11. Choose **Restore backup**.

    The **Restore jobs** pane appears. A message at the top of the page provides information about the restore job.

> **Note**
>
> If you only keep one weekly backup, you can only restore to the state of the file system at the time you took that backup. You can't restore to prior incremental backups.

Use the AWS Backup API, CLI, or SDK to restore Amazon EFS recovery points

Use `StartRestoreJob`. When restoring an Amazon EFS instance, you can restore an entire file system or specific files or directories. To restore Amazon EFS resources, you need the following information:

- `file-system-id` — The ID of the Amazon EFS file system that is backed up by AWS Backup. Returned in `GetRecoveryPointRestoreMetadata`. This is not required when a **new** file system is restored (this value is ignored if parameter `newFileSystem` is `True`).
- `Encrypted` — A Boolean value that, if true, specifies that the file system is encrypted. If `KmsKeyId` is specified, `Encrypted` must be set to `true`.
- `KmsKeyId` — Specifies the AWS KMS key that is used to encrypt the restored file system.
- `PerformanceMode` — Specifies the throughput mode of the file system. Valid values are `generalPurpose` (default) and `maxIO`. The `generalPurpose` mode provides the lowest latency per operation and can achieve up to 7,000 file operations per second. The `maxIO` mode can scale to higher levels of aggregate throughput and operations per second with a slightly higher latency for file operations.
- `CreationToken` — A user-supplied value that ensures the uniqueness (idempotency) of the request.
- `newFileSystem` — A Boolean value that, if true, specifies that the recovery point is restored to a new Amazon EFS file system.
- `ItemsToRestore` — An array of up to five strings where each string is a file path. Use `ItemsToRestore` to restore specific files or directories rather than the entire file system. This parameter is optional.

You may also include `aws:backup:request-id`.

One Zone restores can be performed by including parameters:

```
"singleAzFilesystem": "true"
"availabilityZoneName": "ap-northeast-3"
```

For more information about Amazon EFS configuration values, see create-file-system.

Disabling automatic backups in Amazon EFS

By default, Amazon EFS creates backups of data automatically. These backups are represented as recovery points in AWS Backup. Attempts to remove the recovery point will result in an error message that notes there are insufficient privileges to perform the action.

It is best practice to keep this auto-backup active. Particularly in the case of accidental data deletion, this backup allows restoration of file system content to the date of the last recovery point created.

In the unlikely event you wish to turn these off, the access policy must be changed from `"Effect": "Deny"` to `"Effect": "Allow"`. See the *Amazon EFS User Guide* for more information about turning automatic backups on or off.

Restore an Amazon EKS cluster

You can restore EKS cluster backups using the AWS Backup console or CLI. EKS backups are composite recovery points that include both EKS cluster state and persistent volume backups.

AWS Backup supports multiple restore experiences including granular namespace-level restores. Restores are non-destructive and will not overwrite any existing Kubernetes objects in your target EKS cluster. Restores will also not overwrite the Kubernetes versions of the target EKS cluster.

EKS Backups have to be restored to a target EKS cluster, meaning an Amazon EKS cluster that has been pre-provisioned. As part of the restore workflow, you can opt to create a new EKS cluster which AWS Backup will create on your behalf.

> **Note**
>
> AWS Backup will provide a limited set of options for creating a new EKS cluster as a part of a restore. For all EKS cluster creation functionality, customers can create a new EKS cluster using the EKS Console or APIs and select this as their restore target.

Restore capabilities for Amazon EKS

Restore type	Restore target	Restore behavior
Existing cluster restore	Restore to the source EKS cluster or existing EKS cluster	Restores all Kubernetes resources and persistent volumes to existing EKS clusters. All restores are non-destructives and existing objects are not overwritten. For objects that are skipped, you can subscribe to SNS Notifications

Restore type	Restore target	Restore behavior
New cluster restore	Creates a new Amazon EKS cluster as part of your EKS restore	Restore creates new EKS cluster and restores all Kubernetes resources and persistent volumes to newly created cluster
Namespace restore	Existing Amazon EKS cluster	Restores only specified namespaces, their Kubernetes resources and corresponding persistent storage restores are non-destructive and existing objects are not overwritten. All cluster-scoped resources will also be restored. For objects that are skipped, you can subscribe to SNS Notifications
Peristent Storage Restore	Persistent Storage Dependent	Restore individual persistent storage as standalone restores. See Restore Behavior of Amazon EBS, Amazon S3, Amazon EFS.

Permissions

The permissions required depend on the restore type and target destination.

- AWS Backup's managed policy AWSBackupServiceRolePolicyForRestores contains the required permissions to restore your Amazon EKS cluster and EBS and EFS persistent storage.
- If your EKS Cluster contains an S3 bucket, or you are restoring the child S3 recovery point alone you will need to ensure the following policies or permissions within are assigned to your role AWSBackupServiceRolePolicyForS3Restore.

Considerations before restoring

Before you begin an EKS restore job, review the following. If you are restoring an EKS backup that has been copied across account or region ensure you check these considerations ahead of restores to prevent restore failures.

1. **IAM Roles**: when restoring onto a different cluster, the IAM Roles used in the source cluster (such as Pod identity, IRSA. OIDC provider configs etc) must be present in the account / region as the destination cluster.
2. **Ensure EKS Version and Compatibility**: The API Versions of the objects that you're wanting to restore should be the same version (or as close to as possible) and supported in the new cluster. AWS Backup will perform a best effort restore between EKS versions, though compatibility issues may arise when restoring between significantly different versions.
3. **Matching Storage Classes**: For restores to an existing EKS cluster, ensure that the appropriate CSI Storage Driver add-ons are installed prior to restore
4. **S3 Buckets**: When restoring an EKS cluster with S3 Buckets, ensure your S3 bucket are versioned and accessible in the destination account or region.
5. **Image Repository**: When restoring an EKS cluster ensure that the destination EKS cluster's account or region have access to the images that are being referenced as part of the restore. Check that your registry has the sufficient cross-region / account policy permissions.
6. **Security Groups**: Security groups should be pre-created for ALB, Pod Identities, EKS Node Groups etc. in the target account and region if creating a new EKS cluster as part of your restore
7. **EBS Availability Zones and Nodes**: The Availability Zones where you recover your EBS volumes should be mapped to the Availability Zone of an existing EKS node
8. **Non-destructive restores**: All EKS restores will be non-destructive and not overwrite Kubernetes objects of the target restore.
9. **Enable EKS Audit Logs**: Enable EKS Audit Logs for additional logging and troubleshooting prior to restore. You can also subscribe to SNS notifications to notify of skipped or failed objects on restore.
10. **New EKS Cluster Creation Restore Buffer**: When creating a new EKS cluster during restore, AWS Backup introduces a 15-minute buffer after the EKS cluster reaches an available state but before creating any additional EKS resources. This buffer ensures that all underlying EKS components are fully initialized before dependent resources are created.
11. **User data in launch template**: When creating a node group using a launch template, do not specify `spec.cluster` in the user data section of the launch template. Amazon EKS automatically injects the cluster identity parameters (`apiServerEndpoint`,

`certificateAuthority`, and `serviceIpv4Cidr`) and merges them with any additional configurations defined in the user data.

12 **Amazon EKS Auto Mode and restored Amazon EBS volumes**: If the cluster has Amazon EKS Auto Mode enabled, Auto Mode only attaches an Amazon EBS volume whose `eks:eks-cluster-name` tag matches the cluster's name. AWS Backup does not update this tag on restored volumes, so pods referencing restored volumes might remain in `ContainerCreating` with `FailedAttachVolume` errors when restoring to a different cluster, even though the restore job succeeds. To resolve this, manually tag each restored Amazon EBS volume with the destination cluster's name with the following command: `aws ec2 create-tags --resources vol-123456abcdef --tags Key=eks:eks-cluster-name,Value=destination-cluster-name`

EKS Configurations

When you restore the composite Amazon AWS Backup, you choose the restore type and target destination. You can choose to restore to the source EKS cluster, an existing EKS cluster or create a new EKS cluster as the restore target. For new EKS clusters, you can choose to use the same existing infrastructure settings (for example, VPC, subnets) as the backed up cluster or configure new ones. AWS Backup is designed to perform a non-destructive restore that doesn't overwrite existing resources.

For namespace restores, you can specify up to 5 namespaces to restore selectively. We will also restore all cluster-scoped resources as well.

As an advanced setting you can opt to change the restore order of the Kubernetes Objects. By Default, AWS Backup will restore all Kubernetes objects in the following order:

Cluster Scoped Kubernetes Resources

1. Custom Resource Definitions
2. Namespaces (the namespace itself, not the resources within that namespace)
3. StorageClasses
4. PersistentVolumes

Namespace Scoped Kubernetes Resources

1. PersistentVolumeClaims

2. Secrets
3. ConfigMaps
4. ServiceAccounts
5. LimitRanges
6. Pods
7. ReplicaSets

Persistent Storage Configurations

As part of the composite Amazon EKS backup restore, the second step will be to configure your Persistent Storage configurations. This will vary based on the persistent storage backed up as part of your EKS cluster.

For Amazon EBS Snapshots you are required to provide an Availability Zone, where the Amazon EBS volume will be restored and created. AWS Backup will then attempt to create the EKS pod in the same availability zone as selected so your volume can be remounted to your EKS cluster as part of restore.

As part of the restore, AWS Backup will remount your Amazon EBS volumes and Amazon S3 buckets to your restored EKS cluster. Amazon EFS filesystems restore to random prefixes and require manual access point creation after restore to remount to your EKS cluster. AWS Backup does not create access points or mount targets on your behalf, refer to guidance here for access points and mount targets.

Restore customizations

AWS Backup can customize your Kubernetes objects during an Amazon EKS restore. You define *resource modifiers*, a set of rules that transform objects before they are restored to the target cluster. These customizations are applied in memory during the restore and do not alter the original backup data.

Resource modifiers are particularly useful when a restored cluster contains references to the source environment that must be updated for the destination, such as container image registry URIs, database endpoints, or storage class names. This is common in cross-Region and cross-account disaster recovery, where the restored objects would otherwise point at resources that do not exist in the destination. You can use resource modifiers for any change you need to make to your objects as they are restored.

Providing resource modifiers

1. Create a resource modifiers file in YAML format.
2. Upload the file to an Amazon S3 bucket in the same account and Region from which you start the restore job.
3. Pass the Amazon S3 URI of the file in the `resourceModifiers` restore metadata key when you start the restore job.

AWS Backup reads the file once, when you start the restore job, and uses that copy for the whole restore. Changing or deleting the file after you start the restore job has no effect on that restore.

The principal that calls `StartRestoreJob` must have `s3:GetObject` permission on the resource modifiers file. This is the identity making the API call, not the role you pass in `iam-role-arn`.

The file must be valid UTF-8 YAML, must set `version: v1`, and must not exceed 1 MB. It must not use YAML anchors or aliases, and it must not contain duplicate keys. AWS Backup validates the file when you submit the restore job and rejects the job before it starts if the file is missing, inaccessible, or invalid.

The file must be in an Amazon S3 bucket in the same AWS account and Region as the restore job. If the bucket has a bucket policy, make sure that it does not deny the calling principal. AWS Backup reads the file on your behalf rather than from your network, so a policy that restricts access to a specific VPC endpoint, IP address range, or principal list blocks the read even when your IAM permissions are correct.

The `resourceModifiers` restore metadata key is optional. If you omit it, the restore proceeds with no customer-defined modifications.

Resource modifiers file schema

```
version: v1
resourceModifierRules:
  - conditions:
      groupResource: group/version/resource   # ex: apps/v1/deployments, or v1/pods for
 core resources
      resourceNameRegex: string               # ex: my-pod.*
      namespaces:
        - namespace1
        - namespace2
```

```
    labelSelector:
      matchLabels:
        foo: bar
        test: label
  patches:
    - operation: replace
      path: "/spec/storageClassName"
      value: "new-sc"
    - operation: remove
      path: "/metadata/uid"
```

Each rule consists of conditions, which decide which objects are modified, and patches, which decide what is modified in those objects.

Conditions

Conditions decide which objects a rule applies to. All conditions are optional, but each rule **must specify** at least one. When a rule specifies more than one condition, an object must match all of them.

- `groupResource` - The Kubernetes group, version, and plural resource name that the patches apply to, in the form `group/version/resource`, or `version/resource` for core resources that have no group. The last segment is the plural resource name reported by the Kubernetes API, such as `deployments` or `pods`, not the object kind. A rule whose `groupResource` does not correspond to a resource in the backup matches no objects, and the restore job completes without applying it.
- `resourceNameRegex` - A regular expression matched against the object's name, using RE2 syntax. RE2 does not support lookarounds or backreferences. The expression must match the entire name, not part of it, so `my-pod` matches only an object named exactly `my-pod`. To match a family of names, use `my-pod.*`. Omit this condition to match objects with any name.
- `namespaces` - A list of namespaces that the objects must be in. If omitted or empty, the rule matches objects in all namespaces. This is the namespace name as stored in the backed-up object, not the namespace the object is restored into.
- `labelSelector` - A set of key and value label pairs. An object matches only if it has all of the specified labels with exactly the specified values. An object with no labels never matches a rule that specifies a label selector.

Patches

Patches decide what is modified in the matched objects. Each rule must specify at least one patch. AWS Backup applies each patch as a JSON Patch `add`, `remove`, or `replace` operation, and `path` uses JSON Pointer syntax, including array indexes. Each patch supports the following fields:

- `operation` - The modification to perform, one of `add`, `remove`, or `replace`.
- `path` - The JSON Pointer to the field to modify. For example, `/spec/storageClassName` targets the value at `spec.storageClassName`.
- `value` - The value to apply. Required for `add` and `replace`, and not allowed for `remove`. AWS Backup treats a null value as no value, so it rejects the restore job if an `add` or `replace` patch omits `value` or sets it to null, or if a `remove` patch specifies `value`.

The parent path of a patch must already exist on the object. `add` creates only the final field in the path, so adding `/metadata/labels/team` to an object that has no `labels` map fails.

Unlike a standalone JSON Patch document, which is applied all-or-nothing, AWS Backup continues applying the remaining operations after one of them fails.

Patching arrays

To target an element of an array, use its zero-based index in the path. For example, `/spec/template/spec/containers/0/image` is the image of the first container.

- `add` with an index inserts the value at that position and moves later elements back by one. Use `-` as the index to append to the end of the array, for example `/spec/template/spec/containers/-`.
- `replace` and `remove` with an index change or delete the element at that position.

The following patches fail at run time:

- An `add` whose index is greater than the length of the array.
- A `replace` or `remove` whose index does not exist in the array.
- A `replace` or `remove` that uses `-` as the index.

Paths do not support wildcards. To change the same field in several array elements, such as the image of every container, specify one patch for each index.

Order of application

- You can specify multiple patches for a single rule. Patches are applied in the order listed. If multiple patches target the same `path`, the last one takes effect.
- You can specify multiple rules in the file. Rules are applied in the order listed. AWS Backup does not give precedence to the most specific match, so place your most specific rule last to give it the final say.

Modification results and notifications

A patch can fail at run time if the `path` it targets does not exist on the matched object, for example targeting `/spec/containers/3/image` on an object that has only two containers. When a patch fails at run time:

- Only the failed operation is skipped. Operations listed before it in the same rule stay applied, operations after it are still attempted, and changes made by earlier rules are kept. AWS Backup still attempts to restore the object with whatever modifications succeeded, so if it restores, it can be partially modified rather than unmodified.
- The modification failure does not change the job status. If any modification failed, the job's status message includes "Some Kubernetes Objects were restored without their requested modifications." Status messages are additive, so this message can appear alongside other messages from the same restore.

A patch that applies successfully can still produce an object that your cluster does not accept, for example by removing a field the object requires or by setting a field to a value of the wrong type. AWS Backup does not check whether a modified object is still valid, so your cluster rejects it during the restore and that object fails to restore. Other objects in the same restore are not affected. Test your resource modifiers with a non-production restore before you rely on them.

To get details about which objects were modified or failed to be modified, subscribe to the following SNS event notifications:

- `EKS_RESTORE_MODIFICATION_SUCCESS` - One or more rules matched an object and every patch was applied successfully.
- `EKS_RESTORE_MODIFICATION_FAILED` - A rule matched an object but one of its patches failed to apply at run time. AWS Backup still attempts to restore the object, with whatever other modifications succeeded.

AWS Backup reports these events once per object. For more information, see Notification options with AWS Backup.

Example: cross-account, cross-Region restore

The following resource modifiers file is used when restoring from account `111122223333` in `us-east-1` to account `444455556666` in `us-west-2`. It updates the container image and database endpoint of the `payment` Deployment in the `production` namespace, and adds an annotation recording the backup it was restored from. The Deployment must already have an `annotations` map for the third patch to apply.

```
version: v1
resourceModifierRules:
  - conditions:
      groupResource: apps/v1/deployments
      namespaces:
        - production
      labelSelector:
        matchLabels:
          app: payment
    patches:
      - operation: replace
        path: "/spec/template/spec/containers/0/image"
        value: "444455556666.dkr.ecr.us-west-2.amazonaws.com/payment-service:v2.1.0"
      - operation: replace
        path: "/spec/template/spec/containers/0/env/0/value"
        value: "payment-db.cluster-dst789.us-west-2.rds.amazonaws.com"
      - operation: add
        path: "/metadata/annotations/backup.aws.com~1restored-from"
        value: "arn:aws:backup:us-east-1:111122223333:recovery-point:rp-abc123"
```

> **Note**
>
> In JSON Pointer paths, ~1 is the escape sequence for a literal / in a key name, so `backup.aws.com~1restored-from` refers to the annotation key `backup.aws.com/restored-from`.

The Deployment before and after these rules

These excerpts show only the fields that the three patches change.

In the recovery point

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: payment
  namespace: production
  labels:
    app: payment
  annotations:
    team: payments
spec:
  template:
    spec:
      containers:
        - name: payment-service
          image: 111122223333.dkr.ecr.us-east-1.amazonaws.com/payment-service:v2.1.0
          env:
            - name: DB_HOST
              value: payment-db.cluster-src123.us-east-1.rds.amazonaws.com
```

Restored to the cluster

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: payment
  namespace: production
  labels:
    app: payment
  annotations:
    team: payments
    backup.aws.com/restored-from: arn:aws:backup:us-east-1:111122223333:recovery-
point:rp-abc123
spec:
  template:
    spec:
      containers:
        - name: payment-service
          image: 444455556666.dkr.ecr.us-west-2.amazonaws.com/payment-service:v2.1.0
          env:
            - name: DB_HOST
              value: payment-db.cluster-dst789.us-west-2.rds.amazonaws.com
```

The first two `replace` patches changed the container image and the database endpoint, and the `add` patch added the `restored-from` annotation to the existing `annotations` map. The recovery point is not changed.

Amazon EKS restore procedure

Follow these steps to restore Amazon EKS backups using the AWS Backup console or AWS CLI:

Console

To restore your Amazon EKS cluster

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Backup vaults**.
3. Choose the backup vault that contains your Amazon EKS backup, then select the recovery point for your Amazon EKS backup.
4. Choose **Restore**.
5. In the **Restore options** pane, choose your restore type:
   - **Restore full EKS cluster** - Restores the entire Amazon EKS composite recovery point
   - **Select namespaces to restore** - Restores up to five specific namespaces
6. Configure the target destination:
   - For cluster restore, choose to create a new cluster or use an existing cluster
   - For new clusters, specify cluster name, Kubernetes version, VPC configuration, IAM roles, subnets, Additional security groups, node group settings, fargate profiles and Pod identity IAM roles
   - For existing clusters, select the target cluster from the dropdown
   - For namespace restore, specify the target cluster and namespace names
7. Optionally, configure advanced settings for custom restore order for Kubernetes resources.
8. Choose the IAM restore role for the job. If not using the default role, ensure the selected role includes the iam:PassRole permission.

> **Note**
>
> If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

9. Choose **Restore backup**.

AWS CLI

Use the `aws backup start-restore-job` command with Amazon EKS-specific metadata.

The required metadata depends on your restore type. All restore operations require the `clusterName` parameter.

Restore Amazon EKS recovery points through AWS CLI

Use StartRestoreJob. You can specify the following metadata during Amazon EKS restores:

Mandatory metadata:

- `clusterName` - Name of cluster to restore to

Optional metadata:

- `newCluster` - (true/false) If we should create a new EKS cluster during restore. If newCluster is "true", the following metadata fields apply:
  - `eksClusterVersion` - Desired K8s version of cluster if wanting to increase cluster version during restore
  - `clusterRole` - The IAM Role ARN to attach to the created EKS cluster
  - `controlPlaneScalingTier` - The control plane scaling tier for the new cluster. If you omit this field, the restore uses the tier captured in the backup, or the Standard control plane if no tier was captured. To restore a provisioned control plane as Standard, specify `standard`. For available tier values, see Amazon EKS Provisioned Control Plane.
  - `encryptionConfigProviderKeyArn` - Specify the KMS key ARN to encrypt the destination cluster. This can be either the KMS key from the source cluster, or a different

KMS key. A different KMS key must be provided when performing cross-region or cross-account restore. Omit this metadata entirely if the source cluster is not encrypted.

- `clusterVpcConfig` - VPC/Networking configuration for the created EKS cluster. This field has the following nested fields:
  - `vpcId` - The VPC associated with your cluster
  - `subnetIds [Required]` - The subnets associated with your cluster
  - `securityGroupIds [Required]` - The additional security groups associated with your cluster
- `computeConfig` - The compute configuration for clusters created with Amazon EKS Auto Mode. This field has the following nested field:
  - `nodeRoleArn` - The IAM role ARN that Amazon EKS Auto Mode attaches to the nodes. If omitted, the restore uses the node role from the source cluster. When performing a cross-account restore, you must provide a node role in the destination account. Otherwise, the restore fails.
- `nodeGroups` - The Managed Node Groups to be created on the EKS Cluster. The NodeGroups for restore must have all of the same node groups from backup time and have matching nodeGroupId.
  - `nodeGroupId [Required]` - The ID of the node group
  - `subnetIds [Required]` - The subnets that were specified for the Auto Scaling group that is associated with your node group
  - `instanceTypes` - If the node group wasn't deployed with a launch template, then this is the instance type that is associated with the node group
  - `nodeRole [Required]` - The IAM role associated with your node group
  - `securityGroupIds` - The security group IDs that are allowed SSH access to the nodes
  - `remoteAccessEc2SshKey` - The Amazon EC2 SSH key name that provides access for SSH communication with the nodes in the managed node group
  - `launchTemplateId` - Specify the launch template ID to create the node group. This can be either the launch template ID from the source cluster, or a different launch template ID. If the source cluster's launch template contains hard-coded endpoint that points to the source cluster itself, you must provide a different launch template ID. Omit this metadata entirely if the source cluster does not use a launch template.
  - `launchTemplateVersion` - Launch template version associated with the specified launch template ID.

- `fargateProfiles` - The Fargate Profiles to be created on the EKS Cluster. The Fargate Profiles for restore must have all the same Fargate Profiles from backup time and have matching name.
  - `name [Required]` - The name of the Fargate profile
  - `subnetIds` - The IDs of subnets to launch a Pod into
  - `podExecutionRoleArn [Required]` - The IAM Role ARN of the Pod execution role to use for a Pod that matches the selectors in the Fargate profile
- `podIdentityAssociations` - The Pod Identity Associations to be created on the EKS Cluster
  - `associationId` - The ID of the Pod Identity Association
  - `roleArn` - The IAM Role ARN for the Pod Identity Association
- `kubernetesRestoreOrder` - Override the order the Kubernetes manifests are restored in. This order will take precedence over the default service restore order. This follow the format: group/version/kind or version/kind

  Ex: `["v1/persistentvolumes","v1/pods","customresource/v2/custom"]`
- `namespaceLevelRestore` - (true/false) If you would like to perform a namespace level restore
- `namespaces` - A list of namespaces to restore if namespaceLevelRestore is "true". Can provide up to 5 namespaces to restore.

  Ex: `["ns-1","ns-2","ns-3","ns-4","ns-5"]`
- `restoreKubernetesManifestsOnly` - (true/false) If you would like to only restore the Kubernetes manifest files and no persistent storage systems (EBS, S3, EFS, etc.)
- `nestedRestoreJobs` - Restore Metadata configuration of all of the nested Recovery Points for the PersistentVolume storage systems in the composite Recovery Point. This is a map of RecoveryPointArn: RestoreMetadata of that Recovery Point
- `resourceModifiers` - The Amazon S3 URI of a YAML file containing resource modifier rules that customize Kubernetes objects during restore. The file must be in an Amazon S3 bucket in the same account and Region as the restore job, must not exceed 1 MB, and the principal that calls `StartRestoreJob` must have `s3:GetObject` permission on it. AWS Backup validates the file when the job is submitted and rejects the job if the file is missing, inaccessible, or invalid. For the file schema and examples, see **Restore customizations**.

  Ex: `"s3://my-backup-bucket/restore-configs/account-modifiers.yaml"`

Restore to existing cluster

```
aws backup start-restore-job \
    --recovery-point-arn "arn:aws:backup:us-west-2:123456789012:recovery-
point:composite:eks/my-cluster-20240115" \
    --iam-role-arn "arn:aws:iam::123456789012:role/AWSBackupDefaultServiceRole" \
    --metadata '{"clusterName":"existing-cluster","newCluster":"false"}' \
    --resource-type "EKS"
```

Restore specific namespaces to an existing cluster:

```
aws backup start-restore-job \
    --recovery-point-arn "arn:aws:backup:us-west-2:123456789012:recovery-
point:composite:eks/my-cluster-20240115" \
    --iam-role-arn "arn:aws:iam::123456789012:role/AWSBackupDefaultServiceRole" \
    --metadata '{"clusterName":"existing-
cluster","newCluster":"false","namespaceLevelRestore":"true","namespaces":"[\"ns-1\",
\"ns-2\",\"ns-3\",\"ns-4\",\"ns-5\"]"}' \
    --resource-type "EKS"
```

Restore nested persistent volumes to an existing cluster:

```
aws backup start-restore-job \
    --recovery-point-arn "arn:aws:backup:us-west-2:123456789012:recovery-
point:composite:eks/my-cluster-20240115" \
    --iam-role-arn "arn:aws:iam::123456789012:role/AWSBackupDefaultServiceRole" \
    --metadata '{"clusterName":"existing-
cluster","newCluster":"false","namespaceLevelRestore":"true","nestedrestorejobs":"{\"arn:aws
west-2::snapshot/snap-abc123\":\"{\\\"AvailabilityZone\\\":\\\"us-west-2a\\\"}\",
\"arn:aws:backup:us-west-2:123456789012:recovery-point:fa71a304-2555-4c37-8128-
f154b9578032\":\"{\\\"DestinationBucketName\\\":\\\"bucket-name\\\"}\"}"}' \
    --resource-type "EKS"
```

Restore to new cluster

```
aws backup start-restore-job \
    --recovery-point-arn "arn:aws:backup:us-west-2:123456789012:recovery-
point:composite:eks/my-cluster-20240115" \
    --iam-role-arn "arn:aws:iam::123456789012:role/AWSBackupDefaultServiceRole" \
    --metadata '{"clusterName":"new-
cluster","newCluster":"true","clusterRole":"arn:aws:iam::123456789012:role/
EKSClusterRole","eksClusterVersion":"1.33","controlPlaneScalingTier":"tier-2xl","encryptionC
```

```
west-2:123456789012:key/
ecb2b326-784d-4ec0-8d07-20ab826b5a13","clusterVpcConfig":"{\"vpcId\":
\"vpc-1234\",\"subnetIds\":[\"subnet-1\",\"subnet-2\",\"subnet-3\"],
\"securityGroupIds\":[\"sg-123\"]}","nodeGroups":"[{\"nodeGroupId\":
\"nodegroup-1\",\"subnetIds\":[\"subnet-1\",\"subnet-2\",\"subnet-3\"],
\"nodeRole\":\"arn:aws:iam::123456789012:role/EKSNodeGroupRole\",
\"instanceTypes\":[\"t3.small\"],\"launchTemplateId\":\"lt-0b13949aae3f2b867\",
\"launchTemplateVersion\":\"1\"}]","fargateProfiles":"[{\"name\":\"fargate-
profile-1\",\"subnetIds\":[\"subnet-1\",\"subnet-2\",\"subnet-3\"],
\"podExecutionRoleArn\":\"arn:aws:iam::123456789012:role/EKSFargateProfileRole
\"}]","computeConfig":"{\"nodeRoleArn\":\"arn:aws:iam::123456789012:role/
EKSAutoModeNodeRole\"}"}' \
    --resource-type "EKS"
```

After starting the restore job, use `describe-restore-job` to monitor progress:

```
aws backup describe-restore-job --restore-job-id restore-job-id
```

You can subscribe to **Notification Events** for failed and skipped objects for restore. For more information, see Notification options with AWS Backup.

Amazon EKS restore status messages

When a restore job completes, you might see the following status messages. The table shows the possible scenarios and their corresponding job status values:

Scenario	Job Status	Example message
All objects restored successfully	COMPLETED	—
One or more objects failed to be restored	COMPLETED	"One or more Kubernetes objects failed to restore. To get notified of these failures, enable SNS event notifications."
One or more resource modifications failed to apply	COMPLETED	"Some Kubernetes Objects were restored without their requested modifications. To get notified of these failures, enable SNS event notifications."

Scenario	Job Status	Example message
Restore could not complete	FAILED	(error details)

Skipped objects during restore

The following Kubernetes objects are restored on a best-effort basis. If they fail to restore, they are moved to the skipped list. These objects are system-managed and recreated by Kubernetes or Amazon EKS:

- FlowSchemas with the `apf.kubernetes.io/autoupdate-spec: "true"` annotation
- PriorityLevelConfigurations with the `apf.kubernetes.io/autoupdate-spec: "true"` annotation
- The `eks-exempt` FlowSchema (EKS-managed, references the protected exempt PriorityLevelConfiguration)
- The `kubernetes` Service in the `default` namespace (API server endpoint)
- The `kube-dns` Service in the `kube-system` namespace (CoreDNS)

The following Kubernetes objects are always skipped during a restore. These objects include node infrastructure and networking components. They reference cluster-specific state. Kubernetes controllers automatically recreate these objects when nodes and services become active in the target cluster. If you restore these objects from a backup, they might cause conflicts or errors. These conflicts occur because the restored objects reference resources that don't exist on the target cluster.

- Endpoints (`v1/endpoints`) - Network endpoints that define IP addresses for Services. The Endpoints controller automatically manages these based on Service selectors.
- EndpointSlices (`discovery.k8s.io/v1/endpointslices`) - The EndpointSlice controller automatically manages these objects.
- IPAddresses (`networking.k8s.io/v1/ipaddresses`) - Kubernetes networking manages these cluster-internal IP allocations.
- CSINodes (`storage.k8s.io/v1/csinodes`) - The kubelet automatically creates these objects when CSI drivers register on a node.

- VolumeAttachments (`storage.k8s.io/v1/volumeattachments`) - The attach/detach controller automatically creates these objects when pods require volumes.

Objects that already exist on the target cluster are also skipped. EKS restores are non-destructive — existing objects are never overwritten or deleted.

To receive notifications about skipped or failed objects during restore, subscribe to SNS event notifications. For more information, see Notification options with AWS Backup.

OIDC and IRSA considerations for cross-cluster restore

When restoring an Amazon EKS backup to a different cluster than the source, Kubernetes objects referencing IAM Roles for Service Accounts (IRSA) retain the source cluster's OIDC provider endpoint. These references exist in service account annotations and IAM trust policies. AWS Backup does not automatically update them during restore.

If your workloads use IRSA, a cross-cluster restore requires:

- The destination cluster must have an OIDC provider configured.
- All IAM roles referenced by Kubernetes service accounts must have their trust policies updated to include the destination cluster's OIDC provider endpoint.
- Service account annotations (`eks.amazonaws.com/role-arn`) will still reference the original roles — verify these match valid roles in the destination account.

If these dependencies are not satisfied, pods will fail to assume IAM roles after restore, and restore jobs might report failures during validation.

Recommendation: For workloads that require cross-cluster or cross-account restore portability, we recommend using EKS Pod Identity instead of IRSA. EKS Pod Identity does not depend on cluster-specific OIDC providers and simplifies cross-cluster restore scenarios.

Restore an FSx file system

The restore options that are available when you use AWS Backup to restore Amazon FSx file systems are the same as using the native Amazon FSx backup. You can use a backup's recovery point to create a new file system and restore a point-in-time snapshot of another file system.

AWS Backup supports restoring file systems that use Intelligent Tiering storage for both FSx for Lustre and FSx for OpenZFS file systems. Intelligent Tiering file systems have specific configuration requirements during restore operations.

When restoring Amazon FSx file systems, AWS Backup creates a new file system and populates it with the data (Amazon FSx for NetApp ONTAP allows restoring a volume to an existing file system). This is similar to how native Amazon FSx backs up and restores file systems. Restoring a backup to a new file system takes the same amount of time as creating a new file system. The data restored from the backup is lazy-loaded onto the file system. You might therefore experience slightly higher latency during the process.

> **Note**
>
> You can't restore to an existing Amazon FSx file system, and you can't restore individual files or folders.
> FSx for ONTAP doesn't support backing up certain volume types, including DP (data-protection) volumes, LS (load-sharing) volumes, full volumes, or volumes on file systems that are full. For more information, please see FSx for ONTAP Working with backups.
> AWS Backup vaults that contain recovery points of Amazon FSx file systems are visible outside of AWS Backup. You can restore the recovery points using Amazon FSx but you can't delete them.

You can see backups created by the built-in Amazon FSx automatic backup functionality from the AWS Backup console. You can also recover these backups using AWS Backup. However, you can't delete these backups or change the automatic backup schedules of your Amazon FSx file systems using AWS Backup.

Use the AWS Backup console to restore Amazon FSx recovery points

You can restore most Amazon FSx backups created by AWS Backup using the AWS Backup console, API, or AWS CLI.

This section shows you how to use the AWS Backup console to restore Amazon FSx file systems.

Restoring an FSx for Windows File Server file system

To restore an FSx for Windows File Server file system

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.

2. In the navigation pane, choose **Protected resources**, and then choose the Amazon FSx resource ID that you want to restore.
3. On the **Resource details** page, a list of recovery points for the selected resource ID is shown. Choose the recovery point ID of the resource.
4. In the upper-right corner of the pane, choose **Restore** to open the **Restore backup** page.
5. In the **File system details** section, the ID of your backup is shown under **Backup ID**, and the file system type is shown under **File system type**. You can restore both FSx for Windows File Server and FSx for Lustre file systems.
6. For **Deployment type**, accept the default. You can't change the deployment type of a file system during restore.
7. Choose the **Storage type** to use. If the storage capacity of your file system is less than 2,000 GiB, you can't use the **HDD** storage type.
8. For **Throughput capacity**, choose **Recommended throughput capacity** to use the recommended 16 MB per second (MBps) rate, or choose **Specify throughput capacity** and enter a new rate.
9. In the **Network and security** section, provide the required information.
10. If you are restoring an FSx for Windows File Server file system, provide the **Windows authentication** information used to access the file system, or you can create a new one.

    > **Note**
    >
    > When restoring a backup, you can't change the type of Active Directory on the file system.

    For more information about Microsoft Active Directory, see Working with Active Directory in Amazon FSx for Windows File Server in the *Amazon FSx for Windows File Server User Guide*.
11. (Optional) In the **Backup and maintenance** section, provide the information to set your backup preferences.
12. In the **Restore role** section, choose the IAM role that AWS Backup will use to create and manage your backups on your behalf. We recommend that you choose the **Default role**. If there is no default role, one is created for you with the correct permissions. You can also provide your own IAM role.

> **Note**
>
> If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

13. Verify all your entries, and choose **Restore Backup**.

Restoring an Amazon FSx for Lustre file system

AWS Backup supports Amazon FSx for Lustre file systems that have persistent storage deployment type and are not linked to a data repository like Amazon S3.

> **Note**
>
> You can only restore your backup to a file system of the same deployment type, storage class, throughput capacity, storage capacity, data compression type, and AWS Region as the original. You can increase your restored file system's storage capacity after it becomes available.

To restore an Amazon FSx for Lustre file system

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Protected resources**, and then choose the Amazon FSx resource ID that you want to restore.
3. On the **Resource details** page, a list of recovery points for the selected resource ID is shown. Choose the recovery point ID of the resource.
4. In the upper-right corner of the pane, choose **Restore** to open the **Restore backup to new file system** page.
5. In the **Settings** section, the ID of your backup is shown under **Backup ID**, and the file system type is shown under **File system type**. **File system type** should be **Lustre**.
6. Choose a **Deployment type**. AWS Backup only supports the persistent deployment type. You can't change the deployment type of a file system during restore.

Persistent deployment type is for long-term storage. For detailed information about FSx for Lustre deployment options, see Using Available Deployment Options for Amazon FSx for Lustre File Systems in the *Amazon FSx for Lustre User Guide*.

7. Choose the **Throughput per unit storage** that you want to use.

> **Note**
>
> Throughput per unit storage cannot be configured for file systems using Intelligent-Tiering storage class.

8. (Optional) For file systems using Intelligent-Tiering storage class, choose the SSD read cache sizing mode and capacity. For more information, see the FSx documentation for managing provisioned SSD read cache.
9. (Optional) For file systems using Intelligent-Tiering storage class, choose whether to enable EFA (Elastic Fabric Adapter). To enable EFA, make sure that your security group allows all inbound and outbound traffic within the security group.
10. Specify the **Storage capacity** to use. Enter a capacity between 32 GiB and 64,436 GiB.

> **Note**
>
> For Intelligent Tiering file systems, storage capacity is elastic and cannot be specified during restore. The capacity will automatically scale based on your data usage.

11. In the **Network and security** section, provide the required information.
12. (Optional) In the **Backup and maintenance** section, provide the information to set your backup preferences.
13. In the **Restore role** section, choose the IAM role that AWS Backup will use to create and manage your backups on your behalf. We recommend that you choose the **Default role**. If there is no default role, one is created for you with the correct permissions. You can also provide your IAM role.

> **Note**
>
> If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

14. Verify all your entries, and choose **Restore Backup**.

Restoring Amazon FSx for NetApp ONTAP volumes

To restore Amazon FSx for NetApp ONTAP volumes:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Protected resources**, and then choose the Amazon FSx resource ID that you want to restore.
3. On the **Resource details** page, a list of recovery points for the selected resource ID is shown. Choose the recovery point ID of the resource.
4. In the upper-right corner of the pane, choose **Restore** to open the **Restore** page.

   The first section, **File system details**, displays the recovery point ID, the file system ID, and the file system type.
5. Under **Restore options**, there are several selections. First, choose the **File system** from the dropdown menu.
6. Next, choose the preferred **Storage virtual machine** from the dropdown menu.
7. Enter a name for your volume.
8. Specify the **Junction Path**, which is location within your file system where your volume will be mounted.
9. Specify the **Volume size** in megabytes (MB) that you are creating.
10. (*Optional*) You can choose to **Enable storage efficiency** by checking the box. This will allow deduplication, compression, and compaction.
11. In the **Capacity pool tiering policy** dropdown menu, select the tiering preference.
12. In the **Restore permissions**, choose the IAM role that AWS Backup will use to restore backups.
13. Verify all your entries, and choose **Restore Backup**.

Restoring an Amazon FSx for OpenZFS file system

> **Note**
>
> Restoring from a backup with a given storage class to a file system with a different storage class is not supported.

To restore an FSx for OpenZFS file system

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Protected resources**, and then choose the Amazon FSx resource ID that you want to restore.
3. On the **Resource details** page, a list of recovery points for the selected resource ID is shown. Choose the recovery point ID of the resource.
4. In the upper-right corner of the pane, choose **Restore** to open the **Restore backup** page.

   In the **File system details** section, the ID of your backup is shown under **Backup ID**, and the file system type is shown under **File system type**. File system type should be **FSx for OpenZFS**.
5. Under **Restore options**, you may select **Quick restore** or **Standard restore**. Quick restore will use the default settings of the source file system. If you are doing Quick Restore, skip to Step 7.

   If you choose Standard restore, specify the additional following configurations:

   a. **Provisioned SSD IOPS**: You can choose the **Automatic radio button** or you can choose the **User-provisioned option** if available.

   > **Note**
   >
   > SSD IOPS cannot be set for file systems using Intelling-Tiering storage class

   b. **Throughput capacity**: You can choose the **Recommended throughput capacity** of 64 MB/sec (for SSD storage class), and 160 MB/sec (for Intelligent-Tiering storage class), or you can choose to **Specify throughput capacity**.

   c. (*Optional*) **VPC security groups**: You can specify VPC security groups to associate with your file system's network interface.

d. **Encryption key**: Specify the AWS Key Management Service key to protect the restored file system data at rest.

e. (*Optional*) **Root Volume configuration**: This configuration is collapsed by default. You may expand it by clicking the down-pointing carat (arrow). Creating a file system from a backup will create a new file system; the volumes and snapshots will retain their source configurations.

f. (*Optional*) **Backup and maintenance**: To set a scheduled backup, click the down-pointing carat (arrow) to expand the section. You may choose the backup window, hour and minute, retention period, and weekly maintenance window.

6. The **SSD Storage capacity** will display the file system's storage capacity.

> **Note**
>
> For Intelligent Tiering file systems, storage capacity is elastic and cannot be specified during restore. The capacity will automatically scale based on your data usage.

7. (Optional) For file systems using Intelligent-Tiering storage class, choose the SSD read cache sizing mode and capacity. For more information, see the FSx documentation for managing provisioned SSD read cache.

8. Choose the **Virtual Private Cloud** (VPC) from which your file system can be accessed.

9. In the **Subnet** dropdown menu, choose the subnet in which your file system's network interface resides.

10. In the **Restore role** section, choose the IAM role that AWS Backup will use to create and manage your backups on your behalf. We recommend that you choose the **Default role**. If there is no default role, one is created for you with the correct permissions. You can also choose an IAM role.

> **Note**
>
> If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

11. Verify all your entries, and choose **Restore Backup**.

Use the AWS Backup API, CLI, or SDK to restore Amazon FSx recovery points

To restore Amazon FSx using the API or CLI, use `StartRestoreJob`. You can specify the following metadata during an Amazon FSx restore:

```
StorageCapacity
StorageType
VpcId
KmsKeyId
SecurityGroupIds
SubnetIds
DeploymentType
WeeklyMaintenanceStartTime
DailyAutomaticBackupStartTime
AutomaticBackupRetentionDays
CopyTagsToBackups
WindowsConfiguration
LustreConfiguration
OntapConfiguration
OpenZFSConfiguration
aws:backup:request-id
```

> **Note**
>
> Storage capacity cannot be specified for Intelligent Tiering file systems as they use elastic storage that scales automatically based on data usage.

FSx for Windows File Server restore metadata

You can specify the following metadata during an FSx for Windows File Server restore:

- ThroughputCapacity
- PreferredSubnetId
- ActiveDirectoryId

FSx for Lustre restore metadata

You can specify the following subfields of `LustreConfiguration` in the metadata during an FSx for Lustre restore:

- `PerUnitStorageThroughput` - Specifies the throughput capacity per unit of storage provisioned, measured in MB/s per TiB of storage.
- `DriveCacheType` - The type of drive cache used by `PERSISTENT_1` file systems that are provisioned with HDD storage devices. This parameter is required when `StorageType` is set to HDD.
- `DataReadCacheConfiguration` - Specifies the provisioned SSD read cache for Intelligent Tiering file systems. Required when `StorageType` is set to `INTELLIGENT_TIERING`. See LustreReadCacheConfiguration for more details.
- `EfaEnabled` - Specifies whether Elastic Fabric Adapter (EFA) and GPUDirect Storage (GDS) support is enabled for the FSx for Lustre file system.

For complete details about all available parameters in `LustreConfiguration`, please see CreateFileSystemLustreConfiguration in the *Amazon FSx API Reference*.

FSx for ONTAP restore metadata

You can specify the following metadata during an FSx for ONTAP restore:

- Name #name of volume to be created
- OntapConfiguration: # ontap configuration
- `junctionPath`
- `sizeInMegabytes`
- `storageEfficiencyEnabled`
- `storageVirtualMachineId`
- `tieringPolicy`

FSx for OpenZFS restore metadata

You can specify the following subfields of `OpenZFSConfiguration` in the metadata during an FSx for OpenZFS restore:

- `ThroughputCapacity` - Specifies the throughput capacity of the restored file system, measured in MB/s.
- `DiskIopsConfiguration` - When specifying Iops for SSD storage class, use a value between 0 and 160,000. Do not include Mode when Iops is specified.

- `ReadCacheConfiguration` - Specifies the provisioned SSD read cache for Intelligent Tiering file systems. Required when `StorageType` is set to `INTELLIGENT_TIERING`. See OpenZFSReadCacheConfiguration for more details.

For complete details about all available parameters in `OpenZFSConfiguration`, please see CreateFileSystemOpenZFSConfiguration in the *Amazon FSx API Reference*.

Example CLI restore command:

```
aws backup start-restore-job --recovery-point-arn "arn:aws:fsx:us-west-2:1234:backup/
backup-1234" --iam-role-arn "arn:aws:iam::1234:role/Role" --resource-
type "FSx" --region us-west-2 --metadata 'SubnetIds="[\"subnet-1234\",
\"subnet-5678\"]",StorageType=HDD,SecurityGroupIds="[\"sg-bb5efdc4\",
\"sg-0faa52\"]",WindowsConfiguration="{\"DeploymentType\": \"MULTI_AZ_1\",
\"PreferredSubnetId\": \"subnet-1234\",\"ThroughputCapacity\": \"32\"}"'
```

Example restore metadata:

```
"restoreMetadata":  "{\"StorageType\":\"SSD\",\"KmsKeyId\":\"arn:aws:kms:us-
east-1:123456789012:key/123456a-123b-123c-defg-1h2i2345678\",\"StorageCapacity
\":\"1200\",\"VpcId\":\"vpc-0ab0979fa431ad326\",\"FileSystemType\":\"LUSTRE\",
\"LustreConfiguration\":\"{\\\"WeeklyMaintenanceStartTime\\\":\\\"4:10:30\\\",\\
\"DeploymentType\\\":\\\"PERSISTENT_1\\\",\\\"PerUnitStorageThroughput\\\":50,\\
\"CopyTagsToBackups\\\":true}\",\"FileSystemId\":\"fs-0ca11fb3d218a35c2\",\"SubnetIds
\":\"[\\\"subnet-0e66e94eb43235351\\\"]\"}"
```

Restore a Neptune cluster

Use the AWS Backup console to restore Amazon Neptune recovery points

Restoring an Amazon Neptune database requires that you specify multiple restore options. For information about these options, see Restoring from a DB Cluster Snapshot in the *Neptune User Guide*.

To restore an Neptune database

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Protected resources** and the Neptune resource ID that you want to restore.

3. On the **Resource details** page, a list of recovery points for the selected resource ID is shown. To restore a resource, in the **Backups** pane, choose the radio button next to the recovery point ID of the resource. In the upper-right corner of the pane, choose **Restore**.
4. In the **Instance specifications** pane, accept the defaults or specify the **DB engine** and **Version**.
5. In the **Settings** pane, specify a name that is unique for all DB cluster instances owned by your AWS account in the current Region. The DB cluster identifier is case insensitive, but it is stored as all lowercase, as in "`mydbclusterinstance`". This is a required field.
6. In the **Database options** pane, accept the defaults or specify the options for **Database port** and **DB cluster parameter group**.
7. In the **Encryption** pane, accept the default or specify the options for the **Enable encryption** or **Disable encryption** settings.
8. In the **Log exports** pane, choose the log types to publish to Amazon CloudWatch Logs. The **IAM role** is already defined.
9. In the **Restore role** pane, choose the IAM role that AWS Backup will assume for this restore.

   > **Note**
   >
   > If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

10. After specifying all your settings, choose **Restore backup**.

    The **Restore jobs** pane appears. A message at the top of the page provides information about the restore job.
11. After your restore finishes, attach your restored Neptune cluster to an Amazon RDS instance.

Use the AWS Backup API, CLI, or SDK to restore Neptune recovery points

First, restore your cluster. Use `StartRestoreJob`. You can specify the following metadata during Amazon DocumentDB restores:

```
availabilityZones
backtrackWindow
copyTagsToSnapshot // Boolean
databaseName // string
dbClusterIdentifier // string
```

```
dbClusterParameterGroupName // string
dbSubnetGroupName // string
enableCloudwatchLogsExports // string
enableIAMDatabaseAuthentication // Boolean
engine // string
engineMode // string
engineVersion // string
kmsKeyId // string
port // integer
optionGroupName // string
scalingConfiguration
vpcSecurityGroupIds // string
```

Then, attach your restored Neptune cluster to an Amazon RDS instance using `create-db-instance`.

- For Linux, macOS, or Unix:

```
aws neptune create-db-instance --db-instance-identifier sample-instance \
                  --db-instance-class db.r5.large --engine neptune --engine-
version 1.0.5.0 --db-cluster-identifier sample-cluster --region us-east-1
```

- For Windows:

```
aws neptune create-db-instance --db-instance-identifier sample-instance ^
                  --db-instance-class db.r5.large --engine neptune --engine-
version 1.0.5.0 --db-cluster-identifier sample-cluster --region us-east-1
```

For more information, see `RestoreDBClusterFromSnapshot` in the *Neptune Management API reference* and `restore-db-cluster-from-snapshot` in the *Neptune CLI guide*.

Restore an RDS database

Restoring an Amazon RDS database requires specifying multiple restore options. For more information about these options, see Backing Up and Restoring an Amazon RDS DB Instance in the *Amazon RDS User Guide*.

Use the AWS Backup console to restore Amazon RDS recovery points

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.

2. In the navigation pane, choose **Protected resources** and the Amazon RDS resource ID you want to restore.
3. On the **Resource details** page, a list of recovery points for the selected resource ID is shown. To restore a resource, in the **Backups** pane, choose the radio button next to the recovery point ID of the resource. In the upper-right corner of the pane, choose **Restore**.
4. In the **Instance specifications** pane, accept the defaults or specify the options that you need.
5. In the **Settings** pane, specify a name that is unique for all DB instances and clusters owned by your AWS account in the current Region. The DB instance identifier is case insensitive, but it is stored as all lowercase, as in "`mydbinstance`". This is a required field.
6. In the **Network & Security** pane, accept the defaults or specify the options that you need.
7. In the **Database options** pane, accept the defaults or specify the options that you need.
8. In the **Encryption** pane, use the default settings. If the source database instance for the snapshot was encrypted, the restored database instance will also be encrypted. This encryption cannot be removed.
9. In the **Log exports** pane, choose the log types to publish to Amazon CloudWatch Logs. The **IAM role** is already defined.
10. In the **Maintenance** pane, accept the default or specify the option for **Auto minor version upgrade**.
11. In the **Restore role** pane, choose the IAM role that AWS Backup will assume for this restore.

    > **Note**
    >
    > If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

12. Choose **Restore backup**.

    The **Restore jobs** pane appears. A message at the top of the page provides information about the restore job.

Use the AWS Backup API, CLI, or SDK to restore Amazon RDS recovery points

Use `StartRestoreJob`. For information on accepted metadata and values, see `RestoreDBInstanceFromDBSnapshot` and `RestoreDBInstanceToPointInTime` in the

Amazon RDS API Reference. Additionally, AWS Backup accepts the following information-only attributes. However, including them will not affect the restore:

```
EngineVersion
KmsKeyId
Encrypted
vpcId
```

Restore an Amazon Redshift cluster

You can restore automated and manual snapshots in the AWS Backup console or through CLI.

When you restore a Amazon Redshift cluster, the original cluster settings are input into the console by default. You can specify different settings for the configurations below. When restoring a table, you must specify the source and target databases. For more information on these configurations, see Restoring a cluster from a snapshot in the *Amazon Redshift Management Guide*.

- **Single table or cluster**: You can choose to restore an entire cluster or a single table. If you choose to restore a single table, the source database, source schema, and source table name are needed, as well as the target cluster, schema, and new table name.
- **Node type**: Each Amazon Redshift cluster consists of a leader node and at least one compute node. When you restore a cluster, you need to specify the node type that meets your requirements for CPU, RAM, storage capacity, and drive type.
- **Number of nodes**: When restoring a cluster, you need to specify the number of nodes needed.
- **Configuration summary**
- **Cluster Permissions**

Additional permissions for clusters with managed passwords

If the cluster uses Amazon Redshift managed passwords, you must set `ManageMasterPassword` to `true` in the restore metadata.
You must also add the required permissions to the IAM role for the restore operation. For more information about Amazon Redshift managed passwords, see Managing Amazon Redshift admin passwords with AWS Secrets Manager in the *Amazon Redshift Management Guide*.
Add the following AWS Secrets Manager permissions to the restore role:

- `secretsmanager:CreateSecret`
- `secretsmanager:DeleteSecret`
- `secretsmanager:DescribeSecret`
- `secretsmanager:GetRandomPassword`
- `secretsmanager:RotateSecret`
- `secretsmanager:TagResource`
- `secretsmanager:UpdateSecret`

If the cluster uses a customer managed key for the managed password secret, you must also grant the following AWS KMS permissions on the AWS KMS key:

- `kms:CreateGrant`
- `kms:Decrypt`
- `kms:GenerateDataKey`
- `kms:RetireGrant`

To restore an Amazon Redshift cluster or table using the AWS Backup console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Settings** and the Amazon Redshift resource ID that you want to restore.
3. On the **Resource details** page, a list of recovery points for the selected resource ID is shown. To restore a resource, in the **Recovery Points** pane, choose the radio button next to the recovery point ID of the resource. In the upper-right corner of the pane, choose **Restore**.
4. Restore Options
    a. Restore cluster from snapshot, or
    b. Restore single table within a snapshot to new cluster. If you choose this options, then you must configure the following:
        i. Toggle on or off case-sensitive names.
        ii. Input the source table values, including the database, the schema, and the table. The source table information can be found in the Amazon Redshift console.

iii. Input the target table values, including the database, the schema, and the new table name.

5. Specify your new cluster configuration settings.

   a. For cluster restore: choose Cluster identifier, Node type, and number of nodes.

   b. Specify availability zone and maintenance windows.

   c. You can associate additional roles by clicking **Associate IAM roles**.

6. *Optional:* Additional configurations:

   a. **Use defaults** is toggled on by default.

   b. Use the dropdown menus to select settings for Networking and security, VPC security groups, Cluster subnet group, and Availability zone.

   c. Toggle **Enhanced VPC routing** on or off.

   d. Determine if you want to make your cluster endpoint **publicly accessible**. If it is, instances and devices outside the VPC can connect to your database through the cluster endpoint. If this is toggled on, input the elastic IP address.

7. *Optional:* Database configuration. You may choose to input

   a. Database port (by typing into the text field)

   b. Parameter groups

8. Maintenance: You can choose the

   a. Maintenance window

   b. Maintenance track, from among current, trailing, or preview. This controls which cluster version is applied during a maintenance window.

9. Automated snapshot is set to default.

   a. Automated snapshot retention period. Retention period must be 0 to 35 days. Choose 0 to not create automated snapshots.

   b. The manual snapshot retention period is 1 to 3653 days.

   c. There is an optional checkbox for cluster relocation. If this is checked, it permits the ability to relocate your cluster in another Availability Zone. After you enable relocation, you can use the VPC endpoint.

10. Monitoring: After a cluster is restored, you can set up monitoring through CloudWatch or Amazon Redshift.

11. Choose IAM role to be passed to perform restores. You can use the default role, or you can specify a different one.

Your restore jobs will be visible under **Jobs**. You can see the current status of your restore job by clicking the refresh button or CTRL-R.

Restore an Amazon Redshift cluster using API, CLI, or SDK

Use `StartRestoreJob` to restore an Amazon Redshift cluster.

To restore a Amazon Redshift using the AWS CLI, use the command `start-restore-job` and specify the following metadata:

```
ClusterIdentifier // required string
AdditionalInfo // optional string
AllowVersionUpgrade // optional Boolean
AquaConfigurationStatus // optional string
AutomatedSnapshotRetentionPeriod // optional integer 0 to 35
AvailabilityZone // optional string
AvailabilityZoneRelocation // optional Boolean
ClusterParameterGroupName // optional string
ClusterSecurityGroups // optional array of strings
ClusterSubnetGroupName // optional strings
DefaultIamRoleArn // optional string
ElasticIp // optional string
Encrypted // Optional TRUE or FALSE
EnhancedVpcRouting // optional Boolean
HsmClientCertificateIdentifier // optional string
HsmConfigurationIdentifier // optional string
IamRoles // optional array of strings
KmsKeyId // optional string
MaintenanceTrackName // optional string
ManageMasterPassword // optional Boolean
ManualSnapshotRetentionPeriod // optional integer
MasterPasswordSecretKmsKeyId // optional string
NodeType // optional string
NumberOfNodes // optional integer
OwnerAccount // optional string
Port // optional integer
PreferredMaintenanceWindow // optional string
PubliclyAccessible // optional Boolean
ReservedNodeId // optional string
```

```
SnapshotClusterIdentifier // optional string
SnapshotScheduleIdentifier // optional string
TargetReservedNodeOfferingId // optional string
VpcSecurityGroupIds // optional array of strings
RestoreType // CLUSTER_RESTORE or TABLE_RESTORE or NAMESPACE_RESTORE
```

For more information, see RestoreFromClusterSnapshot in the *Amazon Redshift API Reference* and restore-from-cluster-snapshot in the *AWS CLI guide*.

Here is an example template:

```
aws backup start-restore-job \
-\-recovery-point-arn "arn:aws:redshift:region:account-id:snapshot:cluster-
name/snapshot-name" \
-\-iam-role-arn "arn:aws:iam::account:role/role-name" \
-\-metadata 'metadata' \
-\-resource-type Redshift \
-\-region AWS Region \
-\-endpoint-url URL
```

Here is an example:

```
aws backup start-restore-job \
-\-recovery-point-arn "arn:aws:redshift:us-west-2:123456789012:snapshot:redshift-
cluster-1/awsbackup:job-c40dda3c-fdcc-b1ba-fa56-234d23209a40" \
-\-iam-role-arn "arn:aws:iam::974288443796:role/Backup-Redshift-Role" \
-\-metadata 'RestoreType=CLUSTER_RESTORE,ClusterIdentifier=redshift-cluster-
restore-78,Encrypted=true,KmsKeyId=45e261e4-075a-46c7-9261-dfb91e1c739c' \
-\-resource-type Redshift \
-\-region us-west-2 \
```

You can also use DescribeRestoreJob to assist with restore information.

In the AWS CLI, use the operation `describe-restore-job` and use the following metadata:

```
Region
```

Here is an example template:

```
aws backup describe-restore-job -\-restore-job-id restore job ID \
-\-region AWS Region
```

Here is an example:

```
aws backup describe-restore-job -\-restore-job-id BEA3B353-576C-22C0-9E99-09632F262620
 \
-\-region us-west-2 \
```

Amazon Redshift Serverless restore

You can restore manual snapshots of databases or tables using the AWS Backup console or AWS CLI.

Redshift Serverless and AWS Backup support *interchangeable restore* for data warehouse snapshots. This means you can restore Redshift Serverless backups to Amazon Redshift provisioned clusters or restore provisioned backups to Redshift Serverless namespaces. This applies only to full database restore, not single table restore.

Restore capabilities for Redshift Serverless

Restore capabilities	Namespace	Single table
Type of snapshot	Manual	Manual
Information needed	• Source snapshot • Target namespace • Workgroup	• Source snapshot • Source database • Source table name • Target database • New table name
Restore target effect	Restores to an existing namespace through a destructive restore that overwrites existing data	Restores to a new table
Interchangeable restore?	Yes. • Redshift Serverless backups can be restored to Amazon Redshift provisioned clusters.	Not supported.

Restore capabilities	Namespace	Single table
	• Amazon Redshift provisioned backups can be restored to Redshift Serverless clusters.	

For more information about configurations, see Snapshots and recovery points in the *Amazon Redshift Management Guide*.

Considerations before restoring

Before you begin a restore job, review the following:

Configurations

When you restore an Redshift Serverless snapshot, you choose the target namespace to where you want to restore all the databases or a single table.

When you restore the databases in a snapshot to a Serverless namespace, it is a destructive restore. This means all previously extant data in the target restore namespace is overwritten when you restore to that namespace.

When you restore a single table, it is not a destructive restore. To restore a table, specify the workgroup, snapshot, source database, source table, target restore namespace, and the new table name.

Permissions

The permissions required are determined by the target data warehouse (that is, the namespace or provisioned cluster where you will restore the databases or table). The following table can help you determine the permissions, role, and policy to use. For more information on managing IAM policies, see Identity and access management in Amazon Redshift.

Required permissions and roles for restore operations

Restore target	Needed permissio n(s)	IAM role and policy
Amazon Redshift provisioned cluster	`redshift: RestoreFr omCluster Snapshot`	`AWSBackupServiceRolePolicyF orRestores` contains this permission; it can be used for **aws backup start-restore-job**.
Redshift Serverless namespace	`redshift- serverles s:Restore FromSnapshot`	You must add this permission to the role and policy you will use to call **aws backup start-restore-job**. Since this is a destructive restore job, the service role policy for restores cannot be used.

Redshift Serverless restore procedure

Follow these steps to restore Redshift Serverless backups using the AWS Backup console or AWS CLI:

Console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Settings** and select the Redshift Serverless resource ID to restore.
3. On the **Resource details** page, select the recovery point ID in the **Recovery Points** pane, then choose **Restore**.
4. In the **Restore options** pane, choose to restore the entire data warehouse or a single table.
5. Select the destination target in the **Target data warehouse configuration** pane.
   - For a full data warehouse restore, choose between Amazon Redshift provisioned cluster or Redshift Serverless namespace.
   - For a single table restore, specify the source snapshot, database, schema, table name, and target details.

6. Choose the IAM restore role for the job. If not using the default role, ensure the selected role includes the `iam:PassRole` permission.

> **Note**
>
> If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

AWS CLI

Use the **aws backup start-restore-job** command.

AWS Backup works with Redshift Serverless to orchestrate the restore job. The CLI command will be prepended with `aws backup` but will also contain metadata relevant to Redshift Serverless or Amazon Redshift.

The required and optional metadata depends on whether you're restoring a whole data warehouse or a single table.

- For single table restore, see restore-table-from-snapshot in the *AWS CLI Command Reference*.
- For namespace restore, see restore-from-snapshot in the *AWS CLI Command Reference*.
- To restore to a Amazon Redshift provisioned cluster, see restore-from-cluster-snapshot in the *AWS CLI Command Reference*.

Example template for `start-restore-job` to restore to a Serverless namespace:

```
aws backup start-restore-job \
--recovery-point-arn "arn:aws:backup:region:account:snapshot:name--iam-role-arn
 "arn:aws:iam:account:role/role-name" \
--metadata \
--resource-type Redshift Serverless \
--region Region \
--endpoint-url URL
```

Example for start-restore-job to restore to a Serverless namespace:

```
aws backup start-restore-job \
```

```
--recovery-point-arn "arn:aws:redshift-serverless:us-east-1:123456789012:snapshot/
a12bc34d-567e-890f-123g-h4ijk56l78m9" \
--iam-role-arn "arn:aws:iam::974288443796:role/Backup-Redshift-Role" \
--metadata 'RestoreType=NAMESPACE_RESTORE,NamespaceIdentifier=redshift-namespace-1-
restore' \
--resource-type "RedshiftServerless" \
--region us-west-2
```

After starting the restore job, use **describe-restore-job** to monitor progress.

Restore an SAP HANA database on an Amazon EC2 instance

SAP HANA databases on EC2 instances can be restored using the AWS Backup console, using API, or using AWS CLI.

Topics

- Restore an SAP HANA database with the AWS Backup console
- StartRestoreJob API for SAP HANA on EC2
- CLI for SAP HANA on EC2
- SAP HANA High Availability (HA) restore
- Troubleshooting

Restore an SAP HANA database with the AWS Backup console

Note that backup jobs and restore jobs involving the same database cannot occur concurrently. When an SAP HANA database restore job is occurring, attempts to back up the same database will likely result in an error: "Database cannot be backed up while it is stopped."

1. Access the AWS Backup console using the credentials from prerequisites.
2. Under the **Target restore location** dropdown menu, choose a database to overwrite with the recovery point you are using to restore (note that the instance hosting the restore target database must also have the permissions from the prerequisites).

> **Important**
>
> SAP HANA database restores are destructive. Restoring a database will overwrite the database at the specified target restore location.

3. Complete this step only if you are performing a system copy restore; otherwise, skip to step 4.

   System copy restores are restore jobs which restore to a target database different from the source database which generated the recovery point. For system copy restores, notice the `aws ssm-sap put-resource-permission` command provided for you on the console. This command must be copied, pasted, and executed on the machine that completed the prerequisites. When running the command, use the credentials from the role in the prerequisite where you set up the required permissions for registering applications.

   ```
   // Example command
   aws ssm-sap put-resource-permission \
   --region us-east-1 \
   --action-type RESTORE \
   --source-resource-arn arn:aws:ssm-sap-east-1:112233445566:HANA/Foo/DB/HDB \
   --resource-arn arn:aws:ssm-sap:us-east-1:112233445566:HANA/Bar/DB/HDB
   ```

4. Once you choose the restore location, you can see the target database's **Resource ID**, **Application name**, **Database type**, and the **EC2 instance**.
5. *Optionally*, you may expand **Advanced restore settings** to change your catalog restore option. Available options vary based on selected restore settings.
6. Click **Restore backup**.
7. The target location will be overwritten during restore (**"destructive restore"**), so you must provide confirmation that you permit this in the next pop-up dialog box.

   a. To proceed, you must understand that the existing database will be overwritten by the one you are restoring.

   b. Once this is understood, you must acknowledge the existing data will be overwritten. To acknowledge this and to proceed, type **overwrite** into the text input field.

8. Click **Restore backup**.

If the procedure was successful, a blue banner will appear at the top of the console. This signifies that the restore job is in progress. You will be automatically redirected to the Jobs page where your restore job will appear in the list of restore jobs. This most recent job will have a status of `Pending`. You can search for and then click on the restore job ID too see details of each restore job. You can refresh the restore jobs list by clicking the refresh button to view changes to the restore job status.

StartRestoreJob API for SAP HANA on EC2

This action recovers the saved resource identified by an Amazon Resource Name (ARN).

Request Syntax

```
PUT /restore-jobs HTTP/1.1
Content-type: application/json
{
   "IdempotencyToken": "string",
   "Metadata": {
      "string" : "string"
   },
   "RecoveryPointArn": "string",
   "ResourceType": "string"
}
```

URI Request Parameters: The request does not use any URI parameters.

Request Body: The request accepts the following data in JSON format:

IdempotencyTokenA customer-chosen string that you can use to distinguish between otherwise identical calls to `StartRestoreJob`. Retrying a successful request with the same idempotency token results in a success message with no action taken.

Type: String

Required: No

Metadata

A set of metadata key-value pairs. Contains information, such as a resource name, required to restore a recovery point. You can get configuration metadata about a resource at the time it was backed up by calling `GetRecoveryPointRestoreMetadata`. However, values in addition to those provided by `GetRecoveryPointRestoreMetadata` might be required to restore a resource. For example, you might need to provide a new resource name if the original already exists.

You need to include specific metadata to restore an SAP HANA on Amazon EC2 instance. See StartRestoreJob metadata for SAP HANA-specific items.

To retrieve the relevant metadata, you can use the call GetRecoveryPointRestoreMetadata.

Example of a standard SAP HANA database recovery point:

```
"RestoreMetadata": {
        "BackupSize": "1660948480",
        "DatabaseName": "DATABASENAME",
        "DatabaseType": "SYSTEM",
        "HanaBackupEndTime": "1674838362",
        "HanaBackupId": "1234567890123",
        "HanaBackupPrefix": "1234567890123_SYSTEMDB_FULL",
        "HanaBackupStartTime": "1674838349",
        "HanaVersion": "2.00.040.00.1553674765",
        "IsCompressedBySap": "FALSE",
        "IsEncryptedBySap": "FALSE",
        "SourceDatabaseArn": "arn:aws:ssm-sap:region:accountID:HANA/applicationID/
DB/DATABASENAME",
        "SystemDatabaseSid": "HDB",
        "aws:backup:request-id": "46bbtt4q-7unr-2897-m486-yn378k2mrw9c"
    }
```

Example of a continuous SAP HANA database recovery point:

```
"RestoreMetadata": {
        "AvailableRestoreBases":
 "[1234567890123,9876543210987,1472583691472,7418529637418,1678942598761]",
        "BackupSize": "1711284224",
        "DatabaseName": "DATABASENAME",
        "DatabaseType": "TENANT",
        "EarliestRestorablePitrTimestamp": "1674764799789",
        "HanaBackupEndTime": "1668032687",
        "HanaBackupId": "1234567890123",
        "HanaBackupPrefix": "1234567890123_HDB_FULL",
        "HanaBackupStartTime": "1668032667",
        "HanaVersion": "2.00.040.00.1553674765",
        "IsCompressedBySap": "FALSE",
        "IsEncryptedBySap": "FALSE",
        "LatestRestorablePitrTimestamp": "1674850299789",
        "SourceDatabaseArn": "arn:aws:ssm-sap:region:accountID:HANA/applicationID/
DB/SystemDatabaseSid",
        "SystemDatabaseSid": "HDB",
        "aws:backup:request-id": "46bbtt4q-7unr-2897-m486-yn378k2mrw9d"
    }
```

CLI for SAP HANA on EC2

The command `start-restore-job` recovers the saved resource identified by an Amazon Resource Name (ARN). CLI will follow the API guideline above.

Synopsis:

```
start-restore-job
--recovery-point-arn value
--metadata value
--aws:backup:request-id value
[--idempotency-token value]
[--resource-type value]
[--cli-input-json value]
[--generate-cli-skeleton value]
[--debug]
[--endpoint-url value]
[--no-verify-ssl]
[--no-paginate]
[--output value]
[--query value]
[--profile value]
[--region value]
[--version value]
[--color value]
[--no-sign-request]
[--ca-bundle value]
[--cli-read-timeout value]
[--cli-connect-timeout value]
```

Options

`--recovery-point-arn` (string) is a string in the form of an Amazon Resource Number (ARN) that uniquely identifies a recovery point; for example `arn:aws:backup:region:123456789012:recovery-point:46bbtt4q-7unr-2897-m486-yn378k2mrw9d`

`--metadata` (map): A set of metadata key-value pairs. Contains information, such as a resource name, required to restore a recovery point. You can get configuration metadata about a resource at the time it was backed up by calling `GetRecoveryPointRestoreMetadata` . However, values in addition to those provided by `GetRecoveryPointRestoreMetadata` might be required to

restore a resource. You need to specify specific metadata to restore an SAP HANA on Amazon EC2 instance:

- `aws:backup:request-id`: This is any UUID string used for idempotency. It does not alter your restore experience in any way.
- `aws:backup:TargetDatabaseArn`: Specify the database to which you want to restore. This is the SAP HANA on Amazon EC2 database ARN.
- `CatalogRestoreOption`: Specify where to restore your catalog from. One of `NO_CATALOG`, `LATEST_CATALOG_FROM_AWS_BACKUP`, `CATALOG_FROM_LOCAL_PATH`
- `LocalCatalogPath`: If CatalogRestoreOption metadata value is `CATALOG_FROM_LOCAL_PATH`, then specify the path to local catalog on your EC2 instance. This should be a valid file path in your EC2 instance.
- `RecoveryType`: Currently, `FULL_DATA_BACKUP_RECOVERY`, `POINT_IN_TIME_RECOVERY`, and `MOST_RECENT_TIME_RECOVERY` recovery types are supported.

key = (string); value = (string). Shorthand syntax:

```
KeyName1=string,KeyName2=string
```

JSON syntax:

```
{"string": "string"
  ...}
```

`--idempotency-token` is a user-chosen string that you can use to distinguish between otherwise identical calls to `StartRestoreJob`. Retrying a successful request with the same idempotency token results in a success message with no action taken.

`--resource-type` is a string that starts a job to restore a recovery point for one of the following resources: `SAP HANA on Amazon EC2` for SAP HANA on Amazon EC2. *Optionally*, SAP HANA resources can be tagged using the command `aws ssm-sap tag-resource`

Output: `RestoreJobId` is a string that uniquely identifies the job that restores a recovery point.

SAP HANA High Availability (HA) restore

There are important considerations and additional steps to include when you are restoring a high availability (HA) system of SAP HANA. Expand the section below that best aligns your use case.

Restore scenario:

System database to an SAP HANA HA target

Before you restore to the target (destination) SAP HANA HA system,

1. If a cluster is installed, put all cluster notes in Maintenance mode.
2. Stop the SAP HANA database on all nodes, including primary and secondary.
3. *(Recommended)* Disable any backup plans to ensure they don't interfere with the restore operation.

After the restore job completes, go to the restored SAP HANA HA system, then:

1. Start the SAP HANA database on the primary mode.
2. Manually start any tenant database in which the system database was restored but its tenants were not restored.
3. Re-establish SAP HANA system replication (HSR) between the primary and secondary nodes.
4. Start the SAP HANA database on the secondary node.
5. If a cluster is installed, ensure all cluster nodes are online.
6. Enable any backup plans you disabled prior to the restore operation.

(Optional) You can keep the application in sync on AWS Systems Manager for SAP by calling `StartApplicationRefresh`, or you can wait for the scheduled application refresh that will bring the latest SAP metadata.

System database to an SAP HANA single-node target

Before you begin a restore job, go to the target single-node SAP HANA system, then:

1. Stop the SAP HANA database on the target SAP HANA system.
2. *(Recommended)* Disable any backup plans to ensure they don't interfere with the restore operation.

After the restore job completes, go to the target single-node SAP HANA system, then:

1. Start SAP HANA on the target SAP HANA system.

2. Manually start each tenant database on the target node.
3. Enable any backup plans you disabled prior to the restore operation.

(Optional) You can keep the application in sync on AWS Systems Manager for SAP by calling `StartApplicationRefresh`, or you can wait for the scheduled application refresh that will bring the latest SAP metadata.

Tenant database (in place or system copy)

Before you start a restore job, go to the target SAP HANA system, then:

1. *(Optional, but recommended)* Put any installed clusters into maintenance mode to avoid an unexpected takeover during the restore operation.
2. Ensure the system database is running on the target SAP HANA system.
3. *(Recommended)* Disable any backup plans to ensure they don't interfere with the restore operation.

After the restore job completes:

- Enable any backup plans you disabled prior to the restore operation.

Troubleshooting

If any of the following errors occur while attempting a backup operation, see the associated resolution.

- **Error:** Continuous backup log error

  To maintain recovery points for continuous backups, logs are created by SAP HANA for all changes. When the logs are unavailable, the status of each of these continuous recovery points is `STOPPED`. The last certain viable recovery point that can be used to restore is one that has the status of `AVAILABLE`. If the log data is missing for the time between recovery points with a `STOPPED` status and points with `AVAILABLE`, these times cannot be guaranteed to have a successful restore. If you input a date and time within this range, AWS Backup will attempt the backup, but will use the closest available restorable time. This error will be shown by the message "`Encountered an issue with log backups. Please check SAP HANA for details.`"

Resolution: In the console, the most recent restorable time, based on the logs, is displayed. You can input a time more recent than the time shown. However, if the data for this time is unavailable from the logs, AWS Backup will use the most recent restorable time.

- **Error:** `Internal error`

  Resolution: Create a support case from your console or contact Support with the details of your restore such as the restore job ID.

- **Error:** `The provided role arn:aws:iam::ACCOUNT_ID:role/ServiceLinkedRole cannot be assumed by AWS Backup`

  Resolution: Ensure that the role assumed when calling the restore has the required permissions to create service linked roles.

- **Error:** `User: arn:aws:sts::ACCOUNT_ID:assumed-role/ServiceLinkedRole/AWSBackup-ServiceLinkedRole is not authorized to perform: ssm-sap:GetOperation on resource: arn:aws:ssm-sap:us-east-1:ACCOUNT_ID:...`

  Resolution: Ensure that the role assumed when calling the restore permissions outlined in the prerequisites is entered correctly.

- **Error:** `b* 449: recovery strategy could not be determined: [111014] The backup with backup id '1660627536506' cannot be used for recovery SQLSTATE: HY000\n`

  Resolution: Ensure that Backint agent was properly installed. Check all the prerequisites, particularly Install AWS BackInt Agent and AWS Systems Manager for SAP on your SAP application server and then retry installing the BackInt Agent again.

- **Error:** `IllegalArgumentException: Restore job provided is not ready to return chunks, current restore job status is: CANCELLED`

  Resolution: Restore job was cancelled by the service workflow. Retry restore job.

- **Error:** Encountered an issue restore a tenant database on an SAP HANA High Availability system: `b* -10709: Connection failed (RTE:[89006] System call 'connect' failed, rc=111:Connection refused ([::1]:40404 # localhost:30013))\n`

  Resolution: Check SAP HANA to ensure that the SYSTEMDB is up and running.

- **Error:** `b'* 448: recovery could not be completed: [301102] exception 301153: Sending root key to secondary failed: connection refused. This may`

```
be caused by a stopped system replication secondary. Please keep the
secondary online to receive the restored root key. Alternatively you
could unregister the secondary site in case of an urgent recovery.\n
SQLSTATE: HY000\n'
```

Resolution: On a SAP HANA High Availability system, SAP HANA may not be running on the secondary node while an active restore operation is running. Start SAP HANA on the secondary node, then retry the restore job again.

- **Error:** `RequestError: send request failed\ncaused by: read tcp 10.0.131.4:40482->35.84.99.47:443: read: connection timed out"`

  Resolution: Transient network instability is occurring on the instance. Retry the restore. If this issue happens consistently, try adding `ForceRetry: "true"` to agent config file at `/hana/shared/aws-backint-agent/aws-backint-agent-config.yaml`.

For any other AWS Backint agent related issue, refer to Troubleshoot AWS Backint Agent For SAP HANA.

Restore S3 data using AWS Backup

You can restore the S3 data that you backed up using AWS Backup to the S3 Standard storage class. You can restore all the objects in a bucket or specific objects. You can restore them to an existing or new bucket.

Amazon S3 restore permissions

Before you begin restoring resources, ensure the role you're using has sufficient permissions.

For more information, see the following entries on policies:

1. AWSBackupServiceRolePolicyForS3Restore
2. AWSBackupServiceRolePolicyForRestores
3. Managed policies for AWS Backup

Amazon S3 restore considerations

- If ACLs were enabled during the backup, the destination bucket for the restore must have ACLs enabled; otherwise, the restore job fails.

- If Block Public Access is enabled on the destination bucket, the restore job completes successfully, but objects with public ACLs are not restored.
- Restores of objects are skipped if the destination bucket has an object with the same name or version ID.
- When you restore to the original S3 bucket,
  - AWS Backup does not perform a destructive restore, which means AWS Backup will not put an object into a bucket in place of an object that already exists, regardless of version.
  - A delete marker in the current version is treated as the object as nonexistent, so a restore can occur.
  - AWS Backup does not delete objects (without delete markers) from a bucket during a restore (example: keys currently in the bucket which were not present during the backup will remain).
- **Restoring cross-Region copies**
  - While S3 backups can be copied cross-Region, restore jobs only occur in the same Region in which the original backup or copy is located.

    Example

    Example: An S3 bucket created in US East (N. Virginia) Region can be copied to Canada (Central) Region. The restore job can be initiated using the original bucket in US East (N. Virginia) Region and restored to that Region, or the restore job can be initiated using the copy in Canada (Central) Region and restored to that Region.
  - The original encryption method cannot be used to restore a recovery point (backup) copied from another Region. Cross-Region copy AWS KMS encryption is not available for Amazon S3 resources; instead, use a different encryption type for a restore job.

Restoring ACLs and object tags

When restoring Amazon S3 data, you choose whether ACLs are part of the restore.

If ACLs are available in the recovery point, you choose to restore or exclude ACLs using the Restore ACLs setting; if ACLs were not in the backup, they cannot be restored regardless of the setting. If you try to create a restore job with ACls enabled but they were not part of the backup, you may see an error such as `Unable to restore Access Control Lists (ACLs) for bucket because backup was created with the 'BackupACLs' option disabled. Please proceed with restoring without ACLs`.

Object tags are automatically restored if they were included in the original backup

> **Note**
>
> Restoring recovery points without ACLs
> If through AWS CLI you attempt to restore ACLs from a backup that excluded ACLs, the restore operation will fail with an error message indicating invalid restore parameters.

Restore multiple versions

By default, AWS Backup restores only the latest version of your objects. You have the choice to restore additional or all versions of the objects.

See step 6 in the following section for how to restore up the 10 latest versions or all versions using the AWS Backup console.

See ??? later on this page for metadata details to include when restoring programmatically.

Restore through the AWS Backup console

To restore your Amazon S3 data using the AWS Backup console:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Protected resources**, and select the Amazon S3 resource ID that you want to restore.
3. On the **Resource details** page, you will see a list of recovery points for the selected resource ID. To restore a resource:

   a. In the **Backups** pane, choose the recovery point ID of the resource.

   b. In the upper-right corner of the pane, choose **Restore**.

   (Alternatively, you can go to the backup vault, find the recovery point, and then click **Actions** then click **Restore**.)
4. If you are restoring a continuous backup, in the **Restore time** pane, select either option:

   a. Accept the default to restore to the **Latest restorable time**.

   b. **Specify date and time** to restore.
5. In the **Settings** pane, specify whether to **Restore entire bucket** or perform **Item level restore**.

a. If you choose **Item level restore**, you restore up to 5 items (objects or folders in a bucket) per restore job by specifying each item's S3 URI that uniquely identifies that object.

   (For more information about S3 bucket URIs, see Methods for accessing a bucket in the *Amazon Simple Storage Service User Guide*.)

b. Choose **Add item** to specify another item to restore.

6. By default, only the latest version of an object is restored. You can restore up to the 10 latest versions or restore all versions of the objects. Select your preference from the drop-down menu.

7. Choose your **Restore destination**. You can either **Restore to source bucket**, **Use existing bucket**, or **Create new bucket**.

> **Note**
>
> Your restore destination bucket must have versioning turned on. AWS Backup notifies you if the bucket you select does not meet this requirement.

a. If you choose **Use existing bucket**, select the destination S3 bucket from the menu which shows all existing buckets within your current AWS Region.

b. If you choose **Create new bucket**, type in the **new bucket name**. After the bucket is created, you can modify the BPA (Block Public Access) and S3 versioning default settings.

8. For the encryption of objects in your S3 bucket, you can choose your **Restored object encryption**. Use **original encryption keys** (default), **Amazon S3 key (SSE-S3)**, or **AWS Key Management Service key (SSE-KMS)**.

   These settings only apply to encryption of the objects in the S3 bucket. This does not affect the encryption for the bucket itself.

a. **Use original encryption keys (default)** restores objects with the same encryption keys used by the source object. If a source object was unencrypted, this method restores the object without encryption.

   This restore option allows you to optionally choose a substitute encryption key to encrypt the restore object(s) if the original key is unavailable.

b. If you choose **Amazon S3 key (SSE-S3)**, you do not need to specify any other options.

c. If you choose **AWS Key Management Service key (SSE-KMS)**, you can make the following choices: **AWS managed key (aws/s3)**, **Choose from your AWS KMS keys**, or **Enter AWS KMS key ARN**.

   i. If you choose **AWS managed key (aws/s3)**, you do not need to specify any other options.

   ii. If you **Choose from your AWS KMS keys**, select a AWS KMS key from the dropdown menu. Alternatively, choose **Create key**.

   iii. If you **Enter AWS KMS key ARN**, type in the ARN into the text box. Alternatively, choose **Create key**.

9. In the **Restore role** pane, choose the IAM role that AWS Backup will assume for this restore.

> **Note**
>
> If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

10. Choose **Restore backup**. The **Restore jobs** pane appears. A message at the top of the page provides information about the restore job.

Restore Amazon S3 recovery points through AWS CLI

Use `StartRestoreJob`. You can specify the following metadata during Amazon S3 restores:

```
// Mandatory metadata:
DestinationBucketName // The destination bucket for your restore.

// Optional metadata:
RestoreACLs // Boolean. If ACLs were part of the backup, include and set to TRUE. If
 the backup
does not include ACLs and this parameter is included, set to FALSE.
EncryptionType // The type of encryption to encrypt your restored objects. Options
 are original (same encryption as the original object), SSE-S3, or SSE-KMS).
ItemsToRestore // A list of up to five paths of individual objects to restore. Only
 required for item-level restore.
KMSKey // Specifies the SSE-KMS key to use. Only needed if encryption is SSE-KMS.
RestoreLatestVersionsUpTo // Include this optional parameter to multiple versions.
```

```
RestoreTime // The restore time (only valid for continuous recovery points where it is
 required, in format 2021-11-27T03:30:27Z).
```

`RestoreLatestVersionsUpTo` is an optional metadata key-value pair. By default, or if this is omitted, the latest version is restored. Include this metadata to restore additional versions of your objects. Accepted values are:

- 1 (to restore the latest version)
- *n* , where *n* is any positive integer greater than 1. The latest *n* versions of your objects will be restored. If the actual version count of an object is less than *n*, that number of versions will be restored for that object.
- `all` (to restore all versions)

Recovery point status

Recovery points will have a status indicating their state.

`EXPIRED` status indicates that the recovery point has exceeded its retention period, but AWS Backup lacks permission or is otherwise unable to delete it. To manually delete these recovery points, see Step 3: Delete the recovery points in the *Clean up resources* section of *Getting started*.

`STOPPED` status occurs on a continuous backup where a user has taken some action that causes the continuous backup to be disabled. This can be caused by the removal of permissions, turning off versioning, turning off events being sent to Amazon EventBridge, or disabling the EventBridge rules that are put in place by AWS Backup.

To resolve `STOPPED` status, ensure that all requested permissions are in place and that versioning is enabled on the S3 bucket. Once these conditions are met, the next instance of a backup rule running will result in a new continuous recovery point being created. The recovery points with STOPPED status do not need to be deleted.

S3 restore messages

When a restore job completes or fails, you may see the following message. The following table can help you determine the possible cause of the status message.

Scenario	Job Status	Message	Example
All objects failed to be restored.	FAILED	"No objects were restored from *RecoveryPointARN* to *bucket*. To get notified of these failures, enable SNS event notifications."	The role used to start the restore job does not have permission to put objects in the destination bucket. The restore role does not have permissio n to verify if object version exists in the destination bucket.
One or more (but not all) objects failed to be restored.	COMPLETED	"One or more objects failed to be restored from *RecoveryPointARN* to *bucket*. To get notified of these failures, enable SNS event notifications."	The role used to start the restore job does not have access to the KMS key used by one or more of the original objects.
There are no objects to restore.	COMPLETED	"There are no objects that match the restore request for *RecoveryPointARN* ."	The recovery point (backup) of source bucket to be restored has no objects. The prefix used for the restore job does not correspond with any object.

Restore a Storage Gateway volume

If you are restoring an AWS Storage Gateway volume snapshot, you can choose to restore the snapshot as an Storage Gateway volume or as an Amazon EBS volume. This is because AWS Backup

integrates with both services, and any Storage Gateway snapshot can be restored to either an Storage Gateway volume or an Amazon EBS volume.

Restore Storage Gateway through the AWS Backup console

To restore an Storage Gateway volume

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Protected resources** and then choose the Storage Gateway resource ID you want to restore.
3. On the **Resource details** page, a list of recovery points for the selected resource ID is shown. To restore a resource, in the **Backups** pane, choose the radio button next to the recovery point ID of the resource. In the upper-right corner of the pane, choose **Restore**.
4. Specify the restore parameters for your resource. The restore parameters you enter are specific to the resource type that you selected.

   For **Resource type**, choose the AWS resource to create when restoring this backup.
5. If you choose **Storage Gateway volume**, choose a **Gateway** in a reachable state. Also choose your **iSCSI target name**.

   1. For "Volume stored" gateways, choose a **Disk Id**.
   2. For "Volume cached" gateways, choose a capacity that is at least as large as your protected resource.

   If you choose **EBS volume**, provide the values for **Volume type**, **Size (GiB)**, and choose an **Availability zone**.
6. For **Restore role**, choose the IAM role that AWS Backup will assume for this restore.

> **Note**
>
> If the AWS Backup default role is not present in your account, a **Default role** is created for you with the correct permissions. You can delete this default role or make it unusable.

> **Note**
>
> If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

7. Choose **Restore backup**.

   The **Restore jobs** pane appears. A message at the top of the page provides information about the restore job.

Restore Storage Gateway with AWS CLI

In the command line interface, `start-restore-job` allows you to restore a Storage Gateway volume.

The following list is the accepted metadata.

```
gatewayArn // The Amazon Resource Name (ARN) of the gateway. Use the ListGateways
 operation to return a list of gateways for your account and AWS Region.
gatewayType // The type of created gateway. Valid value is BACKUP_VM
targetName
kmsKey
volumeSize
volumeSizeInBytes
diskId
```

Restore an Amazon Timestream table

When you restore a Amazon Timestream table, there are several options to configure, including the new table name, the destination database, your storage allocation preferences (memory and magnetic storage), and which role you'll use to complete the restore job. You can also choose an Amazon S3 bucket in which to store error logs. Magnetic storage writes are asynchronous, so you may wish you log the errors.

Timestream data storage has two tiers: a memory store and a magnetic store. Memory store is required, but you have the option of transferring your restored table to magnetic storage after the specified memory time is finished. Memory store is optimized for high throughput data writes and

fast point-in-time queries. The magnetic store is optimized for lower throughput late-arrival data writes, long-term data storage, and fast analytical queries.

When you restore a Timestream table, you determine how long you want the table to remain in each storage tier. Using the console or API, you can set the storage time for both. Note that the storage is linear and sequential. Timestream will store your restored table in memory storage first, then automatically transition it to magnetic storage when the memory storage time has been reached.

> **Note**
>
> The magnetic store retention period must be equal or greater than the original retention period (shown at the top-right of the console), or data will be lost.

Example: You set the memory store allocation to hold data for one week and set the magnetic store allocation to hold the same data for one year. When the data in the memory store becomes a week old, it is automatically moved to the magnetic store. It is then retained in the magnetic store for a year. At the end of that time, it is deleted from Timestream and from AWS Backup.

To restore a Amazon Timestream table using the AWS Backup console

You can restore Timestream tables in the AWS Backup console that were created by AWS Backup.

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Protected resources** and the Amazon Timestream resource ID that you want to restore.
3. On the **Resource details** page, a list of recovery points for the selected resource ID is shown. To restore a resource, in the **Backups** pane, choose the radio button next to the recovery point ID of the resource. In the upper-right corner of the pane, choose **Restore**.
4. Specify your new table configuration settings, including:
    a. **New table name**, consisting of 2 to 256 characters (letters, numbers, dashes, periods, and underscores).
    b. **Destination database**, chosen from the drop down menu.
5. **Storage allocation**: Set the amount of time the restored table will first reside in memory storage, and set the amount of time the restored table will then reside in magnetic storage.

Memory storage can be set to hours, days, weeks, or months. Magnetic storage can be set to days, weeks, months, or years.

6. *(Optional)* **Enable magnetic storage writes**: You have the option of allowing magnetic storage writes. With this option checked, late-arriving data, which is data with a timestamp outside the memory storage retention period, will be written directly into the magnetic store.
7. *(Optional)* **Amazon S3 error logs location**: You can specify an S3 location in which your error logs will be stored. Browse your S3 files or copy and paste the S3 file path.

> **Note**
>
> If you choose to specify an S3 error log location, the role you use for this restore must have permission to write to an S3 bucket or it must contain a policy with that permission.

8. Choose the IAM role to be passed to perform restores. You can use the default IAM role or specify a different one.
9. Click **Restore backup**.

Your restore jobs will be visible under protected resources. You can see the current status of your restore job by clicking the refresh button or CTRL-R.

To restore a Amazon Timestream table using API, CLI, or SDK

Use `StartRestoreJob` to restore a Timestream table via API..

To restore a Timestream using the AWS CLI, use the operation `start-restore-job.` and specify the following metadata:

```
TableName: string;
DatabaseName: string;
MemoryStoreRetentionPeriodInHours: value: number unit: 'hours' | 'days' | 'weeks' |
 'months'
MagneticStoreRetentionPeriodInDays: value: number unit: 'days' | 'weeks' | 'months' |
 'years'
EnableMagneticStoreWrites?: boolean;
aws:backup:request-id
```

Here is an example template:

```
aws backup start-restore-job \
--recovery-point-arn "arn:aws:backup:us-west-2:accountnumber:recovery-point:1a2b3cde-
f405-6789-012g-3456hi789012_beta" \
--iam-role-arn "arn:aws:iam::accountnumber:role/rolename" \
--metadata
 'TableName=tablename,DatabaseName=databasename,MagneticStoreRetentionPeriodInDays=1,MemoryStor
\":true,\"MagneticStoreRejectedDataLocation\":{\"S3Configuration\":{\"BucketName\":
\"bucketname\",\"EncryptionOption\":\"SSE_S3\"}}}"' \
--region us-west-2 \
--endpoint-url url
```

You can also use GetRecoveryPointRestoreMetadata to assist with restore information.

In the AWS CLI, use the operation `get-recovery-point-restore-metadata` and use the following metadata:

```
TableName: string;
DatabaseName: string;
MemoryStoreRetentionPeriodInHours: value: number unit: 'hours' | 'days' | 'weeks' |
 'months'
MagneticStoreRetentionPeriodInDays: value: number unit: 'days' | 'weeks' | 'months' |
 'years'
EnableMagneticStoreWrites?: boolean;
```

Here is an example template:

```
aws backup describe-restore-job \
--restore-job-id restore job ID \
--region awsregion \
--endpoint-url url
```

Restore a virtual machine using AWS Backup

You can restore a virtual machine to VMware, VMware Cloud on AWS, VMware Cloud on AWS Outposts, an Amazon EBS volume, or to an Amazon EC2 instance. Restoring (or migrating) a virtual machine to EC2 requires a license. By default, AWS will include a license (charges apply). For more information, see Licensing options in the *VM Import/Export User Guide*.

You can restore a VMware virtual machine using the AWS Backup console or through the AWS CLI. When a virtual machine is restored, the VMware Tools folder is not included. See VMware documentation to reinstall VMware Tools.

AWS Backup restores of virtual machines are non-destructive, meaning AWS Backup does not overwrite existing virtual machines during a restore. Instead, the restore job deploys a new virtual machine.

Tasks

- Considerations when restoring a VM to an Amazon EC2 instance
- Use the AWS Backup console to restore virtual machine recovery points
- Use AWS CLI to restore virtual machine recovery points

Considerations when restoring a VM to an Amazon EC2 instance

- Restoring (or migrating) a virtual machine to EC2 requires a license. By default, an AWS will include a license (charges apply). For more information, see Licensing options in the *VM Import/Export User Guide*.
- There is a maximum limit of 5 TB (terabytes) for each virtual machine disk.
- You can't specify a key pair when you restore the virtual machine to an instance. You can add a key pair to `authorized_keys` during launch (through instance user data) or after launch (as described in this troubleshooting section in the Amazon EC2 User Guide).
- Confirm your operating system is supported for import to and export from Amazon EC2 in the *VM Import/Export User Guide*.
- Review limitations involved with Importing VMs to Amazon EC2 in the *VM Import/Export User Guide*.
- When you restore to an Amazon EC2 instance using AWS CLI, you must specify `"RestoreTo":"EC2Instance"`. All other attributes have default values.
- Amazon EC2 offers EC2 Allowed AMIs. If this setting is enabled in your account, add the alias `aws-backup-vault` to your allowlist. Otherwise, restore operations of VM recovery points to EC2 instances will fail with an error message, such as "Source AMI not found in Region".
- VMware restores to EC2 involving more than 21 disks are not supported. As a workaround, use VMware Restores to EBS to restore each disk individually as an EBS volume, then attach the EBS volumes to an EC2 instance.

Use the AWS Backup console to restore virtual machine recovery points

You can restore a virtual machine from multiple locations in the left navigation pane of the AWS Backup console:

- Choose **Hypervisors** to view recovery points for virtual machines managed by a hypervisor that is connected to AWS Backup.
- Choose **Virtual machines** to view recovery points for virtual machines across all your hypervisors that are connected to AWS Backup.
- Choose **Backup vaults** to view recovery points stored in a specific AWS Backup vault.
- Choose **Protected resources** to view recovery points across all your AWS Backup protected resources.

If you need to restore a virtual machine that no longer has a connection with Backup gateway, choose **Backup vaults** or **Protected resources** to locate your recovery point.

Options

- Restore to VMware
- Restore to an Amazon EBS volume
- Restore to an Amazon EC2 instance

To restore a virtual machine to VMware, VMware Cloud on AWS, and VMware Cloud on AWS Outposts

1. In the **Hypervisors** or **Virtual machines** views, choose the **VM name** to restore. In the **Protected resources** view, choose the virtual machine **Resource ID** to restore.
2. Choose the radial button next to the **Recovery point ID** to restore.
3. Choose **Restore**.
4. Choose the **Restore type**.
    a. **Full restore** restores all the virtual machine's disks.
    b. **Disk-level restore** restores a user-defined selection of one or more disks. Use the drop-down menu to select which disks to restore.
5. Choose the **Restore location**. The options are **VMware**, **VMware Cloud on AWS**, and **VMware Cloud on AWS Outposts**.
6. If you are doing a full restore, skip to the next step. If you are performing a disk-level restore, there will be a drop-down menu under **VM disks**. Choose one or more bootable volumes to restore.

7. Select a **Hypervisor** from the dropdown menu to manage the restored virtual machine
8. For the restored virtual machine, use your organization's virtual machine best practices to specify its:

   a. **Name**

   b. **Path** (such as `/datacenter/vm`)

   c. **Compute resource name** (such as VMHost or Cluster)

      If a host is part of a cluster then you cannot restore to the host but only to the given cluster.

   d. **Datastore**
9. For **Restore role,** select either the **Default role** (recommended) or **Choose an IAM role** using the dropdown menu.

   > **Note**
   >
   > If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.
10. Choose **Restore backup**.
11. *Optional*: Check when your restore job has the status `Completed`. In the left navigation menu, choose **Jobs**.

> **Note**
>
> Restoring to an Amazon EBS volume is a data-level (disk) restore. AWS Backup writes the selected disk's contents to a new Amazon EBS volume, but it does not perform the guest operating system conversion required to boot on Amazon EC2. An Amazon EBS volume restored from a VM recovery point is not bootable and cannot be used on its own to launch a working Amazon EC2 instance. To restore a virtual machine as a bootable Amazon EC2 instance, use the procedure to restore to an Amazon EC2 instance instead, which performs a full restore and prepares the instance to boot on Amazon EC2. Use a restore to Amazon EBS when you need to access or recover data from individual disks.

To restore a virtual machine to an Amazon EBS volume

1. In the **Hypervisors** or **Virtual machines** views, choose the **VM name** to restore. In the **Protected resources** view, choose the virtual machine **Resource ID** to restore.
2. Choose the radial button next to the **Recovery point ID** to restore.
3. Choose **Restore**.
4. Choose the **Restore type**.
   - **Disk restore** restores a user-defined selection of one disk. Use the drop-down menu to select which disk to restore.
5. Choose the **Restore location** as **Amazon EBS**.
6. Under the **VM disk** drop-down menu, choose the disk to restore. The restored Amazon EBS volume contains the data from this disk but is not bootable on Amazon EC2.
7. Under **EBS Volume type**, choose the volume type.
8. Choose your Availability Zone.
9. Encryption (optional). Check the box if you choose to encrypt the EBS volume.
10. Select your KMS key from the menu.
11. For **Restore role,** select either the **Default role** (recommended) or **Choose an IAM role**.

    > **Note**
    >
    > If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

12. Choose **Restore backup**.
13. *Optional*: Check when your restore job has the status `Completed`. In the left navigation menu, choose **Jobs**.
14. *Optional*: Visit How do I use LVM to create a logical volume on an Amazon EBS volume's partition? to learn more on how to mount managed volumes and access data on the restored Amazon EBS volume.

To restore a virtual machine to an Amazon EC2 instance

1. In the **Hypervisors** or **Virtual machines** views, choose the **VM name** to restore. In the **Protected resources** view, choose the virtual machine **Resource ID** to restore.
2. Choose the radial button next to the **Recovery point ID** to restore.
3. Choose **Restore**.
4. Choose the **Restore type**.
   - **Full restore** restores the file system completely, including the root-level folder and files.
5. Choose the **Restore location** as **Amazon EC2**.
6. For **Instance type**, choose the combination of compute and memory required to run your application on your new instance.

   > **Tip**
   >
   > Choose an instance type that matches or exceeds the specifications of the original virtual machine. For more information, see the Amazon EC2 Instance Types Guide.

7. For **Virtual Private Cloud (VPC)**, choose a virtual private cloud (VPC), which defines the networking environment for the instance.
8. For **Subnet**, choose one of the subnets in the VPC. Your instance receives a private IP address from the subnet address range.
9. For **security groups**, choose a security group, which acts as a firewall for traffic to your instance.
10. For **Restore role,** select either the **Default role** (recommended) or **Choose an IAM role**.

    > **Note**
    >
    > If role manager is enabled in your account, AWS Backup selects the default service role for you, and a **Customize** option is available. For more information, see IAM role creation in the *IAM User Guide*.

11. *Optional*: To run a script on your instance at launch, expand **Advanced settings** and enter the script in **User data**.
12. Choose **Restore backup**.

13. *Optional*: Check when your restore job has the status `Completed`. In the left navigation menu, choose **Jobs**.

Use AWS CLI to restore virtual machine recovery points

Use `StartRestoreJob`.

You can specify the following metadata for a virtual machine restore to Amazon EC2 and Amazon EBS:

```
RestoreTo
InstanceType
VpcId
SubnetId
SecurityGroupIds
IamInstanceProfileName
InstanceInitiatedShutdownBehavior
HibernationOptions
DisableApiTermination
Placement
CreditSpecification
RamdiskId
KernelId
UserData
EbsOptimized
LicenseSpecifications
KmsKeyId
AvailabilityZone
EbsVolumeType
IsEncrypted
ItemsToRestore
RequireIMDSv2
NetworkInterfaces
```

AWS Backup supports both partial restores to Amazon EBS and full restores to Amazon EC2. A restore to Amazon EBS is a data-level restore: it writes the selected disk contents to an Amazon EBS volume for data recovery, and the resulting volume is not bootable on Amazon EC2. A restore to Amazon EC2 is a full restore that also prepares the guest operating system to boot on Amazon EC2, producing a launchable instance. For partial restores, use `ItemsToRestore` to specify which disk to restore to the specified EBS volume. When restoring to Amazon EC2, the parameter `ItemsToRestore` can be left blank because it is ignored and the full list of disks is restored.

You can specify the following metadata for a virtual machine restore to VMware, VMware Cloud on AWS, and VMware cloud on AWS Outpost:

```
RestoreTo
HypervisorArn
VMName
VMPath
ComputeResourceName
VMDatastore
DisksToRestore
ItemsToRestore
```

AWS Backup supports both partial and full restores to an on-premises virtual machine. You can choose to restore all disks or only a subset of disks. When performing a partial restore, specify your disk selection in `ItemsToRestore`. When performing a full restore, you must either omit both `DisksToRestore` and `ItemsToRestore`, or specify all the disks in `DisksToRestore`. The `DisksToRestore` parameter does not support subsets of disks.

This example shows how to conduct a full restore to VMware:

```
'{"RestoreTo":"VMware","HypervisorArn":"arn:aws:backup-gateway:us-
east-1:209870788375:hypervisor/hype-9B1AB1F1","VMName":"name","VMPath":"/Labster/
vm","ComputeResourceName":"Cluster","VMDatastore":"vsanDatastore","DisksToRestore":"[{\"DiskId
\":\"2000\",\"Label\":\"Hard disk 1\"}]","vmId":"vm-101"}'
```

Restore testing

Restore testing, a feature offered by AWS Backup, provides automated and periodic evaluation of restore viability, as well as the ability to monitor restore job duration times.

Contents

- Overview
- Restore testing compared with restore process
- Restore testing management
- Create a restore testing plan
- Update a restore testing plan
- View existing restore testing plans

- View restore testing jobs
- Delete a restore testing plan
- Audit restore testing
- Restore testing quotas and parameters
- Restore testing failure troubleshooting
- Restore testing inferred metadata
- Restore testing validation

Overview

First, you create a restore testing plan where you provide a name for your plan, the frequency for your restore tests, and the target start time. Then, you assign the resources you want to include in your plan. You then choose to include specific or random recovery points in your test. AWS Backup backup intelligently infers the metadata that will be needed for your restore job to be successful.

When the scheduled time in your plan arrives, AWS Backup starts restore jobs based on your plan and monitors the time taken to complete the restore.

After the restore test plan completes its run, you can use the results to show compliance for organizational or governance requirements such as the successful completion of restore test scenarios or the restore job completion time.

Optionally, you can use Restore testing validation to confirm the restore test results.

Once the optional validation completes or the validation window closes, AWS Backup deletes the resources involved with the restore test, and the resources will be deleted in accordance with service SLAs.

At the end of the testing process, you can view the results and the completion time of the tests.

Restore testing compared with restore process

Restore testing runs restore jobs in the same way as on-demand restores and uses the same recovery points (backups) as an on-demand restore. You will see calls to `StartRestoreJob` in CloudTrail (if opted-in) for each job started by restore testing

However, there are a few differences between the operation of a schedule restore test and an on-demand restore operation:

	Restore Testing	Restore
Account	Recommended best practice is to designate an account to be used for restore tests	You can restore resources from an account
AWS Backup Audit Manager	Can turn on a control to confirm if a restore test meets specified restore objectives	
Cadence	Periodically as part of a scheduled plan.	On demand
Resources	The resource types you can assign to your testing plan include: Aurora, Amazon DocumentDB, Amazon DynamoDB, Amazon EBS, Amazon EC2, Amazon EFS, Amazon FSx (Lustre, ONTAP, OpenZFS, Windows), Amazon Neptune, Amazon RDS, Amazon Redshift, and Amazon S3.	All resources can be restored.
Results	After the restore testing job is completed, the restored resource is deleted after the Restore testing validation window finishes.	Once the restore job is completed, the restored version of the resource remains.
Tags	For resource types which support tag on restore, testing applies tags on restore.	Tags are optional for supported resources.

Restore testing management

You can create, view, update, or delete a restore testing plan in the AWS Backup console.

You can use AWS CLI to programmatically carry out operations for restore testing plans. Each CLI is specific to the AWS service in which it originates. Commands should be prepended with `aws backup`.

Data deletion

When a restore test is finished, AWS Backup begins deleting the resources involved in the test. This deletion is not instantaneous. Each resource has an underlying configuration that determines how those resources are stored and lifecycled. For example, if Amazon S3 buckets are part of the restore test, lifecycle rules are added to the bucket. It can take up to several days for the rules to execute and for the bucket and its objects to be fully deleted, but charges will only occur for these resources until the day when the lifecycle rule initiates (by default this is 1 days). Speed of deletion will depend upon the resource type.

Resources that are restored by a restore testing plan and support tag-on-restore have a tag of `awsbackup-restore-test`. If a user removes this tag, AWS Backup cannot delete the resource at the end of the testing period and the user will have to delete it manually instead. For resources that do not support tag-on-restore, AWS Backup deletes the resource based on its name.

> **Note**
>
> Restores of DynamoDB, Amazon S3, SAP HANA on Amazon EC2 instances, virtual machines, and Amazon Timestream resources do not currently support tag-on-restore. For these resources, AWS Backup does not apply the `awsbackup-restore-test` tag and instead deletes the resource based on its name. For more information, see Copy tags during a restore.
>
> The deletion of a restore-tested Amazon S3 bucket takes longer than other resource types, because it can take a few days for lifecycle policies to delete all objects within the bucket.

To check why resources may not have been deleted as expected, you can search through failed jobs in the console or use the command line interface to call the API request `DescribeRestoreJob` to retrieve deletion status messages.

Backup plans (non-restore testing plans) ignore resources created by restore testing (those with tag `awsbackup-restore-test` or a name starting with `awsbackup-restore-test`).

Cost control

Restore testing has a cost per restore test. Depending on what resources are included in your restore testing plan, the restore jobs that are part of the plan may also have a cost. See AWS Backup Pricing for full details.

When you set up a restore testing plan for the first time, you may find it beneficial to include a minimum number of resource types and protected resources to familiarize yourself with the feature, the process, and the average costs involved. You can update a plan after its creation to add more resource types and protected resources.

Create a restore testing plan

A restore testing plan has two parts: plan creation and assigning resources.

When you use the console, these parts are sequential. In the first part, you set the name, frequency, and start times. During the second part you assign resources to your testing plan.

When using AWS CLI and API, first use `create-restore-testing-plan`. After you receive a successful response and the plan has been created, then use `create-restore-testing-selection`, for each resource type to include in your plan.

When you create a restore testing plan, we create a service-linked role for you. For more information, see Using roles for restore testing.

Console

Part I: Create a restore testing plan using the console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the left-hand navigation, locate **Restore testing** and select it.
3. Choose **create restore testing plan**.
4. **General**

    a. **Name:** Type in a name for your new restore testing plan. The name cannot be changed after creation. The name must consist of only alphanumeric characters and underscores.

b. **Test frequency:** Choose the frequency at which the restore tests will run.

c. **Start Time:** Set the time (in hour and minute) you prefer the test to begin. You can also set the local time zone in which you want the restore testing plan to operate.

d. **Start within:** This value (in hours) is the period of time in which the restore test is designated to begin. AWS Backup makes a best effort to commence all designated restore jobs during the start within time and randomizes start times within this period.

5. **Recovery point selection:** Here you set the source vaults, the recovery point range, and selection criteria for which recovery points (backups) you want to be part of the plan.

   a. **Source vaults:** Choose whether to include all available vaults or just specific vaults to help filter which recovery points can be in your plan. If you choose **specific vaults**, select from the drop down menu the vaults you wish to include.

   b. **Eligible recovery points:** Specify the time frame from which recovery points will be selected. You can select 1 to 365 days, 1 to 52 weeks, 1 to 12 months, or 1 year.

   c. **Selection criteria:** Once your date range of recovery points is specified, you can choose whether to include the latest one or one at random in your plan. You may wish to choose a random one to gauge the general health of recovery points at more regular frequency in case a restore to an older version is ever warranted.

   d. **Point-in-time recovery points:** If your plan includes resources that have continuous backup (point-in-time-restore/PITR) points, you can check this box to have your testing plan include continuous backups as eligible recovery points (see Feature availability by resource for which resources types have this feature).

6. *(optional)* **Tags added to restore testing plan:** You can choose to add up to 50 tags to your restore testing plan. Each tag must be added separately. To add a new tag, select **Add new tag**.

Part II: Assign resources to the plan using the console

In this section, you choose the resources you have backed up to include in your restore testing plan. You will choose the name of the resource assignment, choose the role you use for the restore test, and set the retention period before cleanup. Then, you will select the resource type, select the scope, and optionally refine your selection with tags.

> **Tip**
>
> To navigate back to the restore testing plan to which you want to add resources, you can go to the AWS Backup console, select **Restore testing**, then find your preferred testing plan and select it.

1. **General**

   a. **Resource assignment name:** Input a name for this resource assignment using a string of alphanumeric characters and underscores, with no white spaces.

   b. **Restore IAM role:** The test must use an Identity and Access Management (IAM) role you designate. You can choose the AWS Backup default role or a different one. If the AWS Backup default does not yet exist as you finish this process, AWS Backup will create it for you automatically with the necessary permissions. The IAM role you choose for restore testing must contain the permissions found in AWSBackupServiceRolePolicyForRestores.

   c. **Retention period before cleanup:** During a restore test, backup data is temporarily restored. By default, this data is deleted after the test is complete. You have the option to delay deletion of this data if you wish to run validation on the restore.

      If you plan to run validation, select **retain for a specific number of hours** and input a value from 1 to 168 hours, inclusive. Note that validation can be run programmatically but not from the AWS Backup console.

2. **Protected resources:**

   a. **Select resource type:** Select which resource types and the scope of which backups of those types to include in the resource testing plan. Each plan can contain multiple resource types, but each type of resource must be assigned to the plan individually.

   b. **Resource selection scope:** Once the type is chosen, select if you want to include all available protected resources of that type or if you want to include specific protected resources only.

   c. *(optional)* **Refine resource selection using tags:** If your backups have tags, you can filter by tags to **select specific protected resources**. Enter the tag key, the condition for this key to be or not to be included, and the value for the key. Then, select the **Add tags** button.

Tags on protected resources are evaluated by checking the tags on the latest recovery point within the backup vault containing the protected resource.

3. **Restore parameters:** Certain resources require specifying parameters in preparation for a restore job. In most cases, AWS Backup will infer the values based on the stored backup.

   It is recommended in most cases to maintain these parameters; however, you can change the values by choosing a different selection from the dropdown menu. Examples where changing the values may be optimal can include overriding encryption keys, Amazon FSx settings where data cannot be inferred, and creation of subnets.

   For example, if an RDS database is one of the resource types you assign to your restore testing plan, parameters such as availability zone, database name, database instance class, and VPC security group will appear with inferred values you can change if applicable.

AWS CLI

The CLI command `CreateRestoreTestingPlan` is used to make a restore testing plan.

The testing plan must contain:

- `RestoreTestingPlan`, which must contain a unique `RestoreTestingPlanName`
- `ScheduleExpression` cron expression
- `RecoveryPointSelection`

  Though named similarly, this is **NOT** the same as `RestoreTestingSelection`.

  `RecoveryPointSelection` has five parameters (three required and two optional). The values you specify determine which recovery point is included in the restore test. You must indicate with `Algorithm` if you want the latest recovery point within your `SelectionWindowDays` or if you want a random recovery point, and you must indicate through `IncludeVaults` from which vaults the recovery points can be chosen.

  `SelectionWindowDays`: The selection window is calculated from the actual job execution time, not the plan's scheduled start time. Because jobs can start anywhere within the `StartWindowHours` period, set `SelectionWindowDays` to be greater than your backup frequency interval plus your `StartWindowHours` value to avoid edge-case recovery point exclusions.

A selection can have one or more protected resource ARNs or can have one or more conditions, but it cannot have not both.

You can also include:

- `ScheduleExpressionTimezone`
- `Tags`
- `CreatorRequestId`
- `StartWindowHours`

Use CLI command `create-restore-testing-plan.`

Once the plan has been created successfully, you need to assign resources to it using `create-restore-testing-selection`.

This consists of `RestoreTestingSelectionName`, `ProtectedResourceType`, and one of the following:

> **Note**
>
> **RestoreTestingSelectionName naming requirements:**
>
> - Must be 1-256 characters in length
> - Can contain letters (a-z, A-Z), numbers (0-9), hyphens (-), and underscores (_)
> - Must start with a letter or number
> - Cannot end with a hyphen or underscore
> - Must be unique within the restore testing plan

- `ProtectedResourceArns`

- `ProtectedResourceConditions`

Each protected resource type can have one single value. A restore testing selection can include a wildcard value ("*") for `ProtectedResourceArns` along with `ProtectedResourceConditions`. Alternatively, you can include up to 30 specific protected resource ARNs in `ProtectedResourceArns`.

Restore test frequency

AWS Backup evaluates cron expressions between 00:00 and 23:59. If you create a restore testing plan for "every 12 hours" but provide a start time of later than 11:59, it will only run once per day.

Recovery point determination

Each time a restore testing plan runs, according to the frequency and start time you specified, AWS Backup first selects protected resources to test. Then, for each selected protected resource, AWS Backup restores at most one recovery point.

Protected resource selection

AWS Backup includes a protected resource when the restore testing selection specifies its resource type and either of the following is true:

- The protected resource ARN is specified in that selection.
- The tag conditions on that selection match the tags on the latest recovery point for the protected resource.

Recovery point selection

For each selected protected resource, AWS Backup restores at most one eligible recovery point. A recovery point is eligible if it falls within the specified time frame and included vaults in the restore testing plan. From the eligible recovery points, AWS Backup chooses one using the latest-or-random algorithm in `RecoveryPointSelection`. If no recovery point for a protected resource is eligible, AWS Backup does not include that protected resource in the test.

> **Tag conditions don't apply to recovery point selection**
>
> Tag conditions apply only to protected resource selection, matched against the protected resource's latest recovery point. AWS Backup does not use them to choose which recovery point to restore. As a result, the restored recovery point might not match the tag conditions used for protected resource selection.

Update a restore testing plan

You can update parts of your restore testing plan and the resource selections within it through the console or AWS CLI.

Console

Update restore testing plans and selections in the console

When you view the restore testing plan details page in the console, you can edit (update) many of the settings of your plan. To do this,

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the left-hand navigation, locate **Restore testing** and select it.
3. Select the **Edit** button.
4. Adjust the frequency, the start time, and the time the test will begin within which the test will begin after the chosen start time.
5. Save your changes.

AWS CLI

Update restore testing plans and selections through AWS CLI

Requests UpdateRestoreTestingPlan and UpdateRestoreTestingSelection can be used to send partial updates to a specified plan or selection. The names cannot be changed, but you can update other parameters. Include only parameters you wish to change in each request.

Before sending an update request, use GetRestoreTestingPlan and GetRestoreTestingSelection to determine if your RestoreTestingSelection contains specific ARNs or if it uses the wildcard and conditions.

If your restore testing selection has specified ARNs (instead of wildcard) and you wish to change it to a wildcard with conditions, the update request must include both the ARN wildcard and the conditions. A selection can have either protected resource ARNs or use the wildcard with conditions, but it cannot have both.

- `get-restore-testing-plan`
- `get-restore-testing-selection`
- `update-restore-testing-plan`
- `update-restore-testing-selection`

View existing restore testing plans

Console

View details about an existing restore testing plan and assigned resources in the console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Select **Restore testing** from the left-hand navigation. The display shows your restore testing plans. The plans are displayed by default by last runtime.
3. Select the link from a plan to see its details, including a summary of the plan, its name, frequency, start time, and start within value.

You can also view the protected resources within this plan, the restore testing jobs from the most recent 30 days included in this plan, and any tags you can created to be part of this testing plan.

AWS CLI

Get details about an existing restore testing plan and testing selection using the command line

- `list-restore-testing-plan`
- `list-restore-testing-selections`
- `get-restore-testing-plan`
- `get-restore-testing-selection`

View restore testing jobs

Console

View existing restore testing jobs in the console

Restore testing jobs are included on the restore jobs page.

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Navigate to **Jobs** page.

   Alternatively, you can select **Restore testing**, then select a restore testing plan to see its details and the jobs associated with the plan.

3. Select the **Restore jobs** tab.

   On this page, you can view the status, restore time, restore type, resource ID, resource type, restore testing plan to which the job belongs, the creation time, and the recovery point ID of the restore job.

   Jobs included in a restore testing plan have the restore type **Test**.

Restore testing jobs have several status categories:

- A **status** type that requires attention is underlined; hover over the status to see additional details if they are available.
- A **validation status** will display if Restore testing validation has been initiated on the test (not available within the console).
- Deletion status notes the status of the data generated by a restore test. There are three possible deletion statuses: **Successful**, **Deleting**, and **Failed**.

  If a restore test job deletion failed, you will need to remove the resource manually since the restore testing flow could not complete it automatically. Often, a failed deletion is triggered if the tag `awsbackup-restore-test` is removed from the resource.

AWS CLI

View existing restore testing jobs from the command line

- `list-restore-jobs-by-protected-resource`

Delete a restore testing plan

Console

Delete restore testing plan in the console

1. Go to View existing restore testing plans to see your current restore testing plans.
2. On the restore testing plan details page, delete a plan by selecting **Delete**.
3. After you select delete, a pop-up confirmation screen will appear to ensure you want to delete your plan. On this screen, the name of your specific restore testing plan will be

displayed in bold. To proceed, type in the exact case-sensitive name of the testing plan, including any underscores, dashes, and periods.

If the option for **Delete restore testing plan** is not selectable, re-enter the name until it matches the displayed name. Once it is an exact match, the option to delete the restore testing plan will become selectable.

AWS CLI

Delete restore testing plan through the command line

The CLI command DeleteRestoreTestingSelection can be used to delete a restore testing selection. Include `RestoreTestingPlanName` and `RestoreTestingSelectionName` in the request.

All testing selections associated with a testing plan need to be deleted before you delete the testing plan. Once all testing selections have been deleted, you can use the API request DeleteRestoreTestingPlan to delete a restore testing plan. You need to include `RestoreTestingPlanName`.

- `delete-restore-testing-selection`
- `delete-restore-testing-plan`

Audit restore testing

Restore testing integrations with AWS Backup Audit manager to help you evaluate if a restored resource completed within your target restore time.

For more information, see Restore time for resources meet target control in AWS Backup Audit Manager controls and remediation.

Restore testing quotas and parameters

- 100 restore testing plans
- 50 tags can be added to each restore testing plan
- 30 selections per plan
- 30 protected resource ARNs per selection

- 30 protected resource conditions per selection (including those within both `StringEquals` and `StringNotEquals`)
- 30 vault selectors per selection
- Max selection window days: 365 days
- Start window hours: Min: 1 hour; Max: 168 hours (7 days)
- Max plan name length: 50 characters
- Max selection name length: 50 characters

Additional information regarding limits can be viewed at AWS Backup quotas.

Restore testing failure troubleshooting

If you have restore testing jobs with a restore status of `Failed`, the following reasons can help you determine the cause and remedy.

Error message(s) can be viewed in the AWS Backup console in the job status details page or by using the CLI commands `list-restore-jobs-by-protected-resource` or `list-restore-jobs`.

1. ***Error:*** *`No default VPC for this user. GroupName is only supported for EC2-Classic and default VPC.`*

   Solution 1: Update your restore testing selection and override the parameter `SubnetId`. The AWS Backup console displays this parameter as "Subnet".

   Solution 2: Recreate the default VPC.

   Resource types affected: Amazon EC2

2. ***Error:*** *`No subnets found for the default VPC [vpc]. Please specify a subnet.`*

   Solution 1: Update your restore testing selection and override the `SubnetId` restore parameter. The AWS Backup console displays this parameter as "Subnet".

   Solution 2: Create a default subnet in the default VPC.

   Resource types affected: Amazon EC2

3. ***Error:*** *No default subnet detected in VPC. Please contact AWS Support to recreate default Subnets.*

   Solution 1: Update your restore testing selection and override the `DBSubnetGroupName` restore parameter. The AWS Backup console displays this parameter as Subnet group.

   Solution 2: Create a default subnet in the default VPC.

   Resource types affected: Amazon Aurora, Amazon DocumentDB, Amazon RDS, Neptune

4. ***Error:*** *IAM Role cannot be assumed by AWS Backup.*

   Solution: The restore role must be assumable by AWS Backup. Either update the role's trust policy in IAM to allow it to be assumed by `"backup.amazonaws.com"` or update your restore testing selection to use a role that is assumable by AWS Backup.

   Resource types affected: all

5. ***Error:*** *Access denied to KMS key.* or *The specified AWS KMS key* ARN *does not exist, is not enabled or you do not have permissions to access it.*

   Solution: Verify the following:

   a. The restore role has access to the AWS KMS key used to encrypt your backups and, if applicable, the KMS key used to encrypt the restored resource.

   b. The resource policies on the above KMS key(s) allow the restore role to access them.

   If the above conditions are not yet met, configure the restore role and the resource policies for appropriate access. Then, run the restore testing job again.

   Resource types affected: all

6. ***Errors:*** *User ARN is not authorized to perform action on resource because no identity based policy allows the action.* or *Access denied performing s3:CreateBucket on awsbackup-restore-test-xxxxxx.*

   Solution: The restore role does not have adequate permissions. Update the permissions in IAM for the restore role.

   Resource types affected: all

7. ***Errors:*** *User ARN is not authorized to perform action on resource because no resource-based policy allows the action.* or *User ARN is not authorized to perform action on resource with an explicit deny in a resource based policy.*

   Solution: The restore role does not have adequate access to the resource specified in the message. Update the resource policy on the resource mentioned.

   Resource types affected: all

Restore testing inferred metadata

Restoring a recovery point requires restore metadata. To perform restore tests, AWS Backup automatically infers metadata that is likely to result in a successful restore. The command `get-restore-testing-inferred-metadata` can be used to preview what AWS Backup will infer. The command `get-restore-job-metadata` returns the set of metadata inferred by AWS Backup. Note that for some resource types (Amazon FSx), AWS Backup is not able to infer a complete set of metadata.

Inferred restore metadata is determined during the restore testing process. You can override certain restore metadata keys by including the parameter `RestoreMetadataOverrides` in the body of `RestoreTestingSelection`. Some metadata overrides are not available in the AWS Backup console.

Each supported resource has both inferred restore metadata keys and values, and overridable restore metadata keys. Only `RestoreMetadataOverrides` key value pairs or nested key value pairs marked with *`required for successful restore`* are necessary to include; the others are optional. Note that key values are not case sensitive.

> **Important**
>
> AWS Backup can infer that a resource should be restored to the default setting, such as an Amazon EC2 instance or Amazon RDS cluster restored to the default VPC. However, if the default is not present, for example the default VPC or subnet has been deleted and no metadata override has been input, the restore will not be successful.

Resource type	Inferred restore metadata keys and values	Overridable metadata

DynamoDB	`deletionProtection` , where value is set to `false`
	
	`encryptionType` is set to `Default`
	
	`targetTableName` , where value is set to random value starting with `awsbackup-restore-test-`
		`encryptionType`
		
		`kmsMasterKeyArn`

Amazon EBS	`availabilityZone` , whose value is set to a random availability zone
	
	`encrypted` , whose value is set to `true`
		`availabilityZone`
		
		`iops`
		
		`kmsKeyId`
		
		`throughput`
		
		`volumesize`
		
		`volumetype`

Amazon EC2	`disableApiTermination` value is set to `false`
	
	`instanceType` value is set to the instanceType of the recovery point being restored
	
	`requiredImdsV2` value is set to `true`
		`iamInstanceProfileName` (null or `UseBackedUpValue`)
		
		`instanceType`
		
		`requireImdsV2`
		
		`securityGroupIds`
		
		`subnetId`

Amazon EFS	`encrypted` value is set to `true`	`kmsKeyId`
		
		`performanceMode`

Resource type	Inferred restore metadata keys and values	Overridable metadata

	`file-system-id` value is set to the file system ID of the recovery point being restored
	
	`kmsKeyId` value is set to `alias/aws/elasticfilesystem`
	
	`newFileSystem` value is set to `true`
	
	`performanceMode` value is set to `generalPurpose`	

Amazon FSx for Lustre	`lustreConfiguration` has nested keys. One nested key is `automaticBackupRetentionDays`, the value of which is set to `0`	`kmsKeyId`
		
		`lustreConfiguration` has nested key `logConfiguration`
		
		`securityGroupIds`
		
		`subnetIds`, *required for successful restore*

Resource type	Inferred restore metadata keys and values	Overridable metadata
Amazon FSx for NetApp ONTAP	`name` is set to a random value starting with `awsbackup_restore_test_` `ontapConfiguration` has nested keys, including: • `junctionPath` where `/name` is the name of the volume being restored • `sizeInMegabytes`, the value of which is set to the size in megabytes of the recovery point being restored • `snapshotPolicy` where the value is set to none **Note** For FSx for ONTAP resources, the Storage Virtual Machine (SVM) that you specify in `storageVirtualMachineId` determines the subnet and Availability Zone of the restored volume. If you specify the `SubnetIds` parameter in the	`ontapConfiguration` has specific overrideable nested keys, including: • `junctionPath` • `ontapVolumeType` • `securityStyle` • `sizeInMegabytes` • `storageEfficiencyEnabled` • `storageVirtualMachineId`, *required for successful restore* • `tieringPolicy`

Resource type	Inferred restore metadata keys and values	Overridable metadata
	console, AWS Backup accepts but ignores it during the restore operation. To restore to a different file system or Availability Zone, specify an SVM that belongs to the target file system in your `RestoreMetadataOverrides` . If you specify an SVM in the same file system, the volume is restored to that file system regardless of any subnet configuration that you specify in the console.	

Resource type	Inferred restore metadata keys and values	Overridable metadata

Amazon FSx for OpenZFS	`openZfzConfiguration`, which has nested keys, including:
	• `automaticBackupRetentionDays` with value set to `0`
	• `deploymentType` with value set to the deployment type of the recovery point being restored
	• `throughputCapacity`, whose value is based on the `deploymentType`. If `deploymentType` is `SINGLE_AZ_1`, the value is set to 64; if the `deploymentType` is `SINGLE_AZ_2` or `MULTI_AZ_1`, the value is set to `160`
		`kmsKeyId`
		`openZfsConfiguration` has specific overridable nested keys, including:
		• `deploymentType`
		• `throughputCapacity`
		• `diskiopsConfiguration`
		`securityGroupIds`
		`subnetIds`

Resource type	Inferred restore metadata keys and values	Overridable metadata

Amazon FSx for Windows File Server	`windowsConfiguration`, which has nested keys including:
	• `automaticBackupRetentionDays` with value set to `0`
	• `deploymentType` with value set to the deployment type of the recovery point being restored
	• `throughputCapacity` with value set to 8
		`kmsKeyId`
		`securityGroupIds`
		`subnetIds` *required for successful restore*
		`windowsConfiguration`, with specific overridable nested keys
		• `throughputCapacity`
		• `activeDirectoryId` *required for successful restore*
		• `preferredSubnetId`

Resource type	Inferred restore metadata keys and values	Overridable metadata

Amazon RDS, Aurora, Amazon DocumentDB, Amazon Neptune clusters	`availabilityZones` with value set to a list of up to three random availability zones
	
	`dbClusterIdentifier` with a random value starting with `awsbackup-restore-test`
	
	`engine` with value set to the engine of the recovery point being restored
		`availabilityZones`
		
		`databaseName`
		
		`dbClusterParameterGroupName`
		
		`dbSubnetGroupName`
		
		`enableCloudwatchLogsExports`
		
		`enableIamDatabaseAuthentication`
		
		`engine`
		
		`engineMode`
		
		`engineVersion`
		
		`kmskeyId`
		
		`port`
		
		`optionGroupName`
		
		`scalingConfiguration`
		
		`vpcSecurityGroupIds`

Resource type	Inferred restore metadata keys and values	Overridable metadata

Amazon RDS instances	`dbInstanceIdentifier` with a random value starting with `awsbackup-restore-test-`
	
	`deletionProtection` with value set to `false`
	
	`multiAz` with value set to `false`
	
	`publiclyAccessible` with value set to false
		`allocatedStorage`
		
		`availabilityZones`
		
		`dbInstanceClass`
		
		`dbName`
		
		`dbParameterGroupName`
		
		`dbSubnetGroupName`
		
		`domain`
		
		`domainIamRoleName`
		
		`enableCloudwatchLogsExports`
		
		`enableIamDatabaseAuthentication`
		
		`iops`
		
		`licensemodel`
		
		`multiAz`
		
		`optionGroupName`
		
		`port`
		
		`processorFeatures`
		
		`publiclyAccessible`
		
		`storageType`
		
		`vpcSecurityGroupIds`

Resource type	Inferred restore metadata keys and values	Overridable metadata

Amazon Redshift	`clusterIdentifier` , where value is set to a random value starting with `awsbackup-restore-test-`
	
	`encrypted` , where value is set to `true`
	
	`publiclyAccessible` , where value is set to `false`
	
	`restoreType` , where value is set to `CLUSTER_RESTORE`
		`availabilityZone`
		
		`clusterParameterGroupName`
		
		`clusterSecurityGroups`
		
		`clusterSubnetGroupName`
		
		`defaultIamRoleArn`
		
		`elasticIp`
		
		`enhancedVpcRouting`
		
		`hsmClientCertificateIdentifier`
		
		`hsmConfigurationIdentifier`
		
		`iamRoles`
		
		`kmsKeyId`
		
		`manageMasterPassword`
		
		`masterPasswordSecretKmsKeyId`
		
		`nodeType`
		
		`numberOfNodes`
		
		`port`

Resource type	Inferred restore metadata keys and values	Overridable metadata

		`reservedNodeId`
		
		`snapshotClusterIdentifier`
		
		`targetReservedNodeOfferingId`
		
		`vpcSecurityGroupIds`

Amazon Simple Storage Service (Amazon S3)	`destinationBucketName` with a random value starting with `awsbackup-restore-test-`
	
	`encrypted` with value set to `true`
	
	`encryptionType` with value set to `SSE-S3`
	
	`newBucket` with value set to `true`
	
	`restoreACLs` specifies whether to restore ACLs. The value is `false` when you created the backup with `BackupACLs` disabled. Otherwise, the value is `true`.
		`encryptionType`
		
		`kmsKey`
		
		`restoreACLs`

Format for list-typed override values

Because `RestoreMetadataOverrides` is a map of string key-value pairs, override values that represent lists (such as `vpcSecurityGroupIds`, `availabilityZones`,

`clusterSecurityGroups`, `enableCloudwatchLogsExports`, `iamRoles`, `securityGroupIds`, and `subnetIds`) must be specified as JSON-array-encoded strings.

For example, to specify a single security group:

```
"vpcSecurityGroupIds": "[\"sg-012d52c68c6e88f00\"]"
```

To specify multiple security groups:

```
"vpcSecurityGroupIds": "[\"sg-012d52c68c6e88f00\",\"sg-abcdef01234567890\"]"
```

In AWS CloudFormation templates, you can use the `!Sub` intrinsic function to reference resource attributes within the JSON-array string:

```
RestoreMetadataOverrides:
  vpcSecurityGroupIds: !Sub '["${DBSecurityGroup.GroupId}"]'
```

For multiple values in CloudFormation:

```
RestoreMetadataOverrides:
  vpcSecurityGroupIds: !Sub
 '["${DBSecurityGroup.GroupId}","${AnotherSecurityGroup.GroupId}"]'
```

This JSON-array-string format applies to all list-typed overrides in the preceding table. For more examples of this format, see the section called “RDS restore”.

Restore testing validation

You have the option of creating an event-driven validation that runs when a restore testing job completes.

First, create a validation workflow with any target supported by Amazon EventBridge, such as AWS Lambda. Second, add an EventBridge rule that listens for the restore job reaching the status `COMPLETED`. Third, create a restore testing plan (or let an existing one run as scheduled). Finally, after the restore test has finished, monitor the logs of the validation workflow to ensure it ran as expected (once validation has run, a validation status will display in the AWS Backup console).

1. **Set up validation workflow**

   You can set up a validation workflow using Lambda or any other target supported by EventBridge. For example, if you are validating a restore test containing an Amazon EC2 instance, you may include code that pings a healthcheck endpoint.

   You can use the details in the event to determine which resource(s) to validate.

   You can use Lambda layers to use the latest SDK (because `PutRestoreValidationResult` is not available through the Lambda SDK).

   Here is a sample:

```
import { Backup } from "@aws-sdk/client-backup";

export const handler = async (event) => {
  console.log("Handling event: ", event);

  const restoreTestingPlanArn = event.detail.restoreTestingPlanArn;
  const resourceType = event.detail.resourceType;
  const createdResourceArn = event.detail.createdResourceArn;

  // TODO: Validate the resource

  const backup = new Backup();
  const response = await backup.putRestoreValidationResult({
    RestoreJobId: event.detail.restoreJobId,
    ValidationStatus: "SUCCESSFUL", // TODO
    ValidationStatusMessage: "" // TODO
  });

  console.log("PutRestoreValidationResult: ", response);
  console.log("Finished");
};
```

2. **Add an EventBridge rule**

   Create an EventBridge rule that listens for the restore job `COMPLETED` event.

   Optionally, you can filter events by resource type or restore testing plan ARN. Set the target of this rule to invoke the validation workflow you defined in Step 1. Here is an example:

```
{
  "source":[
    "aws.backup"
  ],
  "detail-type":[
    "Restore Job State Change"
  ],
  "detail":{
    "resourceType":[
      "..."
    ],
    "restoreTestingPlanArn":[
      "..."
    ],
    "status":[
      "COMPLETED"
    ]
  }
}
```

3. **Let the restore testing plan run and complete**

   The restore testing plan will run according to the schedule you have configured.

   See Create a restore testing plan if you do not yet have one or Update a restore testing plan if you wish to change the settings.

4. **Monitor the results**

   Once a restore testing plan has run as scheduled, you can check the logs of your validation workflow to ensure it ran correctly.

   You can call the API `PutRestoreValidationResult` to post the results, which will then be viewable in the AWS Backup console and through AWS Backup API calls that describe and list restore jobs, such as `DescribeRestoreJob` or `ListRestoreJob`.

   Once a validation status is set, it cannot be changed.

Stop a backup job

You can stop a backup job in AWS Backup after it has been initiated. When you do this, the backup is not created, and the backup job record is retained with the status of **aborted**.

To stop a backup job using the AWS Backup console

1. Sign in to the AWS Management Console, and open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane on the left, choose **Jobs**.
3. Choose the backup job that you want to stop.
4. In the backup job details pane, choose **Stop**.

View existing backups

You can view a list of your backups using the AWS Backup console or programmatically.

Topics

- Listing backups by protected resource in the console
- Listing backups by backup vault in the console
- Listing backups programmatically

Listing backups by protected resource in the console

Follow these steps to view a list of backups of a particular resource on the AWS Backup console.

1. Sign in to the AWS Management Console, and open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Protected resources**.
3. Choose a protected resource in the list to view the list of backups. Only resources that have been backed up by AWS Backup are listed under **Protected resources**.

You can view the backups for the resource. From this view, you can also choose a backup and restore it.

Listing backups by backup vault in the console

Follow these steps to view a list of backups organized in a backup vault.

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the navigation pane, choose **Backup vaults**.
3. In the **Backups** section, view the list of all the backups organized in this backup vault. In this view, you can sort backups by any of the column headers (including status), as well as select a backup to restore it, edit it, or delete it.

Listing backups programmatically

You can list backups programmatically using the `ListRecoveryPoint` API operations:

- `ListRecoveryPointsByBackupVault`
- `ListRecoveryPointsByResource`

For example, the following AWS Command Line Interface (AWS CLI) command lists all your backups with the `EXPIRED` status:

```
aws backup list-recovery-points-by-backup-vault \
  --backup-vault-name sample-vault \
  --query 'RecoveryPoints[?Status == `EXPIRED`]'
```

AWS Backup Audit Manager

You can use AWS Backup Audit Manager to audit the compliance of your AWS Backup policies against controls that you define. A *control* is a procedure designed to audit the compliance of a backup requirement, such as the backup frequency or the backup retention period.

AWS Backup Audit Manager helps you answer questions such as:

- "Am I backing up all my resources?"
- "Are all of my backups encrypted?"
- "Are my backups taking place daily?"

You can use AWS Backup Audit Manager to find backup activity and resources that are not yet compliant with the controls that you defined. Note that only active resources will be included when controls evaluate resources for compliance. For example, an Amazon EC2 instance in a running state will be evaluated. An EC2 instance in a stopped state will not be included in the compliance evaluation.

You can also use it to automatically generate an audit trail of daily and on-demand reports for your backup governance purposes.

The following steps provide an overview of how to use AWS Backup Audit Manager. For detailed walkthroughs, choose one of the topics at the end of this page.

1. Create frameworks that contain one or more governance control templates. The preceding questions are examples of three governance control templates. You can customize the parameters of some governance control templates. For example, you can customize the last control to ask, “Are my backups taking place weekly?” instead of daily.
2. View your framework to see how many of your resources are compliant (or non-compliant) with the controls you defined in that framework.
3. Create reports of your backup and compliance status. Store these reports as demonstrable evidence of your compliance practices, or to identify individual backup activities and resources that are not yet in compliance.

   AWS Backup Audit Manager automatically generates a new report for you every 24 hours and publishes it to Amazon S3. You can also generate on-demand reports.

> **Note**
>
> Before you create your first compliance-related framework, you must turn on resource tracking. Doing so allows AWS Config to track your AWS Backup resources. For technical documentation about how to manage resource tracking, see Setting up AWS Config with the console in the *AWS Config Developer Guide*.
> Charges apply when you turn on resource tracking. For information about resource tracking pricing and billing for AWS Backup Audit Manager, see Metering, costs, and billing.

Topics

- Working with audit frameworks
- Working with audit reports
- Using AWS Backup Audit Manager with CloudFormation
- Using AWS Backup Audit Manager with AWS Audit Manager
- Controls and remediation

Working with audit frameworks

A *framework* is a collection of controls that helps you to evaluate your backup practices. You can use pre-built, customizable controls to define your policies and evaluate whether your backup practices comply with your policies. You can also set up automatic daily reports to gain insights into the compliance status of your frameworks.

Each framework applies to a single account and AWS Region. You can deploy a maximum of 15 frameworks per account per Region. You cannot deploy duplicate frameworks (frameworks that contain the same controls and parameters).

There are two different types of frameworks:

- The **AWS Backup framework** (recommended) – Use the AWS Backup framework to deploy all available controls to monitor your backup activity, coverage, and resources against the best practices that we recommend.
- A **custom framework** that you define – Use a custom framework to choose one or more specific controls and to customize control parameters.

Topics

- Choosing your controls
- Turning on resource tracking
- Creating frameworks using the AWS Backup console
- Creating frameworks using the AWS Backup API
- Viewing framework compliance status
- Finding non-compliant resources
- Updating audit frameworks
- Deleting audit frameworks

Choosing your controls

The following table lists the AWS Backup Audit Manager controls, their customizable parameters, and their AWS Config recording resource types.

Available controls

Control name	Control description	Customizable parameters	AWS Config recording resource type
Backup resources are included in at least one backup plan	Evaluates if resources are included in at least one backup plan.	None	`AWS Backup: backup selection`
Backup plan has minimum frequency and minimum retention	Evaluates if backup frequency is at least [1 day] and retention period is at least [35 days].	Backup frequency; retention period	`AWS Backup: backup plans`
Vaults prevent manual deletion of recovery points	Evaluates if backup vaults do not allow manual deletion of recovery points	Up to 5 IAM roles that allow manual deletion of recovery points	`AWS Backup: backup vaults`

Control name	Control description	Customizable parameters	AWS Config recording resource type
	except by certain AWS Identity and Access Managemen t (IAM) roles. By default, there are no IAM role exception s. There are also no IAM role exceptions when you deploy this control with the AWS Backup framework.		
Recovery points are encrypted	Evaluates if the recovery points are encrypted.	None	`AWS Backup: recovery points`
Minimum retention established for recovery point	Evaluates if the recovery point retention period is at least [35 days].	Recovery point retention period	`AWS Backup: recovery points`
Cross-Region backup copy is scheduled	Evaluates if a resource is configured to create copies of its backups to another AWS Region.	AWS Region	`AWS Backup: backup selection`
Cross-account backup copy is scheduled	Evaluates if a resource has a cross-account backup copy configured.	AWS account ID	`AWS Backup: backup selection`

Control name	Control description	Customizable parameters	AWS Config recording resource type
Resources are in a backup plan with an AWS Backup Vault Lock	Evaluates if a resource has a backup plan configured to store backups in a locked backup vault.	Min Retention Days; Max Retention Days	`AWS Backup: backup selection`
Last recovery point was created	Evaluates if a recovery point was created within specified time frame.	Value in hours [1 to 744] or days [1 to 31].	`AWS Backup recovery points`
Restore time for resources meet target	Evaluates if restore testing job completed within target restore time	Value in minutes	None
Resources are inside a logically air-gapped vault	Evaluates if resources have at least one recovery point copied to a logically air-gapped vault within the specified value and timeframe.	Value in minutes, hours, or days	`AWS Backup: recovery points`

For detailed information about these controls, see Controls and remediation.

For a list of AWS Backup-supported resources that don't support all controls, see the AWS Backup Audit Manager section of the Feature availability by resource table.

> **Note**
>
> If you don't want to use any of the preceding controls, you can still use AWS Backup Audit Manager to create daily reports of your backup, copy, and restore jobs. See Working with audit reports.

Turning on resource tracking

Before you create your first compliance-related framework, you must turn on resource tracking. Doing so allows AWS Config to track your AWS Backup resources. For technical documentation about how to manage resource tracking, see Setting up AWS Config with the console in the *AWS Config Developer Guide*.

Charges apply when you turn on resource tracking. For information about resource tracking pricing and billing for AWS Backup Audit Manager, see Metering, costs, and billing.

Topics

- Turning on resource tracking using the console
- Turning on resource tracking using the AWS Command Line Interface (AWS CLI)
- Turning on resource tracking using a CloudFormation template

Turning on resource tracking using the console

To turn on resource tracking using the console:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the left navigation pane, under **Audit Manager**, choose **Frameworks**.
3. Turn on resource tracking by choosing **Manage resource tracking**.
4. Choose **Go to AWS Config Settings**.
5. Choose **Enable or disable recording**.
6. Choose **Enable** recording for all of the following resource types, or choose to enable recording for some resource types. Refer to AWS Backup Audit Manager controls and remediation for which resource types are required for your controls.

    - `AWS Backup: backup plans`

- `AWS Backup: backup vaults`
- `AWS Backup: recovery points`
- `AWS Backup: backup selection`

7. Choose **Close**.
8. Wait for the blue banner with the text **Turning on resource tracking** to transition to the green banner with the text **Resource tracking is on**.

You can check whether you have turned on resource tracking and, if so, which resource types you are recording, in two places in the AWS Backup console. In the left navigation pane, either:

- Choose **Frameworks**, then choose the text under **AWS Config recorder status**.
- Choose **Settings**, then choose the text under **AWS Config recorder status**.

Turning on resource tracking using the AWS Command Line Interface (AWS CLI)

If you have not yet onboarded to AWS Config, it might be faster to onboard using the AWS CLI.

To turn on resource tracking using the AWS CLI:

1. Type the following command to determine if you already enabled your AWS Config recorder.

```
$ aws configservice describe-configuration-recorders
```

   a. If your `ConfigurationRecorders` list is empty like this:

```
{
  "ConfigurationRecorders": []
}
```

      Your recorder is not enabled. Continue to step 2 to create your recorder.

   b. If you already enabled recording for all resources, your `ConfigurationRecorders` output will look like this:

```
{
  "ConfigurationRecorders":[
    {
      "recordingGroup":{
```

```
        "allSupported":true,
        "resourceTypes":[

        ],
        "includeGlobalResourceTypes":true
      },
      "roleARN":"arn:aws:iam::[account]:role/[roleName]",
      "name":"default"
    }
  ]
}
```

Because you enabled all resources you already turned on resource tracking. You do not need to complete the rest of this procedure to use AWS Backup Audit Manager.

2. Create a AWS Config recorder with the AWS Backup Audit Manager resource types

```
$ aws configservice put-configuration-recorder --configuration-recorder
 name=default, \
roleARN=arn:aws:iam::accountId:role/aws-service-role/config.amazonaws.com/
AWSServiceRoleForConfig \
--recording-group
 resourceTypes="['AWS::Backup::BackupPlan','AWS::Backup::BackupSelection', \
'AWS::Backup::BackupVault','AWS::Backup::RecoveryPoint']"
```

3. Describe your AWS Config recorder.

```
$ aws configservice describe-configuration-recorders
```

Verify that it has the AWS Backup Audit Manager resource types by comparing your output with the following expected output.

```
{
  "ConfigurationRecorders":[
    {
      "name":"default",
      "roleARN":"arn:aws:iam::accountId:role/AWSServiceRoleForConfig",
      "recordingGroup":{
        "allSupported":false,
        "includeGlobalResourceTypes":false,
        "resourceTypes":[
          "AWS::Backup::BackupPlan",
```

```
          "AWS::Backup::BackupSelection",
          "AWS::Backup::BackupVault",
          "AWS::Backup::RecoveryPoint"
        ]
      }
    }
  ]
}
```

4. Create an Amazon S3 bucket as the destination to store the AWS Config configuration files.

```
$ aws s3api create-bucket --bucket amzn-s3-demo-bucket –region us-east-1
```

5. Use *policy.json* to grant AWS Config permission to access your bucket. See the following sample *policy.json*.

```
$ aws s3api put-bucket-policy --bucket amzn-s3-demo-bucket --policy file://
policy.json
```

JSON

```
{
  "Version":"2012-10-17",
  "Statement":[
    {
      "Sid":"AWSConfigBucketPermissionsCheck",
      "Effect":"Allow",
      "Principal":{
        "Service":"config.amazonaws.com"
      },
      "Action":"s3:GetBucketAcl",
      "Resource":"arn:aws:s3:::amzn-s3-demo-bucket",
      "Condition":{
        "StringEquals":{
          "aws:SourceAccount":"111122223333"
        }
      }
    },
    {
      "Sid":"AWSConfigBucketExistenceCheck",
      "Effect":"Allow",
      "Principal":{
```

```
        "Service":"config.amazonaws.com"
      },
      "Action":"s3:ListBucket",
      "Resource":"arn:aws:s3:::amzn-s3-demo-bucket",
      "Condition":{
        "StringEquals":{
          "aws:SourceAccount":"111122223333"
        }
      }
    },
    {
      "Sid":"AWSConfigBucketDelivery",
      "Effect":"Allow",
      "Principal":{
        "Service":"config.amazonaws.com"
      },
      "Action":"s3:PutObject",
      "Resource":"arn:aws:s3:::amzn-s3-demo-bucket/*",
      "Condition":{
        "StringEquals":{
          "aws:SourceAccount":"111122223333"
        }
      }
    }
  ]
}
```

6. Configure your bucket as an AWS Config delivery channel

```
$ aws configservice put-delivery-channel --delivery-channel
 name=default,s3BucketName=amzn-s3-demo-bucket
```

7. Enable AWS Config recording

```
$ aws configservice start-configuration-recorder --configuration-recorder-
name default
```

8. Verify that `"FrameworkStatus":"ACTIVE"` in the last line of your `DescribeFramework` output as follows.

```
$ aws backup describe-framework --framework-name test --region us-east-1
```

```
{
  "FrameworkName":"test",
 "FrameworkArn":"arn:aws:backup:us-east-1:accountId:framework:test-
f0001b0a-0000-1111-ad3d-4444f5cc6666",
  "FrameworkDescription":"",
  "FrameworkControls":[
    {
      "ControlName":"BACKUP_RECOVERY_POINT_MINIMUM_RETENTION_CHECK",
      "ControlInputParameters":[
        {
          "ParameterName":"requiredRetentionDays",
          "ParameterValue":"1"
        }
      ],
      "ControlScope":{

      }
    },
    {
      "ControlName":"BACKUP_PLAN_MIN_FREQUENCY_AND_MIN_RETENTION_CHECK",
      "ControlInputParameters":[
        {
          "ParameterName":"requiredFrequencyUnit",
          "ParameterValue":"hours"
        },
        {
          "ParameterName":"requiredRetentionDays",
          "ParameterValue":"35"
        },
        {
          "ParameterName":"requiredFrequencyValue",
          "ParameterValue":"1"
        }
      ],
      "ControlScope":{

      }
    },
    {
      "ControlName":"BACKUP_RESOURCES_PROTECTED_BY_BACKUP_PLAN",
      "ControlInputParameters":[

      ],
```

```
      "ControlScope":{

      }
    },
    {
      "ControlName":"BACKUP_RECOVERY_POINT_ENCRYPTED",
      "ControlInputParameters":[

      ],
      "ControlScope":{

      }
    },
    {
      "ControlName":"BACKUP_RECOVERY_POINT_MANUAL_DELETION_DISABLED",
      "ControlInputParameters":[

      ],
      "ControlScope":{

      }
    }
  ],
  "CreationTime":1633463605.233,
  "DeploymentStatus":"COMPLETED",
  "FrameworkStatus":"ACTIVE"
}
```

Turning on resource tracking using a CloudFormation template

For a CloudFormation template that turns on resource tracking, see Using AWS Backup Audit Manager with CloudFormation.

Creating frameworks using the AWS Backup console

After turning on resource tracking, create a framework using the following steps.

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the left navigation pane, choose **Frameworks**.
3. Choose **Create Framework**.

4. For **Framework name**, enter a unique name. The framework name must be between 1 and 256 characters, starting with a letter, and consisting of letters (a-z, A-Z), numbers (0-9), and underscores (_).
5. (Optional) Enter a **Framework description**.
6. In **Controls**, your active controls will be displayed. By default, all controls eligible for a resource are listed.

   To change which controls are active, click **Edit controls**.

   a. The first check box indicates if the control is turned on. To turn off a control, uncheck the box.

   b. Under **Choose resources to evaluate**, you can select how to choose resources, either by type, by tags, or by a single resource.

   The list of AWS Backup Audit Manager controls describes the customization options for each control.
7. (Optional) Tag your framework by choosing **Add new tag**. You can use tags to search and filter your frameworks or track your costs.
8. Choose **Create framework**.

AWS Backup Audit Manager might take several minutes to create the framework.

If the error `AlreadyExists` occurs, a framework with the same controls and parameters already exists. To successfully create a new framework, at least one control or parameter must be different from existing frameworks.

Creating frameworks using the AWS Backup API

The following table contains sample API requests to CreateFramework for each control, along with sample API responses to the corresponding DescribeFramework requests. To work with AWS Backup Audit Manager programmatically, you can refer to these code snippets.

Control	CreateFramework request	DescribeFramework response
Backup resources are included in at least one backup plan	{"FrameworkName": "Control1", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "BACKUP_RESOURCES_PROTECTED_BY_BACKUP_PLAN", "ControlInputParameters":[], "ControlScope": {"ComplianceResourceTypes": ["RDS"] // Evaluate only RDS instances } }], "IdempotencyToken": "Control1", "FrameworkTags": {"key1": "foo"} }	{"FrameworkName": "Control1", "FrameworkArn": "arn:aws:backup:us-east-1:123456789012:framework/Control1-ce7655ae-1e31-45cb-96a0-4f43d8c19642", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "BACKUP_RESOURCES_PROTECTED_BY_BACKUP_PLAN", "ControlInputParameters":[], "ControlScope": {"ComplianceResourceTypes": ["RDS"] } }], "CreationTime": 1516925490, "DeploymentStatus": "Active", "FrameworkStatus": "Completed", "IdempotencyToken": "Control1", "FrameworkTags": {"key1": "foo"} }

Control	CreateFramework request	DescribeFramework response
Backup plan minimum frequency and minimum retention	{"FrameworkName": "Control2", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "BACKUP_PLAN_MIN_FREQUENCY_AND_MIN_RETENTION_CHECK", "ControlInputParameters": [{"ParameterName": "requiredRetentionDays", "ParameterValue": "35"}, {"ParameterName": "requiredFrequencyUnit", "ParameterValue": "hours"}, {"ParameterName": "requiredFrequencyValue", "ParameterValue": "24"}], "ControlScope": { "Tags": {"key1": "prod"} // Evaluate backup plans that tagged with "key1": "prod". } }],	{"FrameworkName": "Control2", "FrameworkArn": "arn:aws:backup:us-east-1:123456789012:framework/Control2-de7655ae-1e31-45cb-96a0-4f43d8c1969d", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "BACKUP_PLAN_MIN_FREQUENCY_AND_MIN_RETENTION_CHECK", "ControlInputParameters": [{"ParameterName": "requiredRetentionDays", "ParameterValue": "35"}, {"ParameterName": "requiredFrequencyUnit", "ParameterValue": "hours"}, {"ParameterName": "requiredFrequencyValue", "ParameterValue": "24"}], "ControlScope": {

Control	CreateFramework request	DescribeFramework response
	"IdempotencyToken": "Control2", "FrameworkTags": {"key1": "foo"} }	"Tags": {"key1": "prod"} } }], "CreationTime": 1516925490, "DeploymentStatus": "Active", "FrameworkStatus": "Completed", "IdempotencyToken": "Control2", "FrameworkTags": {"key1": "foo"} }

Control	CreateFramework request	DescribeFramework response
Vaults prevent manual deletion of recovery points	{"FrameworkName": "Control3", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "BACKUP_RECOVERY_POINT_MANUAL_DELETION_DISABLED", "ControlInputParameters": [{"ParameterName": "principalArnList", "ParameterValue": "arn:aws:iam::123456789012:role/application_abc/component_xyz/RDSAccess, arn:aws:iam::123456789012:role/aws-service-role/access-analyzer.amazonaws.com/AWSServiceRoleForAccessAnalyzer, arn:aws:iam::123456789012:role/service-role/QuickSightAction"}], "ControlScope": {"ComplianceResourceIds":["default"],	{"FrameworkName": "Control3", "FrameworkArn": "arn:aws:backup:us-east-1:123456789012:framework/Control2-de7655ae-1e31-45cb-96a0-4f43d8c1969d", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "BACKUP_RECOVERY_POINT_MANUAL_DELETION_DISABLED", "ControlInputParameters": [{"ParameterName": "principalArnList", "ParameterValue": "arn:aws:iam::123456789012:role/application_abc/component_xyz/RDSAccess, arn:aws:iam::123456789012:role/aws-service-role/access-analyzer.amazonaws.com/AWSServiceRoleForAccessAnalyzer, arn:aws:iam::123456789012:r

Control	CreateFramework request	DescribeFramework response
	"Complian ceResourceTypes": ["AWS::Backup::Bac kupVault"] } }], "IdempotencyToken": "Control3", "FrameworkTags": {"key1": "foo"} }	ole/service-role/Q uickSightAction"}], "ControlScope": {"Complia nceResourceIds":[" default"], "Complian ceResourceTypes": ["AWS::Backup::Bac kupVault"] } }], "CreationTime": 1516925490, "DeploymentStatus": "Active", "FrameworkStatus": "Completed", "IdempotencyToken": "Control3", "FrameworkTags": {"key1": "foo"} }

Control	CreateFramework request	DescribeFramework response
Minimum retention established for recovery point	{"FrameworkName": "Control4", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "BACKUP_RECOVERY_POINT_MINIMUM_RETENTION_CHECK", "ControlInputParameters": [{"ParameterName": "requiredRetentionDays", "ParameterValue": "35"}], "ControlScope": {} // Default scope (no scope input) sets scope to all recovery points. }], "IdempotencyToken": "Control4", "FrameworkTags": {"key1": "foo"} }	{"FrameworkName": "Control4", "FrameworkArn": "arn:aws:backup:us-east-1:123456789012:framework/Control6-6e7655ae-1e31-45cb-96a0-4f43d8c19642", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "BACKUP_RECOVERY_POINT_MINIMUM_RETENTION_CHECK", "ControlInputParameters": [{"ParameterName": "requiredRetentionDays", "ParameterValue": "35"}], "ControlScope": {} }], "CreationTime": 1516925490, "DeploymentStatus": "Active", "FrameworkStatus": "Completed", "IdempotencyToken": "Control4", "FrameworkTags":

Control	CreateFramework request	DescribeFramework response
		<pre> {"key1": "foo"} }</pre>

Backup recovery points are encrypted	<pre>{"FrameworkName":
	 "Control5",
	 "FrameworkDescript
	ion": "This is a test
	 framework",
	 "FrameworkControls":
	  [
	  {"ControlName":
	 "BACKUP_RECOVERY_P
	OINT_ENCRYPTED",
	  "ControlInputParam
	eters":
	  [],
	  "ControlScope":
	 {} // Default scope
	 (no scope input) is all
	 recovery points
	  }
	  ],
	 "IdempotencyToken":
	 "Control5",
	 "FrameworkTags":
	  {"key1": "foo"}
	}</pre>
		<pre>{"FrameworkName":
		 "Control5",
		"FrameworkArn":
		 "arn:aws:backup:us
		-east-1:1234567890
		12:framework/Contr
		ol7-7e7655ae-1e31-
		45cb-96a0-4f43d8c1
		9642",
		 "FrameworkDescript
		ion": "This is a test
		 framework",
		  "FrameworkControls
		":
		  [
		  {"ControlName":
		 "BACKUP_RECOVERY_P
		OINT_ENCRYPTED",
		  "ControlInputParam
		eters":
		  [],
		  "ControlScope": {}
		  }
		  ],
		 "CreationTime":
		 1516925490,
		 "DeploymentStatus":
		 "Active",
		 "FrameworkStatus":
		 "Completed",
		 "IdempotencyToken":
		 "Control5",
		 "FrameworkTags":
		  {"key1": "foo"}
		}</pre>

Control	CreateFramework request	DescribeFramework response
Cross-Region backup copy is scheduled	{"FrameworkName": "Control6", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "BACKUP_RESOURCES_PROTECTED_BY_CROSS_REGION", "ControlInputParameters":[], "ControlScope": {"ComplianceResourceTypes": ["EC2"] // Evaluate only EC2 instances } }], "IdempotencyToken": "Control6", "FrameworkTags": {"key1": "foo"} }	{"FrameworkName": "Control6", "FrameworkArn": "arn:aws:backup:us-east-1:123456789012:framework/Control6-ce7655ae-1e31-45cb-96a0-4f43d8c19642", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "BACKUP_RESOURCES_PROTECTED_BY_CROSS_REGION", "ControlInputParameters":[], "ControlScope": {"ComplianceResourceTypes": ["EC2"] } }], "CreationTime": 1516925490, "DeploymentStatus": "Active", "FrameworkStatus": "Completed", "IdempotencyToken": "Control6", "FrameworkTags": {"key1": "foo"} }

Control	CreateFramework request	DescribeFramework response
Cross-account backup copy is scheduled	{"FrameworkName": "Control7", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "BACKUP_RESOURCES_PROTECTED_BY_CROSS_ACCOUNT", "ControlInputParameters":[], "ControlScope": {"ComplianceResourceTypes": ["EC2"] // Evaluate only EC2 instances } }], "IdempotencyToken": "Control7", "FrameworkTags": {"key1": "foo"} }	{"FrameworkName": "Control7", "FrameworkArn": "arn:aws:backup:us-east-1:123456789012:framework/Control7-ce7655ae-1e31-45cb-96a0-4f43d8c19642", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "BACKUP_RESOURCES_PROTECTED_BY_CROSS_ACCOUNT", "ControlInputParameters":[], "ControlScope": {"ComplianceResourceTypes": ["EC2"] } }], "CreationTime": 1516925490, "DeploymentStatus": "Active", "FrameworkStatus": "Completed", "IdempotencyToken": "Control7", "FrameworkTags": {"key1": "foo"} }

Control	CreateFramework request	DescribeFramework response
Resources are in a backup plan with an AWS Backup Vault Lock	{"FrameworkName": "Control8", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "BACKUP_RESOURCES_PROTECTED_BY_BACKUP_VAULT_LOCK", "ControlInputParameters":[], "ControlScope": {"ComplianceResourceTypes": ["EC2"] // Evaluate only EC2 instances } }], "IdempotencyToken": "Control8", "FrameworkTags": {"key1": "foo"} }	{"FrameworkName": "Control8", "FrameworkArn": "arn:aws:backup:us-east-1:123456789012:framework/Control8-ce7655ae-1e31-45cb-96a0-4f43d8c19642", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "BACKUP_RESOURCES_PROTECTED_BY_BACKUP_VAULT_LOCK", "ControlInputParameters":[], "ControlScope": {"ComplianceResourceTypes": ["EC2"] } }], "CreationTime": 1516925490, "DeploymentStatus": "Active", "FrameworkStatus": "Completed", "IdempotencyToken": "Control8", "FrameworkTags": {"key1": "foo"} }

Control	CreateFramework request	DescribeFramework response
Last recovery point was created	{"FrameworkName": "Control9", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "BACKUP_LAST_RECOVERY_POINT_CREATED", "ControlInputParameters":[], "ControlScope": {"ComplianceResourceTypes": ["EC2"] // Evaluate only EC2 instances } }], "IdempotencyToken": "Control9", "FrameworkTags": {"key1": "foo"} }	{"FrameworkName": "Control9", "FrameworkArn": "arn:aws:backup:us-east-1:123456789012:framework/Control9-ce7655ae-1e31-45cb-96a0-4f43d8c19642", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "BACKUP_LAST_RECOVERY_POINT_CREATED", "ControlInputParameters":[], "ControlScope": {"ComplianceResourceTypes": ["EC2"] } }], "CreationTime": 1516925490, "DeploymentStatus": "Active", "FrameworkStatus": "Completed", "IdempotencyToken": "Control9", "FrameworkTags": {"key1": "foo"} }

Control	CreateFramework request	DescribeFramework response
Restore time for resources meet target	{"FrameworkName":"Control10", "FrameworkDescription":"This is a test framework", "FrameworkControls":[{ "ControlName":"RESTORE_TIME_FOR_RESOURCES_MEET_TARGET", "ControlInputParameters":[{ "ParameterName":"maxRestoreTime", "ParameterValue":"720" }], "ControlScope":{ "ComplianceResourceIds":[ ], "ComplianceResourceTypes":["DynamoDB" // Evaluates only DynamoDB databases] } }]"IdempotencyToken":"Control10", "FrameworkTags":{ "key1":"foo" }	{"FrameworkName": "Control10", "FrameworkArn": "arn:aws:backup:us-east-1:123456789012:framework/Control9-ce7655ae-1e31-45cb-96a0-4f43d8c19642", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "RESTORE_TIME_FOR_RESOURCES_MEET_TARGET", "ControlInputParameters":[], "ControlScope": {"ComplianceResourceTypes": ["EC2"] } }], "CreationTime": 1516925490, "DeploymentStatus": "Active", "FrameworkStatus": "Completed", "IdempotencyToken": "Control10", "FrameworkTags": {"key1": "foo"} }

Control	CreateFramework request	DescribeFramework response
	}	

Control	CreateFramework request	DescribeFramework response
RESOURCES_IN_LOGICALLY_AIR_GAPPED_VAULT	{"FrameworkName":"Control11", "FrameworkDescription":"This is a test framework", "FrameworkControls":[{ "ControlName":"RESOURCES_IN_LOGICALLY_AIR_GAPPED_VAULT", "ControlInputParameters":[{ "ParameterName":"recoveryPointAgeValue", "ParameterValue":"10" } { "ParameterName":"recoveryPointAgeUnit", "ParameterValue":"days" }], "ControlScope":{ "ComplianceResourceTypes":["EC2"] } }]"IdempotencyToken":"Control11", "FrameworkTags":{	{"FrameworkName": "Control11", "FrameworkArn": "arn:aws:backup:us-east-1:123456789012:framework/Control11-ab1234cd-5e67-89fg-06a0-4f43d8c19642", "FrameworkDescription": "This is a test framework", "FrameworkControls": [{"ControlName": "", "ControlInputParameters":[], "ControlScope": {"ComplianceResourceTypes": ["EC2","EBS"] } }], "CreationTime": 1726087776.316, "DeploymentStatus": "COMPLETED", "FrameworkStatus": "ACTIVE", "IdempotencyToken": "Control11", "FrameworkTags": {"key1": "foo"} }

Control	CreateFramework request	DescribeFramework response

	<pre>"key1":"foo"
	}
	}</pre>	

Viewing framework compliance status

Once you create an audit framework, it appears in your **Frameworks** table. You can view this table by choosing **Frameworks** in the left navigation pane of the AWS Backup console. To view the audit results for your framework, choose its **Framework name**. Doing so takes you to the **Framework detail** page, which has two sections: **Summary** and **Controls**.

The **Summary** section lists the following statuses from left to right:

- **Compliance status** is your audit framework's overall compliance status as determined by the compliance status of each of its controls. Each control's compliance status is determined by the compliance status of each resource it evaluates.

  Framework compliance status is `Compliant` only if all resources in the scope of your control evaluations have passed those evaluations. If one or more resources failed a control evaluation, the compliance status will be `Non-Compliant`. For information on how to find your non-compliant resources, see Finding non-compliant resources. For information on how to bring your resources into compliance, see the remediation section of AWS Backup Audit Manager controls and remediation.
- **Framework status** refers to whether you have turned on resource tracking for all of your resources. The possible statuses are:
  - `Active` when recording is turned on for all resources the framework evaluates.
  - `Partially active` when recording is turned off for at least one resource the framework evaluates.
  - `Inactive` when recording is turned off for all resources that the framework evaluates.
  - `Unavailable` when AWS Backup Audit Manager is unable to validate recording status at this time.

To correct a `Partially active` or `Inactive` status

1. Choose **Frameworks** from the left navigation pane.
2. Choose **Manage resource tracking**.
3. Follow the instructions in the pop-up to enable recording that were previously not enabled for your resource types.

For more information about which resource types require resource tracking based on the controls you included in your frameworks, see the resource component of AWS Backup Audit Manager controls and remediation.

- **Deployment status** refers to your framework's deployment status. This status should most often be `Completed`, but can also be `Create in progress`, `Update in progress`, `Delete in progress`, and `Failed`.
  - A status of `Failed` means the framework didn't deploy correctly. Delete the framework, then recreate the framework through the AWS Backup console or through AWS Backup API.
- **Compliant controls** show a count of framework controls with all evaluations passing.
- **Non-compliant controls** show a count of framework controls with at least one evaluation not passing.

The **Controls** section shows you the following information:

- **Control status** refers to each control's compliance status. A control can be `Compliant`, meaning all resources pass that evaluation; `Non-compliant`, meaning that at least one resource did not pass that evaluation, or `Insufficient data`, meaning the control found no resources within the evaluation scope to evaluate.
- **Evaluation scope** might limit each control to one or more **Resource types**, one **Resource ID**, or one **Tag key** and **Tag value**, based on how you customized your control when creating your audit framework. If all fields are empty (as shown by a dash, "-"), then the control evaluates all applicable resources.

Finding non-compliant resources

AWS Backup Audit Manager helps you find which resources are non-compliant in two ways.

- When Viewing framework compliance status, choose the control name in the **Details section**. Doing so takes you to the AWS Config console, where you can view a list of your of your `Non-Compliant` resources.
- After you Create a report plan with the resource compliance template that includes your framework, you can View your report to identify all your `Non-Compliant` resources across all your controls.

  Furthermore, your `Resource compliance report` shows the last time AWS Backup Audit Manager last evaluated each of your controls.

Updating audit frameworks

You can update the description, controls, and parameters of an existing audit framework.

To update an existing framework

1. In the AWS Backup console left navigation pane, choose **Frameworks**.
2. Choose the framework you want to edit by its **Framework name**.
3. Choose **Edit**.

Deleting audit frameworks

To delete an existing framework

1. In the AWS Backup console left navigation pane, choose **Frameworks**.
2. Choose the framework you want to delete by its **Framework name**.
3. Choose **Delete**.
4. Type the name of your framework and choose **Delete framework**.

Working with audit reports

AWS Backup Audit Manager reports are automatically generated evidence of your AWS Backup activity, such as:

- Which backup jobs finished and when
- Which resources you backed up

There are two types of reports. When you create a report, you choose which type is created.

One type is a **jobs report**, which shows jobs finished in the last 24 hours and all active jobs with comprehensive context about vault properties, backup plan configurations, and lifecycle settings. Jobs reports do not display a status of `completed with issues`. To find this status, you can filter for `Completed` jobs with one or more status messages. AWS Backup will only include a status message as part of a `Completed` job's status if the message requires attention or action.

The second type of report is a **compliance report**. Compliance reports can monitor resource levels or the different controls that are in effect.

AWS Backup Audit Manager delivers a daily report in to your Amazon S3 bucket. If the report is for the current region and current account, you can choose to receive the report in either CSV or JSON format. Otherwise, the report is available in CSV format. The timing of the daily report might fluctuate over several hours because AWS Backup Audit Manager performs randomization to maintain its performance. You can also run an on-demand report anytime.

All account holders can create cross-Region reports; management and delegated administrator account holders can also create cross-account reports.

> **Tip**
>
> To ensure reports generated by delegated administrator accounts show all member account data, create frameworks in each of those member accounts.

> **Important**
>
> AWS Backup Audit Manager delivers reports only to Amazon S3 buckets that belong to the account in which the report plan or on-demand report is created. Cross-account Amazon S3 bucket delivery is not supported.

Enhanced job report context

AWS Backup Audit Manager job reports now include expanded context to help customers better understand backup operations, especially for delegated administrator accounts monitoring across organizations.

The enhanced reports now include:

- *Vault information*: Vault type, lock status, and encryption details
- *Backup plan context*: Plan names, rule names, schedules, and timezones
- *Lifecycle settings*: Retention periods and cold storage transition settings
- *Resource details*: Resource names and enhanced job metadata

This expanded context eliminates the need for additional API calls to DescribeBackupVault, DescribeBackupPlan, and DescribeRecoveryPoint, providing comprehensive job information in a single report.

> **Note**
>
> The expanded context provides information that previously required separate API calls with additional permissions. Ensure that users with access to these reports have appropriate permissions for the enhanced information being provided.

You can have a maximum of 20 report plans per AWS account.

> **Note**
>
> Resources such as RDS that do not have the capability to show incremental bytes of data of a specific backup will display the value `backupSizeInBytes` as 0.

> **Note**
>
> Restore testing validation and deletion status after your restore jobs report has been generated, will be reported in subsequent restore jobs report on completion with the same restore job id.

To allow AWS Backup Audit Manager to create daily or on-demand reports, you must first create a *report plan* from a *report template*.

Topics

- Choosing your report template
- Creating report plans using the AWS Backup console
- Creating report plans using the AWS Backup API
- Creating on-demand reports
- Viewing audit reports
- Updating report plans
- Deleting report plans

Choosing your report template

A report template defines the information that your report plan includes in your report. When you automate your reports using a *report plan*, AWS Backup Audit Manager provides you reports for the previous 24 hours. AWS Backup Audit Manager creates these reports between the hours of 1 and 5 AM UTC. It offers the following report templates.

Backup report templates

Backup report templates. These templates give you daily updates on your backup, restore, or copy jobs with expanded context including vault properties, backup plan details, and lifecycle settings. You can use these reports to monitor your operational posture, understand backup configurations, and identify any failures that might need further action. The following table lists each backup report template name and its sample output.

Backup report template	Sample report in JSON format
BACKUP_JOB_REPORT	{ "reportItems": [{ "reportTimePeriodStart": "2021-07-14T00:00:00Z", "reportTimePeriodEnd": "2021-07-15T00:00:00Z", "accountId": "112233445566", "region": "us-west-2", "backupJobId": "FCCB040A -9426-2A49-2EA9-5EAFFAC656AC", "jobStatus": "COMPLETED",

Backup report template	Sample report in JSON format


```
      "resourceType": "EC2",
      "resourceArn": "arn:aws:ec2:us-
west-2:112233445566:instance/
i-0bc877aee7782ba75",
      "resourceName": "MyEC2Ins
tance",
      "backupPlanArn": "arn:aws:
backup:us-west-2:1122334455
66:backup-plan:349f2247-b48
9-4301-83ac-4b7dd724db9a",
      "backupRuleId": "ab88bbf8-
ff4e-4f1b-92e7-e13d3e65dcfb",
      "initiationDate": "2021-07-
14T23:53:47.229Z",
      "backupPlanName": "MyBackup
Plan",
      "backupRuleName": "DailyBac
kups",
      "backupRuleSchedule": "cron(0
 5 ? * * *)",
      "backupRuleTimezone": "UTC",
      "startWindowEnd": "2021-07-
14T23:53:47.229Z",
      "backupOptions": {},
      "isParentJob": false,
      "parentJobId": null,
      "creationDate": "2021-07-
14T23:53:47.229Z",
      "completionDate": "2021-07-
15T00:16:07.282Z",
      "recoveryPointArn": "arn:aws:
ec2:us-west-2::image/ami-03
0cafb98e5a6dcdf",
      "jobRunTime": "00:22:20",
      "backupSizeInBytes": 858993459
2,
      "backupVaultName": "Default",
      "backupVaultArn": "arn:aws:
backup:us-west-2:1122334455
66:backup-vault:Default",
      "vaultType": "BACKUP_VAULT",
      "vaultLockStatus": "UNLOCKED",
```

Backup report template	Sample report in JSON format

```
      "isEncrypted": true,
      "encryptionKeyArn": "arn:aws:
kms:us-west-2:112233445566:
key/12345678-1234-1234-1234
-123456789012",
      "deleteAfterDays": 30,
      "moveToColdAfterDays": 7,
      "enableArchive": false,
      "iamRoleArn": "arn:aws:
iam::112233445566:role/service-
role/AWSBackupDefaultServiceRole"
    }
  ]
}
```

Note

The `vaultType` field is not included in the API response in regions where logically air-gapped vaults are not available.

Backup report template	Sample report in JSON format

COPY_JOB_REPORT	{
	  "reportItems": [
	    {
	      "reportTimePeriodStart":
	 "2021-07-14T15:48:31Z",
	      "reportTimePeriodEnd":
	 "2021-07-15T15:48:31Z",
	      "accountId": "112233445566",
	      "region": "us-west-2",
	      "copyJobId": "E0AD48A9-0560-B66
	8-3EF0-941FDC0AD6B1",
	      "jobStatus": "RUNNING",
	      "jobRunTime": "2021-07-
	16T00:00:00.010Z",
	      "resourceType": "EC2",
	      "resourceArn": "arn:aws:ec2:us-
	west-2:112233445566:instance/
	i-0bc877aee7782ba75",
	      "resourceName": "string",
	      "initiationDate": "2021-07-
	14T15:48:31Z",
	      "backupPlanName": "string",
	      "backupRuleName": "string",
	      "backupRuleSchedule":
	 "string",
	      "backupRuleTimezone":
	 "string",
	      "startWindowEnd": "2021-07-
	14T15:48:31Z",
	      "backupOptions": {},
	      "isParentJob": false,
	      "parentJobId": null,
	      "creationDate": "2021-07-
	15T15:42:04.771Z",
	      "completionDate": "2021-07-
	16T00:16:07.282Z",
	      "backupSizeInBytes": 858993459
	2,
	      "sourceRecoveryPointArn":
	 "arn:aws:ec2:us-west-2::image/
	ami-007b3819f25697299",

Backup report template	Sample report in JSON format

```
      "sourceBackupVaultArn":
 "arn:aws:backup:us-west-2:1
12233445566:backup-vault:Default",
      "destinationRecoveryPointAr
n": "arn:aws:ec2:us-east-2::image/
ami-0eba2199a0bcece3c",
      "destinationBackupVaultArn"
: "arn:aws:backup:us-east-2:1
12233445566:backup-vault:Default",
      "vaultType": "BACKUP_VAULT",
      "vaultLockStatus": "string",
      "isEncrypted": true,
      "encryptionKeyArn": "arn:aws:
kms:us-west-2:112233445566:
key/...",
      "deleteAfterDays": 30,
      "moveToColdAfterDays": 7,
      "enableArchive": false,
      "iamRoleArn": "arn:aws:
iam::112233445566:role/service-
role/AWSBackupDefaultServiceRole"
    }
  ]
}
```

Note

The `vaultType` field is not included in the API response in regions where logically air-gapped vaults are not available.

Backup report template	Sample report in JSON format

RESTORE_JOB_REPORT	{
	  "reportItems": [
	    {
	      "reportTimePeriod": "2021-07-14T15:53:30Z - 2021-07-15T15:53:30Z",
	      "accountId": "112233445566",
	      "region": "us-west-2",
	      "restoreJobId": "4CACA67D-4E12-DC05-6C2B-0E97D01FA41E",
	      "jobStatus": "RUNNING",
	      "recoveryPointArn": "arn:aws:ec2:us-west-2::image/ami-00201ecb57a5271ae",
	      "resourceName": "string",
	      "initiationDate": "2021-07-14T15:53:30Z",
	      "backupPlanName": "string",
	      "backupRuleName": "string",
	      "backupRuleSchedule": "string",
	      "backupRuleTimezone": "string",
	      "startWindowEnd": "2021-07-14T15:53:30Z",
	      "backupOptions": {},
	      "isParentJob": false,
	      "parentJobId": null,
	      "vaultType": "BACKUP_VAULT",
	      "vaultLockStatus": "string",
	      "isEncrypted": true,
	      "encryptionKeyArn": "arn:aws:kms:us-west-2:112233445566:key/...",
	      "deleteAfterDays": 30,
	      "moveToColdAfterDays": 7,
	      "enableArchive": false,
	      "sourceResourceArn": "arn:aws:ec2:us-west-2:112233445566:instance/i-0bc877aee7782ba75",

Backup report template	Sample report in JSON format
	"backupVaultArn": "arn:aws:backup:us-west-2:112233445566:backup-vault:Default", "creationDate": "2021-07-15T15:52:49.797Z", "backupSizeInBytes": 8589934592, "percentDone": "0.00%", "iamRoleArn": "arn:aws:iam::112233445566:role/service-role/AWSBackupDefaultServiceRole" }] } **Note** The `vaultType` field is not included in the API response in regions where logically air-gapped vaults are not available.

Compliance report templates

Compliance report templates give you daily reports on the compliance of your backup activity and resources against the controls you defined in one or more frameworks. If the compliance status of one of your frameworks is `Non-compliant`, review a compliance report to identify the non-compliant resources.

Types of compliance report templates

- `Control compliance report` helps you track the compliance status of the controls you have defined in your frameworks.
- `Resource compliance report` helps you track the compliance status of your resources against the controls you defined in your frameworks. These reports include detailed evaluation

results, including identifying information on non-compliant resources that you can use to identify and correct those resources.

The following table shows sample output from a compliance report.

Compliance report template	Sample report in JSON format

CONTROL_COMPLIANCE_REPORT	{
	  "reportItems": [
	    {
	      "accountId": "112233445566",
	      "region": "me-south-1",
	      "frameworkName": "TestFram
	ework7",
	      "frameworkDescription": "A test
	 framework",
	      "controlName": "BACKUP_R
	ESOURCES_PROTECTED_BY_BACKUP_PLAN",
	      "controlComplianceStatus":
	 "NON_COMPLIANT",
	      "lastEvaluationTime": "2021-08-
	17T03:21:56.002Z",
	      "numResourcesCompliant": 91,
	      "numResourcesNonCompliant":
	 205,
	      "controlFrequency": "Twelve_H
	ours",
	      "controlScope": "",
	      "controlParameters": ""
	    },
	    {
	      "accountId": "112233445566",
	      "region": "me-south-1",
	      "frameworkName": "TestFram
	ework7",
	      "frameworkDescription": "A test
	 framework",
	      "controlName": "BACKUP_P
	LAN_MIN_FREQUENCY_AND_MIN_R
	ETENTION_CHECK",
	      "controlComplianceStatus":
	 "NON_COMPLIANT",

Compliance report template	Sample report in JSON format


```
      "lastEvaluationTime": "2021-08-
17T03:21:19.995Z",
      "numResourcesCompliant": 0,
      "numResourcesNonCompliant": 25,
      "controlScope": "{Complia
nceResourceTypes: [],}",
      "controlParameters": "{\"requi
redFrequencyValue\":\"1\",\
"requiredRetentionDays\":\"35\",
\"requiredFrequencyUnit\":\"hours
\"}"
    }
  ]
}
```

Compliance report template	Sample report in JSON format

RESOURCE_COMPLIANCE_REPORT	{
	  "reportItems": [
	    {
	      "accountId": "112233445566",
	      "region": "us-west-2",
	      "frameworkName": "MyTestFr
	amework",
	      "frameworkDescription": "",
	      "controlName": "BACKUP_L
	AST_RECOVERY_POINT_CREATED",
	      "resourceId": "AWS::EFS
	::FileSystem/fs-63c74e66",
	      "resourceType": "AWS::EFS
	::FileSystem",
	      "resourceComplianceStatus":
	 "NON_COMPLIANT",
	      "lastEvaluationTime": "2021-07-
	07T18:55:40.963Z"
	    },
	    {
	      "accountId": "112233445566",
	      "region": "us-west-2",
	      "frameworkName": "MyTestFr
	amework",
	      "frameworkDescription": "",
	      "controlName": "BACKUP_L
	AST_RECOVERY_POINT_CREATED",
	      "resourceId": "AWS::EFS
	::FileSystem/fs-b3d7c218",
	      "resourceType": "AWS::EFS
	::FileSystem",
	      "resourceComplianceStatus":
	 "NON_COMPLIANT",
	      "lastEvaluationTime": "2021-07-
	07T18:55:40.961Z"
	    }
	  ]
	}

Scanning report templates

Scanning report templates. These templates give you daily updates on your scanning jobs with expanded context including vault properties, and backup plan details. You can use these reports to monitor your scan job statuses, reports, and identify any failures that might need further action. The following table lists a scanning report template name and its sample output.

Scanning report template	Sample report in JSON format
MALWARE_JOB_REPORT	{ "reportTimePeriodStart": "2025-11-09T00:00:00Z", "reportTimePeriodEnd": "2025-11-10T00:00:00Z", "accountId": "025066259999", "region": "us-east-1", "scanJobId": "489abba3-0a57-4207-93ff-d3947d85a8d3", "scanId": "9ddd3144f68ea3ee6e388c66a0b55467", "malwareScanner": "GUARDDUTY", "jobStatus": "RUNNING", "scanResultStatus":"", "statusMessage": "", "resourceType": "EBS", "resourceArn": "arn:aws:ec2:us-east-1:025066259999:volume/vol-0f1c480a6a9b33cb7", "backupPlanArn": "arn:aws:backup:us-east-1:025066259999:backup-plan:orgs/4232272a-ed54-3a88-b1d0-ace894b6c24c", "creationDate": "2025-10-28T17:30:50.820Z", "recoveryPointArn": "arn:aws:ec2:us-east-1::snapshot/snap-03fef858bad24c7a6", "backupVaultName": "Default", "backupVaultArn": "arn:aws:backup:us-east-1:025066259999:backup-vault:Default",

Scanning report template	Sample report in JSON format
	"iamRoleArn": "arn:aws:iam::025066259999:role/service-role/AWSBackupDefaultServiceRole", "scannerRoleArn": "arn:aws:iam::025066259999:role/AWSBackupGuardDutyRolePolicyForScans" }

Creating report plans using the AWS Backup console

There are two types of reports. One type is a **jobs report**, which shows jobs finished in the last 24 hours and all active jobs. The second type of report is a **compliance report**. Compliance reports can monitor resource levels or the different controls that are in effect. When you create a report, you choose which type of report to create.

Depending on your type of account, the console display may vary. Only management accounts will see multi-account functionality.

Similar to a *backup plan*, you create a *report plan* to automate the creation of your reports and define their destination Amazon S3 bucket. A report plan requires that you have an S3 bucket to receive your reports. You can't use a bucket from another account. For instructions on setting up a new S3 bucket, see Step 1: Create your first S3 bucket in the *Amazon Simple Storage Service User Guide*.

To create your report plan in the AWS Backup console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the left navigation pane, choose **Reports**.
3. Choose **Create report plan**.
4. Choose one of the report templates from the list.
5. Enter a unique **Report plan name**. The name must be between 1 and 256 characters, starting with a letter, and consisting of letters (a-z, A-Z), numbers (0-9), and underscores (_).
6. (Optional) Enter a **Report plan description**.

7. *Compliance report templates for one account only*. Choose one or more frameworks on which to report. You can add a maximum 1,000 frameworks to a report plan.

   1. Choose your AWS Region.
   2. Choose a framework from that Region.
   3. Choose **Add framework**.

8. (Optional) To add tags to your report plan, choose **Add tags to the report plan**.
9. If you are using a management account, you can specify which accounts you want to include in this report plan. You can select **Only my account**, which will generate reports on just the account to which you're currently logged in. Or, you can select **One or more accounts in my organization** (*available to management and delegated administrator accounts*).
10. (*If you are creating a compliance report for one Region only, skip this step*). You can select which Regions to include in your report. Click the drop down menu to show Regions available to you. Select *All available Regions* or the Regions you prefer.

    - The **Include new Regions when they are incorporated into Backup Audit Manager** check box will trigger new Regions to be included in your reports when they become available.

11. Choose the **File format** of your report. All reports can be exported in CSV format. Additionally, reports for a single Region can be exported in JSON format.
12. For **S3 bucket name**, choose a bucket from your account.
13. (Optional) Enter a bucket prefix.

    AWS Backup delivers your *current account, current Region* reports to `s3://amzn-s3-demo-bucket/prefix/Backup/accountID/Region/year/month/day/report-name`.

    AWS Backup delivers your *cross-account* reports to `s3://amzn-s3-demo-bucket/prefix/Backup/crossaccount/Region/year/month/day/report-name`

    AWS Backup delivers your *cross-Region* reports to `s3://amzn-s3-demo-bucket/prefix/Backup/accountID/crossregion/year/month/day/report-name`

14. Choose **Create report plan**.

Next, you must allow your S3 bucket to receive reports from AWS Backup. After you create a report plan, AWS Backup Audit Manager automatically generates an S3 bucket access policy for you to apply.

If you encrypt your bucket using a customer managed KMS key, the KMS key policy must meet the following requirements:

- The `Principal` attribute must include the Backup Audit Manager service-linked role `AWSServiceRolePolicyForBackupReports` ARN.

- The `Action` attribute must include `kms:GenerateDataKey` at minimum.

The policy AWSServiceRolePolicyForBackupReports has these permissions.

To view and apply this access policy to your S3 bucket

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the left navigation pane, choose **Reports**.
3. Under **Report plan name**, select a report plan by choosing its name.
4. Choose **Edit**.
5. Choose **View access policy for S3 bucket**. You can also use the policy at the end of this procedure.
6. Choose **Copy permissions**.
7. Choose **Edit bucket policy**. Note that until the backup report is created the first time, the service-linked role referred to in the S3 bucket policy will not yet exist, resulting in the error "Invalid principal".
8. Copy the permissions to the **Policy**.

Sample bucket policy

JSON

```
{
  "Version":"2012-10-17",
  "Statement":[
    {
      "Effect":"Allow",
      "Principal":{
        "AWS":"arn:aws:iam::123456789012:role/aws-service-role/
reports.backup.amazonaws.com/AWSServiceRoleForBackupReports"
```

```
      },
      "Action":"s3:PutObject",
      "Resource":[
        "arn:aws:s3:::BucketName/*"
      ],
      "Condition":{
        "StringEquals":{
          "s3:x-amz-acl":"bucket-owner-full-control"
        }
      }
    }
  ]
}
```

If you use a custom AWS Key Management Service to encrypt your target S3 bucket that stores the reports, include the following action in your policy:

```
"Action":[
  "kms:GenerateDataKey"
],
"Resource":[
  "*"
],
```

Using an S3 bucket encrypted with a cross-account KMS key

If your S3 bucket (in Account A) is encrypted with a customer managed KMS key from a different account (Account B), complete the following additional steps to allow AWS Backup to deliver reports to the bucket.

To configure a cross-account KMS key for Backup reports

1. Add the S3 bucket policy as described in the preceding section.
2. Update the KMS key policy in Account B to grant permissions to the `AWSServiceRoleForBackupReports` service-linked role in Account A and to the IAM entity that creates the report plan:

```
{
    "Version": "2012-10-17",
```

```
    "Statement": [
        {
            "Sid": "Allow Backup Reports service-linked role to use the key",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::AccountA-ID:role/aws-service-role/
reports.backup.amazonaws.com/AWSServiceRoleForBackupReports"
            },
            "Action": [
                "kms:GenerateDataKey",
                "kms:Encrypt"
            ],
            "Resource": "*"
        },
        {
            "Sid": "Allow IAM entity to create grants on the key",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::AccountA-ID:role/IAM-role-name"
            },
            "Action": "kms:CreateGrant",
            "Resource": "*"
        }
    ]
}
```

Replace *AccountA-ID* with the AWS account ID where the S3 bucket and AWS Backup report plan reside. Replace *IAM-role-name* with the IAM role you use to create the report plan.

3. From Account A, create a KMS grant for the service-linked role by running the following AWS CLI command with an IAM entity that has `kms:CreateGrant` permission on the key:

```
aws kms create-grant \
    --key-id arn:aws:kms:region:AccountB-ID:key/key-id \
    --grantee-principal arn:aws:iam::AccountA-ID:role/aws-service-role/
reports.backup.amazonaws.com/AWSServiceRoleForBackupReports \
    --operations GenerateDataKey
```

Replace *region*, *AccountB-ID*, *key-id*, and *AccountA-ID* with your values.

4. Create your on-demand or scheduled report plan. Reports will be delivered to the S3 bucket encrypted with the cross-account KMS key.

Creating report plans using the AWS Backup API

You can also work with report plans programmatically.

There are two types of reports. One type is a **jobs report**, which shows jobs finished in the last 24 hours and all active jobs. The second type of report is a **compliance report**. Compliance reports can monitor resource levels or the different controls that are in effect. When you create a report, you choose which type of report to create.

Similar to a *backup plan*, you create a *report plan* to automate the creation of your reports and define their destination Amazon S3 bucket. A report plan requires that you have an S3 bucket to receive your reports. For instructions on setting up a new S3 bucket, see Step 1: Create your first S3 bucket in the *Amazon Simple Storage Service User Guide*.

If you encrypt your bucket using a custom KMS key, the KMS key policy must meet the following requirements:

- The `Principal` attribute must include the Backup Audit Manager service-linked role `AWSServiceRoleForBackupReports` ARN.

- The `Action` attribute must include `kms:GenerateDataKey` at minimum.

The policy AWSServiceRolePolicyForBackupReports has these permissions.

For single-account, single-Region reports, use the following syntax to call CreateReportPlan.

```
{
   "ReportPlanName": "string",
   "ReportPlanDescription": "string",
   "ReportSetting": {
        "ReportTemplate": enum, // Can be RESOURCE_COMPLIANCE_REPORT,
 CONTROL_COMPLIANCE_REPORT, BACKUP_JOB_REPORT, COPY_JOB_REPORT, or RESTORE_JOB_REPORT.
 Only include "ReportCoverageList" if your report is a COMPLIANCE_REPORT.
   "ReportDeliveryChannel": {
       "S3BucketName": "string",
       "S3KeyPrefix": "string",
       "Formats": [ enum ] // Optional. Can be either CSV, JSON, or both. Default is
 CSV if left blank.
   },
   "ReportPlanTags": {
       "string" : "string" // Optional.
```

```
   },
   "IdempotencyToken": "string"
}
```

When you call DescribeReportPlan with the unique name of a report plan, the AWS Backup API responds with the following information.

```
{
    "ReportPlanArn": "string",
    "ReportPlanName": "string",
    "ReportPlanDescription": "string",
    "ReportSetting": {
        "ReportTemplate": enum,
    },
    "ReportDeliveryChannel": {
        "S3BucketName": "string",
        "S3KeyPrefix": "string",
        "Formats": [ enum ]
    },
    "DeploymentStatus": enum
    "CreationTime": timestamp,
    "LastAttemptExecutionTime": timestamp,
    "LastSuccessfulExecutionTime": timestamp
}
```

For multi-account, multi-Region reports, use the following syntax to call CreateReportPlan.

```
{
   "IdempotencyToken": "string",
   "ReportDeliveryChannel": {
      "Formats": [ "string" ], *//Organization report only support CSV file*
      "S3BucketName": "string",
      "S3KeyPrefix": "string"
   },
   "ReportPlanDescription": "string",
   "ReportPlanName": "string",
   "ReportPlanTags": {
      "string" : "string"
   },
   "ReportSetting": {
      "Accounts": [ "string" ], // Use string value of "ROOT" to include all
 organizational units
      "OrganizationUnits": [ "string" ],
```

```
      "Regions": ["string"], // Use wildcard value in string to include all Regions
      "FrameworkArns": [ "string" ],
      "NumberOfFrameworks": number,
      "ReportTemplate": "string"
   }
}
```

When you call DescribeReportPlan with the unique name of a report plan, the AWS Backup API responds with the following information for multi-account, multi-Region plans:

```
{
   "ReportPlan": {
      "CreationTime": number,
      "DeploymentStatus": "string",
      "LastAttemptedExecutionTime": number,
      "LastSuccessfulExecutionTime": number,
      "ReportDeliveryChannel": {
         "Formats": [ "string" ],
         "S3BucketName": "string",
         "S3KeyPrefix": "string"
      },
      "ReportPlanArn": "string",
      "ReportPlanDescription": "string",
      "ReportPlanName": "string",
      "ReportSetting": {
         "Accounts":[ "string" ],
         "OrganizationUnits":[ "string" ],
         "Regions": [ "string" ],
         "FrameworkArns": [ "string" ],
         "NumberOfFrameworks": number,
         "ReportTemplate": "string"
      }
   }
}
```

Creating on-demand reports

You can generate new reports at your convenience by creating an on-demand report with the following steps. AWS Backup Audit Manager delivers your on-demand report to the Amazon S3 bucket that you specified in your report plan.

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.

2. In the left navigation pane, choose **Reports**.
3. Under **Report plan name**, select a report plan by choosing its name.
4. Choose **Create on-demand report**.

You can generate an on-demand report for an existing report plan.

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the left navigation pane, choose **Reports**.
3. Under **Report plans**, select a report plan by clicking on the radio button next to the report plan name.
4. Click **Actions**, then click **Create on-demand report**.

You can do this for multiple reports, even while reports are being generated.

Viewing audit reports

You can open, view, and analyze AWS Backup Audit Manager reports using the programs that you ordinarily use to work with CSV or JSON files. Note that reports for multiple regions or multiple accounts are only available in CSV format.

Large files are broken up into multiple reports if the total file size exceeds 50 MB. If the resulting files are over 50 MB, AWS Backup Audit Manager will create additional CSV files with the remainder of the report.

To view a report

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the left navigation pane, choose **Reports**.
3. Under **Report plan name**, select a report plan by choosing its name.
4. Under **Report jobs**, click on the report link to view the report.
5. If your report's **Report status** has a dotted underline, choose it for information about your report.
6. Choose which report to view by its **Completion time**.
7. Choose the **S3 link**. This opens your destination S3 bucket.
8. Under **Name**, choose the name of the report that you want to view.

9. To save the report to your computer, choose **Download**.

Updating report plans

You can update an existing report plan's description, its delivery destination, and format. If applicable, you can also add or remove frameworks from the report plan.

To update an existing report plan

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the left navigation pane, choose **Reports**.
3. Under **Report plan name**, select a report plan by choosing its name.
4. Choose **Edit**.
5. You can edit the report plan details, including the report name and description, as well as which accounts and Regions are included in the report.

Deleting report plans

You can delete an existing report plan. When you delete a report plan, any reports already created by that report plan will remain in their destination Amazon S3 bucket.

To delete an existing report plan

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the left navigation pane, choose **Reports**.
3. Under **Report plan name**, select a report plan by choosing its name.
4. Choose **Delete**.
5. Enter the name of your report plan, and then choose **Delete report plan**.

Using AWS Backup Audit Manager with CloudFormation

We provide the following sample CloudFormation templates for your reference:

Topics

- Turn on resource tracking

- Deploy default controls
- Exempt IAM roles from control evaluation
- Create a report plan

Turn on resource tracking

The following template turns on resource tracking as described in Turning on resource tracking.

```
AWSTemplateFormatVersion: 2010-09-09
Description: Enable AWS Config

Metadata:
  AWS::CloudFormation::Interface:
    ParameterGroups:
      - Label:
          default: Recorder Configuration
        Parameters:
          - AllSupported
          - IncludeGlobalResourceTypes
          - ResourceTypes
      - Label:
          default: Delivery Channel Configuration
        Parameters:
          - DeliveryChannelName
          - Frequency
      - Label:
          default: Delivery Notifications
        Parameters:
          - TopicArn
          - NotificationEmail
    ParameterLabels:
      AllSupported:
        default: Support all resource types
      IncludeGlobalResourceTypes:
        default: Include global resource types
      ResourceTypes:
        default: List of resource types if not all supported
      DeliveryChannelName:
        default: Configuration delivery channel name
      Frequency:
        default: Snapshot delivery frequency
      TopicArn:
```

```
        default: SNS topic name
      NotificationEmail:
        default: Notification Email (optional)

Parameters:
  AllSupported:
    Type: String
    Default: True
    Description: Indicates whether to record all supported resource types.
    AllowedValues:
      - True
      - False

  IncludeGlobalResourceTypes:
    Type: String
    Default: True
    Description: Indicates whether AWS Config records all supported global resource
 types.
    AllowedValues:
      - True
      - False

  ResourceTypes:
    Type: List<String>
    Description: A list of valid AWS resource types to include in this recording group,
 such as AWS::EC2::Instance or AWS::CloudTrail::Trail.
    Default: <All>

  DeliveryChannelName:
    Type: String
    Default: <Generated>
    Description: The name of the delivery channel.

  Frequency:
    Type: String
    Default: 24hours
    Description: The frequency with which AWS Config delivers configuration snapshots.
    AllowedValues:
      - 1hour
      - 3hours
      - 6hours
      - 12hours
      - 24hours
```

```
  TopicArn:
    Type: String
    Default: <New Topic>
    Description: The Amazon Resource Name (ARN) of the Amazon Simple Notification
 Service (Amazon SNS) topic that AWS Config delivers notifications to.

  NotificationEmail:
    Type: String
    Default: <None>
    Description: Email address for AWS Config notifications (for new topics).

Conditions:
  IsAllSupported: !Equals
    - !Ref AllSupported
    - True
  IsGeneratedDeliveryChannelName: !Equals
    - !Ref DeliveryChannelName
    - <Generated>
  CreateTopic: !Equals
    - !Ref TopicArn
    - <New Topic>
  CreateSubscription: !And
    - !Condition CreateTopic
    - !Not
      - !Equals
        - !Ref NotificationEmail
        - <None>

Mappings:
  Settings:
    FrequencyMap:
      1hour   : One_Hour
      3hours  : Three_Hours
      6hours  : Six_Hours
      12hours : Twelve_Hours
      24hours : TwentyFour_Hours

Resources:

  ConfigBucket:
    DeletionPolicy: Retain
    Type: AWS::S3::Bucket
    Properties:
      BucketEncryption:
```

```
          ServerSideEncryptionConfiguration:
            - ServerSideEncryptionByDefault:
                SSEAlgorithm: AES256

  ConfigBucketPolicy:
    Type: AWS::S3::BucketPolicy
    Properties:
      Bucket: !Ref ConfigBucket
      PolicyDocument:
        Version: 2012-10-17
        Statement:
          - Sid: AWSConfigBucketPermissionsCheck
            Effect: Allow
            Principal:
              Service:
                - config.amazonaws.com
            Action: s3:GetBucketAcl
            Resource:
              - !Sub "arn:${AWS::Partition}:s3:::${ConfigBucket}"
          - Sid: AWSConfigBucketDelivery
            Effect: Allow
            Principal:
              Service:
                - config.amazonaws.com
            Action: s3:PutObject
            Resource:
              - !Sub "arn:${AWS::Partition}:s3:::${ConfigBucket}/AWSLogs/
${AWS::AccountId}/*"
          - Sid: AWSConfigBucketSecureTransport
            Action:
              - s3:*
            Effect: Deny
            Resource:
              - !Sub "arn:${AWS::Partition}:s3:::${ConfigBucket}"
              - !Sub "arn:${AWS::Partition}:s3:::${ConfigBucket}/*"
            Principal: "*"
            Condition:
              Bool:
                aws:SecureTransport:
                  false

  ConfigTopic:
    Condition: CreateTopic
    Type: AWS::SNS::Topic
```

```
    Properties:
      TopicName: !Sub "config-topic-${AWS::AccountId}"
      DisplayName: AWS Config Notification Topic
      KmsMasterKeyId: "alias/aws/sns"

  ConfigTopicPolicy:
    Condition: CreateTopic
    Type: AWS::SNS::TopicPolicy
    Properties:
      Topics:
        - !Ref ConfigTopic
      PolicyDocument:
        Statement:
          - Sid: AWSConfigSNSPolicy
            Action:
              - sns:Publish
            Effect: Allow
            Resource: !Ref ConfigTopic
            Principal:
              Service:
                - config.amazonaws.com

  EmailNotification:
    Condition: CreateSubscription
    Type: AWS::SNS::Subscription
    Properties:
      Endpoint: !Ref NotificationEmail
      Protocol: email
      TopicArn: !Ref ConfigTopic

  ConfigRecorderServiceRole:
    Type: AWS::IAM::ServiceLinkedRole
    Properties:
      AWSServiceName: config.amazonaws.com
      Description: Service Role for AWS Config

  ConfigRecorder:
    Type: AWS::Config::ConfigurationRecorder
    DependsOn:
      - ConfigBucketPolicy
      - ConfigRecorderServiceRole
    Properties:
      RoleARN: !Sub arn:${AWS::Partition}:iam::${AWS::AccountId}:role/aws-service-role/
config.amazonaws.com/AWSServiceRoleForConfig
```

```
      RecordingGroup:
        AllSupported: !Ref AllSupported
        IncludeGlobalResourceTypes: !Ref IncludeGlobalResourceTypes
        ResourceTypes: !If
          - IsAllSupported
          - !Ref AWS::NoValue
          - !Ref ResourceTypes

  ConfigDeliveryChannel:
    Type: AWS::Config::DeliveryChannel
    DependsOn:
      - ConfigBucketPolicy
    Properties:
      Name: !If
        - IsGeneratedDeliveryChannelName
        - !Ref AWS::NoValue
        - !Ref DeliveryChannelName
      ConfigSnapshotDeliveryProperties:
        DeliveryFrequency: !FindInMap
          - Settings
          - FrequencyMap
          - !Ref Frequency
      S3BucketName: !Ref ConfigBucket
      SnsTopicARN: !If
        - CreateTopic
        - !Ref ConfigTopic
        - !Ref TopicArn
```

Deploy default controls

The following template creates a framework with the default controls described in AWS Backup Audit Manager controls and remediation.

```
AWSTemplateFormatVersion: '2010-09-09'
Resources:
  TestFramework:
    Type: AWS::Backup::Framework
    Properties:
      FrameworkControls:
        - ControlName: BACKUP_RESOURCES_PROTECTED_BY_BACKUP_PLAN
        - ControlName: BACKUP_RECOVERY_POINT_MINIMUM_RETENTION_CHECK
          ControlInputParameters:
            - ParameterName: requiredRetentionDays
```

```
              ParameterValue: '35'
        - ControlName: BACKUP_RECOVERY_POINT_MANUAL_DELETION_DISABLED
        - ControlName: BACKUP_PLAN_MIN_FREQUENCY_AND_MIN_RETENTION_CHECK
          ControlInputParameters:
            - ParameterName: requiredRetentionDays
              ParameterValue: '35'
            - ParameterName: requiredFrequencyUnit
              ParameterValue: 'hours'
            - ParameterName: requiredFrequencyValue
              ParameterValue: '24'
          ControlScope:
            Tags:
              - Key: customizedKey
                Value: customizedValue
        - ControlName: BACKUP_RECOVERY_POINT_ENCRYPTED
        - ControlName: BACKUP_RESOURCES_PROTECTED_BY_CROSS_REGION
          ControlInputParameters:
            - ParameterName: crossRegionList
              ParameterValue: 'eu-west-2'
        - ControlName: BACKUP_RESOURCES_PROTECTED_BY_CROSS_ACCOUNT
          ControlInputParameters:
            - ParameterName: crossAccountList
              ParameterValue: '111122223333'
        - ControlName: BACKUP_RESOURCES_PROTECTED_BY_BACKUP_VAULT_LOCK
        - ControlName: BACKUP_LAST_RECOVERY_POINT_CREATED
        - ControlName: RESTORE_TIME_FOR_RESOURCES_MEET_TARGET
          ControlInputParameters:
            - ParameterName: maxRestoreTime
              ParameterValue: '720'

Outputs:
  FrameworkArn:
    Value: !GetAtt TestFramework.FrameworkArn
```

Exempt IAM roles from control evaluation

The control BACKUP_RECOVERY_POINT_MANUAL_DELETION_DISABLED allows you to exempt up to five IAM roles that can still manually delete recovery points. The following template deploys this control and also exempts two IAM roles.

```
AWSTemplateFormatVersion: '2010-09-09'
Resources:
  TestFramework:
```

```
    Type: AWS::Backup::Framework
    Properties:
      FrameworkControls:
        - ControlName: BACKUP_RECOVERY_POINT_MANUAL_DELETION_DISABLED
          ControlInputParameters:
            - ParameterName: "principalArnList"
              ParameterValue: !Sub
 "arn:aws:iam::${AWS::AccountId}:role/AccAdminRole,arn:aws:iam::${AWS::AccountId}:role/
ConfigRole"

Outputs:
  FrameworkArn:
    Value: !GetAtt TestFramework.FrameworkArn
```

Create a report plan

The following template creates a report plan.

```
Description: "Basic AWS::Backup::ReportPlan template"

Parameters:
  ReportPlanDescription:
    Type: String
    Default: "SomeReportPlanDescription"
  S3BucketName:
    Type: String
    Default: "some-s3-bucket-name"
  S3KeyPrefix:
    Type: String
    Default: "some-s3-key-prefix"
  ReportTemplate:
    Type: String
    Default: "BACKUP_JOB_REPORT"

Resources:
  TestReportPlan:
    Type: "AWS::Backup::ReportPlan"
    Properties:
      ReportPlanDescription: !Ref ReportPlanDescription
      ReportDeliveryChannel:
        Formats:
          - "CSV"
        S3BucketName: !Ref S3BucketName
```

```
        S3KeyPrefix: !Ref S3KeyPrefix
      ReportSetting:
        ReportTemplate: !Ref ReportTemplate
        Regions: ['us-west-2', 'eu-west-1', 'us-east-1']
        Accounts: ['123456789098']
        OrganizationUnits: ['ou-abcd-1234wxyz']
      ReportPlanTags:
        - Key: "a"
          Value: "1"
        - Key: "b"
          Value: "2"

Outputs:
  ReportPlanArn:
    Value: !GetAtt TestReportPlan.ReportPlanArn
```

Using AWS Backup Audit Manager with AWS Audit Manager

AWS Backup Audit Manager controls map to prebuilt, standard controls in AWS Audit Manager, allowing you to import your AWS Backup Audit Manager compliance findings to your AWS Audit Manager reports. You might want to do so to help a compliance officer, audit manager, or other colleague who reports on backup activity as part of your organization's overall compliance posture.

You can import the compliance results of your AWS Backup Audit Manager controls to your AWS Audit Manager frameworks. To enable AWS Audit Manager to automatically collect data from your AWS Backup Audit Manager controls, create a custom control in AWS Audit Manager using the instructions for Customizing an existing control in the *AWS Audit Manager User Guide*. As you follow those instructions, note that the **Data source** for AWS Backup controls is **AWS Config**.

For a list of AWS Backup controls, see Choosing your controls.

Controls and remediation

This page lists the available controls for AWS Backup Audit Manager. You can choose the right info pane to see a list of controls and jump to a specific control. To quickly compare controls, see the table in Choosing your controls. To programmatically define controls, see the code snippets in Creating frameworks using the AWS Backup API.

You can use up to 50 controls per account per Region. Using the same control in two different frameworks counts as using two controls of the 50 control limit.

This page lists each control with the following information:

- Description. Values in brackets ("[ ]") are the default parameter values.
- The **resource(s)** the control evaluates.
- The **parameters** of the control.
- Occasion when running of control **occurs**.
- The **scope** of the control, as follows:
  - You can specify **Resources by type** by choosing one or more AWS Backup-supported services.
  - You specify a **Tagged resources** scope with a single tag key and optional value.
  - You can specify a single resource using the **Single resource** dropdown list.
- Remediation steps to bring applicable resources into compliance.

Note that only active resources will be included when controls evaluate resources for compliance. For example, an Amazon EC2 instance in a running state will be evaluated by the control Last recovery point was created. An EC2 instance in a stopped state will not be included in the compliance evaluation.

Backup resources are included in at least one backup plan

Description: Evaluates if resources are included in at least one backup plan.

Resource: `AWS Backup: backup selection`

Parameters: None

Occurs: Automatically every 24 hours

Scope:

- Tagged resources
- Resources by type (default)
- Single resource

Remediation: Assign the resources to a backup plan. AWS Backup automatically protects your resources after you assign them to a backup plan. For more information, see Assigning resources to a backup plan.

Backup plan minimum frequency and minimum retention

Description: Evaluates if backup plans contain at least one backup rule for which the backup frequency is at least [1 day] and retention period is at least [35 days].

Resource: `AWS Backup: backup plans`

Parameters:

- Required backup frequency in number of hours or days.
- Numerical value for required backup frequency. Maximum of 24 for hours, 31 for days.
- Required retention period in number of days, weeks, months, or years. We recommend a warm storage retention of period of at least one week to enable AWS Backup to take incremental backups when possible, avoiding additional charges.

Occurs: Configuration changes

Scope:

- Tagged resources
- Single resource

Remediation: Update a backup plan to change either its backup frequency, retention period, or both. Updating your backup plan changes the retention period for recovery points the plan creates after your update.

Vaults prevent manual deletion of recovery points

Description: Evaluates if backup vaults do not allow manual deletion of recovery points except by certain IAM roles.

Resource: `AWS Backup: backup vaults`

Parameters: The Amazon Resource Names (ARNs) of up to five IAM roles allowed to manually delete recovery points.

Occurs: Configuration changes

Scope:

- Tagged resources
- Single resource

Remediation: Create or modify a resource-based access policy on a backup vault. For an example policy and instructions on how to set a backup vault access policy, see Deny access to delete recovery points in a backup vault.

Recovery points are encrypted

Description: Evaluates if recovery points are encrypted.

Resource: `AWS Backup: recovery points`

Parameters: None

Occurs: Configuration changes

Scope:

- Tagged resources

Remediation: Configure encryption for the recovery points. The way you configure encryption for AWS Backup recovery points differs depending on the resource type.

You can configure encryption for resource types that support full AWS Backup management in using AWS Backup. If the resource type does not support full AWS Backup management, you must configure its backup encryption by following that service's instructions, such as Amazon EBS encryption in the *Amazon Elastic Compute Cloud User Guide*. To see the list of resource types that support full AWS Backup management, see the "Full AWS Backup management" section of the Feature availability by resource table.

Minimum retention established for recovery point

Description: Evaluates if recovery point retention period is at least [35 days].

Resource: `AWS Backup: recovery points`

Parameters: Required recovery point retention period in number of days, weeks, months, or years. We recommend a warm storage retention of period of at least one week to enable AWS Backup to take incremental backups when possible, avoiding additional charges.

Occurs: Configuration changes

Scope:

- Tagged resources

Remediation: Change the retention periods of your recovery points. For more information, see Editing a backup.

Cross-Region backup copy is scheduled

Description: Evaluates if a resource is configured to create copies of its backups to another AWS Region.

Resource: `AWS Backup: backup selection`

Parameters:

- Select the AWS Region(s) where the backup copy should exist (Optional)
- Region

Occurs: Automatically every 24 hours

Scope:

- Tagged resources
- Resources by type
- Single resource

Remediation: Update a backup plan to change the AWS Region where backup copy should exist.

Cross-account backup copy is scheduled

Description: Evaluates if a resource is configured to create copies of its backups to another account. You can add up to 5 accounts for the control to evaluate. The destination account must be in the same organization as the source account in AWS Organizations.

Resource: `AWS Backup: backup selection`

Parameters:

- Select the AWS account ID(s) where the backup copy should exist (Optional)
- Account ID

Occurs: Automatically every 24 hours

Scope:

- Tagged resources
- Resources by type
- Single resource

Remediation: Update a backup plan to change or add the AWS account ID(s) where the copy should exist.

Resources are in a backup plan with an AWS Backup Vault Lock

Description: Evaluates if a resource has immutable backups stored in a locked backup vault.

Resource: `AWS Backup: backup selection`

Parameters:

- Input the minimum and maximum retention days for AWS Backup Vault Lock (optional)
- Minimum retention days
- Maximum retention days

Occurs: Automatically every 24 hours

Scope:

- Tagged resources
- Resources by type
- Single resource

Remediation: Lock a backup vault to set its name, change either its minimum retention days, maximum retention days, or both. Can also include `ChangeableForDays` for a vault lock in compliance mode.

Last recovery point was created

Description: This control evaluates if a recovery point has been created within the specified time frame (in days or hours).

The control is compliant if the resource has had a recovery point created within the time frame specified. The control is non-compliant if a recovery point was not created within the number of days or hours specified.

> **Note**
>
> This control evaluates only recovery points created within the same account and Region. Recovery points copied from another account are not evaluated by this control.

Resource: `AWS Backup: recovery points`

Parameters:

- Input the specified time frame in whole numbers, either in hours or days.
- Values of `hours` can range from `1` to `744`.
- Value of `days` can range from `1` to `31`.

Occurs: Automatically every 24 hours

Scope:

- Tagged resources

- Resources by type
- Single resource

Remediation:

- Update a backup plan to change the specified time frame of recovery point creation.
- Additionally, you can create an on-demand backup.

Restore time for resources meet target

Description: Evaluates if restoring protected resources completed within the target restore time.

This control checks if the restore time of a particular resource meets the target duration. The rule is NON_COMPLIANT if `LatestRestoreExecutionTimeMinutes` of a resource type is greater than `maxRestoreTime` in minutes.

Parameters:

- `maxRestoreTime` (in minutes)

Occurs: Automatically every 24 hours

Scope:

- Tagged resources
- Resources by type
- Single resource

> **Note**
>
> AWS Backup does not provide any service-level agreements (SLAs) for a restore time. Restore times can vary based upon system load and capacity, even for restores containing the same resources.

Resources in a logically air-gapped vault

Description: This control evaluates if resources have at least one recovery point copied to a logically air-gapped vault within the specified value and time frame. This control is NON_COMPLIANT if a recovery point has not been copied to a logically air-gapped vault in the time frame configured for the control.

Resource: `AWS Backup: recovery points`

Parameters:

- `recoveryPointAgeValue`
- `recoveryPointAgeUnit`

Input the time period. Specify the unit in `days` or `hours`. Specify a value for that unit. Values of hours can be within 24 to 2184 inclusive. Values of days can be within 1 to 91 inclusive.

A minimum value of 7 days or 168 hours is recommended. The control value should be no more frequent than the copy creation frequency of your backup plan; otherwise, you may see an unexpected `NON_COMPLIANT` status until your next backup is copied into a logically air-gapped vault and this control is run.

Occurs: Automatically every 24 hours

Scope:

- Resources by type
- Single resource

Managing AWS Backup resources across multiple AWS accounts

> **Note**
>
> Before you manage resources across multiple AWS accounts in AWS Backup, your accounts must belong to the same organization in the AWS Organizations service.

Cross-account management overview

You can use the cross-account management feature in AWS Backup to manage and monitor your backup, restore, and copy jobs across AWS accounts that you configure with AWS Organizations. AWS Organizations is a service that offers policy-based management for multiple AWS accounts from a single management account. It enables you to standardize the way you implement backup policies, minimizing manual errors and effort simultaneously. From a central view, you can easily identify resources in all accounts that meet the criteria that you are interested in.

If you set up AWS Organizations, you can configure AWS Backup to monitor activities in all of your accounts in one place. You can also create a backup policy and apply it to selected accounts that are part of your organization and view the aggregate backup job activities directly from the AWS Backup console. This functionality enables backup administrators to effectively monitor backup job status in hundreds of accounts across their entire enterprise from a single management account. AWS Organizations quotas apply.

For example, you define a backup policy A that takes daily backups of specific resources and keeps them for 7 days. You choose to apply backup policy A to the whole organization. (This means that each account in the organization gets that backup policy, which creates a corresponding backup plan that is visible in that account.) Then, you create an OU named Finance, and you decide to keep its backups for only 30 days. In this case, you define a backup policy B, which overrides the lifecycle value, and attach it to that Finance OU. This means that all the accounts under the Finance OU get a new effective backup plan that takes daily backups of all specified resources and keeps them for 30 days.

In this example, backup policy A and backup policy B were merged into a single backup policy, which defines the protection strategy for all accounts under the OU named Finance. All the other

accounts in the organization remain protected by backup policy A. Merging is done only for backup policies that share the same backup plan name. You can also have policy A and policy B coexist in that account without any merging. You can use advanced merging operators in the JSON view of the console only. For details about merging policies, see Defining policies, policy syntax, and policy inheritance in the *AWS Organizations User Guide*. For additional references and use cases, see the blog Managing backups at scale in your AWS Organizations using AWS Backup and the video tutorial Managing backups at scale in your AWS Organizations using AWS Backup.

Cross-account management consists of cross-account monitoring, cross-account backups, backup policies, and delegated administrator accounts. Not all of these elements are available in all Regions. For information about where cross-account management is available, see Feature availability by AWS Region .

To use cross-account management

1. Create a management account in AWS Organizations and add accounts under the management account.
2. Enable the cross-account management feature in AWS Backup.
3. Create a backup policy to apply to all AWS accounts under your management account.

> **Note**
>
> For backup plans that are managed by Organizations, the resource opt-in settings in the management account override the settings in a member account for that particular backup plan, even if one or more delegated administrator accounts are configured. Delegated administrator accounts are member accounts with enhanced features and cannot override settings like a management account can.

4. Manage backup, restore, and copy jobs in all your AWS accounts.

Cross-account management settings

Cross-account management allows you to enable or disable settings for managing and monitoring activities across multiple accounts within your organization.

Cross-account backup

Allow accounts in your organization to copy backups to other accounts.

Multi-party approval

Multi-party approval integration allows you to opt-in all accounts in your organization to Multi-party approval.

Delegated Administrator

Delegated Administrator allows Backup to automatically synchronize delegated administrator permissions with Organizations.

Creating a management account in Organizations

First, you must create your organization and configure it with AWS member accounts in AWS Organizations. For instructions, see Tutorial: Creating and configuring an organization in the *AWS Organizations User Guide*.

As you add member accounts to your organization, ensure that each account has:

- At least one backup and/or logically air-gapped vault
- An IAM role

The backup policies you create will have a backup plan, but they will also identify the AWS Regions, the vault(s) used in the backup plan, the resources that will be backed up, and the IAM role that will be used to create the backup. Backup policies that reference accounts that lack the required information will not work as expected.

For more information, see Syntax for backup policies in the *AWS Organizations User Guide*.

Enabling cross-account management

Before you can use cross-account management in AWS Backup, the management account must enable the feature (that is, *opt in* to it). After the management account enables cross-account management, you can create backup policies that manage resources in multiple accounts.

To enable cross-account management

1. Open the AWS Backup console at https://console.aws.amazon.com/backup/. You must sign in using the credentials of your management account.

2. In the left navigation pane, choose **Settings** to open the cross-account management page.
3. In the **Backup policies** section, choose **Enable**.

   This gives you access to all the accounts and allows you to create policies that automate management of multiple accounts in your organization simultaneously.
4. In the **Cross-account monitoring** section, choose **Enable**.

   This enables you to monitor the backup, copy, and restore activities of all accounts in your organization from your management account.

Backup policies

You can combine backup plans with the scalability of policies in AWS Organizations to create backup policies to simplify management across your organization.

See the *AWS Organizations User Guide* for information on how to enable backup policies for your organization so you can:

- Create a backup policy
- Update a backup policy; or,
- Delete a backup policy
- Backup policy syntax and examples

See AWS Backup quotas for AWS Backup-specific quotas on elements contained in a policy.

> **Important**
>
> Backup policies support resource selection by resource types (for example, `arn:aws:ec2:*:*:volume/*`) or by tags, but not by individual resource ARNs. Use backup policies for broad application across multiple accounts in an organization or OU, not for targeting individual resources.
> To back up individual resources by ARN, use a local backup plan resource assignment instead.
> For full details on the `selections` syntax, see Backup policy syntax and examples in the AWS Organizations User Guide.

Delegated administrator

Delegated administration provides a convenient way for assigned users in a registered member account to perform most AWS Backup administrative tasks. You can choose to delegate administration of AWS Backup to a member account in AWS Organizations, thereby extending the ability to manage AWS Backup from outside the management account and across the entire organization.

A management account, by default, is the account used to edit and manage policies. Using the delegated administrator feature, you can delegate these management functions to member accounts you designate. In turn, those accounts can manage policies, in addition to the management account.

After a member account has been successfully registered for delegated administration, it is a delegated administrator account. Note that accounts, not users, are designated as delegated administrators.

Enabling delegated administrator accounts allows the option of managing backup policies, it minimizes the number of users with access to the management account, and it permits cross-account monitoring of jobs.

Below is a table showing the functions of the management account, accounts delegated as Backup administrators, and accounts that are members within the AWS Organization.

> **Note**
>
> Delegated administrator accounts are member accounts with enhanced features but cannot override service opt-in settings of other member accounts like a management account can.

PRIVILEGES	MANAGEMENT ACCOUNT	DELEGATED ADMINISTRATOR	MEMBER ACCOUNT
Register/deregister delegated administrator accounts	Yes	No	No
Enable cross-account management	Yes	No	No

PRIVILEGES	MANAGEMENT ACCOUNT	DELEGATED ADMINISTRATOR	MEMBER ACCOUNT
Manage backup policies across accounts in AWS Organizations	Yes	Yes	No
Monitor cross-account jobs	Yes	Yes	No

Prerequisites

Before you can delegate backup administration, you must first register at least one member account in your AWS organization as a **delegated administrator**. Before you can register an account as a delegated administrator, you must first configure the following:

- AWS Organizations must be enabled and configured with at least one member account in addition to your default management account.
- It is recommended that the Organization Management Account have the AWS Backup Service-Linked Role (AWSServiceRoleForBackup) present to automatically sync delegated administrator permissions with Organizations. This ensures delegated administrators have proper access when you enable opt-in regions or change administrator assignment. If the Backup Service-Linked Role is not present or got deleted, customers have a few options to create it:
  - Enabling the Delegated Administrator feature.
  - Access the Backup Vaults Console page.
  - Manually create it following the documentation in Create a service-linked role.
- In the AWS Backup console, ensure **backup policies**, **cross-account monitoring**, **cross-account backup**, and **delegated administrator** features are turned on. These are below the **Delegated administrators** pane in the AWS Backup console.
  - Cross-account monitoring allows you to monitor backup activity across all the accounts in your organization from the management account, as well as from delegated administrator accounts.
  - *Optional:* Cross-account backup, which allows accounts in your organization to copy backups to other accounts (for Backup-supported cross-account resources).

- *Optional:* Delegated administrator, which allows AWS Backup to automatically create the Backup Service-Linked Role to sync delegated administrator permissions with Organizations.
- Enable service access with AWS Backup.

There are two steps involved in setting up delegated administration. The first step is to delegate cross-account jobs monitoring. The second step is to delegate backup policy management.

Register a member account as a delegated administrator account

This is the first section: Using the AWS Backup console to register a delegated administrator account to monitor cross- account jobs. To delegate AWS Backup policies, you will use the Organizations console in the next section.

To register a member account using the AWS Backup Console:

1. Open the AWS Backup console at https://console.aws.amazon.com/backup/. You must sign in using the credentials of your management account.
2. Under **My Account** in the left-hand navigation of the console, choose **Settings**.
3. In the **Delegated administrator** pane, click **Register delegated administrator** or **Add delegated administrator**.
4. On the **Register delegated administrator** page, select the account you want to register, and then choose **Register account**.

This designated account will now be registered as a delegated administrator, with administrative privileges to monitor jobs across accounts within the organization. This member account cannot register or deregister other delegated administrator accounts. You can use the console to register up to 5 accounts as delegated administrators.

> **Important**
>
> Registering an account as a delegated administrator alone does not enable backup policy management. To allow the delegated administrator account to view and edit backup policies, you must also complete the following steps:
>
> - Attach the the section called "AWSBackupOrganizationAdminAccess" managed policy to the IAM roles or users in the delegated administrator account that need to perform organization-level tasks. This policy is not attached automatically.

- Delegate backup policy management through AWS Organizations by creating a resource-based delegation policy in the management account. For steps, see Delegate AWS Backup policies through AWS Organizations.

To register a member account using programmatically:

Use the CLI command `register-delegated-administrator`. You can specify the following parameters in your CLI request:

- `service-principal`
- `account-id`

Below is an example of a CLI request to register a member account programmatically:

```
aws organizations register-delegated-administrator \
--account-id 012345678912 \
--service-principal "backup.amazonaws.com"
```

Deregister a member account

Use the following procedure to remove administrative access from AWS Backup by deregistering a member account in your AWS organization that had previously been designated as a delegated administrator.

To deregister a member account using the Console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup/. You must sign in using the credentials of your management account.
2. Under **My Account** in the left-hand navigation of the console, choose **Settings**.
3. In the **Delegated administrator** section, click **Deregister account**.
4. Choose the account(s) you want to deregister.
5. In the **Deregister account** dialog box, review the security implications, and then type `confirm` to complete the deregistration.
6. Choose `Deregister account`.

To deregister a member account using programmatically:

Use the CLI command `deregister-delegated-administrator` to deregister a delegated administrator account. You can specify the following parameters in your API request:

- `service-principal`
- `account-id`

Below is an example of a CLI request to deregister a member account programmatically:

```
aws organizations deregister-delegated-administrator \
--account-id 012345678912 \
--service-principal "backup.amazonaws.com"
```

Delegate AWS Backup policies through AWS Organizations

Within the AWS Organizations console, you can delegate administration of multiple policies, including Backup policies.

From the management account logged into the AWS Organizations console, you can create, view, or delete a resource-based delegation policy for your organization. For steps to delegate policies, see Create a resource-based delegation policy in the *AWS Organizations User Guide*.

Monitoring activities in multiple AWS accounts

To monitor backup, copy, and restore jobs across accounts, you must enable cross-account monitoring. This lets you monitor backup activities in all accounts from your organizations management account. After you opt in, all the jobs across your organization that were created after the opt-in are visible. When you opt out, AWS Backup keeps the jobs in the aggregated view for 30 days (from reaching a terminus state). Created jobs after the opt-out are not visible and do not show any newly created backup jobs. For opt-in instructions, see Enabling cross-account management.

To monitor multiple accounts

1. Open the AWS Backup console at https://console.aws.amazon.com/backup/. You must sign in using the credentials of your management account.
2. In the left navigation pane, choose **Settings** to open the cross-account management page.
3. In the **Cross-account monitoring** section, choose **Enable**.

This enables you to monitor the backup and restore activities of all accounts in your organization from your management account.

4. In the left navigation pane, choose **Cross-account monitoring**.
5. On the **Cross-account monitoring** page, choose the **Backup jobs**, **Restore jobs**, or **Copy jobs** tab to see all the jobs created in all your accounts. You can see each of these jobs by AWS account ID, and you can see all the jobs in a particular account.
6. In the search box, you can filter the jobs by **Account ID**, **Status**, or **Job ID**.

   For example, you can choose the **Backup jobs** tab and see all backup jobs created in all your accounts. You can filter the list by **Account ID** and see all the backup jobs created in that account.

Resource opt-in rules

If a member account's backup plan was created by an Organizations-level backup policy, the AWS Backup opt-in settings for the Organizations management account will override the opt-in settings in that member account, but only for that backup plan.

If the member account also has local-level backup plans created by users, those backup plans will follow the opt-in settings in the member account, without reference to the Organizations management account's opt-in settings.

Defining policies, policy syntax, and policy inheritance

The following topics are documented in the *AWS Organizations User Guide*.

- **Backup policies** – See Backup policies.
- **Policy syntax** – See Backup policy syntax and examples.
- **Inheritance for management policy types** – See Inheritance for management policy types.

AWS Backup and CloudFormation

In general

With CloudFormation, you can provision and manage your AWS resources in a safe, repeatable manner using templates that you create. You can use CloudFormation templates and StackSets to manage your backup plans, backup resource selections, and backup vaults. For information about using CloudFormation, see How Does CloudFormation Work? in the *AWS CloudFormation User Guide*.

Before you create your CloudFormation template or StackSet, consider the following:

- Create separate templates for your backup plans and your backup vaults. You can only delete backup vaults that are empty. You can't delete a stack that includes backup vaults if they contain recovery points.
- Verify you have a service role available before you create your stack. The AWS Backup default service role is created for you the first time you assign resources to a backup plan. If you haven't assigned resources to your backup plan, do so before creating your stack. You can also specify a custom role that you create. For more information about roles, see IAM service roles.

Deploying a backup vault, backup plan, and resource assignment using CloudFormation

For a sample CloudFormation template that deploys a backup vault, backup plans, and resource assignment, see Assign AWS Backup resources through CloudFormation.

Deploying backup plans using CloudFormation

For sample CloudFormation templates that deploy backup plans, see the section called "CloudFormation templates for backup plans".

Deploying AWS Backup Audit Manager frameworks and report plans using CloudFormation

For sample CloudFormation templates that deploy AWS Backup Audit Manager frameworks and report plans, see the section called "Using CloudFormation to deploy AWS Backup Audit Manager resources".

Deploying backup plans across accounts using CloudFormation

You can use AWS CloudFormation StackSets across multiple accounts in an AWS Organization. Sample templates are available in the CloudFormation User Guide.

An excellent starting point and reference is the publication Automate centralized backup at scale across AWS services using AWS Backup. With Ibukun Oyewumi and Sabith Venkitachalapathy (Jul. 2021).

Learning more about CloudFormation

For information about using CloudFormation with AWS Backup, see AWS Backup Resource Type Reference in the *AWS CloudFormation User Guide*.

For information about controlling access to AWS service resources when using CloudFormation, see Controlling Access with AWS Identity and Access Management in the *AWS CloudFormation User Guide*.

AWS Backup network

AWS Backup endpoints

AWS Backup offers both public and private endpoints for your connectivity needs. For these endpoints, AWS Backup supports both Internet Protocols version 4 (IPv4) and version 6 (IPv6) for resource types that support IPv6.

The newer public endpoint `backup.[Region].api.aws` has dual-stack capabilities and can resolve either or both IPv4 endpoints and IPv6 endpoints. When you make a request to a dual-stack AWS Backup API endpoint, the endpoint will resolve to the address determined by the configuration of the protocol used by your network and client.

The older endpoint `backup.[Region].amazonaws.com` can be used for calls that reference only IPv4.

You can view the public service endpoints for AWS Backup in the Amazon Web Services General Reference. You can view the steps for setting up private endpoints in AWS Backup through VPC.

AWS Backup through VPC endpoints

You can establish a private connection between your virtual private cloud (VPC) and AWS Backup by creating an interface VPC endpoint. Interface endpoints are powered by AWS PrivateLink, a technology that enables you to access the AWS Backup API without using an internet gateway, NAT device, VPN connection, or Direct Connect connection. Instances in your VPC don't need public IP addresses to communicate with AWS Backup API endpoints. Your instances also don't need public IP addresses to use any of the available AWS Backup API and Backup gateway API operations.

For more information, see Access AWS services through AWS PrivateLink in the *AWS PrivateLink Guide*.

Considerations for Amazon VPC endpoints

All AWS Backup operations relevant to managing your resources are available from your VPC using AWS PrivateLink.

VPC endpoint policies are supported for Backup endpoints. By default, full access to Backup operations is allowed through the endpoint. Alternatively, you can associate a security group with the endpoint network interfaces to control traffic to AWS Backup through the interface endpoint.

You can select IPv4, IPv6, or dual stack when created an endpoint. You will receive the same DNS names (which will have both IPv4 and IPv6 addresses if you select dual stack).

Create an AWS Backup VPC endpoint

You can create a VPC endpoint for AWS Backup using either the Amazon VPC console or the AWS Command Line Interface (AWS CLI). For more information, see Create an interface endpoint in the *AWS PrivateLink Guide*.

PrivateLink endpoints use the same name structure of IPv4, though each endpoint can be configured for IPv4, IPv6, or dual stack.

Create a VPC endpoint for AWS Backup using the service name `com.amazonaws.region.backup`.

In China (Beijing) Region and China (Ningxia) Region, the service name should be `cn.com.amazonaws.region.backup`.

For Backup dataplane endpoints, use `com.amazonaws.region.backup-storage`. Only SAP Hana supports dataplane endpoints and may require a BackInt agent update.

For Backup gateway endpoints, use `com.amazonaws.region.backup-gateway`.

The following TCP ports must be allowed in the security group when creating a VPC endpoint for backup Gateway:

- TCP 443
- TCP 1026
- TCP 1027
- TCP 1028
- TCP 1031
- TCP 2222

Protocol	Port	Direction	Source	Destination	Usage
TCP	443 (HTTPS)	Outbound	Backup Gateway	AWS	For communication from Backup

Protocol	Port	Direction	Source	Destination	Usage
					Gateway to the AWS service endpoint

Use a VPC endpoint

If you enable private DNS for the endpoint, you can make API requests to AWS Backup with the VPC endpoint using its default DNS name for the AWS Region, for example `backup.us-east-1.api.aws`.

However, for the China (Beijing) Region and China (Ningxia) Region AWS Regions, API requests should be made with the VPC endpoint using `backup.cn-north-1.amazonaws.com.cn` and `backup.cn-northwest-1.amazonaws.com.cn`, respectively.

Creating a VPC endpoint policy

You can attach an endpoint policy to your VPC endpoint that controls access to the Amazon Backup API. The policy specifies:

- The principal that can perform actions.
- The actions that can be performed.
- The resources on which actions can be performed.

> **Important**
>
> When a non-default policy is applied to an interface VPC endpoint for AWS Backup, certain failed API requests, such as those failing from `RequestLimitExceeded`, might not be logged to AWS CloudTrail or Amazon CloudWatch.

For more information, see Control access to services using endpoint policies in the *AWS PrivateLink Guide*.

Example: VPC endpoint policy for AWS Backup actions

The following is an example of an endpoint policy for AWS Backup. When attached to an endpoint, this policy grants access to the listed AWS Backup actions for all principles on all resources.

```
{
  "Statement":[
    {
      "Action":"backup:*",
      "Effect":"Allow",
      "Principal":"*",
      "Resource":"*"
    }
  ]
}
```

Example: VPC endpoint policy that denies all access from a specified AWS account

The following VPC endpoint policy denies AWS account 123456789012 all access to resources using the endpoint. The policy allows all actions from other accounts.

JSON

```
{
  "Id":"Policy1645236617225",
  "Version":"2012-10-17",
  "Statement":[
    {
      "Sid":"Stmt1645236612384",
      "Action":"backup:*",
      "Effect":"Deny",
      "Resource":"*",
      "Principal":{
        "AWS":[
          "123456789012"
        ]
      }
    }
  ]
}
```

For more information about available API responses, see the API Guide.

Security in AWS Backup

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. To learn about the compliance programs that apply to AWS Backup, see AWS Services in Scope by Compliance Program.
- **Security in the cloud** – Your responsibility for AWS Backup includes, but is not limited to, the following. You are also responsible for other factors including the sensitivity of your data, your organization's requirements, and applicable laws and regulations.
  - Responding to communications you receive from AWS.
  - Managing the credentials you and your team use. For more information, see Identity and access management in AWS Backup.
  - Configuring your backup plans and resource assignments to reflect your organization’s data protection policies. For more information, see Managing backup plans.
  - Regularly testing your ability to find certain recovery points and restore them. For more information, see Working with backups.
  - Incorporating AWS Backup procedures in your organization’s disaster recovery and business continuity written procedures. For a start point, see Getting started with AWS Backup.
  - Ensuring that your employees are familiar with and have practiced using AWS Backup along with your organizational procedures in the event of an emergency. For more information, see the AWS Well-Architected Framework.

This documentation helps you understand how to apply the shared responsibility model when using AWS Backup. The following topics show you how to configure AWS Backup to meet your security and compliance objectives. You also learn how to use other AWS services that help you monitor and secure your AWS Backup resources.

Topics

- Compliance validation for AWS Backup
- Data protection in AWS Backup
- Identity and access management in AWS Backup
- Infrastructure security in AWS Backup
- Integrity of Data in AWS Backup
- Legal holds and AWS Backup
- Malware protection in AWS Backup
- Resilience in AWS Backup

Compliance validation for AWS Backup

Our new AWS sign-up experience is not designed for regulated workloads. If you're using our new AWS sign-up experience, but you want to use AWS for regulated workloads, you can sign up for AWS (advanced) or activate advanced features for your AWS environment.

To learn whether an AWS service is within the scope of specific compliance programs, see AWS services in Scope by Compliance Program and choose the compliance program that you are interested in. 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 AWS services is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. For more information about your compliance responsibility when using AWS services, see AWS Security Documentation.

Data protection in AWS Backup

AWS Backup conforms to the AWS shared responsibility model, which includes regulations and guidelines for data protection. AWS is responsible for protecting the global infrastructure that runs all the AWS services. AWS maintains control over data hosted on this infrastructure, including the security configuration controls for handling customer content and personal data. AWS customers

and AWS Partner Network (APN) partners, acting either as data controllers or data processors, are responsible for any personal data that they put in the AWS Cloud.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual user accounts with AWS Identity and Access Management (IAM). This helps ensure that 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 Secure Sockets Layer (SSL)/Transport Layer Security (TLS) to communicate with AWS resources.
- Use AWS encryption solutions, along with all default security controls within AWS services.

We strongly recommend that you never put sensitive identifying information, such as your customers' account numbers, into free-form fields such as a **Name** field. This includes when you work with AWS Backup or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into AWS Backup or other services might get picked up for inclusion in diagnostic logs. When you provide a URL to an external server, don't include credentials information in the URL to validate your request to that server.

For more information about data protection, see the AWS Shared Responsibility Model and GDPR blog post on the *AWS Security Blog*.

Encryption for backups in AWS Backup

Independent encryption

AWS Backup offers independent encryption for resource types that support full AWS Backup management. Independent encryption means that recovery points (backups) you create through AWS Backup can have an encryption method other than one determined by the source resource's encryption. For example, your backup of an Amazon S3 bucket can have a different encryption method than the source bucket you encrypted with Amazon S3 encryption. This encryption is controlled through the AWS KMS key configuration in the backup vault where your backup in stored.

Backups of resource types that are not fully managed by AWS Backup typically inherit the encryption settings from their source resource. You can configure these encryption settings

according to that service's instructions, such as Amazon EBS encryption in the *Amazon EBS User Guide*.

Your IAM role must have access to the KMS key being used to back up and restore the object. Otherwise the job is successful but the objects are not backed up or restored. The permissions in IAM policy and KMS key policy must be consistent. For more information, see Specifying KMS keys in IAM policy statements in the *AWS Key Management Service Developer Guide*.

The following table lists each supported resource type, how encryption is configured for backups, and whether independent encryption for backups is supported. When AWS Backup independently encrypts a backup, it uses the industry-standard AES-256 encryption algorithm. For more information about encryption in AWS Backup, see cross-Region and cross-account backup.

Resource type	How to configure encryption	Independent AWS Backup encryption
Amazon Simple Storage Service (Amazon S3)	Amazon S3 backups are encrypted using a AWS KMS (AWS Key Management Service) key associated with the backup vault. The AWS KMS key can either be a customer-managed key or an AWS-managed key associated with the AWS Backup service. AWS Backup encrypts all backups even if the source Amazon S3 buckets are not encrypted.	Supported
VMware virtual machines	VM backups are always encrypted. The AWS KMS encryption key for virtual machine backups is configure d in the AWS Backup vault in which the virtual machine backups are stored.	Supported

Resource type	How to configure encryption	Independent AWS Backup encryption
Amazon DynamoDB after enabling Advanced DynamoDB backup	DynamoDB backups are always encrypted. The AWS KMS encryption key for DynamoDB backups is configured in the AWS Backup vault that the DynamoDB backups are stored in.	Supported
Amazon DynamoDB without enabling Advanced DynamoDB backup	DynamoDB backups are automatically encrypted with the same encryption key that was used to encrypt the source DynamoDB table. Snapshots of unencrypted DynamoDB tables are also unencrypted. In order for AWS Backup to create a backup of an encrypted DynamoDB table, you must add the permissions `kms:Decrypt` and `kms:GenerateDataKey` to the IAM role used for backup. Alternately, you can use the AWS Backup default service role.	Not supported

Resource type	How to configure encryption	Independent AWS Backup encryption
Amazon Elastic File System (Amazon EFS)	Amazon EFS backups are always encrypted. The AWS KMS encryption key for Amazon EFS backups is configured in the AWS Backup vault that the Amazon EFS backups are stored in.	Supported
Amazon Elastic Block Store (Amazon EBS)	By default, Amazon EBS backups are either encrypted using the key that was used to encrypt the source volume, or they are unencrypted. During restore, you can choose to override the default encryption method by specifying a KMS key.	Not supported
Amazon Elastic Compute Cloud (Amazon EC2) AMIs	AMIs are unencrypted. EBS snapshots are encrypted by the default encryption rules for EBS backups (see entry for EBS). EBS snapshots of data and root volumes can be encrypted and attached to an AMI.	Not supported

Resource type	How to configure encryption	Independent AWS Backup encryption
Amazon Relational Database Service (Amazon RDS)	Amazon RDS snapshots are automatically encrypted with the same encryption key that was used to encrypt the source Amazon RDS database. Snapshots of unencrypted Amazon RDS databases are also unencrypted.	Not supported
Amazon Aurora	Aurora cluster snapshots are automatically encrypted with the same encryption key that was used to encrypt the source Amazon Aurora cluster. Snapshots of unencrypt ed Aurora clusters are also unencrypted.	Not supported

Resource type	How to configure encryption	Independent AWS Backup encryption
AWS Storage Gateway	Storage Gateway snapshots are automatically encrypted with the same encryptio n key that was used to encrypt the source Storage Gateway volume. Snapshots of unencrypted Storage Gateway volumes are also unencrypted. You don't need to use a customer managed key across all services to enable Storage Gateway. You only need to copy the Storage Gateway backup to a vault that configured a KMS key. This is because Storage Gateway does not have a service-specific AWS KMS managed key.	Not supported
Amazon FSx	Encryption features for Amazon FSx file systems differ based on the underlyin g file system. To learn more about your particular Amazon FSx file system, see the appropriate FSx User Guide.	Not supported

Resource type	How to configure encryption	Independent AWS Backup encryption
Amazon DocumentDB	Amazon DocumentDB cluster snapshots are automatically encrypted with the same encryption key that was used to encrypt the source Amazon DocumentDB cluster. Snapshots of unencrypted Amazon DocumentDB clusters are also unencrypted.	Not supported
Amazon Neptune	Neptune cluster snapshots are automatically encrypted with the same encryption key that was used to encrypt the source Neptune cluster. Snapshots of unencrypted Neptune clusters are also unencrypted.	Not supported
Amazon Timestream	Timestream table snapshot backups are always encrypted . The AWS KMS encryption key for Timestream backups is configured in the backup vault in which the Timestrea m backups are stored.	Supported

Resource type	How to configure encryption	Independent AWS Backup encryption
Amazon Redshift	Amazon Redshift clusters are automatically encrypted with the same encryptio n key that was used to encrypt the source Amazon Redshift cluster. Snapshots of unencrypted Amazon Redshift clusters are also unencrypted.	Not supported
Amazon Redshift Serverless	Redshift Serverless snapshots are automatically encrypted with the same encryption key that was used to encrypt the source.	Not supported
CloudFormation	CloudFormation backups are always encrypted. The AWS KMS encryption key for CloudFormation backups is configured in the AWS Backup vault in which the CloudForm ation backups are stored.	Supported
SAP HANA databases on Amazon EC2 instances	SAP HANA database backups are always encrypted. The AWS KMS encryption key for SAP HANA database backups is configured in the AWS Backup vault in which the database backups are stored.	Supported

> **Tip**
>
> AWS Backup Audit Manager helps you automatically detect unencrypted backups.

Encryption for copies of a backup to a different account or AWS Region

When you copy your backups across accounts or Regions, AWS Backup automatically encrypts those copies for most resource types, even if the original backup is unencrypted. AWS Backup encrypts the copy using the KMS key of the target vault.

Before you copy a backup from one account to another (cross-account copy job) or copy a backup from one Region to another (cross-Region copy job), note the following conditions, many of which depend on whether the resource type in the backup (recovery point) is fully managed by AWS Backup or not fully managed.

- A copy of a backup to another AWS Region is encrypted using the key of the destination vault.
- For a copy of a recovery point (backup) of a resource that is **fully managed by AWS Backup**, you can choose to encrypt it with a customer managed key (CMK) or an AWS Backup managed key (`aws/backup`).

  For a copy of a recovery point of a resource that is **not fully managed** by AWS Backup, the key associated to the destination vault must be a CMK or the managed key of the service that owns the underlying resource. For example, if you are copying an EC2 instance, a Backup managed key cannot be used. Instead, a CMK or Amazon EBS KMS key (`aws/ebs`) must be used to avoid copy job failure.
- Cross-account copy with AWS managed keys isn't supported for resources that aren't fully managed by AWS Backup. The key policy of an AWS managed key is immutable, which prevents copying the key across accounts. If your resources are encrypted with AWS managed keys and you want to perform a cross-account copy, you may change the encryption keys to a customer managed key, which can be used for cross-account copying. Or, you can follow the instructions in Protecting encrypted Amazon RDS instances with cross-account and cross-Region backups to continue using AWS managed keys.
- Copies of unencrypted Amazon Aurora, Amazon DocumentDB, and Amazon Neptune clusters are also unencrypted.

AWS Backup permissions, grants, and deny statements

To help avoid failed jobs, you can examine the AWS KMS key policy to ensure it has required permissions and does not have any deny statements that prevent successful operations.

Failed jobs can occur due to either one or more Deny statements applied to the KMS key or due to a grant revoked for the key.

In an AWS managed access policy such as `AWSBackupFullAccess`, there are Allow actions that permit AWS Backup to interface with AWS KMS to create a grant on a KMS key on a customer's behalf as part backup, copy, and storage operations.

At a minimum, the key policy requires the following permissions:

- `kms:createGrant`
- `kms:generateDataKey`
- `kms:decrypt`

If Deny policies are necessary, you will need to allowlist the required roles for backup and restore operations.

These elements can look like:

```
{
    "Version":"2012-10-17",
    "Statement": [
      {
          "Sid": "KmsPermissions",
          "Effect": "Allow",
          "Principal": {
              "AWS": "arn:aws:iam::123456789012:root"
          },
          "Action": [
              "kms:ListKeys",
              "kms:DescribeKey",
              "kms:GenerateDataKey",
              "kms:Decrypt",
              "kms:ListAliases"
          ],
          "Resource": "*"
      },
```

```
        {
            "Sid": "KmsCreateGrantPermissions",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::123456789012:root"
            },
            "Action": [
                "kms:CreateGrant"
            ],
            "Resource": "*",
            "Condition": {
                "ForAnyValue:StringEquals": {
                    "kms:EncryptionContextKeys": "aws:backup:backup-vault"
                },
                "Bool": {
                    "kms:GrantIsForAWSResource": true
                },
                "StringLike": {
                    "kms:ViaService": "backup.*.amazonaws.com"
                }
            }
        }
    ]
}
```

For customer managed keys, ensure these permissions are included in the key policy. You can view and update the key policy using the steps below.

For AWS managed keys (such as `aws/backup` or `aws/ebs`), these permissions are already included in the key policy by AWS and cannot be modified by the customer.

1. Ensure required permissions are part of KMS key policy
    a. Run KMS CLI `get-key-policy` (`kms:GetKeyPolicy`) to view the key policy attached to the specified KMS key.
    b. Review the returned permissions.
2. Ensure there are no Deny statements that affect operations
    a. Run (or re-run) CLI `get-key-policy` (`kms:GetKeyPolicy`) to view the key policy attached to the specified KMS key.
    b. Review the policy.

c. Remove relevant Deny statements from the KMS key policy.

3. If needed, run `kms:put-key-policy` to replace or update key policy with revised permissions and removed Deny statements.

Additionally, the key associated with the role initiating a cross-Region copy job must have `"kms:ResourceAliases": "alias/aws/backup"` in the `DescribeKey` permission.

Virtual machine hypervisor credential encryption

Virtual machines managed by a hypervisor use AWS Backup Gateway to connect on-premises systems to AWS Backup. It is important that hypervisors have the same robust and reliable security. This security can be achieved by encrypting the hypervisor, either by AWS owned keys or by customer managed keys.

AWS owned and customer managed keys

AWS Backup provides encryption for hypervisor credentials to protect sensitive customer login information using **AWS owned encryption** keys. You have the option of using **customer managed keys** instead.

By default, the keys used to encrypt credentials in your hypervisor are **AWS owned keys**. AWS Backup uses these keys to automatically encrypt hypervisor credentials. You can neither view, manage, or use AWS owned keys, nor can you audit their use. However, you don't have to take any action or change any programs to protect the keys that encrypt your data. For more information, see AWS owned keys in the *AWS KMS Developer Guide*.

Alternatively, credentials can be encrypted using *Customer managed keys*. AWS Backup supports the use of symmetric customer-managed keys that you create, own, and manage to perform your encryption. Because you have full control of this encryption, you can perform tasks such as:

- Establishing and maintaining key policies
- Establishing and maintaining IAM policies and grants
- Enabling and disabling key policies
- Rotating key cryptographic material
- Adding tags
- Creating key aliases
- Scheduling keys for deletion

When you use a customer managed key, AWS Backup validates whether your role has permission to decrypt using this key (prior to a backup or restore job being run). You must add the `kms:Decrypt` action to the role used to start a backup or restore job.

Because the `kms:Decrypt` action cannot be added to the default backup role, you must use a role other than the default backup role to use customer managed keys.

For more information, see customer managed keys in the *AWS Key Management Service Developer Guide*.

Grant required when using customer managed keys

AWS KMS requires a grant to use your customer managed key. When you import a hypervisor configuration encrypted with a customer managed key, AWS Backup creates a grant on your behalf by sending a `CreateGrant` request to AWS KMS. AWS Backup uses grants to access a KMS key in a customer account.

You can revoke access to the grant, or remove AWS Backup's access to the customer managed key at any time. If you do, all your gateways associated with your hypervisor can no longer access the hypervisor's username and password encrypted by the customer managed key, which will affect your backup and restore jobs. Specifically, backup and restore jobs you perform on the virtual machines in this hypervisor will fail.

Backup gateway uses the `RetireGrant` operation to remove a grant when you delete a hypervisor.

Monitoring encryption keys

When you use an AWS KMS customer managed key with your AWS Backup resources, you can use AWS CloudTrail or Amazon CloudWatch Logs to track requests that AWS Backup sends to AWS KMS.

Look for AWS CloudTrail events with the following `"eventName"` fields to for monitor AWS KMS operations called by AWS Backup to access data encrypted by your customer managed key:

- `"eventName": "CreateGrant"`
- `"eventName": "Decrypt"`
- `"eventName": "Encrypt"`
- `"eventName": "DescribeKey"`

Identity and access management in AWS Backup

Access to AWS Backup requires credentials. Those credentials must have permissions to access AWS resources, such as an Amazon DynamoDB database or an Amazon EFS file system. Moreover, recovery points created by AWS Backup for some AWS Backup-supported services cannot be deleted using the source service (such as Amazon EFS). You can delete those recovery points using AWS Backup.

The following sections provide details on how you can use AWS Identity and Access Management (IAM) and AWS Backup to help secure access to your resources.

> **Warning**
>
> AWS Backup uses the same IAM role that you chose when assigning resources to manage your recovery point lifecycle. If you delete or modify that role, AWS Backup cannot manage your recovery point lifecycle. When this occurs, it will attempt to use a service-linked role to manage your lifecycle. In a small percentage of cases, this might also not work, leaving `EXPIRED` recovery points on your storage, which might create unwanted costs. To delete `EXPIRED` recovery points, manually delete them using the procedure in Deleting backups.

Topics

- Authentication
- Access control
- IAM service roles
- Managed policies for AWS Backup
- Using service-linked roles for AWS Backup
- Cross-service confused deputy prevention

Authentication

Access to AWS Backup or the AWS services that you are backing up requires credentials that AWS can use to authenticate your requests. You can access AWS as any of the following types of identities:

- **AWS account root user** – When you sign up for AWS, you provide an email address and password that is associated with your AWS account. This is your *AWS account root user*. Its credentials provide complete access to all of your AWS resources.

  > **Important**
  >
  > For security reasons, we recommend that you use the root user only to create an *administrator*. The administrator is an *IAM user* with full permissions to your AWS account. You can then use this admin user to create other IAM users and roles with limited permissions. For more information, see IAM Best Practices and Creating Your First IAM Admin User and Group in the *IAM User Guide*.

- **IAM user** – An IAM user is an identity within your AWS account that has specific custom permissions (for example, permissions to create a backup vault to store your backups in). You can use an IAM user name and password to sign in to secure AWS webpages like the AWS Management Console, AWS Discussion Forums, or the AWS Support Center.

  In addition to a user name and password, you can also generate access keys for each user. You can use these keys when you access AWS services programmatically, either through one of the several SDKs or by using the AWS Command Line Interface (AWS CLI). The SDK and AWS CLI tools use the access keys to cryptographically sign your request. If you don't use the AWS tools, you must sign the request yourself. For more information about authenticating requests, see Signature Version 4 Signing Process in the *AWS General Reference*.

- **IAM role** – An IAM role is another IAM identity that you can create in your account that has specific permissions. It is similar to an IAM user, but it is not associated with a specific person. An IAM role enables you to obtain temporary access keys that can be used to access AWS services and resources. IAM roles with temporary credentials are useful in the following situations:

  - Federated user access – Instead of creating an IAM user, you can use pre-existing user identities from Directory Service, your enterprise user directory, or a web identity provider. These are known as *federated users*. AWS assigns a role to a federated user when access is requested through an identity provider. For more information about federated users, see Federated Users and Roles in the *IAM User Guide*.
  - Cross-account administration – You can use an IAM role in your account to grant another AWS account permissions to administer your account's resources. For an example, see Tutorial: Delegate Access Across AWS accounts Using IAM Roles in the *IAM User Guide*.

- AWS service access – You can use an IAM role in your account to grant an AWS service permissions to access your account's resources. For more information, see Creating a Role to Delegate Permissions to an AWS Service in the *IAM User Guide*.
- Applications running on Amazon Elastic Compute Cloud (Amazon EC2) – You can use an IAM role to manage temporary credentials for applications running on an Amazon EC2 instance and making AWS API requests. This is preferable to storing access keys within the EC2 instance. To assign an AWS role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and enables programs running on the EC2 instance to get temporary credentials. For more information, see Using an IAM Role to Grant Permissions to Applications Running on Amazon EC2 Instances in the *IAM User Guide*.

Access control

You can have valid credentials to authenticate your requests, but unless you have the appropriate permissions, you can't access AWS Backup resources such as backup vaults. You also can't back up AWS resources such as Amazon Elastic Block Store (Amazon EBS) volumes.

Every AWS resource is owned by an AWS account, and permissions to create or access a resource are governed by permissions policies. An account administrator can attach permissions policies to AWS Identity and Access Management (IAM) identities (that is, users, groups, and roles). And some services also support attaching permissions policies to resources.

An *account administrator* (or administrator user) is a user with administrator permissions. For more information, see IAM Best Practices in the *IAM User Guide*.

When granting permissions, you decide who is getting the permissions, the resources they get permissions for, and the specific actions that you want to allow on those resources.

The following sections cover how access policies work and how you use them to protect your backups.

Topics

- Resources and operations
- Resource ownership
- Specifying policy elements: actions, effects, and principals
- Specifying conditions in a policy

- API permissions: actions, resources, and conditions reference
- Copy tags permissions
- Access policies

Resources and operations

A resource is an object that exists within a service. AWS Backup resources include backup plans, backup vaults, and backups. *Backup* is a general term that refers to the various types of backup resources that exist in AWS. For example, Amazon EBS snapshots, Amazon Relational Database Service (Amazon RDS) snapshots, and Amazon DynamoDB backups are all types of backup resources.

In AWS Backup, backups are also referred to as *recovery points*. When using AWS Backup, you also work with the resources from other AWS services that you are trying to protect, such as Amazon EBS volumes or DynamoDB tables. These resources have unique Amazon Resource Names (ARNs) associated with them. ARNs uniquely identify AWS resources. You must have an ARN when you need to specify a resource unambiguously across all of AWS, such as in IAM policies or API calls.

The following table lists resources, subresources, ARN format, and an example unique ID.

AWS Backup resource ARNs

Resource type	ARN format	Example unique ID
Backup plan	`arn:aws:backup:`*`region`*`:`*`account-id`*`:backup-plan:*`	
Backup vault	`arn:aws:backup:`*`region`*`:`*`account-id`*`:backup-vault:*`	
Recovery point for Amazon EBS	`arn:aws:ec2:`*`region`*`::snapshot/*`	`snapshot/snap-05f426fd8kdjb4224`
Recovery point for Amazon EC2 images	`arn:aws:ec2:`*`region`*`::image/ami-*`	`image/ami-1a2b3e4f5e6f7g890`

Resource type	ARN format	Example unique ID
Recovery point for Amazon RDS	arn:aws:rds: *region*:*account-id* :snapshot:awsbackup:*	awsbackup:job-be59cf2a-2343-4402-bd8b-226993d23453
Recovery point for Aurora	arn:aws:rds: *region*:*account-id* :cluster-snapshot:awsbackup:*	awsbackup:job-be59cf2a-2343-4402-bd8b-226993d23453
Recovery point for Aurora DSQL	arn:*aws-partition* :backup:*region*:*accoid* :recovery-point:*recovery-point-id*	arn:aws:backup:us-east-1:012345678901:recovery-point:8a92c3f1-b475-4d9e-95e6-7c138f2d4b0a
Recovery point for Storage Gateway	arn:aws:ec2: *region*::snapshot/*	snapshot/snap-0d40e49137e31d9e0
Recovery point for DynamoDB without Advanced DynamoDB backup	arn:aws:dynamodb: *region*:*account-id* :table/*/backup/*	table/MyDynamoDBTable/backup/01547087347000-c8b6kdk3
Recovery point for DynamoDB with Advanced DynamoDB backup enabled	arn:aws:backup: *region*:*account-id* :recovery-point:*	12a34a56-7bb8-901c-cd23-4567d8e9ef01
Recovery point for Amazon EFS	arn:aws:backup: *region*:*account-id* :recovery-point:*	d99699e7-e183-477e-bfcd-ccb1c6e5455e
Recovery point for Amazon FSx	arn:aws:fsx: *region*:*account-id* :backup/backup-*	backup/backup-1a20e49137e31d9e0

Resource type	ARN format	Example unique ID
Recovery point for virtual machine	arn:aws:backup: *region*:*account-id* :recovery-point:*	1801234a-5b6b-7dc8-8032-836f7ffc623b
Recovery point for Amazon S3 continuous backup	arn:aws:backup: *region*:*account-id* :recovery-point:*	*amzn-s3-demo-bucket* -5ec207d0
Recovery point for S3 periodic backup	arn:aws:backup: *region*:*account-id* :recovery-point:*	*amzn-s3-demo-bucket* -20211231900000-5ec207d0
Recovery point for Amazon DocumentDB	arn:aws:rds: *region*:*account-id* :cluster-snapshot:awsbackup:*	awsbackup:job-ab12cd3e-4567-8901-fg1h-234567i89012
Recovery point for Neptune	arn:aws:rds: *region*:*account-id* :cluster-snapshot:awsbackup:*	awsbackup:job-ab12cd3e-4567-8901-fg1h-234567i89012
Recovery point for Amazon Redshift	arn:aws:redshift: *region*:*account-id* :snapshot: *resource*/awsbackup:*	awsbackup:job-ab12cd3e-4567-8901-fg1h-234567i89012
Recovery point for Amazon Redshift Serverless	arn:aws:redshift-serverless: *region*:*account-id* :snapshot: *resource*/awsbackup:*	awsbackup:job-ab12cd3e-4567-8901-fg1h-234567i89012

Resource type	ARN format	Example unique ID
Recovery point for Amazon Timestream	arn:aws:backup: *region*:*account-id* :recovery-point:*	recovery-point:1a2b3cde-f405-6789-012g-3456hi789012_beta
Recovery point for AWS CloudFormation template	arn:aws:backup: *region*:*account-id* :recovery-point:*	recovery-point:1a2b3cde-f405-6789-012g-3456hi789012
Recovery point for SAP HANA database on Amazon EC2 instance	arn:aws:backup: *region*:*account-id* :recovery-point:*	recovery-point:1a2b3cde-f405-6789-012g-3456hi789012
Backup access point	arn:aws:backup: *region*:*account-id* :accesspoint/*	accesspoint/my-access-point

Resources that support full AWS Backup management all have recovery points in the format arn:aws:backup:*region*:*account-id:*:recovery-point:*. making it easier for you to apply permissions policies to protect those recovery points. To see which resources support full AWS Backup management, see that section of the Feature availability by resource table.

AWS Backup provides a set of operations to work with AWS Backup resources. For a list of available operations, see AWS Backup Operations.

Resource ownership

The AWS account owns the resources that are created in the account, regardless of who created the resources. Specifically, the resource owner is the AWS account of the principal entity (that is, the AWS account root user, an IAM user, or an IAM role) that authenticates the resource creation request. The following examples illustrate how this works:

- If you use the AWS account root user credentials of your AWS account to create a backup vault, your AWS account is the owner of the vault.

- If you create an IAM user in your AWS account and grant permissions to create a backup vault to that user, the user can create a backup vault. However, your AWS account, to which the user belongs, owns the backup vault resource.
- If you create an IAM role in your AWS account with permissions to create a backup vault, anyone who can assume the role can create a vault. Your AWS account, to which the role belongs, owns the backup vault resource.

Specifying policy elements: actions, effects, and principals

For each AWS Backup resource (see Resources and operations), the service defines a set of API operations (see Operations). To grant permissions for these API operations, AWS Backup defines a set of actions that you can specify in a policy. Performing an API operation can require permissions for more than one action.

The following are the most basic policy elements:

- Resource – In a policy, you use an Amazon Resource Name (ARN) to identify the resource to which the policy applies. For more information, see Resources and operations.
- Action – You use action keywords to identify resource operations that you want to allow or deny.
- Effect – You specify the effect when the user requests the specific action—this can be either allow or deny. If you don't explicitly grant access to (allow) a resource, access is implicitly denied. You can also explicitly deny access to a resource, which you might do to make sure that a user cannot access it, even if a different policy grants access.
- Principal – In identity-based policies (IAM policies), the user that the policy is attached to is the implicit principal. For resource-based policies, you specify the user, account, service, or other entity that you want to receive permissions (applies to resource-based policies only).

To learn more about IAM policy syntax and descriptions, see IAM JSON Policy Reference in the *IAM User Guide*.

For a table showing all of the AWS Backup API actions, see API permissions: actions, resources, and conditions reference.

Specifying conditions in a policy

When you grant permissions, you can use the IAM policy language to specify the conditions when a policy should take effect. For example, you might want a policy to be applied only after a specific

date. For more information about specifying conditions in a policy language, see Condition in the *IAM User Guide*.

AWS supports global condition keys and service-specific condition keys. To see all global condition keys, see AWS global condition context keys in the *IAM User Guide*.

AWS Backup defines its own set of condition keys. To see a list of AWS Backup condition keys, see Condition keys for AWS Backup in the *Service Authorization Reference*.

For restore testing selections, AWS Backup passes the IAM role to the service principal `restore-testing.backup.amazonaws.com`, not `backup.amazonaws.com`. If you use the `iam:PassedToService` condition key for `CreateRestoreTestingSelection` or `UpdateRestoreTestingSelection`, set the value to `restore-testing.backup.amazonaws.com`. For more information, see Using roles for restore testing.

API permissions: actions, resources, and conditions reference

When you are setting up Access control and writing a permissions policy that you can attach to an IAM identity (identity-based policies), you can use the following list as a reference. The list includes each AWS Backup API operation, the corresponding actions for which you can grant permissions to perform the action, and the AWS resource for which you can grant the permissions. You specify the actions in the policy's `Action` field, and you specify the resource value in the policy's `Resource` field. If `Resource` field is blank, you can use the wildcard (*) to include all resources.

You can use AWS-wide condition keys in your AWS Backup policies to express conditions. For a complete list of AWS-wide keys, see Available Keys in the *IAM User Guide*.

[1] Uses the existing vault access policy.

[2] See AWS Backup resource ARNs for resource-specific recovery point ARNs.

[3] `StartRestoreJob` must have the key-value pair in the metadata for the resource. To get the metadata of the resource, call the `GetRecoveryPointRestoreMetadata` API.

For more information, see Actions, resources, and condition keys for AWS Backup in the *Service Authorization Reference*.

Copy tags permissions

When AWS Backup performs a backup or copy job, it attempts to copy the tags from your source resource (or recovery point in the case of copy) to your recovery point.

> **Note**
>
> AWS Backup does **not** natively copy tags during restore jobs. For an event-driven architecture that will copy tags during restore jobs, see How to retain resource tags in AWS Backup restore jobs.

During a backup or copy job, AWS Backup aggregates the tags you specify in your backup plan (or copy plan, or on-demand backup) with the tags from your source resource. However, AWS enforces a limit of 50 tags per resource, which AWS Backup cannot exceed. When a backup or copy job aggregates tags from the plan and the source resource, it might discover more than 50 total tags, it will be unable to complete the job, and will fail the job. This is consistent with AWS-wide tagging best practices.

- Your resource has more than 50 tags after aggregating your backup job tags with your source resource tags. AWS supports up to 50 tags per resource.
- The IAM role you provide to AWS Backup lacks permissions to read the source tags or set the destination tags. For more information and sample IAM role policies, see Managed Policies.

You can use your backup plan (tags added to recovery points) to create tags that contradict your source resource tags. When the two conflict, the tags from your backup plan take precedence. Use this technique if you prefer not to copy a tag value from your source resource. Specify the same tag key, but different or empty value, using your backup plan.

Permissions Required to assign tags to backups

Resource type	Required permission
Amazon EFS file system	`elasticfilesystem:DescribeTags`
Amazon FSx file system	`fsx:ListTagsForResource`
Amazon RDS database and Amazon Aurora cluster	`rds:AddTagsToResource` `rds:ListTagsForResource`
Storage Gateway volume	`storagegateway:ListTagsForResource`

Resource type	Required permission
Amazon EC2 instance and Amazon EBS volume	`EC2:CreateTags` `EC2:DescribeTags`

DynamoDB does not support assigning tags to backups unless you first enable Advanced DynamoDB backup.

When an Amazon EC2 backup creates an Image Recovery Point and a set of snapshots, AWS Backup copies tags to the resulting AMI. AWS Backup also copies the tags from the volumes associated with the Amazon EC2 instance to the resulting snapshots.

Access policies

A *permissions* policy describes who has access to what. Policies attached to an IAM identity are referred to as *identity-based* policies (IAM policies). Policies attached to a resource are referred to as *resource-based* policies. AWS Backup supports both identity-based policies and resource-based policies.

> **Note**
>
> This section discusses using IAM in the context of AWS Backup. It doesn't provide detailed information about the IAM service. For complete IAM documentation, see What Is IAM? in the *IAM User Guide*. For information about IAM policy syntax and descriptions, see IAM JSON Policy Reference in the *IAM User Guide*.

Identity-based policies (IAM policies)

Identity-based policies are policies that you can attach to IAM identities, such as users or roles. For example, you can define a policy that allows a user to view and back up AWS resources, but prevents them from restoring backups.

For more information about users, groups, roles, and permissions, see Identities (Users, Groups, and Roles) in the *IAM User Guide*.

For information about how to use IAM policies to control access to backups, see Managed policies for AWS Backup.

Resource-based policies

AWS Backup supports resource-based access policies for backup vaults. This enables you to define an access policy that can control which users have what kind of access to any of the backups organized in a backup vault. Resource-based access policies for backup vaults provide an easy way to control access to your backups.

Backup vault access policies control user access when you use AWS Backup APIs. Some backup types, such as Amazon Elastic Block Store (Amazon EBS) and Amazon Relational Database Service (Amazon RDS) snapshots, can also be accessed using those services' APIs. You can create separate access policies in IAM that control access to those APIs in order to fully control access to backups.

To learn how to create an access policy for backup vaults, see Vault access policies.

IAM service roles

An AWS Identity and Access Management (IAM) role is similar to a user, in that it is an AWS identity with permissions policies that determine what the identity can and cannot do in AWS. However, instead of being uniquely associated with one person, a role is intended to be assumable by anyone who needs it. A service role is a role that an AWS service assumes to perform actions on your behalf. As a service that performs backup operations on your behalf, AWS Backup requires that you pass it a role to assume when performing backup operations on your behalf. For more information about IAM roles, see IAM Roles in the *IAM User Guide*.

The role that you pass to AWS Backup must have an IAM policy with the permissions that enable AWS Backup to perform actions associated with backup operations, such as creating, restoring, or expiring backups. Different permissions are required for each of the AWS services that AWS Backup supports. The role must also have AWS Backup listed as a trusted entity, which enables AWS Backup to assume the role.

When you assign resources to a backup plan, or if you perform an on-demand backup, copy, or restore, you must pass a service role that has access to perform the underlying operations on the specified resources. AWS Backup uses this role to create, tag, and delete resources in your account.

Using AWS roles to control access to backups

You can use roles to control access to your backups by defining narrowly scoped roles and by specifying who can pass that role to AWS Backup. For example, you could create a role that only grants permissions to back up Amazon Relational Database Service (Amazon RDS) databases and

only grant Amazon RDS database owners permission to pass that role to AWS Backup. AWS Backup provides several predefined managed policies for each of the supported services. You can attach these managed policies to roles that you create. This makes it easier to create service-specific roles that have the correct permissions that AWS Backup needs.

For more information about AWS managed policies for AWS Backup, see Managed policies for AWS Backup.

Default service role for AWS Backup

When using the AWS Backup console for the first time, you can choose to have AWS Backup create a default service role for you. This role has the permissions that AWS Backup needs to create and restore backups on your behalf.

> **Note**
>
> The default role is automatically created when you use the AWS Management Console. You can create the default role using the AWS Command Line Interface (AWS CLI), but it must be done manually.

If you prefer to use custom roles, such as separate roles for different resource types, you can also do that and pass your custom roles to AWS Backup. To view examples of roles that enable backup and restore for individual resource types, see the Customer managed policies table.

The default service role is named `AWSBackupDefaultServiceRole`. This service role contains two managed policies, AWSBackupServiceRolePolicyForBackup and AWSBackupServiceRolePolicyForRestores.

`AWSBackupServiceRolePolicyForBackup` includes an IAM policy that grants AWS Backup permissions to describe the resource being backed up, the ability to create, delete, describe, or add tags to a backup regardless of the AWS KMS key with which it is encrypted.

`AWSBackupServiceRolePolicyForRestores` includes an IAM policy that grants AWS Backup permissions to create, delete, or describe the new resource being created from a backup, regardless of the AWS KMS key with which it is encrypted. It also includes permissions to tag the newly created resource.

To restore an Amazon EC2 instance, you must launch a new instance.

Creating the default service role in the console

Specific actions you take in the AWS Backup Console create the AWS Backup default service role.

To create the AWS Backup default service role in your AWS account

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. To create the role for your account, either assign resources to a backup plan or create an on-demand backup.
   a. Create a backup plan and assign resources to the backup. See Create a backup plan.
   b. Alternatively, create an on-demand backup. See Create an on-demand backup.
3. Verify that you have created the `AWSBackupDefaultServiceRole` in your account by following these steps:
   a. Wait a few minutes. For more information, see Changes that I make are not always immediately visible in the *AWS Identity and Access Management User Guide.*
   b. Sign in to the AWS Management Console and open the IAM console at https://console.aws.amazon.com/iam/.
   c. In the left navigation menu, choose **Roles**.
   d. In the search bar, type `AWSBackupDefaultServiceRole`. If this selection exists, you have created the AWS Backup default role and completed this procedure.
   e. If `AWSBackupDefaultServiceRole` still does not appear, add the following permissions to either the IAM user or IAM role you use to access the console.

      JSON

```
{
  "Version":"2012-10-17",
  "Statement":[
    {
      "Effect":"Allow",
      "Action":[
        "iam:CreateRole",
        "iam:AttachRolePolicy",
        "iam:PassRole"
      ],
```

```
      "Resource":"arn:aws:iam::*:role/service-role/
AWSBackupDefaultServiceRole"
    },
    {
      "Effect":"Allow",
      "Action":[
        "iam:ListRoles"
      ],
      "Resource":"*"
    }
  ]
}
```

For China Regions, replace *aws* with *aws-cn*. For AWS GovCloud (US) Regions, replace *aws* with *aws-us-gov*.

f. If you cannot add permissions to your IAM user or IAM role, ask your administrator to manually create a role with a name *other than* `AWSBackupDefaultServiceRole` and attach that role to these managed policies:

- `AWSBackupServiceRolePolicyForBackup`
- `AWSBackupServiceRolePolicyForRestores`

Managed policies for AWS Backup

Managed policies are standalone identity-based policies that you can attach to multiple users, groups, and roles in your AWS account. When you attach a policy to a principal entity, you give the entity the permissions that are defined in the policy.

AWS managed policies are created and administered by AWS. You can't change the permissions defined in AWS managed policies. If AWS updates the permissions defined in an AWS managed policy, the update affects all principal identities (users, groups, and roles) that the policy is attached to.

Customer managed policies give you fine-grained controls to set access to backups in AWS Backup. For example, you can use them to give your database backup administrator access to Amazon RDS backups but not Amazon EFS ones.

For more information, see Managed policies in the *IAM User Guide*.

AWS managed policies

AWS Backup provides the following AWS managed policies for common use cases. These policies make it easier to define the right permissions and control access to your backups. There are two types of managed policies. One type is designed to be assigned to users to control their access to AWS Backup. The other type of managed policy is designed to be attached to roles that you pass to AWS Backup. The following table lists all the managed policies that AWS Backup provides and describes how they are defined. You can find these managed policies in the **Policies** section of the IAM console.

Policies

- AWSBackupAuditAccess
- AWSBackupAccessPointOperatorAccess
- AWSBackupDataTransferAccess
- AWSBackupFullAccess
- AWSBackupGatewayServiceRolePolicyForVirtualMachineMetadataSync
- AWSBackupGuardDutyRolePolicyForScans
- AWSBackupOperatorAccess
- AWSBackupOrganizationAdminAccess
- AWSBackupRestoreAccessForSAPHANA
- AWSBackupSearchOperatorAccess
- AWSBackupServiceLinkedRolePolicyForBackup
- AWSBackupServiceLinkedRolePolicyForBackupTest
- AWSBackupServiceRolePolicyForBackup
- AWSBackupServiceRolePolicyForItemRestores
- AWSBackupServiceRolePolicyForIndexing
- AWSBackupServiceRolePolicyForRestores
- AWSBackupServiceRolePolicyForS3Backup
- AWSBackupServiceRolePolicyForS3Restore
- AWSBackupServiceRolePolicyForScans
- AWSServiceRolePolicyForBackupReports
- AWSServiceRolePolicyForBackupRestoreTesting

AWSBackupAuditAccess

This policy grants permissions for users to create controls and frameworks that define their expectations for AWS Backup resources and activities, and to audit AWS Backup resources and activities against their defined controls and frameworks. This policy grants permissions to AWS Config and similar services to describe user expectations perform the audits.

This policy also grants permissions to deliver audit reports to Amazon S3 and similar services, and enables users to find and open their audit reports.

To view the permissions for this policy, see AWSBackupAuditAccess in the *AWS Managed Policy Reference*.

AWSBackupAccessPointOperatorAccess

Access point operators are users who need to access backup data in recovery points without performing a full restore. Access point operators have permissions to create, describe, delete, and list backup access points.

This policy contains the necessary permissions for those functions.

To view the permissions for this policy, see AWSBackupAccessPointOperatorAccess in the *AWS Managed Policy Reference*.

AWSBackupDataTransferAccess

This policy provides permissions for the AWS Backup storage plane data transfer APIs, allowing the AWS Backint agent to complete backup data transfer with the AWS Backup storage plane. You can attach this policy to roles assumed by Amazon EC2 instances running SAP HANA with the Backint agent.

To view the permissions for this policy, see AWSBackupDataTransferAccess in the *AWS Managed Policy Reference*.

AWSBackupFullAccess

The backup administrator has full access to AWS Backup operations, including creating or editing backup plans, assigning AWS resources to backup plans, and restoring backups. Backup administrators are responsible for determining and enforcing backup compliance by defining backup plans that meet their organization's business and regulatory requirements. Backup administrators also ensure that their organization's AWS resources are assigned to the appropriate plan.

To view the permissions for this policy, see AWSBackupFullAccess in the *AWS Managed Policy Reference*.

AWSBackupGatewayServiceRolePolicyForVirtualMachineMetadataSync

This policy provides Backup gateway permission to sync the metadata of Virtual Machines on your behalf

To view the permissions for this policy, see AWSBackupGatewayServiceRolePolicyForVirtualMachineMetadataSync in the *AWS Managed Policy Reference*.

AWSBackupGuardDutyRolePolicyForScans

This policy must be added to a new scanning role that grants Amazon GuardDuty permission to read and scan your backups. You'll need to attach this scanning role to your backup plan within the malware protection or scan settings. When AWS Backup initiates a scan, it passes this scanning role to Amazon GuardDuty.

To view the permissions for this policy, see AWSBackupGuardDutyRolePolicyForScans in the *AWS Managed Policy Reference*.

AWSBackupOperatorAccess

Backup operators are users that are responsible for ensuring the resources that they are responsible for are properly backed up. Backup operators have permissions to assign AWS resources to the backup plans that the backup administrator creates. They also have permissions to create on-demand backups of their AWS resources and to configure the retention period of on-demand backups. Backup operators do not have permissions to create or edit backup plans or to delete scheduled backups after they are created. Backup operators can restore backups. You can limit the resource types that a backup operator can assign to a backup plan or restore from a backup. You do this by allowing only certain service roles to be passed to AWS Backup that have permissions for a certain resource type.

To view the permissions for this policy, see AWSBackupOperatorAccess in the *AWS Managed Policy Reference*.

AWSBackupOrganizationAdminAccess

The organization administrator has full access to AWS Organizations operations, including creating, editing, or deleting backup policies, assigning backup policies to accounts and organizational units, and monitoring backup activities within the organization. Organization administrators are

responsible for protecting accounts in their organization by defining and assigning backup policies that meet their organization's business and regulatory requirements.

To view the permissions for this policy, see AWSBackupOrganizationAdminAccess in the *AWS Managed Policy Reference*.

AWSBackupRestoreAccessForSAPHANA

This policy provides AWS Backup permission to restore a backup of SAP HANA on Amazon EC2.

To view the permissions for this policy, see AWSBackupRestoreAccessForSAPHANA in the *AWS Managed Policy Reference*.

AWSBackupSearchOperatorAccess

The search operator role has access to create backup indexes and to create searches of indexed backup metadata.

This policy contains the necessary permissions for those functions.

To view the permissions for this policy, see AWSBackupSearchOperatorAccess in the *AWS Managed Policy Reference*.

AWSBackupServiceLinkedRolePolicyForBackup

This policy is attached to the service-linked role named AWSServiceRoleforBackup to allow AWS Backup to call AWS services on your behalf to manage your backups. For more information, see the section called “Backup and copy”.

To view the permissions for this policy, see AWSBackupServiceLinkedRolePolicyforBackup in the *AWS Managed Policy Reference*.

AWSBackupServiceLinkedRolePolicyForBackupTest

To view the permissions for this policy, see AWSBackupServiceLinkedRolePolicyForBackupTest in the *AWS Managed Policy Reference*.

AWSBackupServiceRolePolicyForBackup

Provides AWS Backup permissions to create backups of all supported resource types on your behalf.

To view the permissions for this policy, see AWSBackupServiceRolePolicyForBackup in the *AWS Managed Policy Reference*.

AWSBackupServiceRolePolicyForItemRestores

Description

This policy grants users permissions to restore individual files and items in a snapshot (periodic backup recovery point) to a new or existing Amazon S3 bucket or new Amazon EBS volume. These permissions include: read permissions to Amazon EBS for snapshots managed by AWS Backup read/write permissions to Amazon S3 buckets, and generate and describe permissions for AWS KMS keys.

Using this policy

You can attach `AWSBackupServiceRolePolicyForItemRestores` to your users, groups, and roles.

Policy details

- **Type:** AWS managed policy
- **Creation time:** 21 November 2024, 22:45 UTC
- **Edited time:** First instance
- **ARN:** `arn:aws:iam::aws:policy/AWSBackupServiceRolePolicyForItemRestores`

Policy version: v1 (default)

This policy's version defines the permissions for the policy. When the user or role with the policy makes a request to access an AWS resource, AWS checks the default version of the policy to determine whether to allow the request or not.

JSON policy document:

AWSBackupServiceRolePolicyForItemRestores JSON

JSON

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Sid": "EBSReadOnlyPermissions",
            "Effect": "Allow",
            "Action": [
```

```
                "ec2:DescribeSnapshots"
            ],
            "Resource": "arn:aws:ec2:*::snapshot/*"
        },
        {
            "Sid": "KMSReadOnlyPermissions",
            "Effect": "Allow",
            "Action": "kms:DescribeKey",
            "Resource": "*"
        },
        {
            "Sid": "EBSDirectReadAPIPermissions",
            "Effect": "Allow",
            "Action": [
                "ebs:ListSnapshotBlocks",
                "ebs:GetSnapshotBlock"
            ],
            "Resource": "arn:aws:ec2:*::snapshot/*"
        },
        {
            "Sid": "S3ReadonlyPermissions",
            "Effect": "Allow",
            "Action": [
                "s3:GetBucketLocation",
                "s3:ListBucket"
            ],
            "Resource": "arn:aws:s3:::*"
        },
        {
            "Sid": "S3PermissionsForFileLevelRestore",
            "Effect": "Allow",
            "Action": [
                "s3:PutObject",
                "s3:AbortMultipartUpload",
                "s3:ListMultipartUploadParts"
            ],
            "Resource": "arn:aws:s3:::*/*"
        },
        {
            "Sid": "KMSDataKeyForS3AndEC2Permissions",
            "Effect": "Allow",
            "Action": [
                "kms:Decrypt",
                "kms:GenerateDataKey"
```

```
            ],
            "Resource": "arn:aws:kms:*:*:key/*",
            "Condition": {
                "StringLike": {
                    "kms:ViaService": [
                        "ec2.*.amazonaws.com",
                        "s3.*.amazonaws.com"
                    ]
                }
            }
        }
    ]
}
```

AWSBackupServiceRolePolicyForIndexing

Description

This policy grants users permissions to index snapshot, also known as periodic, recovery points. These permissions include: read permissions to Amazon EBS for snapshots managed by AWS Backup read/write permissions to Amazon S3 buckets, and generate and describe permissions for AWS KMS keys.

Using this policy

You can attach `AWSBackupServiceRolePolicyForIndexing` to your users, groups, and roles.

Policy details

- **Type:** AWS managed policy
- **Edited time:** First instance
- **ARN:** `arn:aws:iam::aws:policy/AWSBackupServiceRolePolicyForIndexing`

Policy version: v1 (default)

This policy's version defines the permissions for the policy. When the user or or role with the policy makes a request to access an AWS resource, AWS checks the default version of the policy to determine whether to allow the request or not.

JSON policy document:

AWSBackupServiceRolePolicyForIndexing JSON

JSON

```
{
    "Version":"2012-10-17",
    "Statement": [
        {
            "Sid": "EBSReadOnlyPermissions",
            "Effect": "Allow",
            "Action": [
                "ec2:DescribeSnapshots"
            ],
            "Resource": "arn:aws:ec2:*::snapshot/*"
        },
        {
            "Sid": "KMSReadOnlyPermissions",
            "Effect": "Allow",
            "Action": "kms:DescribeKey",
            "Resource": "*"
        },
        {
            "Sid": "EBSDirectReadAPIPermissions",
            "Effect": "Allow",
            "Action": [
                "ebs:ListSnapshotBlocks",
                "ebs:GetSnapshotBlock"
            ],
            "Resource": "arn:aws:ec2:*::snapshot/*"
        },
        {
            "Sid": "KMSDataKeyForEC2Permissions",
            "Effect": "Allow",
            "Action": "kms:Decrypt",
            "Resource": "arn:aws:kms:*:*:key/*",
            "Condition": {
                "StringLike": {
                    "kms:ViaService": [
                        "ec2.*.amazonaws.com"
                    ]
                }
            }
        }
```

```
    ]
}
```

AWSBackupServiceRolePolicyForRestores

Provides AWS Backup permissions to restore backups of all supported resource types on your behalf.

To view the permissions for this policy, see AWSBackupServiceRolePolicyForRestores in the *AWS Managed Policy Reference*.

For EC2 instance restores, you must also include the following permissions to launch the EC2 instance:

JSON

```
{
  "Version":"2012-10-17",
  "Statement": [
    {
      "Sid": "AllowPassRole",
      "Action": "iam:PassRole",
      "Resource": "arn:aws:iam::123456789012:role/role-name",
      "Effect": "Allow"
    }
  ]
}
```

AWSBackupServiceRolePolicyForS3Backup

This policy contains the permissions necessary for AWS Backup to back up any S3 bucket. This includes access to all objects in a bucket and any associated AWS KMS key.

To view the permissions for this policy, see AWSBackupServiceRolePolicyForS3Backup in the *AWS Managed Policy Reference*.

AWSBackupServiceRolePolicyForS3Restore

This policy contains permissions necessary for AWS Backup to restore an S3 backup to a bucket. This includes read and write permissions to the buckets and the usage of any AWS KMS key in regards to S3 operations.

To view the permissions for this policy, see AWSBackupServiceRolePolicyForS3Restore in the *AWS Managed Policy Reference*.

AWSBackupServiceRolePolicyForScans

The policy should be attached to the IAM role that you use in your backup plan's resource selection. This role grants AWS Backup permission to initiate scans in Amazon GuardDuty.

To view the permissions for this policy, see AWSBackupServiceRolePolicyForScans in the *AWS Managed Policy Reference*.

AWSServiceRolePolicyForBackupReports

AWS Backup uses this policy for the AWSServiceRoleForBackupReports service-linked role. This service-linked role gives AWS Backup permissions to monitor and report on the compliance of your backup settings, jobs, and resources with your frameworks.

To view the permissions for this policy, see AWSServiceRolePolicyForBackupReports in the *AWS Managed Policy Reference*.

AWSServiceRolePolicyForBackupRestoreTesting

To view the permissions for this policy, see AWSServiceRolePolicyForBackupRestoreTesting in the *AWS Managed Policy Reference*.

Customer managed policies

The following sections describe the recommended backup and restore permissions for the AWS services and third-party application supported by AWS Backup. You can use the existing AWS managed policies as a model as you create your own policy documents, and then customize them to further restrict access to your AWS resources.

How to build a customer managed policy from managed policy statements

Each service section below lists the specific statement names (Sids) you need from the AWS Backup managed policies. To build a least-privilege customer managed policy for a specific service, follow these steps:

1. Open the managed policy reference for AWSBackupServiceRolePolicyForBackup (for backup permissions) or AWSBackupServiceRolePolicyForRestores (for restore permissions).
2. Locate each statement by its `Sid` value as listed in the service section below.
3. Copy the full JSON block for each required statement into your custom policy document.
4. Wrap the statements in a valid IAM policy structure. Include the `Version` element to avoid IAM Access Analyzer warnings:

```
{
  "Version": "2012-10-17",
  "Statement": [
    // Paste your extracted statements here
  ]
}
```

5. (Optional) Narrow the `Resource` field in each statement to restrict access to specific resources rather than using `"*"`.

Alternatively, you can retrieve the full policy JSON using the AWS CLI:

```
aws iam get-policy-version \
    --policy-arn arn:aws:iam::aws:policy/service-role/
AWSBackupServiceRolePolicyForBackup \
    --version-id $(aws iam get-policy --policy-arn arn:aws:iam::aws:policy/service-
role/AWSBackupServiceRolePolicyForBackup --query "Policy.DefaultVersionId" --output
 text) \
    --query "PolicyVersion.Document" \
    --output json
```

This returns the complete policy document, from which you can extract the statements you need by matching the `Sid` values listed in the service sections below.

Important

When using custom IAM roles for AWS Backup, you must include resource-specific permissions in addition to AWS Backup permissions. For example, when calling `backup:ListTags` on an Amazon RDS resource, your custom IAM role must also include `rds:ListTagsForResource` permission. While these permissions are included in the default AWS Backup service role, they must be explicitly added to customer-managed

policies. The underlying resource permissions required depend on the specific AWS service and operation being performed.

Amazon Aurora

Backup

Start with the following statements from AWSBackupServiceRolePolicyForBackup:

- `DynamoDBBackupPermissions`
- `RDSClusterModifyPermissions`
- `GetResourcesPermissions`
- `BackupVaultPermissions`
- `KMSPermissions`

Restore

Start with the `RDSPermissions` statement from AWSBackupServiceRolePolicyForRestores.

Amazon Aurora DSQL

Backup

Start with the following statements from AWSBackupServiceRolePolicyForBackup:

- `DSQLBackupPermissions`
- `GetResourcesPermissions`
- `BackupVaultPermissions`
- `KMSPermissions`

Restore

Start with the `DSQLRestorePermissions` statement from AWSBackupServiceRolePolicyForRestores.

Amazon DynamoDB

Backup

Start with the following statements from AWSBackupServiceRolePolicyForBackup:

- `DynamoDBPermissions`
- `DynamoDBBackupResourcePermissions`
- `DynamodbBackupPermissions`
- `KMSDynamoDBPermissions`

Restore

Start with the following statements from AWSBackupServiceRolePolicyForRestores:

- `DynamoDBPermissions`
- `DynamoDBBackupResourcePermissions`
- `DynamoDBRestorePermissions`
- `KMSPermissions`

Amazon EBS

Backup

Start with the following statements from AWSBackupServiceRolePolicyForBackup:

- `EBSResourcePermissions`
- `EBSTagAndDeletePermissions`
- `EBSCopyPermissions`
- `EBSSnapshotTierPermissions`
- `GetResourcesPermissions`
- `BackupVaultPermissions`

Restore

Start with the `EBSPermissions` statement from AWSBackupServiceRolePolicyForRestores.

Add the following statement.

```
{
      "Effect":"Allow",
      "Action": [
        "ec2:DescribeSnapshots",
        "ec2:DescribeVolumes"
      ],
      "Resource":"*"
},
```

Amazon EC2

Backup

Start with the following statements from AWSBackupServiceRolePolicyForBackup:

- `EBSCopyPermissions`
- `EC2CopyPermissions`
- `EC2Permissions`
- `EC2TagPermissions`
- `EC2ModifyPermissions`
- `EBSResourcePermissions`
- `GetResourcesPermissions`
- `BackupVaultPermissions`

Restore

Start with the following statements from AWSBackupServiceRolePolicyForRestores:

- `EBSPermissions`
- `EC2DescribePermissions`
- `EC2RunInstancesPermissions`
- `EC2TerminateInstancesPermissions`
- `EC2CreateTagsPermissions`

Add the following statement.

```
{
      "Effect": "Allow",
      "Action": "iam:PassRole",
      "Resource": "arn:aws:iam::account-id:role/role-name"
},
```

Replace *role-name* with the name of the EC2 instance profile role that will be attached to the restored EC2 instance. This is not the AWS Backup service role, but rather the IAM role that provides permissions to applications running on the EC2 instance.

Amazon EFS

Backup

Start with the following statements from AWSBackupServiceRolePolicyForBackup:

- EFSPermissions
- GetResourcesPermissions
- BackupVaultPermissions

Restore

Start with the EFSPermissions statement from AWSBackupServiceRolePolicyForRestores.

Amazon FSx

Backup

Start with the following statements from AWSBackupServiceRolePolicyForBackup:

- FsxBackupPermissions
- FsxCreateBackupPermissions
- FsxPermissions
- FsxVolumePermissions
- FsxListTagsPermissions
- FsxDeletePermissions
- FsxResourcePermissions
- KMSPermissions

Restore

Start with the following statements from AWSBackupServiceRolePolicyForRestores:

- `FsxPermissions`
- `FsxTagPermissions`
- `FsxBackupPermissions`
- `FsxDeletePermissions`
- `FsxDescribePermissions`
- `FsxVolumeTagPermissions`
- `FsxBackupTagPermissions`
- `FsxVolumePermissions`
- `DSPermissions`
- `KMSDescribePermissions`

Amazon Neptune

Backup

Start with the following statements from AWSBackupServiceRolePolicyForBackup:

- `DynamoDBBackupPermissions`
- `RDSClusterModifyPermissions`
- `GetResourcesPermissions`
- `BackupVaultPermissions`
- `KMSPermissions`

Restore

Start with the `RDSPermissions` statement from AWSBackupServiceRolePolicyForRestores.

Amazon RDS

Backup

Start with the following statements from AWSBackupServiceRolePolicyForBackup:

- `DynamoDBBackupPermissions`
- `RDSBackupPermissions`
- `RDSClusterModifyPermissions`
- `GetResourcesPermissions`
- `BackupVaultPermissions`
- `KMSPermissions`

Restore

Start with the `RDSPermissions` statement from AWSBackupServiceRolePolicyForRestores.

Amazon S3

Backup

Start with AWSBackupServiceRolePolicyForS3Backup.

Add the `BackupVaultPermissions` and `BackupVaultCopyPermissions` statements if you need to copy backups to a different account.

Restore

Start with AWSBackupServiceRolePolicyForS3Restore.

AWS Storage Gateway

Backup

Start with the following statements from AWSBackupServiceRolePolicyForBackup:

- `StorageGatewayPermissions`
- `EBSTagAndDeletePermissions`
- `GetResourcesPermissions`
- `BackupVaultPermissions`

Add the following statement.

```
{
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeSnapshots"
      ],
      "Resource":"*"
},
```

Restore

Start with the following statements from AWSBackupServiceRolePolicyForRestores:

- `StorageGatewayVolumePermissions`
- `StorageGatewayGatewayPermissions`
- `StorageGatewayListPermissions`

Virtual machine

Backup

Start with the `BackupGatewayBackupPermissions` statement from AWSBackupServiceRolePolicyForBackup.

Restore

Start with the `GatewayRestorePermissions` statement from AWSBackupServiceRolePolicyForRestores.

Encrypted backup

To restore an encrypted backup, do one of the following

- Add your role to the allowlist for the AWS KMS key policy
- Add the following statements from AWSBackupServiceRolePolicyForRestores to your IAM role for restores:
  - `KMSDescribePermissions`
  - `KMSPermissions`
  - `KMSCreateGrantPermissions`

Policy updates for AWS Backup

View details about updates to AWS managed policies for AWS Backup since this service began tracking these changes.

Change	Description	Date
AWSBackupFullAccess – Update to an existing policy	AWS Backup added the following permission to this policy: • `s3:GetObject` This permission allows AWS Backup to read the resource modifier file that you specify when you restore an Amazon EKS cluster. It is scoped to calls made through AWS Backup using the `aws:CalledVia` condition key.	November 12, 2026
AWSBackupOperatorAccess – Update to an existing policy	AWS Backup added the following permission to this policy: • `s3:GetObject` This permission allows AWS Backup to read the resource modifier file that you specify when you restore an Amazon EKS cluster. It is scoped to calls made through AWS Backup using the `aws:CalledVia` condition key.	November 12, 2026

Change	Description	Date
AWSServiceRolePolicyForBackupRestoreTesting – Update to an existing policy	AWS Backup added the following permission to this policy: • `redshift:DescribeTags` This permission allows AWS Backup Restore Testing to view tags on Amazon Redshift resources. It was added to support onboarding Amazon Redshift to restore testing.	September 24, 2026
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	AWS Backup added the following permission to this policy: • `cloudformation:UntagResource` This permission allows AWS Backup to untag CloudFormation resources.	August 27, 2026

Change	Description	Date
AWSBackupServiceRolePolicyForS3Backup – Update to an existing policy	AWS Backup added the following permission to this policy: • `s3:ListTagsForResource` With this permission, AWS Backup can verify the Amazon S3 bucket system tag `aws:backup:enabled` when configuring Amazon EventBridge event notifications for Amazon S3 backups.	August 24, 2026
AWSBackupAccessPointOperatorAccess – New policy	Use this new managed policy to create and manage backup access points and access backup data in recovery points through Amazon S3 access points.	August 6, 2026
AWSServiceRolePolicyForBackupRestoreTesting – Update to an existing policy	AWS Backup added the following permission to this policy: • `rds:DeleteTenantDatabase` These permissions allow AWS Backup Restore Testing to delete RDS Tenant Databases after restore test completion.	March 18, 2026

Change	Description	Date
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	AWS Backup added the following permissions to this policy: • `guardduty:StartMalwareScan` • `guardduty:GetMalwareScan` • `iam:PassRole` These permissions allow AWS Backup to initiate malware scans on your recovery points.	February 23, 2026
AWSBackupGuardDutyRolePolicyForScans – New policy	AWS Backup added a new AWS managed policy that provides Amazon GuardDuty permission to read and scan customer backups. AWS Backup passes a role with this policy to GuardDuty when initiating the operations `StartMalwareScan` . This is necessary to provide all necessary permissions needed for malware scans on recovery points of Amazon EC2,Amazon EBS, and Amazon S3 resources. For more information, see the managed policy AWSBackupGuardDutyRolePolicyForScans .	November 19, 2025

Change	Description	Date
AWSBackupServiceRolePolicyForScans – New policy	AWS Backup added a new AWS managed policy that provides AWS Backup permission to initiate malware scans on your recovery points. This is necessary to provide all necessary permissions needed for malware scans on recovery points of Amazon EC2,Amazon EBS, and Amazon S3 resources. For more information, see the managed policy AWSBackupServiceRolePolicyForScans.	November 19, 2025
AWSBackupFullAccess – Update to an existing policy	Added `malware-protection.guardduty.amazonaws.com` to `IamPassRolePermissions`, which is necessary to initiate malware scan jobs.	November 19, 2025

Change	Description	Date
AWSBackupOperatorAccess – Update to an existing policy	AWS Backup added the following permissions to this policy: • `malware-protection.guardduty.amazonaws.com` to `IamPassRolePermissions` • `backup:StartScanJob` These permissions are necessary to initiate malware scan jobs.	November 19, 2025
AWSBackupFullAccess – Update to an existing policy	AWS Backup added the following permissions to this policy: • `eks:ListClusters` • `eks:ListTagsForResource` • `eks:DescribeCluster` These permissions allow AWS Backup to backup and restore Amazon EKS clusters.	November 10, 2025

Change	Description	Date
AWSBackupOperatorAccess – Update to an existing policy	AWS Backup added the following permissions to this policy: • `eks:ListClusters` • `eks:ListTagsForResource` • `eks:DescribeCluster` These permissions allow AWS Backup to backup and restore Amazon EKS clusters.	November 10, 2025

Change	Description	Date
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	AWS Backup added the following permissions to this policy: • `eks:ListClusters` • `eks:ListTagsForResource` • `eks:DescribeCluster` • `eks:ListAddons` • `eks:DescribeAddon` • `eks:ListNodegroups` • `eks:DescribeNodegroup` • `eks:ListPodIdentityAssociations` • `eks:DescribePodIdentityAssociation` • `eks:ListAccessEntries` • `eks:DescribeAccessEntry` • `eks:ListAssociatedAccessPolicies` • `eks:ListFargateProfiles` • `eks:DescribeFargateProfile` • `ec2:DescribeLaunchTemplateVersions` • `eks:CreateAccessEntry`	November 10, 2025

Change	Description	Date
	• `eks:AssociateAccessPolicy` • `eks:DisassociateAccessPolicy` These permissions allow AWS Backup to create backups of Amazon EKS clusters and their associated resources on behalf of customers.	

Change	Description	Date
AWSBackupServiceLinkedRolePolicyForBackup – Update to an existing policy	AWS Backup added the following permissions to this policy: • eks:ListClusters • eks:ListTagsForResource • eks:DescribeCluster • eks:ListAddons • eks:DescribeAddon • eks:ListNodegroups • eks:DescribeNodegroup • eks:ListPodIdentityAssociations • eks:DescribePodIdentityAssociation • eks:ListAccessEntries • eks:DescribeAccessEntry • eks:ListAssociatedAccessPolicies • eks:ListFargateProfiles • eks:DescribeFargateProfile • ec2:DescribeLaunchTemplateVersions • eks:CreateAccessEntry	November 10, 2025

Change	Description	Date
	• `eks:AssociateAccessPolicy` • `eks:DisassociateAccessPolicy` These permissions allow AWS Backup to create backups of Amazon EKS clusters and their associated resources on behalf of customers.	

Change	Description	Date
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	AWS Backup added the following permissions to this policy: • `eks:CreateCluster` • `eks:DescribeCluster` • `eks:CreateAccessEntry` • `eks:DescribeAccessEntry` • `eks:AssociateAccessPolicy` • `eks:ListAssociatedAccessPolicies` • `eks:CreateAddon` • `eks:DescribeAddon` • `eks:CreateNodegroup` • `eks:DescribeNodegroup` • `eks:CreateFargateProfile` • `eks:DescribeFargateProfile` • `eks:CreatePodIdentityAssociation` • `eks:DescribePodIdentityAssociation` • `eks:TagResource` • `eks:DisassociateAccessPolicy` • `iam:PassRole`	November 10, 2025

Change	Description	Date
	• ec2:DescribeLaunchTemplateVersions • ec2:DescribeSubnets • ec2:RunInstances • ec2:CreateTags • iam:GetRole • iam:ListAttachedRolePolicies • backup:StartRestoreJob • backup:ListRestoreJobs • backup:ListRecoveryPointsByBackupVault • backup:DescribeRestoreJob These permissions allow AWS Backup to perform restore operations for Amazon EKS clusters and their associated resources on behalf of customers.	

Change	Description	Date
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	AWS Backup added the following permission to this policy: • `organizations:ListDelegatedAdministrators` This permission allows AWS Backup to synchronize delegated administrator information with Organizations for cross-account management features.	September 9, 2025

Change	Description	Date
AWSBackupGuardDuty RolePolicyForScans – New policy	AWS Backup added a new AWS managed policy that provides Amazon GuardDuty permission to read and scan customer backups. AWS Backup passes a role with this policy to GuardDuty when initiating the operations `StartMalwareScan` . This is necessary to provide all necessary permissio ns needed for malware scans on recovery points of Amazon EC2,Amazon EBS, and Amazon S3 resources. For more information, see the managed policy AWSBackup GuardDutyRolePolicyForScans .	November 24, 2025

Change	Description	Date
AWSBackupServiceRolePolicyForScans – New policy	AWS Backup added a new AWS managed policy that provides AWS Backup permission to initiate malware scans on your recovery points. This is necessary to provide all necessary permissio ns needed for malware scans on recovery points of Amazon EC2,Amazon EBS, and Amazon S3 resources. For more information, see the managed policy AWSBackupServiceRolePolicyForScans.	November 24, 2025
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	AWS Backup added the following permissions to this policy: • `dsql:UpdateCluster` • `dsql:AddPeerCluster` • `dsql:RemovePeerClu ster` • `dsql:GetCluster` These permissions are necessary for AWS Backup to perform orchestrated multi-Region restore operations for DSQL resources on behalf of customers.	July 17, 2025

Change	Description	Date
AWSBackupFullAccess – Update to an existing policy	AWS Backup added the following permissions to this policy: • `mpa:GetApprovalTeam` • `mpa:ListApprovalTeams` • `mpa:StartSession` • `mpa:CancelSession` • `mpa:GetSession` • `mpa:ListSessions` These permissions are necessary for AWS Backup integration with AWS Account Management and AWS Organizations so customers have the option of Multi-party approval (MPA) as part of their logically air-gapped vaults.	June 17, 2025

Change	Description	Date
AWSBackupServiceRolePolicyForRestores – Update to an existing policy:	AWS Backup added the following permissions to this policy: • `ec2:DescribeRouteTables` • `ec2:CreateTags` These permissions are necessary to allow customers to restore Amazon FSx for OpenZFS Multi-availability zone (Multi-AZ) snapshots through AWS Backup.	May 27, 2025
AWSBackupFullAccess – Update to an existing policy	AWS Backup added the following permissions to this policy: • `dsql:GetCluster` • `dsql:ListClusters` • `dsql:ListTagsForResource` These permissions allow AWS Backup to backup and restore Amazon Aurora DSQL resources.	May 21, 2025

Change	Description	Date
AWSBackupOperatorAccess – Update to an existing policy	AWS Backup added the following permissions to this policy: • `dsql:GetCluster` • `dsql:ListClusters` • `dsql:ListTagsForResource` These permissions allow AWS Backup to backup and restore Amazon Aurora DSQL resources.	May 21, 2025

Change	Description	Date
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	AWS Backup added the following permissions to this policy: • `kms:Decrypt` • `kms:ReEncryptTo` • `kms:ReEncryptFrom` • `dsql:StartBackupJob` • `dsql:GetBackupJob` • `dsql:StopBackupJob` • `dsql:GetCluster` • `dsql:ListClusters` • `dsql:ListTagsForResource` These permissions allow AWS Backup to create, delete, retrieve, and manage Amazon Aurora DSQL snapshots on behalf of customers.	May 21, 2025

Change	Description	Date
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	AWS Backup added the following permissions to this policy: • `ec2:DescribeRouteTables` • `ec2:CreateTags` • `dsql:StartRestoreJob` • `dsql:GetRestoreJob` • `dsql:StopRestoreJob` • `dsql:TagResource` • `dsql:CreateCluster` • `dsql:PutMultiRegionProperties` • `dsql:PutWitnessRegion` • `kms:DescribeKey` These permissions allow AWS Backup to create, delete, retrieve, encrypt, decrypt, and manage Amazon Aurora DSQL snapshots on behalf of customers.	May 21, 2025

Change	Description	Date
AWSBackupServiceLinkedRolePolicyForBackup – Update to an existing policy	AWS Backup added the following permissions to this policy: • `dsql:ListClusters` • `dsql:ListTagsForResource` These permissions allow AWS Backup to manage Aurora DSQL backups at customer-specified intervals.	May 21, 2025

Change	Description	Date
AWSBackupFullAccess – Update to an existing policy	AWS Backup added the following permissions to this policy: • `aws:CalledVia` • `redshift-serverless:DeleteSnapshot` • `redshift-serverless:GetNamespace` • `redshift-serverless:GetSnapshot` • `redshift-serverless:GetWorkgroup` • `redshift-serverless:ListNamespaces` • `redshift-serverless:ListSnapshots` • `redshift-serverless:ListWorkgroups` These permissions are necessary for designated customers to have full access to Amazon Redshift Serverless backups, including required read permissions as well as the ability to delete Amazon Redshift Serverless recovery points (snapshot backups).	March 31, 2025

Change	Description	Date
AWSBackupOperatorAccess – Update to an existing policy	AWS Backup added the following permissions to this policy: • `redshift-serverless:GetNamespace` • `redshift-serverless:GetSnapshot` • `redshift-serverless:GetWorkgroup` • `redshift-serverless:ListNamespaces` • `redshift-serverless:ListSnapshots` • `redshift-serverless:ListWorkgroups` These permissions are necessary for designated customers to have all necessary backup permissions to Amazon Redshift Serverless, including required read permissions.	March 31, 2025

Change	Description	Date
AWSBackupServiceLinkedRolePolicyForBackup – Update to an existing policy	AWS Backup added the following permissions to this policy: • `redshift-serverless:DeleteSnapshot` • `redshift-serverless:GetNamespace` • `redshift-serverless:GetSnapshot` • `redshift-serverless:GetWorkgroup` • `redshift-serverless:ListNamespaces` • `redshift-serverless:ListSnapshots` • `redshift-serverless:ListTagsForResource` • `redshift-serverless:ListWorkgroups` These permissions are necessary for AWS Backup to manage Amazon Redshift Serverless snapshots at customer specified intervals.	March 31, 2025

Change	Description	Date
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	AWS Backup added the following permissions to this policy: • `redshift-serverless:CreateSnapshot` • `redshift-serverless:DeleteSnapshot` with a condition to only allow snapshots deletion with the `"aws:backup:source-resource"` tag • `redshift-serverless:GetNamespace` • `redshift-serverless:GetSnapshot` • `redshift-serverless:ListNamespaces` • `redshift-serverless:ListSnapshots` • `redshift-serverless:ListTagsForResource` • `redshift-serverless:TagResource` These permissions are necessary to allow AWS Backup to create, delete, retrieve, and manage Amazon	March 31, 2025

Change	Description	Date
	Redshift Serverless snapshots on behalf of customers.	
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	AWS Backup added the following permissions to this policy: • `redshift-serverless:GetNamespace` • `redshift-serverless:GetTableRestoreStatus` • `redshift-serverless:RestoreTableFromSnapshot` These permissions are necessary to allow AWS Backup to restore Amazon Redshift and Amazon Redshift Serverless snapshots on behalf of the customer.	March 31, 2025
AWSBackupSearchOperatorAccess – Added a new AWS managed policy	AWS Backup added the `AWSBackupSearchOperatorAccess` AWS managed policy.	February 27, 2025

Change	Description	Date
AWSBackupServiceLinkedRolePolicyForBackup – Update to an existing policy	AWS Backup added the permission `rds:AddTagsToResource` to support Amazon RDS multi-tenant snapshot cross-account copy of backups. This permission is necessary to complete operations when a customer chooses to create a cross-account copy of a multi-tenant RDS snapshot.	January 8, 2025
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	AWS Backup added the permissions `rds:CreateTenantDatabase` and `rds:DeleteTenantDatabase` to this policy to support the restore process of Amazon RDS resources. These permissions are necessary to complete customer operations for restoring multi-tenant snapshots.	January 8, 2025
AWSBackupServiceRolePolicyForItemRestores – Added a new AWS managed policy	AWS Backup added the `AWSBackupServiceRolePolicyForItemRestores` AWS managed policy.	November 26, 2024
AWSBackupServiceRolePolicyForIndexing – Added a new AWS managed policy	AWS Backup added the `AWSBackupServiceRolePolicyForIndexing` AWS managed policy.	November 26, 2024

Change	Description	Date
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	AWS Backup added permission `backup:TagResource` to this policy. The permission is necessary to obtain tagging permissions during the creation of a recovery point.	May 17, 2024
AWSBackupServiceRolePolicyForS3Backup – Update to an existing policy	AWS Backup added permission `backup:TagResource` to this policy. The permission is necessary to obtain tagging permissions during the creation of a recovery point.	May 17, 2024
AWSBackupServiceLinkedRolePolicyForBackup – Update to an existing policy	AWS Backup added permission `backup:TagResource` to this policy. The permission is necessary to obtain tagging permissions during the creation of a recovery point.	May 17, 2024
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	Added the permission `rds:DeleteDBInstanceAutomatedBackups` . This permission is necessary for AWS Backup to support continuous backup and point-in-time-restore of Amazon RDS instances.	May 1, 2024

Change	Description	Date
AWSBackupFullAccess – Update to an existing policy	AWS Backup updated the Amazon Resource Name (ARN) in permission `storagegateway:ListVolumes` from `arn:aws:storagegateway:*:*:gateway/*` to * in order to accommodate a change in the Storage Gateway API model.	May 1, 2024
AWSBackupOperatorAccess – Update to an existing policy	AWS Backup updated the Amazon Resource Name (ARN) in permission `storagegateway:ListVolumes` from `arn:aws:storagegateway:*:*:gateway/*` to * in order to accommodate a change in the Storage Gateway API model.	May 1, 2024

Change	Description	Date
AWSServiceRolePolicyForBackupRestoreTesting – Update to an existing policy	Added the following permissions to describe and list recovery points and protected resources in order to conduct restore testing plans: `backup:DescribeRecoveryPoint` , `backup:DescribeProtectedResource` , `backup:ListProtectedResources` , and `backup:ListRecoveryPointsByResource` . Added the permission `ec2:DescribeSnapshotTierStatus` to support Amazon EBS archive tier storage. Added the permission `rds:DescribeDBClusterAutomatedBackups` to support Amazon Aurora continuous backups. Added the following permissions to support restore testing of Amazon Redshift backups: `redshift:DescribeClusters` and `redshift:DeleteCluster` . Added the permission `timestream:DeleteT`	February 14, 2024

Change	Description	Date
	`able` to support restore testing of Amazon Timestream backups.	
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	Added the permissions `ec2:DescribeSnapshotTierStatus` and `ec2:RestoreSnapshotTier` . These permissions are necessary for users to have the option to restore Amazon EBS resources stored with AWS Backup from archive storage. For EC2 instance restores, you must also include permissions as shown in the following policy statement to launch the EC2 instance:	November 27, 2023

Change	Description	Date
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	Added the permissions `ec2:DescribeSnapshotTierStatus` and `ec2:ModifySnapshotTier` to support an additional storage option for backed up Amazon EBS resources to be transitioned to the archive storage tier. These permissions are necessary for users to have the option to transition Amazon EBS resources stored with AWS Backup to archive storage.	November 27, 2023

Change	Description	Date
AWSBackupServiceLinkedRolePolicyForBackup – Update to an existing policy	Added the permissions `ec2:DescribeSnapshotTierStatus` and `ec2:ModifySnapshotTier` to support an additional storage option for backed up Amazon EBS resources to be transitioned to the archive storage tier. These permissions are necessary for users to have the option to transition Amazon EBS resources stored with AWS Backup to archive storage. Added the permissions `rds:DescribeDBClusterSnapshots` and `rds:RestoreDBClusterToPointInTime`, which is necessary for PITR (point-in-time restores) of Aurora clusters.	

Change	Description	Date
AWSServiceRolePolicyForBackupRestoreTesting – New policy	Provides the permissions necessary to conduct restore testing. The permissions include the actions `list, read, and write` for the following services to be included in restore tests: Aurora, DocumentDB, DynamoDB, Amazon EBS, Amazon EC2, Amazon EFS, FSx for Lustre, FSx for Windows File Server, FSx for ONTAP, FSx for OpenZFS, Amazon Neptune, Amazon RDS, and Amazon S3.	November 27, 2023
AWSBackupFullAccess – Update to an existing policy	Added `restore-testing.backup.amazonaws.com` to `IamPassRolePermissions` and `IamCreateServiceLinkedRolePermissions`. This addition is necessary for AWS Backup to conduct restore tests on behalf of customers.	November 27, 2023
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	Added the permissions `rds:DescribeDBClusterSnapshots` and `rds:RestoreDBClusterToPointInTime`, which is necessary for PITR (point-in-time restores) of Aurora clusters.	September 6, 2023

Change	Description	Date
AWSBackupFullAccess – Update to an existing policy	Added the permission `rds:DescribeDBClusterAutomatedBackups` , which is necessary for continuous backup and point-in-time restore of Aurora clusters.	September 6, 2023
AWSBackupOperatorAccess – Update to an existing policy	Added the permission `rds:DescribeDBClusterAutomatedBackups` , which is necessary for continuous backup and point-in-time restore of Aurora clusters.	September 6, 2023

Change	Description	Date
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	Added the permission `rds:DescribeDBClusterAutomatedBackups`. This permission is necessary for AWS Backup support of continuous backup and point-in-time restore of Aurora clusters. Added the permission `rds:DeleteDBClusterAutomatedBackups` to allow AWS Backup lifecycle to delete and disassociate Amazon Aurora continuous recovery points when a retention period finishes. This permission is necessary for the Aurora recovery point to avoid a transition into an `EXIPIRED` state. Added the permission `rds:ModifyDBCluster` which allows AWS Backup to interact with Aurora clusters. This addition allows users the ability to enable or disable continuous backups based on desired configurations.	September 6, 2023

Change	Description	Date
AWSBackupFullAccess – Update to an existing policy	Added the action `ram:GetResourceShareAssociations` to grant the user permission to get resource share associations for new vault type.	August 8, 2023
AWSBackupOperatorAccess – Update to an existing policy	Added the action `ram:GetResourceShareAssociations` to grant the user permission to get resource share associations for new vault type.	August 8, 2023
AWSBackupServiceRolePolicyForS3Backup – Update to an existing policy	Added the permission `s3:PutInventoryConfiguration` to enhance backup performance speeds by using a bucket inventory.	August 1, 2023
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	Added the following actions to grant the user permissions to add tags to restore resources: `storagegateway:AddTagsToResource`, `elasticfilesystem:TagResource`, `ec2:CreateTags` for only `ec2:CreateAction` that includes either `RunInstances` or `CreateVolume`, `fsx:TagResource`, and `cloudformation:TagResource`.	May 22, 2023

Change	Description	Date
AWSBackupAuditAccess – Update to an existing policy	Replaced the resource selection within the API `config:DescribeCompliancyByConfigRule` with a wildcard resource to make it easier for a user to select resources.	April 11, 2023
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	Added the following permission to restore Amazon EFS using a customer managed key: `kms:GenerateDataKeyWithoutPlaintext` . This helps to ensure users have required permissions to restore Amazon EFS resources.	March 27, 2023
AWSServiceRolePolicyForBackupReports – Update to an existing policy	Updated the `config:DescribeConfigRules` and `config:DescribeConfigRuleEvaluationStatus` actions to allow AWS Backup Audit Manager to access AWS Backup Audit Manager-managed AWS Config rules.	March 9, 2023

Change	Description	Date
AWSBackupServiceRolePolicyForS3Restore – Update to an existing policy	Added the following permissions: `kms:Decrypt` , `s3:PutBucketOwnershipControls` , and `s3:GetBucketOwnershipControls` to the policy `AWSBackupServiceRolePolicyForS3Restore` . These permissions are necessary to support restores of objects when KMS encryption is used in the original backup and for restoring objects when object ownership is configured on the original bucket instead of ACL.	February 13, 2023

Change	Description	Date
AWSBackupFullAccess – Update to an existing policy	Added the following permissions to schedule backups using VMware tags of virtual machines and to support schedule-based bandwidth throttling: `backup-gateway:GetHypervisorPropertyMappings`, `backup-gateway:GetVirtualMachine`, `backup-gateway:PutHypervisorPropertyMappings`, `backup-gateway:GetHypervisor`, `backup-gateway:StartVirtualMachinesMetadataSync`, `backup-gateway:GetBandwidthRateLimitSchedule`, and `backup-gateway:PutBandwidthRateLimitSchedule`.	December 15, 2022

Change	Description	Date
AWSBackupOperatorAccess – Update to an existing policy	Added the following permissions to schedule backups using VMware tags of virtual machines and to support schedule-based bandwidth throttling: `backup-gateway:GetHypervisorPropertyMappings` , `backup-gateway:GetVirtualMachine` , `backup-gateway:GetHypervisor` , and `backup-gateway:GetBandwidthRateLimitSchedule` .	December 15, 2022
AWSBackupGatewayServiceRolePolicyForVirtualMachineMetadataSync – New policy	Provides permissions for AWS Backup Gateway to sync the metadata of virtual machines in on-premise networks with Backup Gateway.	December 15, 2022

Change	Description	Date
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	Added the following permissions to support Timestream backup jobs: `timestream:StartAwsBackupJob` , `timestream:GetAwsBackupStatus` , `timestream:ListTables` , `timestream:ListDatabases` , `timestream:ListTagsForResource` , `timestream:DescribeTable` , `timestream:DescribeDatabase` , and `timestream:DescribeEndpoints` .	December 13, 2022

Change	Description	Date
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	Added the following permissions to support Timestream restore jobs: `timestream:StartAwsRestoreJob` , `timestream:GetAwsRestoreStatus` , `timestream:ListTables` , `timestream:ListTagsForResource` , `timestream:ListDatabases` , `timestream:DescribeTable` , `timestream:DescribeDatabase` , `s3:GetBucketAcl` , and `timestream:DescribeEndpoints` .	December 13, 2022
AWSBackupFullAccess – Update to an existing policy	Added the following permissions to support Timestream resources: `timestream:ListTables` , `timestream:ListDatabases` , `s3:ListAllMyBuckets` and `timestream:DescribeEndpoints` .	December 13, 2022

Change	Description	Date
AWSBackupOperatorAccess – Update to an existing policy	Added the following permissions to support Timestream resources: `timestream:ListDatabases`, `timestream:ListTables`, `s3:ListAllMyBuckets`, and `timestream:DescribeEndpoints`.	December 13, 2022
AWSBackupServiceLinkedRolePolicyForBackup – Update to an existing policy	Added the following permissions to support Timestream resources: `timestream:ListDatabases`, `timestream:ListTables`, `timestream:ListTagsForResource`, `timestream:DescribeDatabase`, `timestream:DescribeTable`, `timestream:GetAwsBackupStatus`, `timestream:GetAwsRestoreStatus`, and `timestream:DescribeEndpoints`.	December 13, 2022

Change	Description	Date
AWSBackupFullAccess – Update to an existing policy	Added the following permissions to support Amazon Redshift resources : `redshift:DescribeClusters` , `redshift:DescribeClusterSubnetGroups` , `redshift:DescribeNodeConfigurationOptions` , `redshift:DescribeOrderableClusterOptions` , `redshift:DescribeClusterParameterGroups` , `redshift:DescribeClusterTracks` , `redshift:DescribeSnapshotSchedules` , and `ec2:DescribeAddresses` .	November 27, 2022

Change	Description	Date
AWSBackupOperatorAccess – Update to an existing policy	Added the following permissions to support Amazon Redshift resources : `redshift:DescribeClusters` , `redshift:DescribeClusterSubnetGroups` , `redshift:DescribeNodeConfigurationOptions` , `redshift:DescribeOrderableClusterOptions` , `redshift:DescribeClusterParameterGroups`, , `redshift:DescribeClusterTracks` . `redshift:DescribeSnapshotSchedules` , and `ec2:DescribeAddresses` .	November 27, 2022
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	Added the following permissions to support Amazon Redshift restore jobs: `redshift:RestoreFromCluster Snapshot` , `redshift:RestoreTableFromClusterSnapshot` , `redshift:DescribeClusters` , and `redshift:DescribeTableRestoreStatus` .	November 27, 2022

Change	Description	Date
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	Added the following permissions to support Amazon Redshift backup jobs: `redshift:CreateClusterSnapshot` , `redshift:DescribeClusterSnapshots` , `redshift:DescribeTags` , `redshift:DeleteClusterSnapshot` , `redshift:DescribeClusters` , and `redshift:CreateTags` .	November 27, 2022
AWSBackupFullAccess – Update to an existing policy	Added the following permission to support CloudFormation resources: `cloudformation:ListStacks` .	November 27, 2022
AWSBackupOperatorAccess – Update to an existing policy	Added the following permission to support CloudFormation resources: `cloudformation:ListStacks` .	November 27, 2022

Change	Description	Date
AWSBackupServiceLinkedRolePolicyForBackup – Update to an existing policy	Added the following permissions to support CloudFormation resources : `redshift:DescribeClusterSnapshots` , `redshift:DescribeTags` , `redshift:DeleteClusterSnapshot` , and `redshift:DescribeClusters` .	November 27, 2022
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	Added the following permissions to support CloudFormation application stack backup jobs: `cloudformation:GetTemplate` , `cloudformation:DescribeStacks` , and `cloudformation:ListStackResources` .	November 16, 2022
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	Added the following permissions to support CloudFormation application stack backup jobs: `cloudformation:CreateChangeSet` and `cloudformation:DescribeChangeSet`	November 16, 2022

Change	Description	Date
AWSBackupOrganizationAdminAccess – Update to an existing policy	Added the following permissions to this policy to allow organization administrators to usethe Delegated Administrator feature: `organizations:ListDelegatedAdministrator` , `organizations:RegisterDelegatedAdministrator` , and `organizations:DeregisterDelegatedAdministrator`	November 27, 2022
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	Added the following permissions to support SAP HANA on Amazon EC2 instances: `ssm-sap:GetOperation` , `ssm-sap:ListDatabases` , `ssm-sap:BackupDatabase` , `ssm-sap:UpdateHanaBackupSettings` , `ssm-sap:GetDatabase` , and `ssm-sap:ListTagsForResource` .	November 20, 2022

Change	Description	Date
AWSBackupFullAccess – Update to an existing policy	Added the following permissions to support SAP HANA on Amazon EC2 instances: `ssm-sap:GetOperation` , `ssm-sap:ListDatabases` , `ssm-sap:GetDatabase` , and `ssm-sap:ListTagsForResource` .	November 20, 2022
AWSBackupOperatorAccess – Update to an existing policy	Added the following permissions to support SAP HANA on Amazon EC2 instances: `ssm-sap:GetOperation` , `ssm-sap:ListDatabases` , `ssm-sap:GetDatabase` , and `ssm-sap:ListTagsForResource` .	November 20, 2022
AWSBackupServiceLinkedRolePolicyForBackup – Update to an existing policy	Added the following permission to support SAP HANA on Amazon EC2 instances: `ssm-sap:GetOperation` .	November 20, 2022
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	Added the following permission to support Backup gateway restore jobs to an EC2 instance: `ec2:CreateTags` .	November 20, 2022

Change	Description	Date
AWSBackupDataTransferAccess – Update to an existing policy	Added the following permissions to support secure storage data transfer for SAP HANA On Amazon EC2 resources: `backup-storage:StartObject`, `backup-storage:PutChunk`, `backup-storage:GetChunk`, `backup-storage:ListChunks`, `backup-storage:ListObjects`, `backup-storage:GetObjectMetadata`, and `backup-storage:NotifyObjectComplete`.	November 20, 2022

Change	Description	Date
AWSBackupRestoreAccessForSAPHANA – Update to an existing policy	Added the following permissions for resource owners to perform restore of SAP HANA On Amazon EC2 resources: `backup:Get*` , `backup:List*` , `backup:Describe*` , `backup:StartBackupJob` , `backup:StartRestoreJob` , `ssm-sap:GetOperation` , `ssm-sap:ListDatabases` , `ssm-sap:BackupDatabase` , `ssm-sap:RestoreDatabase` , `ssm-sap:UpdateHanaBackupSettings` , `ssm-sap:GetDatabase` , and `ssm-sap:ListTagsForResource` .	November 20, 2022
AWSBackupServiceRolePolicyForS3Backup – Update to an existing policy	Added the permission `s3:GetBucketAcl` to support backup operations of AWS Backup for Amazon S3.	August 24, 2022
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	Added the following actions to grant access to create a database instance to support multi-Availability Zone (Multi-AZ) functionality: `rds:CreateDBInstance` .	July 20, 2022

Change	Description	Date
AWSBackupServiceLinkedRolePolicyForBackup – Update to an existing policy	Added the `s3:GetBucketTagging` permission to grant the user permission to select buckets to backup with a resource wildcard. Without this permission, users who select which buckets to backup with a resource wildcard are unsuccessful.	May 6, 2022
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	Added volume resources in the scope of existing `fsx:CreateBackup` and `fsx:ListTagsForResource` actions, and added new action `fsx:DescribeVolumes` to support FSx for ONTAP volume level backups.	April 27, 2022
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	Added the following actions to grant the users permissions to restore FSx for ONTAP volumes `fsx:DescribeVolumes` , `fsx:CreateVolumeFromBackup` , `fsx:DeleteVolume` , and `fsx:UntagResource` .	April 27, 2022

Change	Description	Date
AWSBackupServiceRolePolicyForS3Backup – Update to an existing policy	Added the following actions to grant the user permissions to receive notifications of changes to their Amazon S3 buckets during backup operations: `s3:GetBucketNotification` and `s3:PutBucketNotification` .	February 25, 2022

Change	Description	Date
AWSBackupServiceRolePolicyForS3Backup – New policy	Added the following actions to grant the user permissions to back up their Amazon S3 buckets: `s3:GetInventoryConfiguration` , `s3:PutInventoryConfiguration` , `s3:ListBucketVersions` , `s3:ListBucket` , `s3:GetBucketTagging` , `s3:GetBucketVersioning` , `s3:GetBucketNotification` , `s3:GetBucketLocation` , and `s3:ListAllMyBuckets` Added the following actions to grant the user permissions to back up their Amazon S3 objects: `s3:GetObject` , `s3GetObjectAcl` , `s3:GetObjectVersionTagging` , `s3:GetObjectVersionAcl` , `s3:GetObjectTagging` , and `s3:GetObjectVersion` . Added the following actions to grant the user permissions to back up their encrypted Amazon S3 data: `kms:Decrypt` and `kms:DescribeKey` .	February 17, 2022

Change	Description	Date
	Added the following actions to grant the user permissions to take incremental backups of their Amazon S3 data using Amazon EventBridge rules: `events:DescribeRule` ,`events:EnableRule` , `events:PutRule` , `events:DeleteRule` , `events:PutTargets` , `events:RemoveTargets` , `events:ListTargetsByRule` ,`events:DisableRule` ,`cloudwatch:GetMetricData` ,and `events:ListRules` .	

Change	Description	Date
AWSBackupServiceRolePolicyForS3Restore – New policy	Added the following actions to grant the user permissions to restore their Amazon S3 buckets: `s3:CreateBucket` , `s3:ListBucketVersions` , `s3:ListBucket` , `s3:GetBucketVersioning` , `s3:GetBucketLocation` , and `s3:PutBucketVersioning` . Added the following actions to grant the user permissions to restore their Amazon S3 buckets: `s3:GetObject` , `s3:GetObjectVersion` , `s3:DeleteObject` , `s3:PutObjectVersionAcl` , `s3:GetObjectVersionAcl` , `s3:GetObjectTagging` , `s3:PutObjectTagging` , `s3:GetObjectAcl` , `s3:PutObjectAcl` , `s3:PutObject` , and `s3:ListMultipartUploadParts` . Added the following actions to grant the user permissions to encrypt their restored Amazon S3 data: `kms:Decrypt` , `kms:DescribeKey` ,	February 17, 2022

Change	Description	Date
	and `kms:GenerateDataKey` .	
AWSBackupServiceLinkedRolePolicyForBackup – Update to an existing policy	Added `s3:ListAllMyBuckets` to grant the user permissions to view a list of their buckets and choose which ones to assign to a backup plan.	February 14, 2022
AWSBackupServiceLinkedRolePolicyForBackup – Update to an existing policy	Added `backup-gateway:ListVirtualMachines` to grant the user permissions to view a list of their virtual machines and choose which ones to assign to a backup plan. Added `backup-gateway:ListTagsForResource` to grant the user permissions to list the tags for their virtual machines.	November 30, 2021
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	Added `backup-gateway:Backup` to grant the user permissions restore their virtual machine backups. AWS Backup also added `backup-gateway:ListTagsForResource` to grant the user permissions to list the tags assigned to their virtual machine backups.	November 30, 2021

Change	Description	Date
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	Added `backup-gateway:Restore` to grant the user permissions restore their virtual machine backups.	November 30, 2021

Change	Description	Date
AWSBackupFullAccess – Update to an existing policy	Added the following actions to grant the users permissions to use AWS Backup Gateway to back up, restore, and manage their virtual machines: backup-gateway:AssociateGatewayToServer ,backup-gateway:CreateGateway ,backup-gateway:DeleteGateway ,backup-gateway:DeleteHypervisor ,backup-gateway:DisassociateGatewayFromServer ,backup-gateway:ImportHypervisorConfiguration ,backup-gateway:ListGateways ,backup-gateway:ListHypervisors ,backup-gateway:ListTagsForResource ,backup-gateway:ListVirtualMachines ,backup-gateway:PutMaintenanceStartTime ,backup-gateway:TagResource ,backup-gateway:TestHypervisorConfiguration ,backup-gateway:UntagResource ,backup-gateway:UpdateGatewayInformat	November 30, 2021

Change	Description	Date
	`ion` , and `backup-gateway:UpdateHypervisor` .	
AWSBackupOperatorAccess – Update to an existing policy	Added the following actions to grant the user permissions to back up their virtual machines: `backup-gateway:ListGateways` , `backup-gateway:ListHypervisors` , `backup-gateway:ListTagsForResource` , and `backup-gateway:ListVirtualMachines` .	November 30, 2021
AWSBackupServiceLinkedRolePolicyForBackup – Update to an existing policy	Added `dynamodb:ListTagsOfResource` to grant the user permissions to list tags of their DynamoDB tables to back up using AWS Backup's advanced DynamoDB backup features.	November 23, 2021

Change	Description	Date
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	Added `dynamodb:StartAwsBackupJob` to grant the user permissions to back up their DynamoDB tables using advanced backup features. Added `dynamodb:ListTagsOfResource` to grant the user to permissions to copy tags from their source DynamoDB tables to their backups.	November 23, 2021
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	Added `dynamodb:RestoreTableFromAwsBackup` to grant the user permissions restore their DynamoDB tables backed up using AWS Backup's advanced DynamoDB advanced backup features.	November 23, 2021
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	Added `dynamodb:RestoreTableFromAwsBackup` to grant the user permissions restore their DynamoDB tables backed up using AWS Backup's advanced DynamoDB advanced backup features.	November 23, 2021

Change	Description	Date
AWSBackupOperatorAccess – Update to an existing policy	Removed the actions `backup:GetRecoveryPointRestoreMetadata` and `rds:DescribeDBSnapshots` because they were redundant. AWS Backup did not need both `backup:GetRecoveryPointRestoreMetadata` and `backup:Get*` as part of `AWSBackupOperatorAccess`. Also, AWS Backup did not need both `rds:DescribeDBSnapshots` and `rds:describeDBSnapshots` as part of `AWSBackupOperatorAccess`.	November 23, 2021

Change	Description	Date
AWSBackupServiceLinkedRolePolicyForBackup – Update to an existing policy	Added the new actions `elasticfilesystem:DescribeFileSystems`, `dynamodb:ListTables`, `storagegateway:ListVolumes`, `ec2:DescribeVolumes`, `ec2:DescribeInstances`, `rds:DescribeDBInstances`, `rds:DescribeDBClusters`, and `fsx:DescribeFileSystems` to allow customers to view and choose from a list of their AWS Backup-supported resources when selecting which resources to assign to a backup plan.	November 10, 2021
AWSBackupAuditAccess – New policy	Added `AWSBackupAuditAccess` to grant the user permissions to use AWS Backup Audit Manager. Permissions include the ability to configure compliance frameworks and generate reports.	August 24, 2021

Change	Description	Date
AWSServiceRolePolicyForBackupReports – New policy	Added `AWSServiceRolePolicyForBackupReports` to grant permissions for a service-linked role to automate the monitoring of backup settings, jobs, and resources for compliance with frameworks configured by the user.	August 24, 2021
AWSBackupFullAccess – Update to an existing policy	Added `iam:CreateServiceLinkedRole` to create a service-linked role (on a best-effort basis) to automate the deletion of expired recovery points for you. Without this service-linked role, AWS Backup cannot delete expired recovery points after customers delete the original IAM role they used to create their recovery points.	July 5, 2021
AWSBackupServiceLinkedRolePolicyForBackup – Update to an existing policy	Added the new action `dynamodb:DeleteBackup` to grant `DeleteRecoveryPoint` permission to automate the deletion of expired DynamoDB recovery points based on your backup plan lifecycle settings.	July 5, 2021

Change	Description	Date
AWSBackupOperatorAccess – Update to an existing policy	Removed the actions `backup:GetRecoveryPointRestoreMetadata` and `rds:DescribeDBSnapshots` because they were redundant. AWS Backup did not need both `backup:GetRecoveryPointRestoreMetadata` and `backup:Get*` as part of `AWSBackupOperatorAccess` Also, AWS Backup did not need both `rds:DescribeDBSnapshots` and `rds:describeDBSnapshots` as part of `AWSBackupOperatorAccess`	May 25, 2021

Change	Description	Date
AWSBackupOperatorAccess – Update to an existing policy	Removed the actions `backup:GetRecoveryPointRestoreMetadata` and `rds:DescribeDBSnapshots` because they were redundant. AWS Backup did not need both `backup:GetRecoveryPointRestoreMetadata` and `backup:Get*` as part of `AWSBackupOperatorAccess`. Also, AWS Backup did not need both `rds:DescribeDBSnapshots` and `rds:describeDBSnapshots` as part of `AWSBackupOperatorAccess`.	May 25, 2021
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	Added the new action `fsx:TagResource` to grant `StartRestoreJob` permission to allow you to apply tags to Amazon FSx file systems during the restore process.	May 24, 2021

Change	Description	Date
AWSBackupServiceRolePolicyForRestores – Update to an existing policy	Added the new actions `ec2:DescribeImages` and `ec2:DescribeInstances` to grant `StartRestoreJob` permission to allow you to restore Amazon EC2 instances from recovery points.	May 24, 2021
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	Added the new action `fsx:CopyBackup` to grant `StartCopyJob` permission to allow you to copy Amazon FSx recovery points across Regions and accounts.	April 12, 2021
AWSBackupServiceLinkedRolePolicyForBackup – Update to an existing policy	Added the new action `fsx:CopyBackup` to grant `StartCopyJob` permission to allow you to copy Amazon FSx recovery points across Regions and accounts.	April 12, 2021
AWSBackupServiceRolePolicyForBackup – Update to an existing policy	Updated to comply with the following requirement: For AWS Backup to create a backup of an encrypted DynamoDB table, you must add the permissions `kms:Decrypt` and `kms:GenerateDataKey` to the IAM role used for backup.	March 10, 2021

Change	Description	Date
AWSBackupFullAccess – Update to an existing policy	Updated to comply with the following requirements: To use AWS Backup to configure continuous backups for your Amazon RDS database, verify the API permission `rds:ModifyDBInstance` exists in the IAM role defined by your Backup plan configuration. To restore Amazon RDS continuous backups, you must add the permission `rds:RestoreDBInstanceToPointInTime` to the IAM role you submitted for restore job. In the AWS Backup console, to describe the range of times available for point-in-time recovery, you must include the `rds:DescribeDBInstanceAutomatedBackups` API permission in your IAM-managed policy.	March 10, 2021
AWS Backup started tracking changes	AWS Backup started tracking changes for its AWS-managed policies.	March 10, 2021

Using service-linked roles for AWS Backup

AWS Backup uses AWS Identity and Access Management (IAM) service-linked roles. A service-linked role is a unique type of IAM role that is linked directly to AWS Backup. Service-linked roles are predefined by AWS Backup and include all the permissions that the service requires to call other AWS services on your behalf.

Topics

- Using roles to back up and copy
- Using roles for AWS Backup Audit Manager
- Using roles for restore testing

Using roles to back up and copy

AWS Backup uses AWS Identity and Access Management (IAM) service-linked roles. A service-linked role is a unique type of IAM role that is linked directly to AWS Backup. Service-linked roles are predefined by AWS Backup and include all the permissions that the service requires to call other AWS services on your behalf.

A service-linked role makes setting up AWS Backup easier because you don’t have to manually add the necessary permissions. AWS Backup defines the permissions of its service-linked roles, and unless defined otherwise, only AWS Backup can assume its roles. The defined permissions include the trust policy and the permissions policy, and that permissions policy cannot be attached to any other IAM entity.

You can delete a service-linked role only after first deleting its related resources. This protects your AWS Backup resources because you can't inadvertently remove permission to access the resources.

For information about other services that support service-linked roles, see AWS Services That Work with IAM and look for the services that have **Yes** in the **Service-Linked Role** column. Choose a **Yes** with a link to view the service-linked role documentation for that service.

Service-linked role permissions for AWS Backup

AWS Backup uses the service-linked role named **AWSServiceRoleForBackup** to create and delete backups of your AWS resources on your behalf.

The AWSServiceRoleForBackup service-linked role trusts the following services to assume the role:

- `backup.amazonaws.com`

To view the permissions for this policy, see AWSBackupServiceLinkedRolePolicyforBackup in the *AWS Managed Policy Reference*.

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. For more information, see Service-linked role permissions in the *IAM User Guide*.

Creating a service-linked role for AWS Backup

You don't need to manually create a service-linked role. When you list resources to back up, set up cross-account backup, or perform backups in the AWS Management Console, the AWS CLI, or the AWS API, AWS Backup creates the service-linked role for you.

> **Important**
>
> This service-linked role can appear in your account if you completed an action in another service that uses the features supported by this role. To learn more, see A New Role Appeared in My IAM Account.

If you delete this service-linked role, and then need to create it again, you can use the same process to recreate the role in your account. When you list resources to back up, set up cross-account backup, or perform backups, AWS Backup creates the service-linked role for you again.

Editing a service-linked role for AWS Backup

AWS Backup does not allow you to edit the AWSServiceRoleForBackup service-linked role. After you create a service-linked role, you cannot change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see Edit a service-linked role description in the *IAM User Guide*.

Deleting a service-linked role for AWS Backup

If you no longer need to use a feature or service that requires a service-linked role, we recommend that you delete that role. That way you don’t have an unused entity that is not actively monitored or maintained. However, you must clean up your service-linked role before you can manually delete it.

Cleaning up a service-linked role

Before you can use IAM to delete a service-linked role, you must first delete any resources used by the role. First, you must delete all your recovery points. Then, you must delete all your backup vaults.

> **Note**
>
> If the AWS Backup service is using the role when you try to delete the resources, then the deletion might fail. If that happens, wait for a few minutes, then try the operation again.

To delete AWS Backup resources used by the AWSServiceRoleForBackup (console)

1. To delete all your recovery points and backup vaults (except for your default vault), follow the procedure in Delete a vault.
2. To delete your default vault, use the following command in the AWS CLI:

```
aws backup delete-backup-vault --backup-vault-name Default --region us-east-1
```

To delete AWS Backup resources used by the AWSServiceRoleForBackup (AWS CLI)

1. To delete all your recovery points, use delete-recovery-point.
2. To delete all your backup vaults, use delete-backup-vault.

To delete AWS Backup resources used by the AWSServiceRoleForBackup (API)

1. To delete all your recovery points, use `DeleteRecoveryPoint`.
2. To delete all your backup vaults, use `DeleteBackupVault`.

Manually delete the service-linked role

Use the IAM console, the AWS CLI, or the AWS API to delete the AWSServiceRoleForBackup service-linked role. For more information, see Delete a service-linked role in the *IAM User Guide*.

Supported Regions for AWS Backup service-linked roles

AWS Backup supports using service-linked roles in all of the Regions where the service is available. For more information, see AWS Backup supported features and Regions.

Using roles for AWS Backup Audit Manager

AWS Backup uses AWS Identity and Access Management (IAM) service-linked roles. A service-linked role is a unique type of IAM role that is linked directly to AWS Backup. Service-linked roles are predefined by AWS Backup and include all the permissions that the service requires to call other AWS services on your behalf.

A service-linked role makes setting up AWS Backup easier because you don't have to manually add the necessary permissions. AWS Backup defines the permissions of its service-linked roles, and unless defined otherwise, only AWS Backup can assume its roles. The defined permissions include the trust policy and the permissions policy, and that permissions policy cannot be attached to any other IAM entity.

You can delete a service-linked role only after first deleting its related resources. This protects your AWS Backup resources because you can't inadvertently remove permission to access the resources.

For information about other services that support service-linked roles, see AWS Services That Work with IAM and look for the services that have **Yes** in the **Service-Linked Role** column. Choose a **Yes** with a link to view the service-linked role documentation for that service.

Service-linked role permissions for AWS Backup

AWS Backup uses the service-linked role named **AWSServiceRoleForBackupReports** – Provides AWS Backup with permission to create controls, frameworks, and reports.

The AWSServiceRoleForBackupReports service-linked role trusts the following services to assume the role:

- `reports.backup.amazonaws.com`

To view the permissions for this policy, see AWSServiceRolePolicyForBackupReports in the *AWS Managed Policy Reference*.

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. For more information, see Service-linked role permissions in the *IAM User Guide*.

Creating a service-linked role for AWS Backup

You don't need to manually create a service-linked role. When you create a framework or a report plan in the AWS Management Console, the AWS CLI, or the AWS API, AWS Backup creates the service-linked role for you.

> **Important**
>
> This service-linked role can appear in your account if you completed an action in another service that uses the features supported by this role. To learn more, see A New Role Appeared in My IAM Account.

If you delete this service-linked role, and then need to create it again, you can use the same process to recreate the role in your account. When you create a framework or a report plan, AWS Backup creates the service-linked role for you again.

Editing a service-linked role for AWS Backup

AWS Backup does not allow you to edit the AWSServiceRoleForBackupReports service-linked role. After you create a service-linked role, you cannot change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see Edite a service-linked role description in the *IAM User Guide*.

Deleting a service-linked role for AWS Backup

If you no longer need to use a feature or service that requires a service-linked role, we recommend that you delete that role. That way you don’t have an unused entity that is not actively monitored or maintained. However, you must clean up your service-linked role before you can manually delete it.

Cleaning up a service-linked role

Before you can use IAM to delete a service-linked role, you must first delete any resources used by the role. You must delete all frameworks and report plans.

> **Note**
>
> If the AWS Backup service is using the role when you try to delete the resources, then the deletion might fail. If that happens, wait for a few minutes, then try the operation again.

To delete AWS Backup resources used by the AWSServiceRoleForBackupReports (console)

1. To delete all frameworks, see Deleting frameworks.
2. To delete all report plans, see Deleting report plans.

To delete AWS Backup resources used by the AWSServiceRoleForBackupReports (AWS CLI)

1. To delete all frameworks, use delete-framework.
2. To delete all report plans, use delete-report-plan.

To delete AWS Backup resources used by the AWSServiceRoleForBackupReports (API)

1. To delete all frameworks, use DeleteFramework.
2. To delete all report plans, use DeleteReportPlan.

Manually delete the service-linked role

Use the IAM console, the AWS CLI, or the AWS API to delete the AWSServiceRoleForBackupReports service-linked role. For more information, see Delete a service-linked role in the *IAM User Guide*.

Supported Regions for AWS Backup service-linked roles

AWS Backup supports using service-linked roles in all of the Regions where the service is available. For more information, see AWS Backup supported features and Regions.

Using roles for restore testing

AWS Backup uses AWS Identity and Access Management (IAM) service-linked roles. A service-linked role is a unique type of IAM role that is linked directly to AWS Backup. Service-linked roles are predefined by AWS Backup and include all the permissions that the service requires to call other AWS services on your behalf.

A service-linked role makes setting up AWS Backup easier because you don't have to manually add the necessary permissions. AWS Backup defines the permissions of its service-linked roles, and unless defined otherwise, only AWS Backup can assume its roles. The defined permissions include the trust policy and the permissions policy, and that permissions policy cannot be attached to any other IAM entity.

You can delete a service-linked role only after first deleting its related resources. This protects your AWS Backup resources because you can't inadvertently remove permission to access the resources.

For information about other services that support service-linked roles, see AWS Services That Work with IAM and look for the services that have **Yes** in the **Service-Linked Role** column. Choose a **Yes** with a link to view the service-linked role documentation for that service.

Service-linked role permissions for AWS Backup

AWS Backup uses the service-linked role named **AWSServiceRoleForBackupRestoreTesting** – Provides backup permissions to conduct restore testing.

The **AWSServiceRoleForBackupRestoreTesting** service-linked role trusts the following services to assume the role:

- `restore-testing.backup.amazonaws.com`

To view the permissions for this policy, see AWSServiceRolePolicyForBackupRestoreTesting in the *AWS Managed Policy Reference*.

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. For more information, see Service-linked role permissions in the *IAM User Guide*.

Creating a service-linked role for AWS Backup

You don't need to manually create a service-linked role. When you conduct restore testing in the AWS Management Console, the AWS CLI, or the AWS API, AWS Backup creates the service-linked role for you.

> **Important**
>
> This service-linked role can appear in your account if you completed an action in another service that uses the features supported by this role. To learn more, see A New Role Appeared in My IAM Account.

If you delete this service-linked role, and then need to create it again, you can use the same process to recreate the role in your account. When you conduct restore testing, AWS Backup creates the service-linked role for you again.

Editing a service-linked role for AWS Backup

AWS Backup does not allow you to edit the AWSServiceRoleForBackupRestoreTesting service-linked role. After you create a service-linked role, you cannot change the name of the role because various entities might reference the role. However, you can edit the description of the role using IAM. For more information, see Edit a service-linked role description in the *IAM User Guide*.

Deleting a service-linked role for AWS Backup

If you no longer need to use a feature or service that requires a service-linked role, we recommend that you delete that role. That way you don't have an unused entity that is not actively monitored or maintained. However, you must clean up your service-linked role before you can manually delete it.

Cleaning up a service-linked role

Before you can use IAM to delete a service-linked role, you must first delete any resources used by the role. You must delete all restore testing plans.

> **Note**
>
> If the AWS Backup service is using the role when you try to delete the resources, then the deletion might fail. If that happens, wait for a few minutes, then try the operation again.

To delete AWS Backup resources used by the AWSServiceRoleForBackupRestoreTesting (console)

- To delete all restore testing plans, see Restore testing.

To delete AWS Backup resources used by the AWSServiceRoleForBackupRestoreTesting (AWS CLI)

- To delete restore testing plans, use `delete-restore-testing-plan`.

To delete AWS Backup resources used by the AWSServiceRoleForBackupRestoreTesting (API)

- To delete restore testing plans, use `DeleteRestoreTestingPlan`.

Manually delete the service-linked role

Use the IAM console, the AWS CLI, or the AWS API to delete the **AWSServiceRoleForBackupRestoreTesting** service-linked role. For more information, see Delete a service-linked role in the *IAM User Guide*.

Supported Regions for AWS Backup service-linked roles

AWS Backup supports using service-linked roles in all of the Regions where the service is available. For more information, see AWS Backup supported features and Regions.

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 can coerce a more-privileged entity to perform the action. In AWS, cross-service impersonation can result in the confused deputy problem. Cross-service impersonation can occur when one service (the *calling service*) calls another service (the *called service*). The calling service can be manipulated to use its permissions to act on another customer's resources in a way it should not otherwise have permission to access. To prevent this, AWS provides tools that help you protect your data for all services with service principals that have been given access to resources in your account.

We recommend using the `aws:SourceArn` and `aws:SourceAccount` global condition context keys in resource policies to limit the permissions that AWS Backup gives another service to the resource. If you use both global condition context keys, the `aws:SourceAccount` value and the account in the `aws:SourceArn` value must use the same account ID when used in the same policy statement.

The value of `aws:SourceArn` must be a AWS Backup vault when using AWS Backup to publish Amazon SNS topics on your behalf.

The most effective way to protect against the confused deputy problem is to use the `aws:SourceArn` global condition context key with the full ARN of the resource. If you don't know the full ARN of the resource or if you are specifying multiple resources, use the `aws:SourceArn` global context condition key with wildcards (*) for the unknown portions of the ARN. For example, `arn:aws::`*`servicename`*`::`*`123456789012`*`:*`.

The following example shows how you can use the `aws:SourceArn` and `aws:SourceAccount` global condition context keys in AWS Backup to prevent the confused deputy problem.
Add the following statement to your *KMS key policy* to deny the service principal `backup-`

`storage.amazonaws.com` from performing KMS actions unless the request originates from your specified backup vaults and account:

```
{
  "Sid": "Deny Backup Storage confused deputy",
  "Effect": "Deny",
  "Principal": {
    "Service": "backup-storage.amazonaws.com"
  },
  "Action": [
    "kms:Decrypt",
    "kms:RetireGrant",
    "kms:GenerateDataKey"
  ],
  "Resource": "*",
  "Condition": {
    "StringNotEquals": {
      "aws:SourceAccount": "123456789012"
    },
    "ArnNotLike": {
      "aws:SourceArn": "arn:aws::backup:us-east-1:123456789012:backup-vault:*"
    }
  }
}
```

Replace *us-east-1* with your AWS Region and *123456789012* with your AWS account ID. This policy denies the AWS Backup storage service principal from using your KMS key unless the request comes from a backup vault in your specified account and Region.

Infrastructure security in AWS Backup

As a managed service, AWS Backup is protected by AWS global network security. For more 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 AWS Backup through the network. Clients must support Transport Layer Security (TLS) 1.2 or later. Clients must also support cipher suites with perfect forward secrecy (PFS) such as Ephemeral Diffie-Hellman (DHE) or Elliptic Curve Diffie-Hellman Ephemeral (ECDHE). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an IAM principal. Or you can use the AWS Security Token Service (AWS STS) to generate temporary security credentials to sign requests.

Integrity of Data in AWS Backup

AWS Backup data integrity goal

AWS Backup seeks to maintain integrity during transmission, storage, and processing of your data. AWS Backup treats stored resource data as content-agnostic critical information, in that we offer the same high level of security to customers, regardless of the type of data you store. We are vigilant about our customers' security and have implemented sophisticated technical and physical measures against unauthorized access. You retain complete control over how your data is classified, the Regions in which you store your data, and how you control, archive, and protect your data against disclosure.

AWS Backup data integrity implementation

AWS Backup works in concert with other AWS and Amazon services to maintain integrity of the data it stores and with which it interacts. The tools used may vary and can include (but are not limited to):

- Continuous object validation against their checksum to prevent object corruption
- Internal checksums to confirm integrity of data in transit and at rest
- Checksums calculated on data in backups created from the primary store
- Checksums are always verified to be correct before using the corresponding data. If we find data that doesn't match its checksum, we replace it with a correct copy. If we fail to replace the correct copy, we will fail the corresponding jobs
- Automatic attempt to restore normal levels of object storage redundancy in the event of disk corruption or detection of device failure
- Redundant storage of data across multiple physical locations
- Object durability enhancement across multiple availability zones during the initial write, combined with further replication in the event of device unavailability or detected bit-rot
- Checksums on all network traffic to detect corruption of data packets when storing or retrieving data

AWS Backup natively stores data for Amazon DynamoDB with advanced features, Amazon EFS, Amazon S3, Amazon Timestream, and virtual machines running with VMware connected through Backup gateway. AWS Backup facilitates backups of data stored with other services, including Amazon Aurora, Amazon DocumentDB, Amazon DynamoDB, Amazon EBS, Amazon EC2, Amazon FSx for Windows File Server, Amazon FSx for Lustre, Amazon FSx for OpenZFS, Amazon FSx for NetApp ONTAP, Amazon Neptune, Amazon RDS, and Amazon Redshift.

Objective confirmation and audit of AWS Backup data integrity

The data stored directly by AWS Backup and the data stored in partnership with fellow AWS services with which AWS Backup interacts is subjected to the rigorous process of Amazon Simple Storage Service (Amazon S3) underpinning this data integrity. This integrity is confirmed by an independent, third-party auditor through an annual SOC audit report which is available through AWS Artifact.

Legal holds and AWS Backup

Legal hold overview

A legal hold is an administrative tool that helps prevent backups from being deleted while under a hold. While the hold is in place, backups under a hold cannot be deleted and lifecycle policies that would alter the backup status (such as transition to a `Deleted` state) are delayed until the legal hold is removed.

Legal holds can be applied to one or more backups (also known as recovery points) created by AWS Backup if their lifecycles allow it. Legal holds do not apply to continuous backups.

When a legal hold is created, it can take into account specific filtering criteria, such as resource types and resource IDs. Additionally, you can define the creation date range of backups you wish to include in a legal hold.

Legal holds apply only to the original backup on which they are placed. When a backup is copied across Regions or accounts (if the resource supports it), it does not retain or carry its legal hold with it. A legal hold, like other resources, has a unique ARN (Amazon Resource Name) associated with it. Only recovery points created by AWS Backup can be part of a legal hold.

Note that while AWS Backup Vault Lock provides additional protections and immutability to a vault, a legal hold provides additional protection against deletion of individual backups (recovery

points). The legal hold does not expire and retains the data within the backup indefinitely. The hold remains active until it is released by a user with sufficient permissions.

Multiple legal holds

A backup can have more than one legal hold. Legal holds and backups have a many:many relationship, meaning that a backup can have more than legal hold and a legal hold can include more than one backup.

A backup cannot be deleted as long as it has at least one legal hold. After all legal holds on a backup are removed, it is subject to its retention lifecycle properties. Maintain at least one legal hold to prevent backup deletion. Legal holds can be applied to a recovery point retained past its backup lifecycle retention date (due to an existing legal hold).

Each account can have a maximum of 50 legal holds active at one time.

Create a legal hold

A legal hold can be added to an existing backup (recovery point).

Backups (recovery points) with a status of `EXPIRED` or `DELETING` will not be included in the legal hold. Recovery points (backups) with the status of `CREATING` may not be included in the legal hold, depending on the time of completion.

Legal holds can be added by users who have the required IAM permissions.

Create a legal hold using the console

To create a legal hold

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. In the dashboard in the left of the console, find My Account. Choose **Legal holds**.
3. Choose **Add legal hold**.
4. Three panels are shown: **Legal hold details**, **Legal hold scope**, and **Legal hold tags**.

    a. Under **Legal hold details**, enter a legal hold title and a description for the hold in the text boxes provided.

    b. In the panel **Legal hold scope**, choose how you wish to select the resource to include in the hold. When you create a hold, you choose the method used to select the resources that are within the legal hold. You can choose to include one of the following:

- Specific resource types and IDs
- Select backup vaults
- All resources types or all backup vaults within your account

c. Specify the date range of your legal hold. Enter the dates in the YYYY:MM:DD format (dates are inclusive).

d. Optionally, you can add tags for the hold under **Legal hold tags**. Tags can help categorize the hold for future reference and organization. You can add up to 50 tags total.

5. When you are satisfied with the configuration of your new legal hold, click the button **Add new hold**.

Create a legal hold using the AWS CLI

You can create a legal hold using the create-legal-hold command.

```
aws backup create-legal-hold --title "my title" \
    --description "my description" \
    --recovery-point-selection
 "VaultNames=string,DateRange={FromDate=timestamp,ToDate=timestamp}"
```

View legal holds

You can see legal hold details in the AWS Backup console or programmatically.

View legal holds using the console

To view all legal holds within an account using the Backup console,

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Using the left part of the dashboard, under **My account**, click **Legal holds**.
3. The **legal hold** table displays the title, status, description, ID, and creation date of existing holds. Click on the carat (down arrow) next to the table header to filter the table by the selected column.

View legal holds programatically

To view all legal holds programmatically, you can use the following API calls: ListLegalHolds and GetLegalHold.

The following JSON template can be used for `GetLegalHold`.

```
GET /legal-holds/{legalHoldId} HTTP/1.1

Request

empty body

Response

{
    Title: string,
    Status: LegalHoldStatus,
    Description: string, // 280 chars max
    CancelDescription: string, // this is provided during cancel // 280 chars max
    LegalHoldId: string,
    LegalHoldArn: string,
    CreatedTime: number,
    CanceledTime: number,

   ResourceSelection: {
        VaultArns: [ string ]
        Resources: [ string ]
   },
   ResourceFilters: {
        DateRange: {
          FromDate: number,
          ToDate: number
        }
   }
}
```

The following JSON template can be used for `ListLegalHolds`.

```
GET /legal-holds/
  &maxResults=MaxResults
  &nextToken=NextToken


Request

empty body
url params:
```

```
  MaxResults: number  // optional,
  NextToken: string  // optional

status: Valid values: CREATING | ACTIVE | CANCELED | CANCELING
maxResults: 1-1000



Response

{
  NextToken: token,
  LegalHolds: [
    Title: string,
    Status: string,
    Description: string, // 280 chars max
    CancelDescription: string, // this is provided during cancel // 280 chars max
    LegalHoldId: string,
    LegalHoldArn: string,
    CreatedTime: number,
    CanceledTime: number,
  ]

}
```

The following are the possible status values.

Status	Description
CREATING	Requested recovery points are in the process of being held, and delete requests of those recovery points may be successful since the hold hasn't finished being created.
ACTIVE	The legal hold has been created, All recovery points listed under this legal hold are held.
CANCELLING	Legal holds are in the process of being removed, and delete requests of recovery points under the hold may succeed.

Status	Description
CANCELED	Legal hold is fully released and no longer has any effect. Recovery points can be deleted.

View recovery points in legal holds

You can see recovery points in legal holds in the AWS Backup console or programmatically.

View recovery points in legal holds using the console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Using the left part of the dashboard, under **My account**, click **Legal holds**.
3. Select a legal hold, then you can view **Legal hold details**, **Resources selected for this hold**, **Recovery points on hold**, and **Tags** for this hold.

View recovery points in legal holds programmatically

To view recovery points in a legal hold programmatically, you can use the ListRecoveryPointsByLegalHold API call.

The following is an example request and response for `ListRecoveryPointsByLegalHold`.

```
GET /legal-holds/{legalHoldId}/recovery-points HTTP/1.1

Request

empty body

Response

{
  "RecoveryPoints": [
    {
      "RecoveryPointArn": "string",
      "ResourceArn": "string",
      "ResourceType": "string",
      "BackupVaultName": "string"
    }
  ],
```

```
  "NextToken": "string"
}
```

Release a legal hold

Legal holds remain in effect until they are removed by a user with sufficient permissions. Removing a legal hold is also known as cancelling, deleting, or releasing a legal hold. Removing a legal hold eliminates it from all backups to which it was attached. Any backups that expired during the legal hold are deleted within 24 hours after the legal hold is removed.

Release a legal hold using the console

To release a hold using the console

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Enter the description you would like associated with the release.
3. Review the details, then click **Release hold**.
4. When the Release hold dialogue box appears, confirm your intent to release the hold by typing `confirm` into the text box.
    - Check the box that acknowledges you are cancelling the hold.

On the **Legal holds** page you can see all your holds. If the release was successful, the status of that hold will be shown as `Released`.

Release a legal hold programmatically

To remove a hold programmatically, use the API call CancelLegalHold.

Use the following JSON template.

```
DELETE /legal-holds/{legalHoldId}


Request

{
   CancelDescription: String
   DeleteAfterDays: number // optional
}
```

```
DeleteAfterDays: optional.
  Defaults to 180 days. how long to keep legal hold record after canceled.
  This applies to the actual legal hold record only.
  Recovery points are unlocked as soon as cancelation processes and are not subject to
 this date.

Response

Empty body

200 if successful
other standard codes
```

Malware protection in AWS Backup

Malware scanning of your backups is provided by Amazon GuardDuty Malware Protection. Using Amazon GuardDuty Malware Protection for AWS Backup allows you to automate scanning of recovery points through existing backup workflows, or initiate on-demand scans of previously created backups. This AWS native solution helps ensure your backups are clean from potential malware, allowing you to meet compliance requirements and respond to malicious incidents faster by ensuring the recovery of clean data.

To see a list of supported resource types and regions, visit the feature availability page.

Topics

- Integration with Amazon GuardDuty
- How to use malware scanning
- Access
- Incremental vs full scans
- Monitoring your malware scans
- Understanding scan results
- Troubleshooting scan failures
- Metering
- Quotas
- Console and CLI usage steps for malware scan types

Integration with Amazon GuardDuty

AWS Backup integrates with Amazon GuardDuty Malware Protection to provide threat detection for your recovery points. When you start a malware scan, AWS Backup automatically calls Amazon GuardDuty's `StartMalwareScan` API after each backup completes, passing the recovery point details and your scanner role credentials. Amazon GuardDuty then begins reading, decrypting, and scanning all files and objects within the backup.

When Amazon GuardDuty accesses your backup data, that access is logged in AWS CloudTrail for visibility.

For more information about this integration, see the Amazon GuardDuty Malware Protection documentation.

How to use malware scanning

When you use Amazon GuardDuty Malware Protection with AWS Backup, you can automatically scan your backups for malware. This integration helps you detect malicious code in your backups and identify clean recovery points for restore operations.

Amazon GuardDuty Malware Protection supports two primary workflows for scanning your backups:

- **Automatic malware scanning through backup plans** – Enable malware scanning in backup plans to automate malware detection with AWS Backup. When enabled, AWS Backup automatically initiates an Amazon GuardDuty scan after each successful backup completion. You can configure either full or incremental scanning for specific backup plan rules, which determines how frequently your backups are scanned. For more information about scan types, see the section called “Incremental vs full scans” below. AWS Backup recommends enabling automatic malware scanning in backup plans for proactive threat detection after backup creation.
- **On-demand scans** – Run on-demand scans to manually scan existing backups, choosing between full or incremental scan types. AWS Backup recommends using on-demand scans to identify your last clean backup. When scanning before a restore operation, use a full scan to examine the entire backup with the latest threat detection model.

Access

Before you start with malware protection, your account must have required permissions for the operations.

AWS Backup malware scanning requires two IAM roles to scan your recovery points for potential malware:

- First, the AWSBackupServiceRolePolicyForBackup or AWSBackupServiceRolePolicyForScans managed policy must be attached to your existing or new backup role. This is the same role found in the resource assignment for your backup plan in the console, or through the BackupSelection API. This managed policy allows AWS Backup to initiate malware scans with Amazon GuardDuty.
- Second, a new scanner role is required with AWSBackupGuardDutyRolePolicyForScans managed policy that trusts `malware-protection.guardduty.amazonaws.com`. To create this role, you can specify a custom trust policy and attach the required managed policy. This is the same role found in the malware protection portion of your backup plan in the console or through the scan settings in your BackupPlan API. This role is passed by AWS Backup to Amazon GuardDuty when a scan is initiated, providing access to backups.

Incremental vs full scans

With malware scanning, you have the option to choose between incremental and full scans based on your security requirements and cost considerations.

Incremental scans analyze only the data that changed between the target and base recovery point. These scans are faster and more cost-effective for regular scanning, making them ideal for frequent periodic backups where you want to scan newly backed up data.

Even when incremental scanning is selected, AWS Backup performs a full scan in these situations:

- **First-time scans:** The initial scan of a resource is always a full scan, allowing Amazon GuardDuty to establish a baseline of potential threats. Subsequent scans will then be incremental.
- **Expired baseline:** If your baseline recovery point was scanned more than 365 days ago, a full scan occurs. Because Amazon GuardDuty retains finding information for only 365 days, a new baseline must be established to ensure accurate scanning results.
- **Deleted baseline:** If your base recovery point is deleted before your next incremental scan starts, a full scan occurs automatically.

Full scans examine the entire recovery point regardless of previous scans. While these scans provide comprehensive coverage, they take longer to complete and incur higher costs. You can run full scans on-demand or schedule them through your backup plans. AWS Backup recommends

configuring periodic full scans in your backup plans at extended intervals to ensure your entire backup data is regularly scanned with the latest malware signature model.

For optimal security versus cost management, consider your backup frequency when choosing scan types. When you configure malware scanning for Amazon S3 continuous backups, your scan job will run at the frequency of the corresponding backup rule.

> **Note**
>
> On-demand scans for Amazon S3 continuous backups do not support incremental scans. Instead you can scan up to any restorable point in time.

> **Note**
>
> Incremental malware scanning is not supported for Amazon EC2 recovery points in a logically air-gapped vault or copied Amazon EC2 recovery points.

> **Note**
>
> AWS Backup does not support using a recovery point scanned only via on-demand scan from Amazon GuardDuty as a base recovery point for incremental scans. To perform incremental scans via AWS Backup, ensure AWS Backup has previously scanned the base recovery point.

Monitoring your malware scans

After scanning is enabled, both AWS Backup and Amazon GuardDuty provide monitoring and notification mechanisms you can use to track your results:

- **AWS Backup Console:** The AWS Backup console is powered by `ListScanJobs` and `DescribeScanJob` APIs. You can visit the Malware protection section to view the list of scan jobs, representing job status and scan results. AWS Backup also supports a `ListScanJobSummaries` API, though not available in the console.
- **AWS Backup Audit Manager:** You can set up a scanning report to view all AWS Backup initiated malware scan jobs over the last 24 hours.

- **Amazon GuardDuty Console:** If base Amazon GuardDuty is enabled, you can view details in the Malware Scan results and investigate malware on the Amazon GuardDuty findings page. You can view information such as the threat and file name, file path, objects/files scanned, bytes scanned, etc. Note that this detailed threat information is not available through AWS Backup, and you must have the appropriate Amazon GuardDuty permissions to view this information.
- **Amazon EventBridge:** Both AWS Backup and Amazon GuardDuty emit EventBridge events, allowing backup and security administrators to be alerted synchronously. You can set up custom rules to receive notifications when scans complete or malware is detected.
- **AWS CloudTrail:** Both AWS Backup and Amazon GuardDuty emit CloudTrail events, allowing you to monitor API access.

Understanding scan results

Your scan jobs from AWS Backup will have a scan state and scan result.

Scan States

The scan state indicates the job state and can have values of: `CREATED`, `COMPLETED`, `COMPLETED_WITH_ISSUES`, `RUNNING`, `FAILED`, or `CANCELED`.

There are multiple situations in which your scan job will finish with state `COMPLETED_WITH_ISSUES`:

For Amazon S3 backups, there are object size/type limitations which will prevent objects from being scanned. When at least one object is skipped within a scan, the corresponding scan job will be marked as `COMPLETED_WITH_ISSUES`. For Amazon EC2/Amazon EBS backups, there are volume size/quantity limitations which result in volumes being skipped during scanning. These situations will result in an Amazon EC2/Amazon EBS backup job to result in `COMPLETED_WITH_ISSUES`.

If your job finishes with state `COMPLETED_WITH_ISSUES` and you need further information about the reasons, you will need to get those details from the corresponding scan job through Amazon GuardDuty.

> **Note**
>
> Incremental scan jobs only scan the difference in data between two backups. Therefore, if an incremental scan job does not encounter any of the situations described above, it will

finish in state COMPLETE and will not inherit the COMPLETED_WITH_ISSUES from the base recovery point.

In rare cases, Amazon GuardDuty might experience internal issues when scanning files and objects, exhausting all retry attempts. When this happens, AWS Backup displays the scan job as FAILED, while Amazon GuardDuty shows it as COMPLETED_WITH_ISSUES. This status difference enables you to view available scan results in Amazon GuardDuty while indicating that Amazon GuardDuty did not successfully scan all supported files and objects. In these cases, AWS Backup cannot use the recovery point with a FAILED scan as a base recovery point for subsequent incremental scans.

Scan Results

The scan results indicate an aggregated result from Amazon GuardDuty and can have values of: THREATS_FOUND, or NO_THREATS_FOUND.

Scan results indicate whether potential malware was detected in your recovery points. A NO_THREATS_FOUND status means no potential malware was detected, while THREATS_FOUND indicates potential malware was discovered. For detailed threat information, access the full Amazon GuardDuty findings through the Amazon GuardDuty console or APIs. Scan results are also available through EventBridge events, allowing you to build automated workflows that respond to infected backups.

Amazon GuardDuty retains findings for 365 days, tracking files or objects across incremental scans to monitor if threats are removed or malware signatures change. For example, if malware is detected in backup 2, the scan result shows THREATS_FOUND. When you perform an incremental scan on backup 3 using backup 2 as a base, the scan result remains THREATS_FOUND unless the threat has been removed from the data.

Note

For malware scans of Amazon S3 continuous backups, you can find your latest clean restorable point in time using the GetPITRMalwareScanResults API, or natively through the AWS Backup console. If you query a time which has not been scanned, you will receive a scan result of unknown.

Troubleshooting scan failures

Common scan failures include insufficient IAM permissions, service limits, and resource access issues.

Permission errors occur when the backup role lacks `AWSBackupServiceRolePolicyForScans` permissions or the scanner role doesn't have `AWSBackupGuardDutyRolePolicyForScans` with proper trust relationships.

Service limit errors happen when you exceed the 150 concurrent scans per account or 5 concurrent scans per resource type - scan jobs will remain in `CREATED` state until capacity becomes available.

Access denied errors may indicate encrypted recovery points without proper AWS KMS permissions or deleted parent recovery points for incremental scans.

Timeout failures can occur with very large recovery points or during high Amazon GuardDuty load periods.

To troubleshoot, check the scan job status using `DescribeScanJob` API, verify IAM role configurations, ensure recovery points exist and are accessible, and consider switching to full scans if incremental scan parent references are missing.

Monitor your concurrent scan usage and implement jittering in automated workflows to avoid hitting service limits.

Metering

Malware protection is provided and billed by Amazon GuardDuty. You will not see any AWS Backup charges related to using this feature. All usage can be viewed under Amazon GuardDuty's billing. To learn more, visit Amazon GuardDuty pricing.

Quotas

Both AWS Backup and Amazon GuardDuty have quota limits for Amazon GuardDuty Malware Protection for AWS Backup.

For more information, visit AWS Backup quotas and Amazon GuardDuty quotas.

Console and CLI usage steps for malware scan types

The following sections show the steps for configuring different malware scan types using both the console and AWS CLI.

How to set up malware scans

Console

1. Navigate to AWS Backup console → Backup plans
2. Create new backup plan or select existing plan
3. Enable **Malware protection** toggle
4. Select **Scanner role** to choose a new scanner role. Make sure both backup role and scanner role have appropriate permissions as discussed in the section called “Access”.
5. Select **Scannable resource types**. This will filter malware scanning to the resource selection criteria you have chosen. As an example, if your scannable resource type selection is Amazon EBS, but your plan's resource selection includes Amazon EBS and Amazon S3, then only Amazon EBS malware scans will take place.
6. Set **Scan type** for each backup rule. You can choose between full, incremental, and no scan. The scan type selection means that scan will occur at the schedule frequency of the associated backup rule.
7. Save backup plan

AWS CLI

CreateBackupPlan

You can create a backup plan with malware scanning enabled using the create-backup-plan command.

```
aws backup create-backup-plan \
    --region us-west-2 \
    --cli-input-json '{
                        "BackupPlan": {
                          "BackupPlanName": "scan-initial-test-demo",
                          "Rules": [
                            {
```

```
            "RuleName": "full",
            "TargetBackupVaultName": "Default",
            "ScheduleExpression": "cron(0 * * * ? *)",
            "StartWindowMinutes": 120,
            "CompletionWindowMinutes": 6000,
            "Lifecycle": {
              "DeleteAfterDays": 3,
              "OptInToArchiveForSupportedResources": true
            },
            "RecoveryPointTags": {
              "key1": "foo",
              "key2": "foo"
            },
            "EnableContinuousBackup": true,
            "ScanActions": [
              {
                "MalwareScanner": "GUARDDUTY",
                "ScanMode": "FULL_SCAN"
              }
            ]
          },
          {
            "RuleName": "incremental",
            "TargetBackupVaultName": "Default",
            "ScheduleExpression": "cron(30 * * * ? *)",
            "StartWindowMinutes": 100,
            "CompletionWindowMinutes": 5000,
            "Lifecycle": {
              "DeleteAfterDays": 2,
              "OptInToArchiveForSupportedResources": true
            },
            "RecoveryPointTags": {
              "key1": "foo",
              "key2": "foo"
            },
            "EnableContinuousBackup": true,
            "ScanActions": [
              {
                "MalwareScanner": "GUARDDUTY",
                "ScanMode": "INCREMENTAL_SCAN"
              }
            ]
          }
        ],
```

```
                        "ScanSettings": [
                          {
                            "MalwareScanner": "GUARDDUTY",
                            "ResourceTypes": ["EBS", "EC2", "S3"],
                            "ScannerRoleArn": "arn:aws:iam::300949271314:role/
TestBackupScannerRole"
                          }
                        ]
                      }
                    }'
```

UpdateBackupPlan

You can update a backup plan with malware scanning enabled using the update-backup-plan command.

```
aws backup update-backup-plan \
    --region us-west-2 \
    --cli-input-json '{
                      "BackupPlanId": "d1391282-68cf-4fce-93ad-e08bc5178bac",
                      "BackupPlan": {
                      "BackupPlanName": "scan-initial-test-demo",
                      "Rules":
                        [
                          {"RuleName": "full",
                          "TargetBackupVaultName": "Default",
                          "ScheduleExpression": "cron(0 * * * ? *)",
                          "StartWindowMinutes": 60,
                          "CompletionWindowMinutes": 3000,
                          "Lifecycle":
                            {
                            "DeleteAfterDays": 6,
                            "OptInToArchiveForSupportedResources": false},
                          "RecoveryPointTags":
                            {"key1": "foo",
                            "key2": "foo"},
                          "EnableContinuousBackup": false,
                          "ScanActions":
                            [
                              {"MalwareScanner": "GUARDDUTY",
                              "ScanMode": "FULL_SCAN"}
                            ]
                          },
```

```
                            {
                              "RuleName": "incremental",
                              "TargetBackupVaultName": "Default",
                              "ScheduleExpression": "cron(30 * * * ? *)",
                              "StartWindowMinutes": 120,
                              "CompletionWindowMinutes": 6000,
                              "Lifecycle":
                                {
                                "DeleteAfterDays": 9,
                                "OptInToArchiveForSupportedResources": false},
                              "RecoveryPointTags":
                                {"key1": "foo",
                                "key2": "foo"},
                              "EnableContinuousBackup": false,
                              "ScanActions":
                                [
                                  {
                                  "MalwareScanner": "GUARDDUTY",
                                  "ScanMode": "INCREMENTAL_SCAN"
                                  }
                                ]
                            }
                          ],
                        "ScanSettings":
                          [
                            {
                            "MalwareScanner": "GUARDDUTY",
                            "ResourceTypes":
                              ["ALL", "EBS"],
                            "ScannerRoleArn": "arn:aws:iam::300949271314:role/
TestBackupScannerRole"
                            }
                          ]
                        }
                      }'
```

Key Notes

- Target ARN entry is required before scan options become enabled (Console)
- Both backup IAM role and scanner IAM role required for all configurations
- Use `aws backup list-scan-jobs` to view all scan jobs (AWS CLI)

- Cost implications vary by scan type (incremental vs full) and frequency

AWS CLI Key Notes

- Use `aws backup list-scan-jobs` to view all scan jobs (AWS CLI)
- Scan results available via `describe-recovery-point` API with ScanResults field
- Both backup IAM role and scanner IAM role required for all configurations
- JSON backup plan structure includes ScanSettings at plan level and ScanActions in rules

Resilience in AWS Backup

AWS Backup takes its resilience — and your data security — extremely seriously.

AWS Backup stores your backups with *at least* as much resilience and durability as your resource's original AWS service would give you, if you backed it up there.

AWS Backup is designed to use the AWS global infrastructure to resiliently and redundantly store your backups across multiple Availability Zones for durability of 99.999999999% (11 nines) in any given year, provided that you adhere to the current AWS Backup documentation.

AWS Backup encrypts your backup plans at rest and continuously backs them up. You can also restrict access to your backup plans using AWS Identity and Access Management (IAM) credentials and policies. For more information, see Authentication, Access Control, and Security Best Practices in IAM.

The AWS global infrastructure is built around AWS Regions and Availability Zones. AWS Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. AWS Backup stores your backups across Availability Zones. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures. For more information, see AWS Backup Service Level Agreement (SLA).

Furthermore, AWS Backup empowers you to copy your backups across Regions for even greater resilience. For more information about the AWS Backup cross-Region copy feature, see Creating a Backup Copy.

For more information about AWS Regions and Availability Zones, see AWS Global Infrastructure.

AWS Backup quotas

The following quotas apply when working with AWS Backup.

Quotas

- Backup
- Backup index and search quotas
- Policy quotas
- Malware scanning quotas
- Amazon Timestream resource quotas
- AWS Backup Audit Manager quotas
- Restore testing plan quotas
- AWS Backup gateway quotas
- Amazon EKS quotas
- Logically air-gapped vault quotas
- Related quotas

Backup

Name	Default	Adjustable
Total vaults (backup and logically air-gapped) per Region per account	300	Yes
Recovery points per backup vault	1,000,000	Yes
Backup plans per Region per account	300	Yes
Versions per backup plan	2,000	Yes

Name	Default	Adjustable
Resource assignments per backup plan	100	No
Concurrent cross-Region or in-Region copy jobs per account in destination Region	100[2]	No
Additional cross-Region copy jobs per vault in a destinati on Region after the limit in row above entry has been reached.[1]	5[2]	No
Concurrent cross-account copies that can be made of the same resource to the same destination Region	30	No
Concurrent backup and copy jobs per resource	1	No
Tags per resource selection in a cross account backup policy	30	No. Include additional tags using multiple resource assignments or backups plans.
Hypervisors	20	No
Legal holds per account	50	No
Nested backup layers of application stacks	10	No

[1]The limit for concurrent copy jobs from one Region to another Region is 100 per account per Region. Once this limit is reached, if a specific vault in the destination Region has fewer than 5 concurrent copy jobs, new copy jobs can begin, up to a maximum of 5 concurrently.

[2]Limit only apply to resource types fully managed by AWS Backup.

Backup index and search quotas

Name	Default	Adjustable
Concurrent indexing jobs in account (most AWS Regions)	40	Yes
Concurrent indexing jobs in account in Asia Pacific (Malaysia), Canada (Central) , Asia Pacific (Thailand), and Mexico (Central) AWS Regions.	10	Yes
Concurrent indexing jobs for each resource	5	No
Concurrent on-demand indexing job	1	No
Concurrent search jobs in account	10	
Concurrent export jobs	5	
Number of recovery points included in search job	20	
Concurrent Amazon EBS file level restore jobs (most AWS Regions)	25	
Concurrent Amazon EBS file level restore jobs in Asia Pacific (Malaysia), Canada (Central), Asia Pacific	5	

Name	Default	Adjustable
(Thailand), and Mexico (Central) AWS Regions.		

Policy quotas

Name	Default	Adjustable
Resource assignments per backup plan	100	No
Tags in a resource selection	30	No
Resource selections that use tags in a plan	10	No
Backup plan rules in a plan	10	No
Tags added to a recovery point	10	No
Copy actions per backup rule	5	No
Conditions in a resource assignment in a backup plan	30	No

Malware scanning quotas

The following table lists quotas for AWS Backup malware scanning with Amazon GuardDuty.

Name	Default	Adjustable
Running scan jobs per resource per account	5	Yes

Name	Default	Adjustable
Running scan jobs per recovery point per account	1	No
Running scan jobs per account	150	Yes
Created scan jobs per resource per account	10	No
Incremental scan base constraint	1	No

When you hit your created scan jobs per resource per account limit, we will fail oldest queued job.

You might also encounter quotas imposed by Amazon GuardDuty. For more information, see Amazon GuardDuty quotas in the *AWS General Reference*.

Amazon Timestream resource quotas

Name	Default	Adjustable
Concurrent Timestream backup jobs per account	4	Yes
Concurrent Timestream restore jobs per account	1	Yes

AWS Backup Audit Manager quotas

Name	Default	Adjustable
Frameworks per account per Region	15	Yes

Name	Default	Adjustable
Controls per account per Region	50	Yes
Report plans per account	20	Yes
Frameworks per report plan	1,000	No
[Number of accounts] multiplied by [number of Regions in a report plan] multiplied by [number of daily jobs plus evaluations in a report plan]	100,000	No

Restore testing plan quotas

Name	Default	Adjustable
Restore testing plans	100	No
Tags per plan	50	No
Selections per plan	30	No
ARNs per restore testing selection	30	No
Conditions per selection (both `StringEquals` and `StringNotEquals`)	30	No
Vault selectors per restore testing selection	30	No
Maximum value (in days) of selection window	365 days	No

Name	Default	Adjustable
Boundaries of start window hours	Minimum: 1 hour; Maximum: 168 hours	No
Maximum character length of restore testing plan name	50 characters (alphanumeric and underscores, no white spaces)	No
Maximum character length of restore testing selection name	50 characters (alphanumeric and underscores, no white spaces)	No

Each resource type has a limit on the number of concurrent restore jobs that can exist at one time for restore jobs that are created through a restore testing plan. Once this limit is reached, no new restore jobs for that resource type will be created until a job in a state of `RUNNING` transitions to `COMPLETED`.

If a scheduled restore job did not start due to this quota, that job will result in a `FAILED` status with the status message `"Restore job was unable to start within the specified start window. Try increasing your start window."`. If you receive a failed job with this status message, the best practice is to first increase your start window with sufficient time to allow jobs to finish. Then, retry the jobs.

Note quotas do not apply to on-demand restore jobs, but to restore jobs created by a restore testing plan. For some resource types, you may request an increase in the quota limit.

Name	Default	Adjustable
Amazon Aurora	40	Yes
Amazon DocumentDB	40	Yes
Amazon DynamoDB	40	No
Amazon EBS	100	Yes
Amazon EC2	100	Yes

Name	Default	Adjustable
Amazon EFS	30	Yes
Amazon FSx	40	Yes
Amazon Neptune	40	Yes
Amazon RDS	40	Yes
Amazon Redshift	30	Yes
Amazon S3	30	Yes

AWS Backup gateway quotas

Name	Default	Adjustable
Backup or restore jobs per gateway	4	No. Create more gateways and connect them to your hypervisor.
VMware tags to resource tags mapping per hypervisor	10	No

Amazon EKS quotas

Name	Default	Adjustable
Namespaces per EKS cluster backup	10,000	Yes
Persistent Storage backups per EKS cluster backup	1,200	Yes

Name	Default	Adjustable
Restore jobs per target EKS cluster	1	No
EKS Restore jobs per account	5	Yes

Logically air-gapped vault quotas

The following quotas apply per source account and per copy destination Region. They are currently not adjustable.

Resource type	Maximum number of concurrent copies	Adjustable
EC2	The EBS concurrent copy limit applies to any snapshots being copied as part of an AMI copy.	No
EBS	20	No
Aurora	20	No
DocumentDB	20	No
Neptune	20	No
Storage Gateway	20	No
FSx	5	No

Related quotas

There are quotas on a single resource assignment in a single backup rule. You can create a backup plan with multiple backup rules.

When you manage backups across multiple accounts using AWS Organizations, you might encounter quotas that AWS Organizations imposes. For these quotas, see Quotas for AWS Organizations in the *AWS Organizations User Guide*.

You might also encounter quotas imposed by a AWS Backup-supported service, including the following:

- Amazon Elastic File System
- Amazon Elastic Block Store
- Amazon RDS
- Amazon Aurora
- Amazon EC2
- AWS Storage Gateway
- Amazon DynamoDB
- Amazon FSx for Lustre
- Amazon FSx for Windows File Server
- Amazon DocumentDB
- Amazon Neptune
- Amazon Simple Storage Service
- Amazon Timestream

Monitoring AWS Backup

AWS Backup works with other AWS tools to empower you to monitor its workloads. These tools include the following:

- AWS Backup console dashboards
  - The jobs dashboard brings job health monitoring, where you can view metrics showing job successes and failures, filtered by reasons, accounts, Region, and resource type.
  - The jobs dashboard is available in Regions where AWS Backup Audit Manager is supported. See Feature availability by AWS Region for those Regions. All other Regions will be able to access the CloudWatch Dashboard.
- **Amazon CloudWatch** and **Amazon EventBridge** to monitor AWS Backup processes.
  - You can use CloudWatch to track metrics, create alarms, and view dashboards. For more information, see AWS Backup metrics with Amazon CloudWatch.
  - You can use EventBridge to view and monitor AWS Backup events. For more information, see Monitoring AWS Backup events using Amazon EventBridge.
- **AWS CloudTrail** to monitor AWS Backup API calls. You can identify the time, source IP, users, and accounts making those calls. For more information, see Logging AWS Backup API calls with CloudTrail.
- **Amazon Simple Notification Service** (Amazon SNS) to subscribe to AWS Backup-related topics such as backup, restore, and copy events. For more information, see Notification options with AWS Backup.

AWS Backup console dashboards

> **Note**
>
> The jobs dashboard is available in all Regions where AWS Backup Audit Manager is supported. See Feature availability by AWS Region for those Regions. All other Regions will be able to access the CloudWatch Dashboard.

Topics

- Backup dashboards overview

- Viewing the jobs dashboard
- Reasons for problematic jobs
- Obtaining dashboard data through AWS CLI

Backup dashboards overview

AWS Backup provides a Jobs dashboard in the console to help you monitor the health of your backup, copy, and restore jobs. The same data that is visually displayed in the console can be retrieved in the command line through AWS CLI.

The jobs dashboard can be used to identify issues with backup, copy, and restore jobs through organization level or member account monitoring. With this information, you can identify and diagnose events and possible issues to help ensure fidelity in your activities.

The jobs dashboard can display two timeframes. By default, data from the latest 14 days are displayed, but you can change the view to show the latest 7 days. If you change the timeframe, the data will update to reflect to new time interval.

Note the dashboard displays data until the most recent 0:00 UTC; that is, the current day's data is not included. The dashboard updates daily during approximately 1:30 - 2:30 UTC.

Viewing the jobs dashboard

To view the jobs dashboard, log into the AWS Backup console and select **Jobs dashboards** in the left navigation bar.

On the jobs dashboard page, you can select from the backup, copy, or restore jobs tab.

The jobs dashboard overview displays the aggregated view over the specified timeframe for job activity, including completed, completed with issues, expired, and failed jobs. By default, data from the latest 14 days are displayed, but you can change the view to show 7 days.

> **Note**
>
> `Completed with issues` is a status of a job displayed in the console that denotes a completed job with a status message.

Job health

The line chart displays the successful and unsuccessful jobs rate lines over time. The successful rate line shows an aggregation of completed and completed with issues jobs. The unsuccessful rate line shows the sum of failed and expired jobs according to the specified time range.

Jobs in a non-completed or non-failed state (jobs with a status of created, pending, running, aborted, aborting, or partial) are not included; percentage totals may not equal 100%.

Job status over time

With the bar chart, you can generate a custom bar chart that shows the number of jobs in each category (Completed, Complete with issues, Failed, and Expired), distributed by days.

With the dropdown menus, choose the status(es), resource types, and AWS Regions you want to see in the chart. If you want to explore your selection further, select **View jobs** to see a pre-filtered portion of the jobs/cross-account monitoring page.

You can hover the mouse over a bar to display a popover that shows detailed job data for the selected date.

Problematic jobs

A **problematic** job is a job that has the status Failed, Expired, or Completed with issues. Each chart displays the corresponding metric that contains either the accounts, resource types, or top reasons that contain the highest number of problematic jobs.

The default display sorts the dashboard widget by the specified metric in descending order, starting with the metric with the highest number of problematic jobs that belong to the metric.

The top problematic accounts display will only be visible in accounts that have access through Organizations, such as administrative accounts and delegated administrator accounts. If visible, you can hover over an account to display the number of problematic jobs that belong to the chosen account.

You can select a bar within the graph to open a popup window. In this window, you can select a job status to open a jobs/cross-account monitoring table filtered by the status selected.

Reasons for problematic jobs

The **Top problematic reasons** widget shows the message code category to which error messages belong. However, the category might not explain the issues a job experiences. Expand the message

code categories below to see more details about the specific messages or errors your jobs could be encountering. All error messages not belonging to the message code categories below will fall under "MISCELLANEOUS".

"VSS_ERROR"

- "Windows VSS Backup attempt failed because either Instance or SSM Agent has invalid state or insufficient privileges."
- "Windows VSS Backup attempt failed because of insufficient privileges to perform this operation"
- "Windows VSS Backup attempt failed because ec2-vss-agent.exe is not installed in the Instance"
- "Windows VSS Backup Job Error encountered, trying for regular backup"
- "Windows VSS Backup attempt failed because of timeout on VSS enabled snapshot creation"
- "Windows VSS Backup attempt failed because of unsupported Windows Server version. Supported Versions are Windows Server 2012 or later."
- "Windows VSS Backup attempt failed because of timeout on VSS enabled snapshot creation"

"LIMIT_EXCEEDED"

- "Subscriber limit exceeded: You have reached the maximum concurrent number of backups, which is 300. Wait until other jobs finish, and try again. You can also reach out to Support to request a quota increase."
- "Maximum allowed in-progress snapshots for a single volume exceeded."
- "Maximum allowed active snapshot limit exceeded."
- "Image creation throttled due to maximum number of concurrent snapshots on an EBS volume."
- "Cannot create more than 20 user snapshots"
- "The resultant tag set must not have more than 50 user tags."
- "You have reached the maximum supported backups for your account/database. See Quotas in the Timestream developer guide for additional information."
- "You have reached your quota of 50,000 for the number of public and private images allowed in this Region. Deregister unused images, or request an increase in your AMI quota."
- "Your backup succeeded, but we were unable to persist NetworkInterfaces metadata as its size exceeded our internal limits."
- "REGEX#subscriber limit exceeded"

- "REGEX#More than 50 tags specified"
- "REGEX#can have at most"

"ACCESS_DENIED"

- "You are not authorized to perform this operation."
- "Access Denied trying to call AWS Backup service"
- "Customer credential miss permission to get bucket versioning for bucket."
- "Images from AWS Marketplace cannot be copied to another AWS account."
- "Copy job failed because the destination Backup vault is encrypted with the default Backup service managed key. The contents of this vault cannot be copied. Only the contents of a Backup vault encrypted by an AWS KMS key may be copied.
- Snapshots encrypted with the AWS managed key can't be shared. Specify another snapshot.
- "Encrypted snapshots with Amazon EBS default key cannot be shared
- "Copy job failed. Both source and destination account must be a member of the same organization."
- "REGEX#access denied"
- "REGEX#not authorized to"
- "REGEX#cannot be assumed by AWS Backup
- "REGEX#does not have permission"
- "REGEX#missing permission"

"CONCURRENT_JOB"

- "Backup job failed because there was a running job for the same resource."
- "There are too many copy jobs running for resource."

"FEATURE_NOT_ENABLED"

- "Copy job failed. Cross-account copy feature is not enabled for the current organization."

"JOB_EXPIRED"

- "Backup job expired before completion."

"INVALID_LIFECYCLE"

- "Copy job failed. The retention specified in the job is not within the range specified for the target Backup Vault."
- "REGEX#could not start because it is either inside or too close to the weekly maintenance window configured"
- "REGEX#could not start because it is either inside or too close to the automated backup window configured"

"INVALID_STATE"

- "REGEX#Instance is not in state"
- "REGEX#not in the available state"
- "REGEX#not in available state"
- "REGEX#Cannot snapshot volume"

"KMS_KEY_ERROR"

- "KMS key is either disabled or pending deletion or access to KMS key is denied"
- "Given key ID is not accessible"
- "AMI snapshot copy failed with error: Given key ID is not accessible. You must have DescribeKey permissions on the default key"
- "REGEX#kms key"

"ACCESS_KEY_ERROR"

- "The AWS Access Key Id needs a subscription for the service"

"HYPERVISOR_OFFLINE"

- "This operation is not valid for the specified hypervisor because it is not online"

"RESOURCE_NOT_FOUND"

- "The specified volume was not found."

- "The virtual machine is not found."
- "Given key ID does not exist"
- "Table not found."
- "REGEX#does not exist"
- "REGEX#Could not find resource"
- "REGEX#Could not find cryopod"
- "REGEX#Cannot find recovery point"
- "REGEX#resource not found"
- "REGEX#no longer available"
- "REGEX#is invalid"

"RESOURCE_NOT_SUPPORTED"

- "REGEX#unsupported resource type"
- "REGEX#Unsupported resource type"

"TAG_COPY_ERROR"

- "We are unable to copy resource tags to your backup because of the Internal Failure."
- "We are unable to copy resource tags to your backup because source or destination recovery point is unavailable"

"TOKEN_EXPIRED"

- "Token expired. Try again."

"UNSUPPORTED_OPERATION"

- "CreateSnapshot method not supported on hypervisor during snapshot creation. Aborted backup job"
- "UnsupportedOperation : Storage Gateway backup copies require a user-created backup vault and key at destination."
- "REGEX#Feature is not supported for provided resource type."

"FATAL_ERROR"

- "An internal error occurred."
- "Copy job encountered a fatal error. Please contact AWS Support for further assistance."
- "Copy job encountered a fatal error."
- "REGEX#Backup job encountered a fatal error"

Obtaining dashboard data through AWS CLI

You can use the command line to retrieve the same data which appears in the console. Use one of the following CLI commands:

- list-backup-job-summaries
- list-copy-job-summaries
- list-restore-job-summaries

There are the valid parameters you can include in each command:

```
BackupJobSummaries (list)
    Region (string),
    Account (string),
    State (string),
    ResourceType (string),
    MessageCategory (string),
AggregationPeriod: (string),
NextToken (string),
MaxResults (number)

CopyJobSummaries (list)
    Region (string),
    Account (string),
    State (string),
    ResourceType (string),
    MessageCategory (string),
AggregationPeriod: (string),
NextToken (string),
MaxResults (number)

RestoreJobSummaries (list)
```

```
    Region (string),
    Account (string),
    State (string),
    ResourceType (string),
AggregationPeriod: (string),
NextToken (string)
```

This example shows a sample request where the a user has input `list-backup-job-summaries` where the request asks to return all available accounts with a state of `FAILED` over the prior 14 days:

```
GET /audit/backup-job-summaries/
    ?accountId=ANY
    &state=FAILED
    &aggregationPeriod=FOURTEEN_DAYS
```

To obtain a job count for jobs with a status of `completed with issues`, subtract the job count of `COMPLETED` jobs with a `MessageCategory` of `SUCCESS` from the total number of `COMPLETED`.

Monitoring AWS Backup events using Amazon EventBridge

AWS Backup sends events to Amazon EventBridge when the state of a backup or copy job changes. You can use EventBridge to monitor AWS Backup events. For example, you can receive an alarm when a backup job fails. AWS Backup emits events to EventBridge in a best-effort manner every 5 minutes.

To track events using EventBridge, see the following:

- Creating a rule that reacts to events (*Amazon EventBridge User Guide*)
- Amazon CloudWatch Events and Metrics for AWS Backup (blog - see *Configure AWS Backup events to send to Amazon EventBridge*)

Some events report `status: COMPLETED` whereas other events report `state: COMPLETED`. This is consistent with the AWS Backup API. Some statuses are specific to the AWS Backup console: the status `Completed with issues` status is a representation of `Completed` jobs with status messages. To monitor `Completed with issues` events, monitor `COMPLETED` jobs that have a status message. Please note this is only specific to backup jobs.

You can alternatively use the AWS Backup notification API to track AWS Backup events with Amazon Simple Notification Service (Amazon SNS). However, EventBridge tracks more changes than the notification API does, including changes to backup vaults, copy job state, Region settings, and the number of cold or warm recovery points.

Events

- Backup Job events
- Backup Plan events
- Backup Vault events
- Copy Job events
- Recovery Point events
- Region Settings events
- Restore Job events
- Recovery point indexing events
- Malware scan Job events
- Backup access point events

Backup Job events

The following are example events.

State

- State: FAILED
- State: COMPLETED
- State: RUNNING
- State: ABORTED
- State: EXPIRED
- State: PENDING
- State: CREATED

State: FAILED

```
{
```

```
  "version": "0",
  "id": "710b0398-d48e-f3c3-afca-cfeb2fdaa656",
  "detail-type": "Backup Job State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-07-29T20:15:26Z",
  "region": "us-west-2",
  "resources": [],
  "detail": {
    "backupJobId": "34176239-e96d-4e1d-9fad-529dbb3c3556",
    "backupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:9ab3e749-82c6-4342-9320-5edbf4918b86",
    "backupVaultName": "9ab3e749-82c6-4342-9320-5edbf4918b86",
    "bytesTransferred": "0",
    "creationDate": "2020-07-29T20:13:07.392Z",
    "iamRoleArn": "arn:aws:iam::1112233445566:role/MockRCBackupTestRole",
    "resourceArn": "arn:aws:service:us-west-2:1112233445566:resource-type/resource-id",
    "resourceType": "type",
    "state": "FAILED",
    "statusMessage": "\"Backup job failed because backup vault arn:aws:backup:us-
west-2:1112233445566:backup-vault:9ab3e749-82c6-4342-9320-5edbf4918b86 does not exist.
\"",
    "startBy": "2020-07-30T04:13:07.392Z",
    "percentDone": 0,
    "retryCount": 3
  }
}
```


State: COMPLETED


```
{
  "version": "0",
  "id": "dafac799-9b88-0134-26b7-fef4d54a134f",
  "detail-type": "Backup Job State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-07-15T21:41:17Z",
  "region": "us-west-2",
  "resources": [
    "arn:aws:backup:us-west-2:1112233445566:recovery-point:f1d966fe-a3bd-410b-
b292-99f442d13b56"
  ],
  "detail": {
```

```
    "backupJobId": "a827233a-d405-4a86-a440-759fa94f34dd",
    "backupSizeInBytes": "36048",
    "backupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:9732c1b4-1091-472a-9d9f-52e0565ee39a",
    "backupVaultName": "9732c1b4-1091-472a-9d9f-52e0565ee39a",
    "bytesTransferred": "36048",
    "creationDate": "2020-07-15T21:40:31.207Z",
    "iamRoleArn": "arn:aws:iam::1112233445566:role/MockRCBackupTestRole",
    "resourceArn": "arn:aws:service:us-west-2:1112233445566:resource-type/resource-id",
    "resourceType": "type",
    "state": "COMPLETED",
    "completionDate": "2020-07-15T21:41:05.921Z",
    "startBy": "2020-07-16T05:40:31.207Z",
    "percentDone": 100,
    "retryCount": 3
  }
}
```


State: RUNNING


```
{
  "version": "0",
  "id": "44946c39-b519-3505-44e6-ba74afeb2e30",
  "detail-type": "Backup Job State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-07-15T21:39:13Z",
  "region": "us-west-2",
  "resources": [],
  "detail": {
    "backupJobId": "B6EC38D2-CB3C-EF0A-F5A4-3CF324EF4945",
    "backupSizeInBytes": "3221225472",
    "backupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:e6625738-0655-4aa9-bd37-6ec1dd183b15",
    "backupVaultName": "e6625738-0655-4aa9-bd37-6ec1dd183b15",
    "bytesTransferred": "0",
    "creationDate": "2020-07-15T21:38:31.152Z",
    "iamRoleArn": "arn:aws:iam::1112233445566:role/FullBackupTestRole",
    "resourceArn": "arn:aws:ec2:us-west-2:1112233445566:volume/vol-0b5ae24f2ee72d926",
    "resourceType": "EBS",
    "state": "RUNNING",
    "startBy": "2020-07-16T05:00:00Z",
    "expectedCompletionDate": "Jul 15, 2020 9:39:07 PM",
```

```
    "percentDone": 99,
    "createdBy": {
      "backupPlanId": "bde0f455-4e24-4668-aeaa-4932a97f5cc5",
      "backupPlanArn": "arn:aws:backup:us-west-2:1112233445566:backup-
plan:bde0f455-4e24-4668-aeaa-4932a97f5cc5",
      "backupPlanVersion": "YTkzNmM0MmUtMWRhNS00Y2RkLThmZGUtNjA5NTc4NGM1YTc5",
      "backupPlanRuleId": "1f97bafa-14d6-4f39-94fd-94b51bd6d0d5"
    }
  }
}
```

State: ABORTED

```
{
  "version": "0",
  "id": "4c91ceb0-b798-da82-6818-c29b3dce7543",
  "detail-type": "Backup Job State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-07-15T21:33:16Z",
  "region": "us-west-2",
  "resources": [],
  "detail": {
    "backupJobId": "58cdef95-7680-4c74-80d5-1b64093999c8",
    "backupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:f59bffcd-2538-4bbe-8343-1c60dae27c27",
    "backupVaultName": "f59bffcd-2538-4bbe-8343-1c60dae27c27",
    "bytesTransferred": "0",
    "creationDate": "2020-07-15T21:33:00.803Z",
    "iamRoleArn": "arn:aws:iam::1112233445566:role/MockRCBackupTestRole",
    "resourceArn": "arn:aws:service:us-west-2:1112233445566:resource-type/resource-id",
    "resourceType": "type",
    "state": "ABORTED",
    "statusMessage": "\"Backup job was stopped by user.\"",
    "completionDate": "2020-07-15T21:33:01.621Z",
    "startBy": "2020-07-16T05:33:00.803Z",
    "percentDone": 0
  }
}
```

State: EXPIRED

```
{
```

```
  "version": "0",
  "id": "1d7bbc04-6120-1145-13b9-49b0af465328",
  "detail-type": "Backup Job State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-07-29T13:04:57Z",
  "region": "us-west-2",
  "resources": [],
  "detail": {
    "backupJobId": "01EE26DC-7107-4D8E-0C54-EAC27C662BA4",
    "backupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-vault:aws/backup/
AutomatedBackupVaultDel2",
    "backupVaultName": "aws/backup/AutomatedBackupVaultDel2",
    "bytesTransferred": "0",
    "creationDate": "2020-07-29T05:10:20.077Z",
    "iamRoleArn": "arn:aws:iam::1112233445566:role/MockRCBackupTestRole",
    "resourceArn": "arn:aws:service:us-west-2:1112233445566:resource-type/resource-id",
    "resourceType": "type",
    "state": "EXPIRED",
    "statusMessage": "\"Backup job failed because there was a running job for the same
 resource.\"",
    "completionDate": "2020-07-29T13:02:15.234Z",
    "startBy": "2020-07-29T13:00:00Z",
    "percentDone": 0,
    "createdBy": {
      "backupPlanId": "aws/efs/414a5bd4-f880-47ad-95f3-f085108a4c3b",
      "backupPlanArn": "arn:aws:backup:us-west-2:1112233445566:backup-plan:aws/
efs/414a5bd4-f880-47ad-95f3-f085108a4c3b",
      "backupPlanVersion": "NjBjOTUzZjYtYzZiNi00NjhlLWIzMTEtNWRjOWY0YTNjN2Vj",
      "backupPlanRuleId": "3eb0017c-f262-4211-a802-302cebb11dc2"
    }
  }
}
```


State: PENDING


```
{
  "version": "0",
  "id": "64dd1897-f863-31a3-9ee5-b05e306d81ff",
  "detail-type": "Backup Job State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-07-29T20:03:30Z",
```

```
  "region": "us-west-2",
  "resources": [],
  "detail": {
    "backupJobId": "2cffdb68-d6ed-485f-9f9b-8b530749f1c2",
    "backupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:ed1f2661-5587-48bf-8a98-fadb977bf975",
    "backupVaultName": "ed1f2661-5587-48bf-8a98-fadb977bf975",
    "bytesTransferred": "0",
    "creationDate": "2020-07-29T20:01:06.224Z",
    "iamRoleArn": "arn:aws:iam::1112233445566:role/MockRCBackupTestRole",
    "resourceArn": "arn:aws:service:us-west-2:1112233445566:resource-type/resource-id",
    "resourceType": "type",
    "state": "PENDING",
    "statusMessage": "",
    "startBy": "2020-07-30T04:01:06.224Z",
    "percentDone": 0
  }
}
```

State: CREATED

```
{
  "version": "0",
  "id": "29af2bf2-eace-58ab-da3a-8c0bf738d692",
  "detail-type": "Backup Job State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-06-22T20:32:53Z",
  "region": "us-west-2",
  "resources": [],
  "detail": {
    "backupJobId": "7e8845b5-ca30-415f-a842-e0152bf4d0ca",
    "state": "CREATED",
    "creationDate": "2020-06-22T20:32:47.466Z"
  }
}
```

Backup Plan events

The following are example events.

State

- State: MODIFIED
- State: DELETED
- State: CREATED

State: MODIFIED


```
{
  "version": "0",
  "id": "2895aefb-dd4a-0a23-6071-2652abd92c3f",
  "detail-type": "Backup Plan State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-06-24T23:18:25Z",
  "region": "us-west-2",
  "resources": [
    "arn:aws:backup:us-west-2:1112233445566:backup-plan:83fcb8ee-2d93-42ac-
b06f-591563f3f8de"
  ],
  "detail": {
    "backupPlanId": "83fcb8ee-2d93-42ac-b06f-591563f3f8de",
    "versionId": "NjIwNDFjMDEtNmZlNC00M2JmLTkzZDgtNzNkZjQyNzkxNDk0",
    "modifiedAt": "2020-06-24T23:18:19.168Z",
    "state": "MODIFIED"
  }
}
```


State: DELETED


```
{
  "version": "0",
  "id": "33fc5c1d-6db2-b3d9-1e70-1c9a2c23645c",
  "detail-type": "Backup Plan State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-06-24T23:18:25Z",
  "region": "us-west-2",
  "resources": [
    "arn:aws:backup:us-west-2:1112233445566:backup-plan:83fcb8ee-2d93-42ac-
b06f-591563f3f8de"
  ],
  "detail": {
```

```
    "backupPlanId": "83fcb8ee-2d93-42ac-b06f-591563f3f8de",
    "versionId": "NjIwNDFjMDEtNmZlNC00M2JmLTkzZDgtNzNkZjQyNzkxNDk0",
    "deletionDate": "2020-06-24T23:18:19.411Z",
    "state": "DELETED"
  }
}
```

State: CREATED

```
{
  "version": "0",
  "id": "b64fb2d0-ae16-ff9a-faf6-0bdd0d4bfdef",
  "detail-type": "Backup Plan State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-06-24T23:18:19Z",
  "region": "us-west-2",
  "resources": [
    "arn:aws:backup:us-west-2:1112233445566:backup-plan:2c103c5f-6d6e-4cac-9147-
d3afa4c84f59"
  ],
  "detail": {
    "backupPlanId": "2c103c5f-6d6e-4cac-9147-d3afa4c84f59",
    "versionId": "N2Q4OTczMzEtZmY1My00N2UwLWE3ODUtMjViYWYyOTUzZWY4",
    "creationDate": "2020-06-24T23:18:15.318Z",
    "state": "CREATED"
  }
}
```

Backup Vault events

The following are example events.

State

- State: CREATED
- State: MODIFIED
- State: DELETED

State: CREATED

```
{
  "version": "0",
  "id": "d415609e-5f35-d9a2-76d1-613683e4e024",
  "detail-type": "Backup Vault State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-06-24T23:18:19Z",
  "region": "us-west-2",
  "resources": [
    "arn:aws:backup:us-west-2:1112233445566:backup-vault:d8864642-155c-4283-a168-
a04f40e12c97"
  ],
  "detail": {
    "backupVaultName": "d8864642-155c-4283-a168-a04f40e12c97",
    "state": "CREATED"
  }
}
```

State: MODIFIED

```
{
  "version": "0",
  "id": "1a2b3cd4-5e6f-7g8h-9i0j-123456k7l890",
  "detail-type": "Backup Vault State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-06-24T23:18:19Z",
  "region": "us-west-2",
  "resources": [
    "arn:aws:backup:us-west-2:1112233445566:backup-vault:nameOfTestBackup"
  ],
  "detail": {
    "backupVaultName": "vaultName",
    "state": "MODIFIED",
    "isLocked": "true"
  }
}
```

State: DELETED


```
{
  "version": "0",
  "id": "344bccc1-6d2e-da93-3adf-b3f82460294d",
  "detail-type": "Backup Vault State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-06-22T02:42:37Z",
  "region": "us-west-2",
  "resources": [
    "arn:aws:backup:us-west-2:1112233445566:backup-vault:e8189629-1f8e-4ed2-af7d-
b32415d04db1"
  ],
  "detail": {
    "backupVaultName": "e8189629-1f8e-4ed2-af7d-b32415d04db1",
    "state": "DELETED"
  }
}
```


Copy Job events

The following are example events.

State

- State: FAILED
- State: RUNNING
- State: COMPLETED
- State: CREATED

State: FAILED


```
{
  "version": "0",
  "id": "4660bc92-a44d-c939-4542-cda503f14855",
  "detail-type": "Copy Job State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-07-15T20:37:34Z",
```

```
  "region": "us-west-2",
  "resources": [
    "arn:aws:ec2:us-west-2::image/ami-00179b33a7a88cac5"
  ],
  "detail": {
    "copyJobId": "47C8EF56-74D8-059D-1301-C5BE1D5C926E",
    "backupSizeInBytes": 22548578304,
    "creationDate": "2020-07-15T20:36:13.239Z",
    "iamRoleArn": "arn:aws:iam::1112233445566:role/
RoleForEc2BackupWithNoDescribeTagsPermissions",
    "resourceArn": "arn:aws:ec2:us-west-2:1112233445566:instance/i-0515aee7de03f58e1",
    "resourceType": "EC2",
    "sourceBackupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:55aa945e-c46a-421b-aa27-f94b074e31b7",
    "state": "FAILED",
    "statusMessage": "Access denied exception while trying to list tags",
    "completionDate": "2020-07-15T20:37:28.704Z",
    "destinationBackupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:55aa945e-c46a-421b-aa27-f94b074e31b7",
    "destinationRecoveryPointArn": {}
  }
}
```


State: RUNNING


```
{
  "version": "0",
  "id": "d17480ae-7042-edb2-0ff5-8b94822c58e4",
  "detail-type": "Copy Job State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-07-15T22:07:48Z",
  "region": "us-west-2",
  "resources": [
    "arn:aws:ec2:us-west-2::snapshot/snap-03886bc8d6ef3a1f9"
  ],
  "detail": {
    "copyJobId": "0175DE71-5784-589F-D8AC-541ACCB4CAC8",
    "backupSizeInBytes": 3221225472,
    "creationDate": "2020-07-15T22:06:27.234Z",
    "iamRoleArn": "arn:aws:iam::1112233445566:role/OrganizationCanaryTestRole",
    "resourceArn": "arn:aws:ec2:us-west-2:1112233445566:volume/vol-050eba21ee4d3c001",
    "resourceType": "EBS",
```

```
    "sourceBackupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:846869de-4589-45c3-ab60-4fbbabcdd3ec",
    "state": "RUNNING",
    "destinationBackupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:846869de-4589-45c3-ab60-4fbbabcdd3ec",
    "destinationRecoveryPointArn": {},
    "createdBy": {
      "backupPlanId": "b58e3621-1c53-4997-ad8a-afc3347a850e",
      "backupPlanArn": "arn:aws:backup:us-west-2:1112233445566:backup-
plan:b58e3621-1c53-4997-ad8a-afc3347a850e",
      "backupPlanVersion": "Mjc4ZTRhMzUtMGE5Ni00NmQ5LWE1YmMtOWMwY2IwMTY4NWQ4",
      "backupPlanRuleId": "78e356d3-1a11-4f61-8585-af5d6b69bb18"
    }
  }
}
```

State: COMPLETED

```
{
  "version": "0",
  "id": "47deb974-6473-aef1-56c2-52c3eaedfceb",
  "detail-type": "Copy Job State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-07-15T22:08:04Z",
  "region": "us-west-2",
  "resources": [
    "arn:aws:ec2:us-west-2::snapshot/snap-03886bc8d6ef3a1f9"
  ],
  "detail": {
    "copyJobId": "0175DE71-5784-589F-D8AC-541ACCB4CAC8",
    "backupSizeInBytes": 3221225472,
    "creationDate": "2020-07-15T22:06:27.234Z",
    "iamRoleArn": "arn:aws:iam::1112233445566:role/OrganizationCanaryTestRole",
    "resourceArn": "arn:aws:ec2:us-west-2:1112233445566:volume/vol-050eba21ee4d3c001",
    "resourceType": "EBS",
    "sourceBackupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:846869de-4589-45c3-ab60-4fbbabcdd3ec",
    "state": "COMPLETED",
    "completionDate": "2020-07-15T22:07:58.111Z",
    "destinationBackupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:846869de-4589-45c3-ab60-4fbbabcdd3ec",
```

```
    "destinationRecoveryPointArn": "arn:aws:ec2:us-west-2::snapshot/
snap-0726fe70935586180",
    "createdBy": {
      "backupPlanId": "b58e3621-1c53-4997-ad8a-afc3347a850e",
      "backupPlanArn": "arn:aws:backup:us-west-2:1112233445566:backup-
plan:b58e3621-1c53-4997-ad8a-afc3347a850e",
      "backupPlanVersion": "Mjc4ZTRhMzUtMGE5Ni00NmQ5LWE1YmMtOWMwY2IwMTY4NWQ4",
      "backupPlanRuleId": "78e356d3-1a11-4f61-8585-af5d6b69bb18"
    }
  }
}
```

State: CREATED

```
{
  "version": "0",
  "id": "8398a4c4-8fe8-2b49-a4b9-fd4fdcd34a4e",
  "detail-type": "Copy Job State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-06-22T21:06:32Z",
  "region": "us-west-2",
  "resources": [
    "arn:aws:ec2:us-west-2::image/ami-0888b126e2170b98e"
  ],
  "detail": {
    "creationDate": "2020-06-22T21:06:25.754Z",
    "state": "CREATED",
    "sourceBackupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:ef09da5a-21a6-461f-a98f-857e9e621a17",
    "destinationBackupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:ef09da5a-21a6-461f-a98f-857e9e621a17"
  }
}
```

Recovery Point events

The following are the events.

State

- COMPLETED

- PARTIAL
- DELETING
- EXPIRED
- AVAILABLE
- STOPPED
- CREATING

State: COMPLETED


```
{
    "version": "0",
    "id": "ab32977c-378d-2122-e985-fgh4596f0709",
    "detail-type": "Recovery Point State Change",
    "source": "aws.backup",
    "account": "1112233445566",
    "time": "2020-07-15T21:39:07Z",
    "region": "us-west-2",
    "resources": [
        "arn:aws:rds:us-west-2:1112233445566:cluster-snapshot:awsbackup:job-4ece7121-
d60e-00c2-5c3b-49960142d03b"
    ],
    "detail": {
        "backupVaultName": "e6625738-0655-4aa9-bd37-6ec1dd183b15",
        "backupVaultArn": "arn:aws:backup:us-west-2:496821122410:backup-
vault:e6625738-0655-4aa9-bd37-6ec1dd183b15",
        "creationDate": "2020-07-15T21:38:31.152Z",
        "iamRoleArn": "arn:aws:iam::1112233445566:role/FullBackupTestRole",
        "resourceType": "Aurora",
        "resourceArn": "arn:aws:rds:us-west-2:1112233445566:cluster:id",
        "status": "COMPLETED",
        "isEncrypted": "false",
        "storageClass": "WARM",
        "completionDate": "2020-07-15T21:39:05.689Z",
        "createdBy": {
            "backupPlanId": "bde0f455-4e24-4668-aeaa-4932a97f5cc5",
            "backupPlanArn": "arn:aws:backup:us-west-2:1112233445566:backup-
plan:bde0f455-4e24-4668-aeaa-4932a97f5cc5",
            "backupPlanVersion": "YTkzNmM0MmUtMWRhNS00Y2RkLThmZGUtNjA5NTc4NGM1YTc5",
            "backupPlanRuleId": "1f97bafa-14d6-4f39-94fd-94b51bd6d0d5"
        },
```

```
        "lifecycle": {
            "deleteAfterDays": 100
        },
        "calculatedLifeCycle": {
            "deleteAt": "2020-10-23T21:38:31.152Z"
        }
    }
}
```

Region Settings events

The following is an example event.

```
{
  "version": "0",
  "id": "e7ed82ba-4955-4de5-10d6-dbafcfb68b4f",
  "detail-type": "Region Setting State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-06-24T22:55:03Z",
  "region": "us-west-2",
  "resources": [],
  "detail": {
    "modifiedAt": "2020-06-24T22:54:57.161Z",
    "ResourceTypeOptInPreference": {
      "Aurora": true
    },
    "state": "MODIFIED"
  }
}
```

Restore Job events

The following are example events. Note that your use case of a restore job will determine the required and optional parameters to include. For example, if your restore job is part of a restore testing plan, the parameter `restoreTestingPlanArn` is included. See `DescribeRestoreJob` for possible parameters.

State

- State: FAILED
- State: RUNNING

- State: COMPLETED
- State: PENDING
- State: CREATED

State: FAILED


```
{
  "version": "0",
  "id": "ab32977c-378d-2122-e985-fgh4596f0709",
  "detail-type": "Restore Job State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-07-15T20:19:29Z",
  "region": "us-west-2",
  "resources": [
    "arn:aws:ec2:us-west-2::image/ami-12b3456dfb7f8cf90"
  ],
  "detail": {
    "restoreJobId": "1B234A56-789B-01CD-2A34-4567A08901FD",
    "backupSizeInBytes": "22548578304",
    "creationDate": "2020-07-15T20:19:07.303Z",
    "iamRoleArn": "arn:aws:iam::1112233445566:role/TestAWSBackupRole",
    "percentDone": 0,
    "resourceType": "EC2",
    "status": "FAILED",
    "statusMessage": "AWS Backup does not permit attaching a new instance profile to an
 EC2 instance. Please restore using the backed up instance profile."
  }
}
```


State: RUNNING


```
{
  "version": "0",
  "id": "ab32977c-378d-2122-e985-fgh4596f0709",
  "detail-type": "Restore Job State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-07-29T20:26:06Z",
  "region": "us-west-2",
  "resources": [
```

```
    "arn:aws:ec2:us-west-2::snapshot/snap-0fe123ca456cfad7c"
  ],
  "detail": {
    "restoreJobId": "1B234A56-789B-01CD-2A34-4567A08901FD",
    "backupSizeInBytes": "3221225472",
    "creationDate": "2020-07-29T20:26:00.098Z",
    "iamRoleArn": "arn:aws:iam::1112233445566:role/RestoreTestRole",
    "percentDone": 0,
    "resourceType": "EBS",
    "status": "RUNNING"
  }
}
```

State: COMPLETED

```
{
  "version":"0",
  "id":"ab32977c-378d-2122-e985-fgh4596f0709",
  "detail-type":"Restore Job State Change",
  "source":"aws.backup",
  "account":"1112233445566",
  "time":"2020-07-15T03:14:58Z",
  "region":"us-west-2",
  "resources":[
    "arn:aws:rds:us-
west-2:1112233445566:snapshot:awsbackup:job-1a2bcd34-567e-8901-23f4-5g6hijkl7890"
  ],
  "detail":{
    "restoreJobId":"AB123456-78C9-0123-456D-789012E34567",
    "backupVaultArn":"arn:aws:backup:us-west-2:1112233445566:backup-
vault:ExampleVault",
    "recoveryPointArn":"arn:aws:backup:us-west-2:1112233445566:recovery-
point:6f7fc1f8-2f2e-40ac-943b-8f8efa9ba99d",
    "sourceResourceArn":"arn:aws:rds:us-west-2:1112233445566:db:example-database",
    "backupSizeInBytes":"0",
    "creationDate":"2020-07-15T03:10:01.742Z",
    "iamRoleArn":"arn:aws:iam::1112233445566:role/RestoreTestRole",
    "percentDone":0,
    "resourceType":"RDS",
    "status":"COMPLETED",
    "createdResourceArn":"arn:aws:rds:us-
west-2:1112233445566:db:testinginstance1a2bcd34-567e-8901-23f4-5g6hijkl7890",
    "completionDate":"2020-07-15T03:14:53.128Z"
```

```
  }
}
```

State: PENDING

```
{
  "version": "0",
  "id": "ab32977c-378d-2122-e985-fgh4596f0709",
  "detail-type": "Restore Job State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-07-29T20:08:26Z",
  "region": "us-west-2",
  "resources": [
    "arn:aws:backup:us-west-2:1112233445566:recovery-point:42bb8260-92cd-46a2-ab8d-
b29f4edb47b1"
  ],
  "detail": {
    "restoreJobId": "123EA45F-C678-EFE9-0123-4D56FC0E789A",
    "backupSizeInBytes": "36048",
    "creationDate": "2020-07-29T20:08:21.083Z",
    "iamRoleArn": "arn:aws:iam::1112233445566:role/RestoreTestRole",
    "percentDone": 0,
    "resourceType": "EC2",
    "status": "PENDING"
  }
}
```

State: CREATED

```
{
  "version": "0",
  "id": "ab32977c-378d-2122-e985-fgh4596f0709",
  "detail-type": "Restore Job State Change",
  "source": "aws.backup",
  "account": "1112233445566",
  "time": "2020-06-22T18:50:49Z",
  "region": "us-west-2",
  "resources": [
    "arn:aws:backup:us-west-2:1112233445566:recovery-point:a6560b33-3660-494c-8d47-
efgh939ij32k"
  ],
```

```
  "detail": {
    "restoreJobId": "123EA45F-C678-EFE9-0123-4D56FC0E789A",
    "creationDate": "2020-06-22T18:50:46.407Z",
    "state": "CREATED"
  }
}
```

Recovery point indexing events

The following are example events.

State

- State: ACTIVE
- State: DELETED
- State: FAILED

State: ACTIVE

```
{
  "version":"0",
  "id":"ab32977c-378d-2122-e985-fgh4596f0709",
  "detail-type":"Recovery Point Index State Change",
  "source":"aws.backup",
  "account":"1112233445566",
  "time":"2023-12-15T21:39:07Z",
  "region":"us-west-2",
  "resources":[
    "arn:aws:backup:us-west-2:1112233445566:recovery-point:abcd1234-5678-abcd-9012-
abcdef123456"
  ],
  "detail":{
    "recoveryPointArn":"arn:aws:backup:us-west-2:1112233445566:recovery-
point:abcd1234-5678-abcd-9012-abcdef123456",
    "accountId":"1112233445566",
    "indexStatus":"ACTIVE",
    "iamRoleArn":"arn:aws:iam::1112233445566:role/BackupIndexRole",
    "resourceType":"EBS",
    "backupVaultArn":"arn:aws:cryo:us-west-2:1112233445566:pod/backup-pod-12345",
    "indexCreationTime":"2025-05-25T21:38:31.152Z",
    "isIndexingContinuous":false,
```

```
    "sourceResourceArn":"arn:aws:ec2:us-west-2:1112233445566:volume/
vol-01234567890abcdef",
    "backupCreationTime":"2023-12-15T21:38:00.000Z",
    "indexStatusMessage":"An AWS Backup recovery point index was
 successfully completed. Indexed recovery point arn : arn:aws:backup:us-
west-2:1112233445566:recovery-point:abcd1234-5678-abcd-9012-abcdef123456",
    "indexCompletionTime":"2025-05-25T21:39:05.689Z"
  }
}
```

State: DELETED

```
{
  "version":"0",
  "id":"ab32977c-378d-2122-e985-fgh4596f0709",
  "detail-type":"Recovery Point Index State Change",
  "source":"aws.backup",
  "account":"1112233445566",
  "time":"2023-12-15T21:39:07Z",
  "region":"us-west-2",
  "resources":[
    "arn:aws:backup:us-west-2:1112233445566:recovery-point:abcd1234-5678-abcd-9012-
abcdef123456"
  ],
  "detail":{
    "recoveryPointArn":"arn:aws:backup:us-west-2:1112233445566:recovery-
point:abcd1234-5678-abcd-9012-abcdef123456",
    "accountId":"1112233445566",
    "indexStatus":"DELETED",
    "iamRoleArn":"arn:aws:iam::1112233445566:role/BackupIndexRole",
    "resourceType":"EBS",
    "backupVaultArn":"arn:aws:cryo:us-west-2:1112233445566:pod/backup-pod-12345",
    "indexCreationTime":"2025-05-25T21:38:31.152Z",
    "isIndexingContinuous":false,
    "sourceResourceArn":"arn:aws:ec2:us-west-2:1112233445566:volume/
vol-01234567890abcdef",
    "backupCreationTime":"2023-12-15T21:38:00.000Z",
    "indexStatusMessage":"An AWS Backup recovery point index was deleted.
 Indexed recovery point arn : arn:aws:backup:us-west-2:1112233445566:recovery-
point:abcd1234-5678-abcd-9012-abcdef123456",
    "indexDeletionTime":"2025-05-27T22:39:05.689Z"
  }
}
```

State: FAILED


```
{
  "version":"0",
  "id":"ab32977c-378d-2122-e985-fgh4596f0709",
  "detail-type":"Recovery Point Index State Change",
  "source":"aws.backup",
  "account":"1112233445566",
  "time":"2023-12-15T21:39:07Z",
  "region":"us-west-2",
  "resources":[
    "arn:aws:backup:us-west-2:1112233445566:recovery-point:abcd1234-5678-abcd-9012-
abcdef123456"
  ],
  "detail":{
    "recoveryPointArn":"arn:aws:backup:us-west-2:1112233445566:recovery-
point:abcd1234-5678-abcd-9012-abcdef123456",
    "accountId":"1112233445566",
    "indexStatus":"FAILED",
    "iamRoleArn":"arn:aws:iam::1112233445566:role/BackupIndexRole",
    "resourceType":"EBS",
    "backupVaultArn":"arn:aws:cryo:us-west-2:1112233445566:pod/backup-pod-12345",
    "indexCreationTime":"2025-05-25T21:38:31.152Z",
    "isIndexingContinuous":false,
    "sourceResourceArn":"arn:aws:ec2:us-west-2:1112233445566:volume/
vol-01234567890abcdef",
    "backupCreationTime":"2023-12-15T21:38:00.000Z",
    "indexStatusMessage":"An AWS Backup recovery point index failed to create.
 Indexed recovery point arn : arn:aws:backup:us-west-2:1112233445566:recovery-
point:abcd1234-5678-abcd-9012-abcdef123456"
  }
}
```


Malware scan Job events

The following are example events.

State

- State: CREATED
- State: RUNNING
- State: COMPLETED

- State: COMPLETED WITH ISSUES
- State: FAILED
- State: CANCELED

State: CREATED


```
{
    "version": "0",
    "id": "60ce181d-67c7-496b-90fb-69636b42daee",
    "detail-type": "Scan Job State Change",
    "source": "aws.backup",
    "account": "1112233445566",
    "time": "2025-12-12T12:12:12Z",
    "region": "us-west-2",
    "resources": [
        "arn:aws:backup:us-west-2:1112233445566:recovery-point:abcd1234-5678-abcd-9012-
abcdef123456"
    ],
    "detail": {
        "accountId": "1112233445566",
        "backupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:9ab3e749-82c6-4342-9320-5edbf4918b86",
        "backupVaultName": "9ab3e749-82c6-4342-9320-5edbf4918b86",
        "createdBy": {
            "backupPlanArn": "arn:aws:backup:us-west-2:1112233445566:backup-
plan:5d5a14cc-4ee2-4b9f-beb6-4afa998bb98f",
            "backupPlanId": "5d5a14cc-4ee2-4b9f-beb6-4afa998bb98f",
            "backupPlanVersion": "Mjc4ZTRhMzUtMGE5Ni00NmQ5LWE1YmMtOWMwY2IwMTY4NWQ4",
            "backupRuleId": "256ad167-f523-4cb9-93f3-f0c933efd97f"
        },
        "creationDate": "2025-12-12T12:12:00Z",
        "iamRoleArn": "arn:aws:iam::1112233445566:role/
RoleForEc2BackupWithNoDescribeTagsPermissions",
        "scannerRoleArn": "arn:aws:iam::1112233445566:role/
RoleForAwsBackupGuarddutyScanner",
        "recoveryPointArn": "arn:aws:backup:us-west-2:1112233445566:recovery-
point:abcd1234-5678-abcd-9012-abcdef123456",
        "resourceArn": "arn:aws:ec2:us-west-2:1112233445566:instance/
i-0515aee7de03f58e1",
        "resourceType": "EC2",
        "scanId": "43bc14a6-03e0-436a-abf5-8825d3aa6835",
        "scanJobId": "4d5fc12e-7e33-4336-a981-bbe43c300298",
```

```
        "scanBaseRecoveryPointArn": "arn:aws:backup:us-west-2:1112233445566:recovery-
point:46bd3edb-a499-4db1-8f8f-269a2ff76fab",
        "malwareScanner": "GUARDDUTY",
        "state": "CREATED"
    }
}
```

State: RUNNING

```
{
    "version": "0",
    "id": "60ce181d-67c7-496b-90fb-69636b42daee",
    "detail-type": "Scan Job State Change",
    "source": "aws.backup",
    "account": "1112233445566",
    "time": "2025-12-12T12:12:12Z",
    "region": "us-west-2",
    "resources": [
        "arn:aws:backup:us-west-2:1112233445566:recovery-point:abcd1234-5678-abcd-9012-
abcdef123456"
    ],
    "detail": {
        "accountId": "1112233445566",
        "backupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:9ab3e749-82c6-4342-9320-5edbf4918b86",
        "backupVaultName": "9ab3e749-82c6-4342-9320-5edbf4918b86",
        "createdBy": {
            "backupPlanArn": "arn:aws:backup:us-west-2:1112233445566:backup-
plan:5d5a14cc-4ee2-4b9f-beb6-4afa998bb98f",
            "backupPlanId": "5d5a14cc-4ee2-4b9f-beb6-4afa998bb98f",
            "backupPlanVersion": "Mjc4ZTRhMzUtMGE5Ni00NmQ5LWE1YmMtOWMwY2IwMTY4NWQ4",
            "backupRuleId": "256ad167-f523-4cb9-93f3-f0c933efd97f"
        },
        "creationDate": "2025-12-12T12:12:00Z",
        "iamRoleArn": "arn:aws:iam::1112233445566:role/
RoleForEc2BackupWithNoDescribeTagsPermissions",
        "scannerRoleArn": "arn:aws:iam::1112233445566:role/
RoleForAwsBackupGuarddutyScanner",
        "recoveryPointArn": "arn:aws:backup:us-west-2:1112233445566:recovery-
point:abcd1234-5678-abcd-9012-abcdef123456",
        "resourceArn": "arn:aws:ec2:us-west-2:1112233445566:instance/
i-0515aee7de03f58e1",
        "resourceType": "EC2",
```

```
        "scanId": "43bc14a6-03e0-436a-abf5-8825d3aa6835",
        "scanJobId": "4d5fc12e-7e33-4336-a981-bbe43c300298",
        "scanBaseRecoveryPointArn": "arn:aws:backup:us-west-2:1112233445566:recovery-
point:46bd3edb-a499-4db1-8f8f-269a2ff76fab",
        "malwareScanner": "GUARDDUTY",
        "state": "RUNNING"
    }
}
```

State: COMPLETED

```
{
    "version": "0",
    "id": "60ce181d-67c7-496b-90fb-69636b42daee",
    "detail-type": "Scan Job State Change",
    "source": "aws.backup",
    "account": "1112233445566",
    "time": "2025-12-12T12:12:12Z",
    "region": "us-west-2",
    "resources": [
        "arn:aws:backup:us-west-2:1112233445566:recovery-point:abcd1234-5678-abcd-9012-
abcdef123456"
    ],
    "detail": {
        "accountId": "1112233445566",
        "backupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:9ab3e749-82c6-4342-9320-5edbf4918b86",
        "backupVaultName": "9ab3e749-82c6-4342-9320-5edbf4918b86",
        "completionDate": "2025-12-12T12:12:12Z",
        "createdBy": {
            "backupPlanArn": "arn:aws:backup:us-west-2:1112233445566:backup-
plan:5d5a14cc-4ee2-4b9f-beb6-4afa998bb98f",
            "backupPlanId": "5d5a14cc-4ee2-4b9f-beb6-4afa998bb98f",
            "backupPlanVersion": "Mjc4ZTRhMzUtMGE5Ni00NmQ5LWE1YmMtOWMwY2IwMTY4NWQ4",
            "backupRuleId": "256ad167-f523-4cb9-93f3-f0c933efd97f"
        },
        "creationDate": "2025-12-12T12:12:00Z",
        "iamRoleArn": "arn:aws:iam::1112233445566:role/
RoleForEc2BackupWithNoDescribeTagsPermissions",
        "scannerRoleArn": "arn:aws:iam::1112233445566:role/
RoleForAwsBackupGuarddutyScanner",
        "recoveryPointArn": "arn:aws:backup:us-west-2:1112233445566:recovery-
point:abcd1234-5678-abcd-9012-abcdef123456",
```

```
        "resourceArn": "arn:aws:ec2:us-west-2:1112233445566:instance/
i-0515aee7de03f58e1",
        "resourceType": "EC2",
        "scanId": "43bc14a6-03e0-436a-abf5-8825d3aa6835",
        "scanJobId": "4d5fc12e-7e33-4336-a981-bbe43c300298",
        "scanBaseRecoveryPointArn": "arn:aws:backup:us-west-2:1112233445566:recovery-
point:46bd3edb-a499-4db1-8f8f-269a2ff76fab",
        "malwareScanner": "GUARDDUTY",
        "scanResult": {
            "scanResultStatus": "THREATS_FOUND"
        },
        "state": "COMPLETED",
        "statusMessage": "An AWS Backup scan job was successful completed."
    }
}
```

State: COMPLETED WITH ISSUES

```
{
    "version": "0",
    "id": "60ce181d-67c7-496b-90fb-69636b42daee",
    "detail-type": "Scan Job State Change",
    "source": "aws.backup",
    "account": "1112233445566",
    "time": "2025-12-12T12:12:12Z",
    "region": "us-west-2",
    "resources": [
        "arn:aws:backup:us-west-2:1112233445566:recovery-point:abcd1234-5678-abcd-9012-
abcdef123456"
    ],
    "detail": {
        "accountId": "1112233445566",
        "backupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:9ab3e749-82c6-4342-9320-5edbf4918b86",
        "backupVaultName": "9ab3e749-82c6-4342-9320-5edbf4918b86",
        "completionDate": "2025-12-12T12:12:12Z",
        "createdBy": {
            "backupPlanArn": "arn:aws:backup:us-west-2:1112233445566:backup-
plan:5d5a14cc-4ee2-4b9f-beb6-4afa998bb98f",
            "backupPlanId": "5d5a14cc-4ee2-4b9f-beb6-4afa998bb98f",
            "backupPlanVersion": "Mjc4ZTRhMzUtMGE5Ni00NmQ5LWE1YmMtOWMwY2IwMTY4NWQ4",
            "backupRuleId": "256ad167-f523-4cb9-93f3-f0c933efd97f"
        },
```

```
        "creationDate": "2025-12-12T12:12:00Z",
        "iamRoleArn": "arn:aws:iam::1112233445566:role/
RoleForEc2BackupWithNoDescribeTagsPermissions",
        "scannerRoleArn": "arn:aws:iam::1112233445566:role/
RoleForAwsBackupGuarddutyScanner",
        "recoveryPointArn": "arn:aws:backup:us-west-2:1112233445566:recovery-
point:abcd1234-5678-abcd-9012-abcdef123456",
        "resourceArn": "arn:aws:ec2:us-west-2:1112233445566:instance/
i-0515aee7de03f58e1",
        "resourceType": "EC2",
        "scanId": "43bc14a6-03e0-436a-abf5-8825d3aa6835",
        "scanJobId": "4d5fc12e-7e33-4336-a981-bbe43c300298",
        "scanBaseRecoveryPointArn": "arn:aws:backup:us-west-2:1112233445566:recovery-
point:46bd3edb-a499-4db1-8f8f-269a2ff76fab",
        "malwareScanner": "GUARDDUTY",
        "scanResult": {
            "scanResultStatus": "NO_THREATS_FOUND"
        },
        "state": "COMPLETED_WITH_ISSUES",
        "statusMessage": "Scan job partially completed. View more details in Amazon
 GuardDuty"
    }
}
```

State: FAILED

```
{
    "version": "0",
    "id": "60ce181d-67c7-496b-90fb-69636b42daee",
    "detail-type": "Scan Job State Change",
    "source": "aws.backup",
    "account": "1112233445566",
    "time": "2025-12-12T12:12:12Z",
    "region": "us-west-2",
    "resources": [
        "arn:aws:backup:us-west-2:1112233445566:recovery-point:abcd1234-5678-abcd-9012-
abcdef123456"
    ],
    "detail": {
        "accountId": "1112233445566",
        "backupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:9ab3e749-82c6-4342-9320-5edbf4918b86",
        "backupVaultName": "9ab3e749-82c6-4342-9320-5edbf4918b86",
```

```
        "completionDate": "2025-12-12T12:12:12Z",
        "createdBy": {
            "backupPlanArn": "arn:aws:backup:us-west-2:1112233445566:backup-
plan:5d5a14cc-4ee2-4b9f-beb6-4afa998bb98f",
            "backupPlanId": "5d5a14cc-4ee2-4b9f-beb6-4afa998bb98f",
            "backupPlanVersion": "Mjc4ZTRhMzUtMGE5Ni00NmQ5LWE1YmMtOWMwY2IwMTY4NWQ4",
            "backupRuleId": "256ad167-f523-4cb9-93f3-f0c933efd97f"
        },
        "creationDate": "2025-12-12T12:12:00Z",
        "iamRoleArn": "arn:aws:iam::1112233445566:role/
RoleForEc2BackupWithNoDescribeTagsPermissions",
        "scannerRoleArn": "arn:aws:iam::1112233445566:role/
RoleForAwsBackupGuarddutyScanner",
        "recoveryPointArn": "arn:aws:backup:us-west-2:1112233445566:recovery-
point:abcd1234-5678-abcd-9012-abcdef123456",
        "resourceArn": "arn:aws:ec2:us-west-2:1112233445566:instance/
i-0515aee7de03f58e1",
        "resourceType": "EC2",
        "scanId": "43bc14a6-03e0-436a-abf5-8825d3aa6835",
        "scanJobId": "4d5fc12e-7e33-4336-a981-bbe43c300298",
        "scanBaseRecoveryPointArn": "arn:aws:backup:us-west-2:1112233445566:recovery-
point:46bd3edb-a499-4db1-8f8f-269a2ff76fab",
        "malwareScanner": "GUARDDUTY",
        "state": "FAILED",
        "statusMessage": "<failure reason>"
    }
}
```

State: CANCELED

```
{
    "version": "0",
    "id": "60ce181d-67c7-496b-90fb-69636b42daee",
    "detail-type": "Scan Job State Change",
    "source": "aws.backup",
    "account": "1112233445566",
    "time": "2025-12-12T12:12:12Z",
    "region": "us-west-2",
    "resources": [
        "arn:aws:backup:us-west-2:1112233445566:recovery-point:abcd1234-5678-abcd-9012-
abcdef123456"
    ],
    "detail": {
```

```
        "accountId": "1112233445566",
        "backupVaultArn": "arn:aws:backup:us-west-2:1112233445566:backup-
vault:9ab3e749-82c6-4342-9320-5edbf4918b86",
        "backupVaultName": "9ab3e749-82c6-4342-9320-5edbf4918b86",
        "completionDate": "2025-12-12T12:12:12Z",
        "createdBy": {
            "backupPlanArn": "arn:aws:backup:us-west-2:1112233445566:backup-
plan:5d5a14cc-4ee2-4b9f-beb6-4afa998bb98f",
            "backupPlanId": "5d5a14cc-4ee2-4b9f-beb6-4afa998bb98f",
            "backupPlanVersion": "Mjc4ZTRhMzUtMGE5Ni00NmQ5LWE1YmMtOWMwY2IwMTY4NWQ4",
            "backupRuleId": "256ad167-f523-4cb9-93f3-f0c933efd97f"
        },
        "creationDate": "2025-12-12T12:12:00Z",
        "iamRoleArn": "arn:aws:iam::1112233445566:role/
RoleForEc2BackupWithNoDescribeTagsPermissions",
        "scannerRoleArn": "arn:aws:iam::1112233445566:role/
RoleForAwsBackupGuarddutyScanner",
        "recoveryPointArn": "arn:aws:backup:us-west-2:1112233445566:recovery-
point:abcd1234-5678-abcd-9012-abcdef123456",
        "resourceArn": "arn:aws:ec2:us-west-2:1112233445566:instance/
i-0515aee7de03f58e1",
        "resourceType": "EC2",
        "scanId": "43bc14a6-03e0-436a-abf5-8825d3aa6835",
        "scanJobId": "4d5fc12e-7e33-4336-a981-bbe43c300298",
        "scanBaseRecoveryPointArn": "arn:aws:backup:us-west-2:1112233445566:recovery-
point:46bd3edb-a499-4db1-8f8f-269a2ff76fab",
        "malwareScanner": "GUARDDUTY",
        "state": "CANCELED",
        "statusMessage": "Scan job was stopped by user."
    }
}
```


Backup access point events

The following are example events.

States:

- AVAILABLE
- FAILED
- DELETED
- EXPIRED

- DISASSOCIATED

State: AVAILABLE

```
{
  "version": "0",
  "id": "12345678-1234-1234-1234-123456789012",
  "detail-type": "Backup Access Point State Change",
  "source": "aws.backup",
  "account": "123456789012",
  "time": "2026-07-28T12:00:00Z",
  "region": "us-east-1",
  "resources": [
    "arn:aws:backup:us-east-1:123456789012:accesspoint/my-access-point"
  ],
  "detail": {
    "accessPointArn": "arn:aws:backup:us-east-1:123456789012:accesspoint/my-access-
point",
    "s3AccessPointArn": "arn:aws:s3:us-east-1:123456789012:accesspoint/my-access-
point",
    "backupAccessPointName": "my-access-point",
    "backupVaultArn": "arn:aws:backup:us-east-1:123456789012:backup-vault:MyVault",
    "backupVaultName": "MyVault",
    "recoveryPointArn": "arn:aws:backup:us-east-1:123456789012:recovery-
point:rp-1234567890abcdef0",
    "resourceArn": "arn:aws:s3:::my-bucket",
    "resourceType": "S3",
    "creationDate": "2026-07-28T11:59:00Z",
    "status": "AVAILABLE"
  }
}
```

State: FAILED

```
{
  "version": "0",
  "id": "12345678-1234-1234-1234-123456789012",
  "detail-type": "Backup Access Point State Change",
  "source": "aws.backup",
  "account": "123456789012",
  "time": "2026-07-28T12:05:00Z",
  "region": "us-east-1",
```

```
  "resources": [
    "arn:aws:backup:us-east-1:123456789012:accesspoint/my-access-point"
  ],
  "detail": {
    "accessPointArn": "arn:aws:backup:us-east-1:123456789012:accesspoint/my-access-
point",
    "backupAccessPointName": "my-access-point",
    "backupVaultArn": "arn:aws:backup:us-east-1:123456789012:backup-vault:MyVault",
    "backupVaultName": "MyVault",
    "recoveryPointArn": "arn:aws:backup:us-east-1:123456789012:recovery-
point:rp-1234567890abcdef0",
    "resourceArn": "arn:aws:s3:::my-bucket",
    "resourceType": "S3",
    "creationDate": "2026-07-28T12:04:00Z",
    "statusMessage": "An S3 access point with the same name already exists in this
 Region and account.",
    "status": "FAILED"
  }
}
```

AWS Backup metrics with Amazon CloudWatch

Topics

- CloudWatch Dashboard
- Metrics with CloudWatch

CloudWatch Dashboard

> **Note**
>
> The console dashboard depends on which Region is accessing the console. See Feature availability by AWS Region to see which Regions have access to the Jobs dashboard. Regions not listed will be able to access the CloudWatch dashboard.

Your AWS Backup console includes a dashboard to see metrics on completed or failed backup, copy, and restore jobs. Within this dashboard, you can view job status by time period, customized to the time frame you desire.

TO ACCESS THE DASHBOARD

1. Open the AWS Backup console at https://console.aws.amazon.com/backup.
2. Select **Dashboard** in the left-hand navigation pane.

VIEW AND UNDERSTAND THE DASHBOARD

The CloudWatch dashboard displays several widgets. Each widget shows job metrics by count. Each widget shows several line graphs. Each line corresponds to a protected resource (if you do not see an expected resource displayed, ensure the resource is turned on in **Settings**). The displays do not show in-progress jobs.

The y-axis (vertical values) shows the count. The x-axis (horizontal values) shows points in time. If there are no data points to visualize in the selected job status, the value will be set to 0 with a horizontal line on the x-axis. The legend showing the resources will still be visible.

The metrics display account-specific and Region-specific information related to the current login. To see other accounts or Regions, you must login under the chosen account.

CUSTOMIZE THE DASHBOARD

By default, the displayed time frame is one week. Along the top menu, there are options for redefining the displayed time frame. You can choose from among 1 hour, 3 hours, 12 hours, 1 day, 3 days, and 1 week. Additionally, you can select **Custom** to specify a different value. Customization will temporarily change the current view to your specifications.

You can hover over a widget, which will display a **Enlarge** button in the top right of the widget. Click on **Enlarge** to open the widget in full-screen view. In full screen, there are more options for customizing the graph display, such as changing the period (the time between every data point). Any changes will not be retained once the full-screen view is closed.

To view only one resource type at a time, click on the label text of the resource type you wish to view in the graph legend. This will deselect other all resource types. To reverse this, click on a resource type color box in the legend. To go back to default view of all resource types with all the labels selected, click again on the label text of any resource type selected.

Clicking the three vertical dots in the top right corner of a widgets opens up a drop down menu with options to refresh, enlarge, view in metrics and view in logs. “View in metrics” opens up the metric used in the widget in CloudWatch console. You can make any changes to the widget there and add the widget to a custom dashboard in CloudWatch dashboard. Any changes you make in

the CloudWatch dashboard will not be reflected on the dashboard in AWS Backup Console. “View as logs” opens up the logs view page in CloudWatch console.

To add widgets displayed to your own custom CloudWatch dashboard, click on the **Add to dashboard** button located on the top right of the dashboard. This will open up the CloudWatch console where you can select in which custom dashboard to add all the six widgets.

For more information, see Using Amazon CloudWatch metrics.

Metrics with CloudWatch

You can use CloudWatch to monitor AWS Backup metrics. The `AWS/Backup` namespace allows you to track the following metrics. AWS Backup emits updated metrics to CloudWatch every 5 minutes.

The purpose of this documentation page is to provide you with the reference materials to use CloudWatch to monitor AWS Backup. To learn how to monitor a metric using CloudWatch, see the blog Amazon CloudWatch Events and Metrics for AWS Backup or Focus on Metrics and Alarms in a Single AWS Service in the *CloudWatch User Guide*. To set alarms, see Using Amazon CloudWatch Alarms in the *CloudWatch User Guide*.

Category	Metrics	Example dimensions	Example use case
Jobs	Number of backup, restore, and copy jobs across each state, including `CREATED`, `PENDING`, `RUNNING`, `ABORTED`, `COMPLETED` , `FAILED`, and `EXPIRED`. Different job types have different available states.	Resource type, vault name. The vault name of copy jobs is that of their destination vault.	Monitor the number of failed backup jobs within one or more specific backup vaults. When there are more than five failed jobs within 1 hour, send an email or SMS using Amazon SNS or open a ticket to the engineering team to investigate. **Reporting criteria**: There is a nonzero value

Category	Metrics	Example dimensions	Example use case
Recovery points	Number of warm and cold recovery points across each state: `MODIFIED`, `COMPLETED` , `PARTIAL`, `EXPIRED`, `DELETED`.	Resource type, vault name.	Track the number of deleted recovery points for your Amazon EBS volumes, and separately track the number of warm and cold recovery points in each backup vault. **Reporting criteria**: There is a nonzero value

> **Note**
>
> The job status of `Completed with issues` is specific to only the AWS Backup console; it cannot be tracked via CloudWatch.

The following table lists all the metrics available to you.

Metric	Description
`NumberOfBackupJobsCreated`	The number of backup jobs that AWS Backup created.
`NumberOfBackupJobsPending`	The number of backup jobs about to run in AWS Backup.
`NumberOfBackupJobsRunning`	The number of backup jobs currently running in AWS Backup.
`NumberOfBackupJobsAborted`	The number of user cancelled backup jobs.

Metric	Description
NumberOfBackupJobsCompleted	The number of backup jobs that AWS Backup finished.
NumberOfBackupJobsFailed	The number of backup jobs with status of Failed. Often caused by scheduling a backup job during or 1 hour before a database resource or 3 hours before or during a Amazon FSx maintenance window or automated backup window and not using AWS Backup to perform continuous backup for point-in-time restores. See Point-in-Time Recovery for a list of supported services and instructions on how to use AWS Backup to take continuous backups, or reschedule your backup jobs.
NumberOfBackupJobsExpired	The number of backup jobs that have a status of EXPIRED. A backup job changes from status CREATED to EXPIRED if a backup cannot begin within the start window time.
NumberOfCopyJobsCreated	The number of cross-account and cross-Region copy jobs that AWS Backup created.
NumberOfCopyJobsRunning	The number of cross-account and cross-Region copy jobs currently running in AWS Backup.
NumberOfCopyJobsCompleted	The number of cross-account and cross-Region copy jobs that AWS Backup finished.
NumberOfCopyJobsFailed	The number of cross-account and cross-Reg ion copy jobs that AWS Backup attempted but could not complete.
NumberOfRestoreJobsPending	The number of restore jobs about to run in AWS Backup.

Metric	Description
NumberOfRestoreJobsRunning	The number of restore jobs currently running in AWS Backup.
NumberOfRestoreJobsCompleted	The number of restore jobs that AWS Backup finished.
NumberOfRestoreJobsFailed	The number of restore jobs that AWS Backup attempted but could not complete.
NumberOfRecoveryPointsCompleted	The number of recovery points that AWS Backup created.
NumberOfRecoveryPointsPartial	The number of recovery points that AWS Backup started to create but could not finish. AWS retries the process later, but because the retry occurs at the later time, it retains the partial recovery point.
NumberOfRecoveryPointsExpired	The number of recovery points that AWS Backup attempted to delete based on your backup retention lifecycle, but could not delete. You are billed for the storage that expired backups consume and should delete them manually.
NumberOfRecoveryPointsDeleting	The number of recovery points that AWS Backup is deleting.
NumberOfRecoveryPointsCold	The number of recovery points that AWS Backup tiered to cold storage.

More dimensions are available beyond those listed in the table. To view all the dimensions of a metric, type the name of that metric into the AWS/Backup namespace of the **Metrics** section of the CloudWatch console.

Logging AWS Backup API calls with CloudTrail

AWS Backup is integrated with AWS CloudTrail a service that provides a record of actions taken by a user, role, or an AWS service service. CloudTrail captures all API calls for AWS Backup as events. The calls captured include calls from the AWS Backup console and code calls to the AWS Backup API operations. Using the information collected by CloudTrail, you can determine the request that was made to AWS Backup, the IP address from which the request was made, when it was made, and additional details.

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 user or user credentials.
- Whether the request was made on behalf of an IAM Identity Center user.
- Whether the request was made with temporary security credentials for a role or federated user.
- Whether the request was made by another AWS service.

CloudTrail is active in your AWS account when you create the account and you automatically have access to the CloudTrail **Event history**. The CloudTrail **Event history** provides a viewable, searchable, downloadable, and immutable record of the past 90 days of recorded management events in an AWS Region. For more information, see Working with CloudTrail Event history in the *AWS CloudTrail User Guide*. There are no CloudTrail charges for viewing the **Event history**.

For an ongoing record of events in your AWS account past 90 days, create a trail or a CloudTrail Lake event data store.

CloudTrail trails

A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. All trails created using the AWS Management Console are multi-Region. You can create a single-Region or a multi-Region trail by using the AWS CLI. Creating a multi-Region trail is recommended because you capture activity in all AWS Regions in your account. If you create a single-Region trail, you can view only the events logged in the trail's AWS Region. For more information about trails, see Creating a trail for your AWS account and Creating a trail for an organization in the *AWS CloudTrail User Guide*.

You can deliver one copy of your ongoing management events to your Amazon S3 bucket at no charge from CloudTrail by creating a trail, however, there are Amazon S3 storage charges. For

more information about CloudTrail pricing, see AWS CloudTrail Pricing. For information about Amazon S3 pricing, see Amazon S3 Pricing.

CloudTrail Lake event data stores

CloudTrail Lake lets you run SQL-based queries on your events. CloudTrail Lake converts existing events in row-based JSON format to Apache ORC format. ORC is a columnar storage format that is optimized for fast retrieval of data. Events are aggregated into *event data stores*, which are immutable collections of events based on criteria that you select by applying advanced event selectors. The selectors that you apply to an event data store control which events persist and are available for you to query. For more information about CloudTrail Lake, see Working with AWS CloudTrail Lake in the *AWS CloudTrail User Guide*.

CloudTrail Lake event data stores and queries incur costs. When you create an event data store, you choose the pricing option you want to use for the event data store. The pricing option determines the cost for ingesting and storing events, and the default and maximum retention period for the event data store. For more information about CloudTrail pricing, see AWS CloudTrail Pricing.

AWS Backup events in CloudTrail

AWS Backup generates these CloudTrail events when it performs backups, restores, copies, scans, notifications, or access point operations. These events are not necessarily generated by use of the AWS Backup public APIs. For more information, see AWS service events in the *AWS CloudTrail User Guide*.

- `AssociateBackupVaultMpaApprovalTeamCompleted`
- `AssociateBackupVaultMpaApprovalTeamFailed`
- `BackupAccessPointCreated`
- `BackupAccessPointCreationFailed`
- `BackupAccessPointDeleted`
- `BackupAccessPointDeletionFailed`
- `BackupAccessPointExpired`
- `BackupAccessPointDisassociated`
- `BackupDeleted`
- `BackupJobCompleted`

- BackupJobStarted
- BackupSelectionDeletedDueToSLRDeletion
- BackupTransitionedToCold
- ContinuousRecoveryPointIntegrityCheckCompleted
- ContinuousRecoveryPointIntegrityCheckFailed
- ContinuousRecoveryPointIntegrityCheckStarted
- CopyJobCompleted
- CopyJobStarted
- CreateRestoreAccessBackupVaultFailed
- DisassociateBackupVaultMpaApprovalTeamFailed
- PutBackupVaultNotifications
- ProcessBackupPlanSelection
- RecoveryPointCreated
- ReportJobCompleted
- ReportJobStarted
- RestoreAccessBackupVaultDeleted
- RestoreCompleted
- RestoreStarted
- RevokeRestoreAccessBackupVaultFailed
- ScanJobCompleted
- ScanJobCreated
- ScanJobFailed
- ScanJobStarted

Understanding AWS Backup 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 `StartBackupJob`, `StartRestoreJob`, and `DeleteRecoveryPoint` actions and also the `BackupJobCompleted` event.


```
{
    "eventVersion": "1.05",
    "userIdentity": {
        "type": "Root",
        "principalId": "123456789012",
        "arn": "arn:aws:iam::123456789012:root",
        "accountId": "123456789012",
        "accessKeyId": "AKIAI44QH8DHBEXAMPLE",
        "sessionContext": {
            "attributes": {
                "mfaAuthenticated": "false",
                "creationDate": "2019-01-10T12:24:50Z"
            }
        }
    },
    "eventTime": "2019-01-10T13:45:24Z",
    "eventSource": "backup.amazonaws.com",
    "eventName": "StartBackupJob",
    "awsRegion": "us-east-1",
    "sourceIPAddress": "12.34.567.89",
    "userAgent": "aws-internal/3 aws-sdk-java/1.11.465
 Linux/4.9.124-0.1.ac.198.73.329.metal1.x86_64 OpenJDK_64-Bit_Server_VM/25.192-b12
 java/1.8.0_192",
    "requestParameters": {
        "backupVaultName": "Default",
        "resourceArn": "arn:aws:ec2:us-east-1:123456789012:volume/
vol-00a422a05b9c6asd3",
        "iamRoleArn": "arn:aws:iam::123456789012:role/AWSBackup",
        "startWindowMinutes": 60
    },
    "responseElements": {
        "backupJobId": "8a3c2a87-b23e-4d56-b045-fa9e88ede4e6",
        "creationDate": "Jan 10, 2019 1:45:24 PM"
    },
    "requestID": "98cf4d59-8c76-49f7-9201-790743931234",
    "eventID": "fe8146a5-7812-4a95-90ad-074498be1234",
    "eventType": "AwsApiCall",
    "recipientAccountId": "account-id"
},
```

```
{
    "eventVersion": "1.05",
    "userIdentity": {
        "type": "Root",
        "principalId": "123456789012",
        "arn": "arn:aws:iam::123456789012:root",
        "accountId": "123456789012",
        "accessKeyId": "AKIAI44QH8DHBEXAMPLE",
        "sessionContext": {
            "attributes": {
                "mfaAuthenticated": "false",
                "creationDate": "2019-01-10T12:24:50Z"
            }
        }
    },
    "eventTime": "2019-01-10T13:49:50Z",
    "eventSource": "backup.amazonaws.com",
    "eventName": "StartRestoreJob",
    "awsRegion": "us-east-1",
    "sourceIPAddress": "12.34.567.89",
    "userAgent": "aws-internal/3 aws-sdk-java/1.11.465
 Linux/4.9.124-0.1.ac.198.73.329.metal1.x86_64 OpenJDK_64-Bit_Server_VM/25.192-b12
 java/1.8.0_192",
    "requestParameters": {
        "recoveryPointArn": "arn:aws:ec2:us-east-1::snapshot/snap-00a129455bdbc9d99",
        "metadata": {
            "volumeType": "gp2",
            "availabilityZone": "us-east-1b",
            "volumeSize": "100"
        },
        "iamRoleArn": "arn:aws:iam::123456789012:role/AWSBackup",
        "idempotencyToken": "a9c8b4fb-d369-4a58-944b-942e442a8fe3",
        "resourceType": "EBS"
    },
    "responseElements": {
        "restoreJobId": "9808E090-8C76-CCB8-4CEA-407CF6AC4C43"
    },
    "requestID": "783ddddc-6d7e-4539-8fab-376aa9668543",
    "eventID": "ff35ddea-7577-4aec-a132-964b7e9dd423",
    "eventType": "AwsApiCall",
    "recipientAccountId": "account-id"
},
{
    "eventVersion": "1.05",
```

```
    "userIdentity": {
        "type": "Root",
        "principalId": "123456789012",
        "arn": "arn:aws:iam::123456789012:root",
        "accountId": "123456789012",
        "accessKeyId": "AKIAI44QH8DHBEXAMPLE",
        "sessionContext": {
            "attributes": {
                "mfaAuthenticated": "false",
                "creationDate": "2019-01-10T12:24:50Z"
            }
        }
    },
    "eventTime": "2019-01-10T14:52:42Z",
    "eventSource": "backup.amazonaws.com",
    "eventName": "DeleteRecoveryPoint",
    "awsRegion": "us-east-1",
    "sourceIPAddress": "12.34.567.89",
    "userAgent": "aws-internal/3 aws-sdk-java/1.11.465
 Linux/4.9.124-0.1.ac.198.73.329.metal1.x86_64 OpenJDK_64-Bit_Server_VM/25.192-b12
 java/1.8.0_192",
    "requestParameters": {
        "backupVaultName": "Default",
        "recoveryPointArn": "arn:aws:ec2:us-east-1::snapshot/snap-05f426fd9daab3433"
    },
    "responseElements": null,
    "requestID": "f1f1b33a-48da-436c-9a8f-7574f1ab5fd7",
    "eventID": "2dd70080-5aba-4a79-9a0f-92647c9f0846",
    "eventType": "AwsApiCall",
    "recipientAccountId": "account-id"
},
{
    "eventVersion": "1.05",
    "userIdentity": {
        "accountId": "123456789012",
        "invokedBy": "backup.amazonaws.com"
    },
    "eventTime": "2019-01-10T08:24:39Z",
    "eventSource": "backup.amazonaws.com",
    "eventName": "BackupJobCompleted",
    "awsRegion": "us-east-1",
    "sourceIPAddress": "backup.amazonaws.com",
    "userAgent": "backup.amazonaws.com",
    "requestParameters": null,
```

```
    "responseElements": null,
    "eventID": "2e7e4fcf-0c52-467f-9fd0-f61c2fcf7d17",
    "eventType": "AwsServiceEvent",
    "recipientAccountId": "account-id",
    "serviceEventDetails": {
        "completionDate": {
            "seconds": 1547108091,
            "nanos": 906000000
        },
        "state": "COMPLETED",
        "percentDone": 100,
        "backupJobId": "8A8E738B-A8C5-E058-8224-90FA323A3C0E",
        "backupVaultName": "BackupVault",
        "backupVaultArn": "arn:aws:backup:us-east-1:123456789012:backup-
vault:BackupVault",
        "recoveryPointArn": "arn:aws:ec2:us-east-1::snapshot/snap-07ce8c3141d361233",
        "resourceArn": "arn:aws:ec2:us-east-1:123456789012:volume/
vol-06692095a6a421233",
        "creationDate": {
            "seconds": 1547101638,
            "nanos": 272000000
        },
        "backupSizeInBytes": 8589934592,
        "iamRoleArn": "arn:aws:iam::123456789012:role/AWSBackup",
        "resourceType": "EBS"
    }
}
```


Logging cross-account management events

With AWS Backup, you can manage your backups across all AWS accounts inside your AWS Organizations structure. AWS Backup generates these CloudTrail events in your member account when you create, update, or delete an AWS Organizations backup policy (that applies backup plans to your member accounts) or when there is an invalid organization backup plan:

- `CreateOrganizationalBackupPlan`
- `UpdateOrganizationalBackupPlan`
- `DeleteOrganizationalBackupPlan`
- `InvalidOrganizationBackupPlan`

Example: AWS Backup log file entries for cross-account management

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 `CreateOrganizationalBackupPlan` action.


```
{
    "eventVersion": "1.05",
    "userIdentity": {
        "accountId": "123456789012",
        "invokedBy": "backup.amazonaws.com"},
    "eventTime": "2020-06-02T00:34:00Z",
    "eventSource": "backup.amazonaws.com",
    "eventName": "CreateOrganizationalBackupPlan",
    "awsRegion": "ca-central-1",
    "sourceIPAddress": "backup.amazonaws.com",
    "userAgent": "backup.amazonaws.com",
    "requestParameters": null,
    "responseElements": null,
    "eventID": "f2642255-af77-4203-8c37-7ca19d898e84",
    "readOnly": false,
    "eventType": "AwsServiceEvent",
    "recipientAccountId": "account-id",
    "serviceEventDetails": {
        "backupPlanId": "orgs/544033d1-b19c-3f2a-9c20-40bcfa82ca68",
        "backupPlanVersionId": "ZTA1Y2ZjZDYtNmRjMy00ZTA1LWIyNTAtM2M1NzQ4OThmNzRj",
        "backupPlanArn": "arn:aws:backup:ca-central-1:123456789012:backup-
plan:orgs/544033d1-b19c-3f2a-9c20-40bcfa82ca68",
        "backupPlanName": "mybackupplan",
        "backupRules": "[{\"id\":\"745fd0ea-7f57-3f35-8a0e-ed4b8c48a8e2\",
\"name\":\"hourly\",\"description\":null,\"cryopodArn\":\"arn:aws:backup:ca-
central-1:123456789012:backup-vault:ControllerCAMTestBackupVault\",\"scheduleExpression
\":\"cron(0 0/1 ? * * *)\",\"startWindow\":\"PT1H\",\"completionWindow\":\"PT2H\",
\"lifecycle\":{\"moveToColdStorageAfterDays\":null,\"deleteAfterDays\":\"7\"},\"tags
\":null,\"copyActions\":[]}]",
        "backupSelections": "[{\"name\":\"selectiondatatype\",\"arn\":
\"arn:aws:backup:ca-central-1:123456789012:selection:8b40c6d9-3641-3d49-926d-
a075ea715686\",\"role\":\"arn:aws:iam::123456789012:role/OrganizationmyRoleTestRole\",
```

```
\"resources\":[],\"notResources\":[],\"conditions\":[{\"type\":\"STRINGEQUALS\",\"key
\":\"dataType\",\"value\":\"PII\"},{\"type\":\"STRINGEQUALS\",\"key\":\"dataType\",
\"value\":\"RED\"}],\"creationDate\":\"2020-06-02T00:34:00.695Z\",\"creatorRequestId
\":null}]",
        "creationDate": {
            "seconds": 1591058040,
            "nanos": 695000000
        },
        "organizationId": "org-id",
        "accountId": "123456789012"
    }
}
```


The following example shows a CloudTrail log entry that demonstrates the `DeleteOrganizationalBackupPlan` action.


```
{
    "eventVersion": "1.05",
    "userIdentity": {
        "accountId": "123456789012",
        "invokedBy": "backup.amazonaws.com"
    },
    "eventTime": "2020-06-02T00:34:25Z",
    "eventSource": "backup.amazonaws.com",
    "eventName": "DeleteOrganizationalBackupPlan",
    "awsRegion": "ca-central-1",
    "sourceIPAddress": "backup.amazonaws.com",
    "userAgent": "backup.amazonaws.com",
    "requestParameters": null,
    "responseElements": null,
    "eventID": "5ce66cd0-b90c-4957-8e00-96ea1077b4fa",
    "readOnly": false,
    "eventType": "AwsServiceEvent",
    "recipientAccountId": "account-id",
    "serviceEventDetails": {
        "backupPlanId": "orgs/544033d1-b19c-3f2a-9c20-40bcfa82ca68",
        "backupPlanVersionId": "ZTA1Y2ZjZDYtNmRjMy00ZTA1LWIyNTAtM2M1NzQ4OThmNzRj",
        "backupPlanArn": "arn:aws:backup:ca-central-1:123456789012:backup-
plan:orgs/544033d1-b19c-3f2a-9c20-40bcfa82ca68",
        "backupPlanName": "mybackupplan",
        "deletionDate": {
            "seconds": 1591058065,
            "nanos": 519000000
```

```
        },
        "organizationId": "org-id",
        "accountId": "123456789012"
    }
}
```

The following example shows a CloudTrail log entry that demonstrates the event `InvalidOrganizationBackupPlan`, which is sent when AWS Backup receives an invalid backup plan from Organizations.

```
{
    "eventVersion": "1.08",
    "userIdentity": {
        "accountId": "123456789012",
        "invokedBy": "backup.amazonaws.com"
    },
    "eventTime": "2022-06-11T13:29:23Z",
    "eventSource": "backup.amazonaws.com",
    "eventName": "InvalidOrganizationBackupPlan",
    "awsRegion": "Region",
    "sourceIPAddress": "backup.amazonaws.com",
    "userAgent": "backup.amazonaws.com",
    "requestParameters": null,
    "responseElements": null,
    "eventID": "ab1de234-fg56-7890-h123-45ij678k9l01",
    "readOnly": false,
    "eventType": "AwsServiceEvent",
    "managementEvent": true,
    "recipientAccountId": "987654321098",
    "serviceEventDetails": {
        "effectivePolicyVersion": 7,
        "effectivePolicyId": "12345678-a9b0-123c-45d6-78e901f23456",
        "lastUpdatedTimestamp": "Jun 11, 2022 1:29:22 PM",
        "policyType": "BACKUP_POLICY",
        "effectiveBackupPlan": {
            "logicalName": "logical-name",
            "regions": [
                "Region"
            ],
            "rules": [
                {
                    "name": "test-orgs",
                    "targetBackupVaultName": "vault-name",
```

```
                    "ruleLifecycle": {
                        "deleteAfterDays": 100
                    },
                    "copyActions": [],
                    "enableContinuousBackup": true
                }
            ],
            "selections": {
                "tagSelections": [
                    {
                        "selectionName": "selection-name",
                        "iamRoleArn": "arn:aws:iam::$account:role/role",
                        "targetedTags": [
                            {
                                "tagKey": "key",
                                "tagValue": "value"
                            }
                        ]
                    }
                ]
            },
            "backupPlanTags": {
                "key": "value"
            }
        },
        "organizationId": "org-id",
        "accountId": "123456789012"
    },
    "eventCategory": "Management"
}
```

Notification options with AWS Backup

There are two ways to receive notifications about AWS Backup:

- User Notifications can send notifications, including Amazon CloudWatch alarms, AWS Support, and other services' notifications.
- Amazon Simple Notification Service can notify you of AWS Backup events.

User Notifications and AWS Backup

AWS Backup supports management of your backup notifications from the User Notifications console. With User Notifications, you can view the progress of your backup, copy, and restore jobs and changes to your backup policies, vaults, recovery points, and settings from the User Notifications Notification Center.

Amazon CloudWatch, Amazon EventBridge alarms, and AWS Support case updates are among other types of notifications you can manage from the console. Additionally, you can set up several delivery options, including email, Amazon Q Developer in chat applications notifications, and AWS Console Mobile Application push notifications.

Amazon SNS and AWS Backup events

AWS Backup takes advantage of the robust notifications delivered by Amazon Simple Notification Service (Amazon SNS). You can configure Amazon SNS to notify you of AWS Backup events from the Amazon SNS console.

Limitations

- While the Amazon SNS service allows cross-account notifications, AWS Backup does not currently support this feature. You must specify your own AWS account ID and the resource ARN of your topic.
- AWS Backup supports Standard topics for SNS best-effort deduplication, but AWS Backup does not currently support SNS FIFO topics for Strict deduplication.

Common use cases

- Set up notifications for failed backup jobs by following the steps in How can I get notifications for AWS Backup jobs that failed? from AWS Premium Support.
- Review sample Amazon SNS notification JSONs for completed, failed, and expired backup jobs in the Examples of events table below.

For more information about Amazon SNS generally, see Getting Started with Amazon SNS in the *Amazon Simple Notification Service Developer Guide*.

AWS Backup notification APIs

After creating your topics using the Amazon SNS console or AWS Command Line Interface (AWS CLI), you can use the following AWS Backup API operations to manage your backup notifications.

- DeleteBackupVaultNotifications — Deletes event notifications for the specified backup vault.
- GetBackupVaultNotifications — Lists all event notifications for the specified backup vault.
- PutBackupVaultNotifications — Turns on notifications for the specified topic and events.

AWS Backup supports the following events:

Job type	Event
Backup job	`BACKUP_JOB_STARTED | BACKUP_JOB_COMPLETED | BACKUP_JOB_SUCCESSFUL | BACKUP_JOB_FAILED | BACKUP_JOB_EXPIRED | CONTINUOUS_BACKUP_INTERRUPTED`
Copy job	`COPY_JOB_STARTED | COPY_JOB_SUCCESSFUL | COPY_JOB_FAILED`
Restore job	`RESTORE_JOB_STARTED | RESTORE_JOB_COMPLETED | RESTORE_JOB_SUCCESSFUL | RESTORE_JOB_FAILED`
Recovery point	`RECOVERY_POINT_MODIFIED`
Recovery point indexing	`RECOVERY_POINT_INDEX_COMPLETED | RECOVERY_POINT_INDEX_DELETED | RECOVERY_POINT_INDEXING_FAILED`
Access point	`ACCESS_POINT_AVAILABLE | ACCESS_POINT_CREATION_FAILED | ACCESS_POINT_DELETED | ACCESS_POINT_DELETION_FAILED | ACCESS_POINT_EXPIRED | ACCESS_POINT_DISASSOCIATED`

AWS Backup for S3 supports two additional events:

- S3_BACKUP_OBJECT_FAILED notifies you of any S3 object that AWS Backup failed to back up during a backup job.
- S3_RESTORE_OBJECT_FAILED notifies you of any S3 object that AWS Backup failed to restore during a restore job.

AWS Backup for EKS supports five additional events:

- EKS_BACKUP_OBJECT_FAILED notifies you of any EKS objects that AWS Backup failed to back up during a backup job.
- EKS_RESTORE_OBJECT_FAILED notifies you of any EKS objects that AWS Backup failed to restore during a restore job.
- EKS_RESTORE_OBJECT_SKIPPED notifies you of any EKS objects that AWS Backup skipped during a restore job.
- EKS_RESTORE_MODIFICATION_SUCCESS notifies you of any EKS objects that AWS Backup modified during a restore job using your resource modifiers.
- EKS_RESTORE_MODIFICATION_FAILED notifies you of any EKS objects where one of your resource modifiers could not be applied during a restore job.

Examples of events

Example: Backup job completed

```
{
    "Records": [{
        "EventSource": "aws:sns",
        "EventVersion": "1.0",
        "EventSubscriptionArn": "arn:aws:sns:...-a3802aa1ed45",
        "Sns": {
            "Type": "Notification",
            "MessageId": "12345678-abcd-123a-def0-abcd1a234567",
            "TopicArn": "arn:aws:sns:us-west-1:123456789012:backup-2sqs-sns-topic",
            "Subject": "Notification from AWS Backup",
            "Message": "An AWS Backup job was completed successfully. Recovery point
 ARN: arn:aws:ec2:us-west-1:123456789012:volume/vol-012f345df6789012d. Resource ARN :
 arn:aws:ec2:us-west-1:123456789012:volume/vol-012f345df6789012e. BackupJob ID :
 1b2345b2-f22c-4dab-5eb6-bbc7890ed123",
```

```
            "Timestamp": "2019-08-02T18:46:02.788Z",
            ...
            "MessageAttributes": {
                "EventType": {"Type":"String","Value":"BACKUP_JOB"},
                "State": {"Type":"String","Value":"COMPLETED"},
                "AccountId": {"Type":"String","Value":"123456789012"},
                "Id": {"Type":"String","Value":"1b2345b2-f22c-4dab-5eb6-bbc7890ed123"},
                "StartTime": {"Type":"String","Value":"2019-09-02T13:48:52.226Z"}
            }
        }
    }]
}
```

Example: Backup job failed

```
{
    "Records": [{
        "EventSource": "aws:sns",
        "EventVersion": "1.0",
        "EventSubscriptionArn": "arn:aws:sns:...-a3802aa1ed45",
        "Sns": {
            "Type": "Notification",
            "MessageId": "12345678-abcd-123a-def0-abcd1a234567",
            "TopicArn": "arn:aws:sns:us-west-1:123456789012:backup-2sqs-sns-topic",
            "Subject": "Notification from AWS Backup",
            "Message": "An AWS Backup job failed. Resource ARN : arn:aws:ec2:us-
west-1:123456789012:volume/vol-012f345df6789012e. BackupJob ID : 1b2345b2-
f22c-4dab-5eb6-bbc7890ed123",
            "Timestamp": "2019-08-02T18:46:02.788Z",
            ...
            "MessageAttributes": {
                "EventType": {"Type":"String","Value":"BACKUP_JOB"},
                "State": {"Type":"String","Value":"FAILED"},
                "AccountId": {"Type":"String","Value":"123456789012"},
                "Id": {"Type":"String","Value":"1b2345b2-f22c-4dab-5eb6-bbc7890ed123"},
                "StartTime": {"Type":"String","Value":"2019-09-02T13:48:52.226Z"}
            }
        }
    }]
}
```

Example: Backup job could not complete during the backup window

```
{
    "Records": [{
        "EventSource": "aws:sns",
        "EventVersion": "1.0",
        "EventSubscriptionArn": "arn:aws:sns:...-a3802aa1ed45",
        "Sns": {
            "Type": "Notification",
            "MessageId": "12345678-abcd-123a-def0-abcd1a234567",
            "TopicArn": "arn:aws:sns:us-west-1:123456789012:backup-2sqs-sns-topic",
            "Subject": "Notification from AWS Backup",
            "Message": "An AWS Backup job failed to complete in time. Resource ARN :
 arn:aws:ec2:us-west-1:123456789012:volume/vol-012f345df6789012e. BackupJob ID :
 1b2345b2-f22c-4dab-5eb6-bbc7890ed123",
            "Timestamp": "2019-08-02T18:46:02.788Z",
            ...
            "MessageAttributes" : {
              "EventType" : {"Type":"String","Value":"BACKUP_JOB"},
              "State" : {"Type":"String","Value":"EXPIRED"},
              "AccountId" : {"Type":"String","Value":"123456789012"},
              "Id" : {"Type":"String","Value":"1b2345b2-f22c-4dab-5eb6-bbc7890ed123"},
              "StartTime" : {"Type":"String","Value":"2019-09-02T13:48:52.226Z"}
            }
        }
    }]
}
```

Example: Recovery point indexing completed

```
{
    "Records": [{
        "EventSource": "aws:sns",
        "EventVersion": "1.0",
        "EventSubscriptionArn": "arn:aws:sns:...-a3802aa1ed45",
        "Sns": {
            "Type": "Notification",
            "MessageId": "12345678-abcd-123a-def0-abcd1a234567",
            "TopicArn": "arn:aws:sns:us-west-1:123456789012:backup-2sqs-sns-topic",
            "Subject": "Notification from AWS Backup",
            "Message": "An AWS Backup backup index job was completed. Indexed recovery
 point arn: arn:aws:backup:us-west-2:1112233445566:recovery-point:abcd1234-5678-
abcd-9012-abcdef123456",
```

```
            "Timestamp": "2025-05-25T18:46:02.788Z",
            ...
            "MessageAttributes" : {
              "EventType" :
 {"Type":"String","Value":"RECOVERY_POINT_INDEXING_COMPLETED"},
              "AccountId" : {"Type":"String","Value":"123456789012"},
              "IndexStatus" : {"Type":"String","Value":"ACTIVE"},
              "IsIndexingContinuous" : {"Type":"String","Value":"false"},
              "RecoveryPointArn" : {"Type":"String","Value":"arn:aws:backup:us-
west-2:1112233445566:recovery-point:abcd1234-5678-abcd-9012-abcdef123456"}
            }
        }
    }]
}
```

Example: EKS backup object failed

```
{
    "Records": [{
        "EventSource": "aws:sns",
        "EventVersion": "1.0",
        "EventSubscriptionArn": "arn:aws:sns:...-a3802aa1ed45",
        "Sns": {
            "Type": "Notification",
            "MessageId": "12345678-abcd-123a-def0-abcd1a234567",
            "TopicArn": "arn:aws:sns:us-west-1:123456789012:backup-2sqs-sns-topic",
            "Subject": "Notification from AWS Backup",
            "Message": "A Kubernetes resource failed to backup from your Amazon
 EKS Backup. Resource: example.resource.io/v1. EKS Cluster Name: eks-cluster-name.
 BackupJob ID: 1b2345b2-f22c-4dab-5eb6-bbc7890ed123",
            "Timestamp": "2025-05-25T18:46:02.788Z",
            ...
            "MessageAttributes" : {
              "eventType" : {"Type":"String","Value":"EKS_BACKUP_OBJECT_FAILED"},
              "backupJobId" : {"Type":"String","Value":"1b2345b2-f22c-4dab-5eb6-
bbc7890ed123"},
              "clusterName" : {"Type":"String","Value":"eks-cluster-name"},
              "reason" : {"Type":"String","Value":"Example failure reason."},
              "resourceName" : {"Type":"String","Value":"example.resource.io/v1"}
            }
        }
    }]
}
```

Example: EKS restore object failed

```
{
    "Records": [{
        "EventSource": "aws:sns",
        "EventVersion": "1.0",
        "EventSubscriptionArn": "arn:aws:sns:...-a3802aa1ed45",
        "Sns": {
            "Type": "Notification",
            "MessageId": "12345678-abcd-123a-def0-abcd1a234567",
            "TopicArn": "arn:aws:sns:us-west-1:123456789012:backup-2sqs-sns-topic",
            "Subject": "Notification from AWS Backup",
            "Message": "A Kubernetes resource failed to restore from your Amazon EKS
 Backup. Resource: apiextensions.k8s.io/v1/customresourcedefinitions. Resource Name:
 exampleresource. Destination EKS Cluster Name: eks-restore-target-cluster-name.
 RestoreJob ID: 1b2345b2-f22c-4dab-5eb6-bbc7890ed123",
            "Timestamp": "2025-05-25T18:46:02.788Z",
            ...
            "MessageAttributes" : {
              "eventType" : {"Type":"String","Value":"EKS_RESTORE_OBJECT_FAILED"},
              "clusterName" : {"Type":"String","Value":"eks-restore-target-cluster-
name"},
              "parentRestoreJobId" : {"Type":"String","Value":"12345678-abcd-123a-def0-
abcd1a234567"},
              "reason" : {"Type":"String","Value":"Example failure reason."},
              "resourceName" : {"Type":"String","Value":"exampleresourceio"},
              "resourceType" : {"Type":"String","Value":"apiextensions.k8s.io/v1/
customresourcedefinitions"},
              "restoreJobId" : {"Type":"String","Value":"1b2345b2-f22c-4dab-5eb6-
bbc7890ed123"}
            }
        }
    }]
}
```

Example: EKS restore object skipped

```
{
    "Records": [{
        "EventSource": "aws:sns",
        "EventVersion": "1.0",
        "EventSubscriptionArn": "arn:aws:sns:...-a3802aa1ed45",
        "Sns": {
```

```
            "Type": "Notification",
            "MessageId": "12345678-abcd-123a-def0-abcd1a234567",
            "TopicArn": "arn:aws:sns:us-west-1:123456789012:backup-2sqs-sns-topic",
            "Subject": "Notification from AWS Backup",
            "Message": "A Kubernetes resource was skipped from restore from your
 Amazon EKS Backup. This Kubernetes resource already exists in your target EKS
 cluster. Resource: apiextensions.k8s.io/v1/customresourcedefinitions. Resource Name:
 exampleresource. Destination EKS Cluster Name: eks-restore-target-cluster-name.
 RestoreJob ID: 1b2345b2-f22c-4dab-5eb6-bbc7890ed123",
            "Timestamp": "2025-05-25T18:46:02.788Z",
            ...
            "MessageAttributes" : {
              "eventType" : {"Type":"String","Value":"EKS_RESTORE_OBJECT_SKIPPED"},
              "clusterName" : {"Type":"String","Value":"eks-restore-target-cluster-
name"},
              "parentRestoreJobId" : {"Type":"String","Value":"12345678-abcd-123a-def0-
abcd1a234567"},
              "reason" : {"Type":"String","Value":"Already exists."},
              "resourceName" : {"Type":"String","Value":"exampleresource"},
              "resourceType" : {"Type":"String","Value":"apiextensions.k8s.io/v1/
customresourcedefinitions"},
              "restoreJobId" : {"Type":"String","Value":"1b2345b2-f22c-4dab-5eb6-
bbc7890ed123"}
            }
        }
    }]
}
```

Example: EKS restore modification succeeded

```
{
    "Records": [{
        "EventSource": "aws:sns",
        "EventVersion": "1.0",
        "EventSubscriptionArn": "arn:aws:sns:...-a3802aa1ed45",
        "Sns": {
            "Type": "Notification",
            "MessageId": "12345678-abcd-123a-def0-abcd1a234567",
            "TopicArn": "arn:aws:sns:us-west-1:123456789012:backup-2sqs-sns-topic",
            "Subject": "Notification from AWS Backup",
            "Message": "A Kubernetes resource was successfully modified by your
 resource modifiers during restore from your Amazon EKS Backup. Resource: apps/v1/
deployments. Resource Name: exampledeployment. Namespace: production. Destination EKS
```

```
 Cluster Name: eks-restore-target-cluster-name. RestoreJob ID: 1b2345b2-f22c-4dab-5eb6-
bbc7890ed123",
            "Timestamp": "2025-05-25T18:46:02.788Z",
            ...
            "MessageAttributes" : {
              "eventType" :
 {"Type":"String","Value":"EKS_RESTORE_MODIFICATION_SUCCESS"},
              "clusterName" : {"Type":"String","Value":"eks-restore-target-cluster-
name"},
              "namespace" : {"Type":"String","Value":"production"},
              "parentRestoreJobId" : {"Type":"String","Value":"12345678-abcd-123a-def0-
abcd1a234567"},
              "resourceName" : {"Type":"String","Value":"exampledeployment"},
              "resourceType" : {"Type":"String","Value":"apps/v1/deployments"},
              "restoreJobId" : {"Type":"String","Value":"1b2345b2-f22c-4dab-5eb6-
bbc7890ed123"}
            }
        }
    }]
}
```

Example: EKS restore modification failed

```
{
    "Records": [{
        "EventSource": "aws:sns",
        "EventVersion": "1.0",
        "EventSubscriptionArn": "arn:aws:sns:...-a3802aa1ed45",
        "Sns": {
            "Type": "Notification",
            "MessageId": "12345678-abcd-123a-def0-abcd1a234567",
            "TopicArn": "arn:aws:sns:us-west-1:123456789012:backup-2sqs-sns-topic",
            "Subject": "Notification from AWS Backup",
            "Message": "One of your resource modifiers could not be applied to
 a Kubernetes resource during restore from your Amazon EKS Backup. The failed
 modification did not take effect; other requested modifications may have been
 applied. Resource: apps/v1/deployments. Resource Name: exampledeployment. Namespace:
 production. Destination EKS Cluster Name: eks-restore-target-cluster-name. RestoreJob
 ID: 1b2345b2-f22c-4dab-5eb6-bbc7890ed123",
            "Timestamp": "2025-05-25T18:46:02.788Z",
            ...
            "MessageAttributes" : {
```

```
            "eventType" :
 {"Type":"String","Value":"EKS_RESTORE_MODIFICATION_FAILED"},
            "clusterName" : {"Type":"String","Value":"eks-restore-target-cluster-
name"},
            "namespace" : {"Type":"String","Value":"production"},
            "parentRestoreJobId" : {"Type":"String","Value":"12345678-abcd-123a-def0-
abcd1a234567"},
            "reason" : {"Type":"String","Value":"Resource modifier operation replace
 on path /spec/replicas could not be applied."},
            "resourceName" : {"Type":"String","Value":"exampledeployment"},
            "resourceType" : {"Type":"String","Value":"apps/v1/deployments"},
            "restoreJobId" : {"Type":"String","Value":"1b2345b2-f22c-4dab-5eb6-
bbc7890ed123"}
          }
        }
    }]
}
```

Example: Access point available

```
{
    "Type": "Notification",
    "MessageId": "dd7579e6-47bb-51af-866c-1b36e07d9aaa",
    "TopicArn": "arn:aws:sns:us-east-1:123456789012:my-backup-topic",
    "Subject": "Notification from AWS Backup",
    "Message": "An AWS Backup access point was created successfully. Access point ARN:
 arn:aws:backup:us-east-1:123456789012:accesspoint/my-access-point. S3 access point
 ARN: arn:aws:s3:us-east-1:123456789012:accesspoint/my-access-point. Recovery point
 ARN: arn:aws:backup:us-east-1:123456789012:recovery-point:rp-1234567890abcdef0.",
    "Timestamp": "2026-07-28T12:00:26.426Z",
    "MessageAttributes": {
        "EventType": {"Type":"String","Value":"ACCESS_POINT_AVAILABLE"},
        "Status": {"Type":"String","Value":"AVAILABLE"},
        "AccountId": {"Type":"String","Value":"123456789012"},
        "AccessPointArn": {"Type":"String","Value":"arn:aws:backup:us-
east-1:123456789012:accesspoint/my-access-point"},
        "CreationTime": {"Type":"String","Value":"2026-07-28T12:00:25.938846787Z"}
    }
}
```

Example: Access point creation failed

```
{
```

```
    "Type": "Notification",
    "MessageId": "abcdef01-2345-6789-abcd-ef0123456789",
    "TopicArn": "arn:aws:sns:us-east-1:123456789012:my-backup-topic",
    "Subject": "Notification from AWS Backup",
    "Message": "An AWS Backup access point failed to create. Access point ARN:
 arn:aws:backup:us-east-1:123456789012:accesspoint/my-access-point. Recovery point ARN:
 arn:aws:backup:us-east-1:123456789012:recovery-point:rp-1234567890abcdef0.",
    "Timestamp": "2026-07-28T12:05:10.123Z",
    "MessageAttributes": {
        "EventType": {"Type":"String","Value":"ACCESS_POINT_CREATION_FAILED"},
        "Status": {"Type":"String","Value":"FAILED"},
        "AccountId": {"Type":"String","Value":"123456789012"},
        "AccessPointArn": {"Type":"String","Value":"arn:aws:backup:us-
east-1:123456789012:accesspoint/my-access-point"},
        "CreationTime": {"Type":"String","Value":"2026-07-28T12:04:55.123456789Z"}
    }
}
```

Example: Access point expired

```
{
    "Type": "Notification",
    "MessageId": "fedcba98-7654-3210-fedc-ba9876543210",
    "TopicArn": "arn:aws:sns:us-east-1:123456789012:my-backup-topic",
    "Subject": "Notification from AWS Backup",
    "Message": "An AWS Backup access point expired. Access point ARN:
 arn:aws:backup:us-east-1:123456789012:accesspoint/my-access-point. Recovery point ARN:
 arn:aws:backup:us-east-1:123456789012:recovery-point:rp-1234567890abcdef0.",
    "Timestamp": "2026-07-28T18:30:05.789Z",
    "MessageAttributes": {
        "EventType": {"Type":"String","Value":"ACCESS_POINT_EXPIRED"},
        "Status": {"Type":"String","Value":"EXPIRED"},
        "AccountId": {"Type":"String","Value":"123456789012"},
        "AccessPointArn": {"Type":"String","Value":"arn:aws:backup:us-
east-1:123456789012:accesspoint/my-access-point"},
        "CreationTime": {"Type":"String","Value":"2026-07-28T12:00:25.938846787Z"}
    }
}
```

AWS Backup notification command examples

You can use AWS CLI commands to subscribe to, list, and delete Amazon SNS notifications for your AWS Backup events.

Example put backup vault notification

The following command subscribes to an Amazon SNS topic for the specified backup vault that notifies you when a restore job is started or completed, or when a recovery point is modified.

```
aws backup put-backup-vault-notifications
    --backup-vault-name myBackupVault
    --sns-topic-arn arn:aws:sns:region:account-id:myBackupTopic
    --backup-vault-events RESTORE_JOB_STARTED RESTORE_JOB_COMPLETED
 RECOVERY_POINT_MODIFIED
```

Example get backup vault notification

The following command lists all events currently subscribed to an Amazon SNS topic for the specified backup vault.

```
aws backup get-backup-vault-notifications
    --backup-vault-name myVault
```

The sample output is as follows:

```
{
    "SNSTopicArn": "arn:aws:sns:region:account-id:myBackupTopic",
    "BackupVaultEvents": [
        "RESTORE_JOB_STARTED",
        "RESTORE_JOB_COMPLETED",
        "RECOVERY_POINT_MODIFIED"
    ],
    "BackupVaultName": "myVault",
    "BackupVaultArn": "arn:aws:backup:region:account-id:backup-vault:myVault"
}
```

Example delete backup vault notification

The following command unsubscribes from an Amazon SNS topic for the specified backup vault.

```
aws backup delete-backup-vault-notifications
    --backup-vault-name myVault
```

Specifying AWS Backup as a service principal

> **Note**
>
> To allow AWS Backup to publish SNS topics on your behalf, you must specify AWS Backup as a service principal.

Include the following JSON in the access policy of the Amazon SNS topic that you use to track AWS Backup events. You must specify the resource Amazon Resource Name (ARN) of your topic.

```
{
      "Sid": "My-statement-id",
      "Effect": "Allow",
      "Principal": {
        "Service": "backup.amazonaws.com"
      },
      "Action": "SNS:Publish",
      "Resource": "arn:aws:sns:region:account-id:myTopic"
}
```

For more information about specifying a service principal in an Amazon SNS access policy, see Allowing Any AWS Resource to Publish to a Topic in the *Amazon Simple Notification Service Developer Guide*.

> **Note**
>
> If your topic is encrypted, you must include additional permissions in your policy to allow AWS Backup to publish to it. For more information about enabling services to publish to encrypted topics, see Enable Compatibility between Event Sources from AWS Services and Encrypted Topics in the *Amazon Simple Notification Service Developer Guide*.

Troubleshooting AWS Backup

When you use AWS Backup, you might encounter issues. The following sections can help you troubleshoot some common issues that might occur.

For general questions about AWS Backup, see the AWS Backup FAQ. You can also search for answers and post questions in AWS re:Post.

Topics

- Troubleshooting general issues
- Troubleshoot creating resources
- Troubleshooting deleting resources
- Troubleshooting restoring resources
- Troubleshooting formatting errors
- Common troubleshooting questions
- Additional important links

Troubleshooting general issues

When you back up and restore resources, you must have permission to use AWS Backup and permission to access the resources that you want to protect. The easiest way to have the proper permissions is to choose the **Default role** when you assign resources to a backup plan. For more information about access control using AWS Identity and Access Management (IAM) with AWS Backup, see Access control.

If you get an `AccessDenied` error when attempting to access a AWS Backup resource, such as a backup vault, either the resource does not exist or you do not have permissions to access the resource.

If you run into issues with backing up and restoring a particular resource type, it can be helpful to review the backup and restore troubleshooting topic for that resource. For more information, see the links under How AWS Backup works with supported AWS services.

If AWS Backup fails to create or delete a resource, you can learn more about the issue by using AWS CloudTrail to view error messages or logs. For more information about using CloudTrail with AWS Backup, see Logging AWS Backup API calls with CloudTrail.

If you have a role with `tag:GetResources` permissions, but you also have a service control policy (SCP) that blocks `tag:GetResources`, then AWS Backup cannot successfully run a backup plan with tag-based selections. In these situations, the CloudTrail event `ProcessBackupPlanSelection` shows a `tag:GetResources` permissions error.

Troubleshoot creating resources

The following information can help you troubleshoot problems with creating backups.

- In general, **AWS database** services cannot start backups 1 hour before or during their maintenance window or automatic backup window. **Amazon FSx** cannot start backups 3 hours before or during the maintenance window or automatic backup window (Amazon Aurora is exempt from both the maintenance window restriction and the automatic backup window restriction). Snapshot backups scheduled during those times will fail. One exception: when you opt in to using AWS Backup for both snapshot and continuous backups for a supported service, you no longer need to worry about those windows because AWS Backup will schedule them for you. See Point-in-Time Recovery for a list of supported services and instructions on how to use AWS Backup to take continuous backups.
- Creating backups for **DynamoDB tables** will fail while tables are being created. Creating a DynamoDB table typically takes a couple of minutes.
- Backing up **Amazon EFS file systems** can take up to 7 days when the file systems are very large. Only one concurrent backup at a time can be queued for an Amazon EFS file system. If a subsequent backup is queued while a previous one is still in progress, the backup window can expire and no backup is created.
- **Amazon EBS** has a soft quota of 100,000 backups per AWS Region per account, and additional backups fail when this quota is reached. If you reach this quota, you can delete excess backups or request a quota increase. For more information about requesting a quota increase, see AWS Service Quotas.
- When creating **Amazon Relational Database Service (RDS) backups**, consider the following:
  - If you do not use AWS Backup to manage both Amazon RDS snapshots and continuous backups with point-in-time recovery, your backups will fail if initiated if scheduled or made on-demand during the daily, user-configurable 30-minute backup window. For more information about automated Amazon RDS backups, see Working With Backups in the *Amazon RDS User Guide*. You can avoid this limitation by using AWS Backup to manage both Amazon RDS snapshots and continuous backups with point-in-time recovery.

- If you initiate a backup job from the Amazon RDS console, this can conflict with an Aurora clusters backup job, causing the error `Backup job expired before completion.` If this occurs, configure a longer backup window in AWS Backup.
- AWS Backup does not currently pass on the TDE option group when a copy job is created. If you intend to use this option group for copy job creation, you must use the Amazon RDS console or Amazon RDS API instead of AWS Backup tools. See Copying an option group in the *Amazon Relational Database Service User Guide* for more information.
- **ERROR:** On-demand backups complete but scheduled backups fail with error "The source snapshot KMS key does not exist, is not enabled or you do not have permissions to access it." The on-demand job is completed because it uses the API call `CopyDBSnapshot`, which doesn't require KMS access.

  REMEDY: Add your IAM role to your KMS key.
- For resources that support full AWS Backup management with recovery points in the format `arn:aws:backup:region:account-id:recovery-point:*` and all continuous backups, ensure your IAM role has permission to perform `backup:TagResource` if your source resources contain tags or you want to add additional tags to your recovery points. Apply the `backup:TagResource` permission to `"Resource": "arn:aws:backup:*:*:recovery-point:*"`.

Troubleshooting deleting resources

Recovery points that are created by AWS Backup cannot be deleted in the console window of the protected resource. You can delete them on the AWS Backup console by selecting them in the vault where they are stored and then choosing **Delete**.

To delete a recovery point or a backup vault, you need the appropriate permissions. For more information about access control using IAM with AWS Backup, see Access control.

Troubleshooting restoring resources

Restoring using API

To restore a backup programmatically, use the StartRestoreJob API operation.

To get the configuration metadata that your backup was created with, you can call GetRecoveryPointRestoreMetadata.

See Restoring a backup for more information.

Restoring using the Console

- Restoring Amazon S3 data
- Restoring a virtual machine
- Restoring an Amazon FSx file system
- Restoring an Amazon EBS volume
- Restoring an Amazon EFS file system
- Restoring an Amazon DynamoDB table
- Restoring an Amazon RDS database
- Restoring an Aurora cluster
- Restoring an Amazon EC2 instance
- Restoring a Storage Gateway volume
- Restoring a Amazon DocumentDB cluster
- Restoring a Neptune cluster

Troubleshooting formatting errors

When a wildcard (*) is included for the value in a parameter, the wildcard is processed to include values other than whitespaces. Values in a key-value pair that contain white spaces will not included as part of the wildcard.

Common troubleshooting questions

The following resources can help you troubleshoot common issues with AWS Backup.

- Why does my AWS Backup job fail?
- How do I troubleshoot common errors for Amazon S3 backups that are failing in AWS Backup?
- How do I troubleshoot common errors for VMware backups that are failing in AWS Backup?
- For additional VMware backup troubleshooting, see the section called “Troubleshoot VM issues”.
- How do I troubleshoot the "You are not authorized to perform this operation" error when I try to restore my Amazon EC2 instance?

- How do I troubleshoot a backup policy that doesn't create jobs in my member accounts in an organization?
- How do I resolve the error "This image is managed by AWS Backup and cannot be deleted via Amazon EC2 APIs" when I try to delete an Amazon EC2 backup?
- How do I troubleshoot Amazon EC2 VSS failures in AWS Backup?
- How can I troubleshoot not receiving notifications from Amazon EventBridge for AWS Backup?
- How do I resolve "Insufficient privileges to perform this action" errors when performing an Amazon EFS restore using AWS Backup?
- How can I turn off automatic backups in Amazon EFS and remove the stored backup data?
- Why is my recovery point in the expired status in AWS Backup?
- How can I stop an Amazon RDS continuous backup in AWS Backup?
- How do I remove an AWS Backup Vault Lock?
- Why is my cross-account copy failing in AWS Backup?
- How can I resolve the "Given key ID not accessible" error when performing a cross-account copy in AWS Backup?
- How do I resolve the error "Access Denied trying to call AWS Backup service" when I try to create a cross-account copy in AWS Backup?
- Why do I get an Access Denied error when I try to create an AWS Backup vault?
- Why is my Aurora cluster endpoint stuck in the Creating status after the AWS Backup restore job completes?
- Why are my scheduled backup plans in AWS Backup not running?

Additional important links

The following resources provide additional guidance for working with AWS Backup.

- Troubleshoot a logically air-gapped vault issue
- Delegated administrator accounts — Delegated administrator accounts are member accounts with enhanced features but cannot override service opt-in settings of other member accounts like a management account can.
- Managing multiple accounts — For backup plans that are managed by Organizations, the resource opt-in settings in the management account override the settings in a member account, even if one or more delegated administrator accounts are configured.

- Backup plan options and configuration — If you have a backup plan with multiple rules and the time frames of the two rules overlap, AWS Backup optimizes the backup and takes a backup for the rule with the longer retention time.
- Metering and billing — To avoid additional charges, we recommend that you configure your backup plan retention policy with a warm storage duration of at least one week.
- AWS Backup pricing
- AWS Backup SLA
- AWS Backup quotas
- Continuous and point-in-time recovery considerations — A resource can only have one continuous backup.
- Prerequisites for Amazon S3 backups and considerations for Amazon S3 backups
- Best practices and cost considerations for Amazon S3 backups
- Feature availability, supported resources, and AWS Regions

AWS Backup API

In addition to using the console, you can use the AWS Backup API actions and data types to programmatically configure and manage AWS Backup and its resources. For a complete list of actions, data types, and errors, see the AWS Backup API Reference.

FAQs on supported protocols

The AWS Backup Gateway 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 AWS Backup

- **API version:** November 19, 2025
- **Latest documentation update:** November 12, 2026

The following table lists all AWS Backup launches since the launch of the service in January 2019 to present. For notifications about updates to this documentation you can subscribe to the RSS feed above.

Change	Description	Date
Logically air-gapped vault support for Amazon EBS cold storage	AWS Backup now supports transitioning EBS backups from warm to cold storage in logically air-gapped vaults. For more information, see Logically air-gapped vault, Amazon EBS Archive Tier for cold storage, and Lifecycle and storage tiers.	November 12, 2026
Cross-account backup copy and cross-account management in the AWS European Sovereign Cloud (Germany) Region	AWS Backup now supports cross-account backup copy and cross-account management in the AWS European Sovereign Cloud (Germany) Region. See Feature availability by AWS Region.	October 29, 2026
Restore testing support for Amazon Redshift	You can now include Amazon Redshift provisioned clusters in a restore testing plan. For more information, see Restore testing and Restore testing inferred metadata.	October 28, 2026

Change	Description	Date
Amazon FSx support in Asia Pacific (New Zealand)	AWS Backup now supports Amazon FSx in Asia Pacific (New Zealand). For more information, see Supported services by AWS Region.	October 13, 2026
Logically air-gapped vault support for FSx for ONTAP	Recovery points (backups) of FSx for ONTAP resources can now be copied to logically air-gapped vaults. For more information, see Logically air-gapped vault and Feature availability by resource.	September 28, 2026
Aurora DSQL multi-Region backup and restore Regional expansion	Aurora DSQL multi-Region backup and restore support is now available in Asia Pacific (Hong Kong). For more information, see Amazon Aurora DSQL backups.	September 24, 2026
Expanded Amazon S3 bucket support	AWS Backup now supports protecting more than 1,000 Amazon S3 buckets per account by integrating with Amazon S3 Event Notificat ions with system-generated tags. See Amazon S3 backups for more details.	September 1, 2026

Change	Description	Date
AWS Backup adds cross-Region and cross-account copy support for FSx for ONTAP	AWS Backup now supports cross-account and cross-Region copy of FSx for ONTAP backups in most AWS Regions where FSx for ONTAP is supported. For more information, see Creating backup copies across AWS Regions.	August 27, 2026
Cross-account backup copy and cross-account management in the China Regions	AWS Backup now supports cross-account backup copy and cross-account management in China (Beijing) and China (Ningxia). See Feature availability by AWS Region.	August 27, 2026
Amazon DocumentDB cross-account and cross-Region copy in opt-in Regions	Cross-account and cross-Region backup copy for Amazon DocumentDB is now supported in AWS Regions where opt-in is required. For more information, see Feature availability by AWS Region.	August 26, 2026

Change	Description	Date
Amazon Neptune backup Regional expansion	Amazon Neptune backup and restore support is now available in Asia Pacific (Malaysia) and South America (São Paulo). Corrected the matrix to show that Asia Pacific (New Zealand) is not yet supported. For more information, see Supported services by AWS Region.	August 7, 2026
AWS Backup extends logically air-gapped vault support for Amazon Neptune	Logically air-gapped vault support for Amazon Neptune is now available in Asia Pacific (Melbourne), Europe (Spain), and Europe (Zurich). To learn more, see Feature availability by Region.	August 6, 2026
Access points for S3 backups	AWS Backup now lets you create read-only access points for Amazon S3 recovery points, enabling you to access backup data using S3 APIs without initiating a restore. See Access points for more details.	August 6, 2026

Change	Description	Date
Aurora DSQL backup and restore Regional expansion	Aurora DSQL multi-Region backup and restore support is now available in Europe (Stockholm), Europe (Spain), Asia Pacific (Mumbai), and Asia Pacific (Singapore). Aurora DSQL single-Region backup and restore support is now available in Asia Pacific (Sydney) and Asia Pacific (Melbourne). For more information, see Amazon Aurora DSQL backups.	July 30, 2026
Restore testing Regional expansion	Restore testing is now available in Asia Pacific (Taipei), Asia Pacific (New Zealand), Asia Pacific (Malaysia), Asia Pacific (Thailand), Canada West (Calgary), Israel (Tel Aviv), and Mexico (Central). For more information, see Feature availability by Region.	July 14, 2026
Aurora DSQL multi-Region backup and restore Regional expansion	Aurora DSQL multi-Region backup and restore support is now available in Canada West (Calgary) and Canada (Central). For more informati on, see Amazon Aurora DSQL backups.	July 10, 2026

Change	Description	Date
Aurora DSQL single-Region backup and restore Regional expansion	Aurora DSQL single-Region backup and restore support is now available in Europe (Spain), Canada (Central), and Canada West (Calgary) . For more information, see Supported services by Region.	July 10, 2026
Aurora DSQL single-Region backup and restore Regional expansion	Aurora DSQL single-Region backup and restore support is now available in Asia Pacific (Hong Kong), Asia Pacific (Mumbai), Asia Pacific (Singapore), Europe (Stockholm), and South America (São Paulo). For more information, see Supported services by Region.	May 4, 2026
Added documentation for tagging existing recovery points	Added a new section describing how to add, modify, and remove tags on existing recovery points using the console and AWS CLI. For more information, see Add or remove tags on existing recovery points.	April 30, 2026
Amazon EKS support for logically air-gapped vaults	AWS Backup now supports backup, copy, and restore functionality for Amazon EKS clusters to logically air-gapped vaults.	March 12, 2026

Change	Description	Date
AWS Backup updates Amazon RDS PITR workflow	AWS Backup uses snapshot-based backups for Amazon RDS instances, when PITR is enabled for the first time, until the next maintenan ce window to avoid an immediate instance restart. PITR is automatically enabled after maintenance completes . When PITR is disabled, the change is scheduled for your maintenance window. You may continue to incur backup storage costs until the maintenance window applies the change. For more information, see Point-in-time recovery and Stopping point-in-time recovery.	February 26, 2026
Cross-Region snapshot copying support for logically air-gapped vault	Amazon Aurora, Amazon DocumentDB, and Amazon Neptune now support cross-Region snapshot copying to logically air-gapped vaults.	February 12, 2026
AWS Backup Regional expansion	AWS Backup, along with many of its features, is now available in AWS European Sovereign Cloud (Germany) . For more information, see Feature availability by Region.	January 15, 2026

Change	Description	Date
Amazon GuardDuty Malware Protection for AWS Backup	Amazon GuardDuty Malware Protection now supports AWS Backup, extending malware detection capabilities to EC2, EBS, and S3 backups. For more information, see the product details page.	November 19, 2025
Added a new AWS managed policy	AWS Backup added the `AWSBackupGuardDutyRolePolicyForScans` AWS managed policy. For more information, see Managed policies for AWS Backup.	November 19, 2025
Added a new AWS managed policy	AWS Backup added the `AWSBackupServiceRolePolicyForScans` AWS managed policy. For more information, see Managed policies for AWS Backup.	November 19, 2025

Change	Description	Date
Managed policy updates	AWS Backup added additional permissions to the following AWS managed policies: • `AWSBackupFullAccess` • `AWSBackupOperatorAccess` For more information, see Managed policy updates for AWS Backup.	November 19, 2025
Backup tiering	AWS Backup introduces a low-cost warm storage tier for Amazon S3 backup data. For more information, see Backup tiering.	November 18, 2025

Change	Description	Date
Primary backups to logically air-gapped vaults	Logically air-gapped vaults are now supported as a primary backup target. This enhancement allows customers to leverage the security and recovery benefits of logically air-gapped vaults while maintaining only a single backup copy, reducing storage and operational costs. Previously, logically air-gapped vaults could only store copies of existing backups. For more information, see Primary backups to logically air-gapped vaults.	November 17, 2025
AWS Backup support for Amazon Elastic Kubernetes Service clusters	AWS Backup now supports Amazon EKS clusters. Customers can create backups and restore recovery points for Amazon EKS clusters and their associated resources. For more information, see Feature availability by resource and Managed policy updates for AWS Backup.	November 10, 2025

Change	Description	Date
Managed policy updates	AWS Backup added additional permissions to the following AWS managed policies: • `AWSBackupFullAccess` • `AWSBackupOperatorAccess` • `AWSBackupServiceRolePolicyForBackup` • `AWSBackupServiceLinkedRolePolicyForBackup` • `AWSBackupServiceRolePolicyForRestores` For more information, see Managed policy updates for AWS Backup.	November 10, 2025
Logically air-gapped vault support for AWS KMS customer-managed keys	Logically air-gapped vaults now support encryption with AWS KMS customer-managed keys. For more information, see Logically air-gapped vault.	November 6, 2025

Change	Description	Date
Cross-Region and cross-account snapshot copying support	Amazon RDS, Aurora, DocumentDB, and Neptune now support cross-Region and cross-account snapshot copying in a single action. For more information, see Feature availability by resource.	October 23, 2025
AWS Backup Regional support expansion for Aurora DSQL	Aurora DSQL backup and restore support is now available in Europe (Frankfurt). For more information, see Supported services by Region and Aurora DSQL backups.	October 22, 2025
AWS Backup Regional expansion	AWS Backup, along with many of its features, is now available in Asia Pacific (New Zealand). For more information, see Feature availability by Region.	September 29, 2025
Amazon S3 optional ACLs and object tags backup support	AWS Backup now supports optional exclusion of Amazon S3 ACLs and object tags, allowing customers to help reduce backup costs. For more information, see Amazon S3 backups.	September 9, 2025

Change	Description	Date
Managed policy updates	AWS Backup added additional permissions to the following AWS managed policies: • AWSBackupServiceLinkedRolePolicyForBackup For more information, see Managed policy updates for AWS Backup.	September 9, 2025
AWS Backup Audit Manager Regional expansion	AWS Backup Audit Manager cross-Region and cross-account reports are now available in additional AWS Regions, including Asia Pacific (Hyderabad), Asia Pacific (Jakarta), Asia Pacific (Melbourne), Europe (Spain), Europe (Zurich), and Middle East (UAE). See Feature availability by AWS Region and Creating report plans for more information.	September 5, 2025

Change	Description	Date
Increased AWS Backup support for Amazon Aurora DSQL multi-Region restore	AWS Backup users are now able to restore Aurora DSQL multi-Region clusters through the AWS Backup console or through AWS Backup CLI. See Restore an Aurora DSQL multi-Region cluster for more information.	July 29, 2025
Managed policy updates	AWS Backup added additional permissions to the following AWS managed policies: • AWSBackupServiceRolePolicyForRestores For more information, see Managed policy updates for AWS Backup.	July 29, 2025
AWS Backup support for Amazon Aurora DSQL multi-Region restore in additional AWS Regions	AWS Backup now supports Aurora DSQL multi-Region restore in additional AWS Regions. For more information, see Restore an Aurora DSQL multi-Region cluster and Supported services by Region.	July 3, 2025

Change	Description	Date
AWS Backup feature Regional expansion	AWS Backup adds support for cross-account and cross-Region copy of Amazon S3 backups within AWS GovCloud (US-East) and AWS GovCloud (US-West) AWS Regions. For more information, see Amazon S3 backups and Backup and tag copy.	June 26, 2025
AWS Backup support for Amazon Aurora DSQL in Asia Pacific (Seoul)	AWS Backup now supports Aurora DSQL in Asia Pacific (Seoul). For more information, see Supported services by Region.	June 26, 2025
AWS Backup support for Multi-party approval	AWS Backup now supports optional integration for logically air-gapped vaults to use Multi-party approval. This capability is currently available for accounts in US East (N. Virginia). For more information, see: • Multi-party approval for logically air-gapped vaults Multi-party approval User Guide	June 17, 2025

Change	Description	Date
AWS managed policy update	AWS Backup added new permissions to AWS managed policy `AWSBackupFullAccess`. For more information and specific permission changes, see policy updates.	June 17, 2025
AWS Backup Audit Manager Regional expansion	AWS Backup Audit Manager is now available in additional AWS Regions, including: • Asia Pacific (Hyderabad) • Asia Pacific (Jakarta) • Asia Pacific (Melbourne) • Europe (Spain) • Europe (Zurich) • Middle East (UAE) For additional information, see Feature availability by Region and Backup Audit Manager.	June 16, 2025
AWS Backup Regional expansion	AWS Backup, along with many of its features, is now available in Asia Pacific (Taipei) AWS Region. For more information, see Feature availability by Region and Supported services by Region.	June 6, 2025

Change	Description	Date
Backup index notification and event support	AWS Backup launches support for Amazon Simple Notification Service notifications and Amazon EventBridge events for backup indexes. AWS Backup support for notification backup indexes is available in all AWS commercial and GovCloud Regions where backup indexes, SNS, and EventBridge are supported. For more information, see the User Notifications console, the EventBridge console, and backup search and index.	June 3, 2025
Additional features for Amazon EC2 restore	Volume configurations can be customized during a restore job. You can specify custom settings for attached Amazon EBS volumes to an Amazon EC2 AMI, including the restore metadata for block device mappings. For additional information, see Restore Amazon EC2 with AWS CLI and the *Amazon Elastic Compute Cloud User Guide*.	May 28, 2025

Change	Description	Date
AWS Backup support for Amazon Aurora DSQL clusters	AWS Backup now offers integration with Aurora DSQL in several AWS Regions. Customers can create backups and restore recovery points for Aurora DSQL single-Region and multi-Region clusters. For more information, see Amazon Aurora DSQL restore, Feature availability by resource, and Manage clusters in Aurora DSQL with the AWS CLI.	May 27, 2025
Managed policy updates	AWS Backup added additional permissions to the following AWS managed policies: `AWSBackupFullAccess`, `AWSBackupOperatorAccess`, `AWSBackupServiceLinkedRolePolicyForBackup`, `AWSBackupServiceRolePolicyForBackup`, and `AWSBackupServiceRolePolicyForRestores`. For more information, see Managed policy updates for AWS Backup.	May 27, 2025

Change	Description	Date
AWS Backup feature Regional expansion	AWS Backup search is now available in Asia Pacific (Thailand), Mexico (Central) , AWS GovCloud (US-East), and AWS GovCloud (US-West) AWS Regions. For more information, see Feature availability by AWS Region and Backup search.	April 22, 2025
AWS Backup support for Amazon Redshift Serverless namespaces	Customers now have the capability to use AWS Backup to create Redshift Serverles s manual snapshot backups and to restore to Redshift Serverless namespaces. For more information, see Amazon Redshift Serverless backups.	March 31, 2025

Change	Description	Date
Managed policy updates	AWS Backup added additional permissions to the following AWS managed policies: • AWSBackupFullAccess • AWSBackupOperatorAccess • AWSBackupServiceLinkedRolePolicyForBackup • AWSBackupServiceRolePolicyForBackup • AWSBackupServiceRolePolicyForRestores For more information, see Managed policy updates for AWS Backup.	March 31, 2025

Change	Description	Date
AWS Backup supported service Regional expansion	AWS Backup support for FSx for OpenZFS is now available in the following AWS Regions: • US West (N. California) • Africa (Cape Town) • Asia Pacific (Hyderabad) • Asia Pacific (Jakarta) • Asia Pacific (Osaka) • Europe (Milan) • Europe (Paris) • Europe (Spain) • Europe (Zurich) • Israel (Tel Aviv) • Middle East (Bahrain) • Middle East (UAE) • South America (São Paulo) For more information, see Supported services by AWS Regions.	March 12, 2025

Change	Description	Date
Logically air-gapped vault resource expansion	Recovery points (backups) of FSx for Lustre, FSx for Windows File Server, and FSx for OpenZFS resources can now be copied into a logically air-gapped vault. For more information, see: • Logically air-gapped vault • Feature availability by resource	March 12, 2025
New AWS managed policy	AWS Backup added a new AWS managed policy: AWSBackupSearchOperatorAccess. For more information, see AWSBackupSearchOperatorAccess in the *AWS Managed Policy Guide*.	February 27, 2025
AWS Backup Regional expansion	AWS Backup, along with many of its features, is now available in Mexico (Central) AWS Region. For more information, see Feature availability by Region.	January 21, 2025

Change	Description	Date
AWS Backup Regional expansion	AWS Backup, along with many of its features, is now available in Asia Pacific (Thailand). For more information, see Feature availability by Region.	January 14, 2025
Updated AWS managed policies	AWS Backup added the permission `rds:AddTagsToResource` to managed policy `AWSBackupServiceLinkedRolePolicyForBackup`. AWS Backup added the permissions `rds:CreateTenantDatabase` and `rds:DeleteTenantDatabase` to policy `AWSBackupServiceRolePolicyForRestores`. For more information, see Policy updates.	January 8, 2025
Support for IPv6	AWS Backup now supports both Internet Protocols version 6 (IPv6) and IPv4. For more information on dual stack capabilities and endpoints, see AWS Backup network.	December 18, 2024

Change	Description	Date
AWS Backup supported service Regional expansion	AWS Backup support for Amazon Timestream is now available in AWS GovCloud (US-West). For more information, see Supported services by AWS Regions and Amazon Timestream backups.	December 18, 2024
Search feature added	AWS Backup introduces the option to search your backup metadata for item level properties to help you find recovery points and other files or objects you want to restore. See Backup search for more details.	December 17, 2024
Added a new AWS managed policy	AWS Backup added the `AWSBackupServiceRolePolicyForItemRestores` AWS managed policy. For more information, see Managed policies for AWS Backup.	December 17, 2024

Change	Description	Date
Added a new AWS managed policy	AWS Backup added the AWSBackupServiceRolePolicyForIndexing AWS managed policy. For more information, see Managed policies for AWS Backup.	December 17, 2024
AWS Backup feature Regional expansion	Cross-account management for monitoring backup, copy, and, restore jobs through AWS Backup and AWS Organizations integration is now available in AWS Regions where opt-in is required. For more information, see Feature availability by Region.	December 12, 2024
AWS Backup supported service Regional expansion	AWS Backup support for Amazon Timestream is now available in Asia Pacific (Mumbai). For more information, see Supported services by AWS Regions and Amazon Timestream backups.	November 22, 2024

Change	Description	Date
Additional features for Amazon S3 restore	AWS Backup adds additiona l flexibility when you restore S3 backups. You can now restore additional versions of objects, including restoring all versions. For more information, see Restore S3 data using AWS Backup.	November 21, 2024
AWS Backup feature Regional expansion	AWS Backup customers can now create cross-Region copies of Amazon Neptune backups in more Regions, including Africa (Cape Town), Asia Pacific (Jakarta), and Asia Pacific (Osaka). For more information, see Feature availability by Region and Creating backup copies across AWS Regions.	November 13. 2024
Cross-Region copy expansion	Cross-Region copy of Amazon S3 backups is now supported in AWS Regions where opt-in is required. For more information, see Feature availability by AWS Region.	November 13, 2024

Change	Description	Date
AWS Backup Audit Manager Regional expansion	AWS Backup Audit Manager is now available in AWS Region Asia Pacific (Osaka). For more information, see Feature availability by Region and Backup Audit Manager.	October 29, 2024
AWS Backup now available in Asia Pacific (Malaysia) Region	Backup and restore for many resource types are now available in AWS Region Asia Pacific (Malaysia). For compatible backup features, see Feature availability by AWS Region. For supported resource types, see Supported services by AWS Region.	October 10, 2024
AWS Backup Audit Manager control for logically air-gapped vault	AWS Backup Audit Manager now offers the control Resources in a logically air-gapped vault to assist with monitoring that resources have been copied to a logically air-gapped vault within the time frame specified. For more information, see Controls and remediation and Audit restore testing.	September 12, 2024

Change	Description	Date
AWS Backup feature Regional expansion	AWS Backup customers can now create cross-Region copies of Amazon Neptune backups in more Regions, including Asia Pacific (Hong Kong), Israel (Tel Aviv), Middle East (Bahrain), and Middle East (UAE). For more information, see feature availability by Region and creating backup copies across AWS Regions.	August 29. 2024
AWS Backup supports cross-Region, cross-account copy for SAP HANA databases	Users who have backed up SAP HANA databases on Amazon EC2 instances can now copy snapshots from full backups to other Regions and other accounts. For more information, see SAP HANA databases on Amazon EC2 instances.	August 20, 2024
AWS Backup now offers a new type of secure vault	AWS Backup has introduce d a secondary type of vault for storing copies of backups. Logically air-gapped vaults offer enhanced security features as well as the ability to share vault access with accounts through AWS RAM. For more information, see Logically air-gapped vault.	August 7, 2024

Change	Description	Date
AWS Backup feature Regional expansion	AWS Backup support for Amazon Simple Storage Service (Amazon S3) is now available in Canada West (Calgary). Customers can now backup and restore resources within the Canada West (Calgary) Region. For more informati on on AWS Backup Amazon S3 Region availability, see Supported Services by AWS Region.	June 27, 2024
Logically air-gapped vault preview closure	Logically air-gapped vault preview has ended.	June 27, 2024
AWS Backup feature Regional expansion	AWS Backup support of Amazon EBS snapshot archive tier is now available in the following Regions: • China (Beijing) • China (Ningxia) • AWS GovCloud (US-West) • AWS GovCloud (US-East)	June 3, 2024

Change	Description	Date
Updated AWS managed policies	AWS Backup added permission `backup:TagResource` to the following managed policies: • AWSBackupServiceRolePolicyForBackup • AWSBackupServiceRolePolicyForS3Backup • AWSBackupServiceLinkedRolePolicyForBackup For more information, see Policy updates.	May 17, 2024
AWS Backup now available in Canada West (Calgary) Region	Backup and restore for many resource types are now available in AWS Region Canada West (Calgary). For compatible backup features, see Feature availability by AWS Region. For supported resource types, see Supported services by AWS Region.	March 14, 2024

Change	Description	Date
Added permissions to managed policy	AWS Backup updated the policy AWSServiceRolePolicyForBackupRestoreTesting by adding permissions to support additional resource types within the restore testing feature. For more information on the specific permissions added, see Policy updates.	February 14, 2024
Backup and restore support for FSx for ONTAP FlexGroup volumes	AWS Backup now supports backup and restore of FSx for ONTAP FlexGroup volumes in most AWS Regions. For more information, see Restoring an Amazon FSx file system .	January 10, 2024
Support for SAP HANA HA backup and restore	AWS Backup now offers support SAP HANA High Availability databases on Amazon EC2 backup and restore. For more information, see SAP HANA on Amazon EC2 backups and Restoring an SAP HANA High Availability system	December 21, 2023

Change	Description	Date
AWS Backup Audit Manager control for restore testing	AWS Backup Audit Manager now offers the control Restore time for resources meet target to assist with monitoring restore times. This control checks if the restore time of a resource meets the target duration. For more information, see Controls and remediation and Audit restore testing.	December 18, 2023
Support for Amazon EBS cold storage	AWS Backup now supports transitioning EBS backups from warm to cold storage. For more information, see Amazon EBS Archive Tier for cold storage, Lifecycle and storage tiers, and Create a backup plan.	November 27, 2023
Introducing restore testing	AWS Backup introduces restore testing, which brings automated and periodic evaluation of restore viability , as well as the ability to monitor restore job duration times. For more information, see Restore testing.	November 27, 2023

Change	Description	Date
Updated AWS managed policies	AWS Backup added the permissions `ec2:DescribeSnapshotTierStatus` and `ec2:ModifySnapshotTier` to the managed policies `AWSBackupServiceRolePolicyForBackups` and `AWSBackupServiceLinkedRolePolicyForBackup`. AWS Backup also added the permissions `ec2:DescribeSnapshotTierStatus` and `ec2:RestoreSnapshotTier` to managed policy `AWSBackupServiceRolePolicyForRestores`. These permissions are necessary for users to have the option to transition Amazon EBS resources stored with AWS Backup to archive storage and to restore resources from the archive storage tier. For more information, see Policy updates.	November 27, 2023

Change	Description	Date
Added pass role permission to support restore testing.	AWS Backup added `restore-testing.backup.amazonaws.com` to `IamPassRolePermissions` and `IamCreateServiceLinkedRolePermissions`. This addition is necessary for AWS Backup to conduct restore tests on behalf of customers.	November 27, 2023
Added new service-linked role	AWS Backup has added the new service-linked role named `AWSServiceRoleForBackupRestoreTesting`, which provides backup permissions to conduct restore testing. This new service-linked role provides AWS Backup with permissions necessary to conduct restore testing. The permissions include the actions `list, read, and write` for the following services to be included in restore tests: Aurora, DocumentDB, DynamoDB, Amazon EBS, Amazon EC2, Amazon EFS, FSx for Lustre, FSx for Windows File Server, FSx for ONTAP, FSx for OpenZFS, Amazon Neptune, Amazon RDS, and Amazon S3.	November 27, 2023

Change	Description	Date
New job metric dashboard in the AWS Backup console	The AWS Backup console now displays a jobs dashboard, simplifying backup health monitoring at scale with a new visual user interface and aggregated backup, copy, and restore metrics for services supported by AWS Backup. The jobs dashboard is available in all Regions where AWS Backup Audit Manager is available. Regions not listed will still be able to access the CloudWatch dashboard. For more information, see AWS Backup console dashboards.	November 15, 2023
Support for nested stack backups	AWS Backup has expanded its support for backups of CloudFormation resources . Your CloudFormation application stacks that have nested stacks within them can be included in your backups. For more information, see CloudFormation stack backups.	November 8, 2023

Change	Description	Date
Support for Amazon S3 in China (Beijing) and China (Ningxia).	AWS Backup support for Amazon S3 is now available in China (Beijing) and China (Ningxia) Regions. For more information, see Feature availability by Region.	October 26, 2023
Support for Amazon Aurora continuous backups and Point-in-time restore	AWS Backup now supports continuous backups and point-in-time restore (PITR) for Aurora resources. For more information, see Continuous backups and Point-in-time recovery.	September 7, 2023
CloudFormation stacks support exclusion of resources	AWS Backup now supports the option to exclude chosen resources from your CloudFormation stack. For more information, see CloudFormation stack backups.	September 6, 2023
Backup plan rules introduce timezone flexibility	AWS Backup plan rules can now have a specified timezone for backup windows. For more information, see Managing backup plans.	August 28, 2023

Change	Description	Date
AWS Backup now available in Israel (Tel Aviv) Region	Many AWS Backup features are now available in the new Israel (Tel Aviv) Region. To see what resources are supported, visit Feature availability by AWS Region.	August 22, 2023
AWS Backup Audit Manager now supports delegated administrator accounts	AWS Backup Audit Manager report generation can now be accessed by delegated administrator accounts. For more information, see Audit backups and create reports with AWS Backup Audit Manager, Working with audit reports, and Delegated Administrator.	August 16, 2023
AWS Backup enhances Amazon S3 backups	AWS Backup has increased performance, size, and speed capabilities for S3 bucket backups. For more information, see Amazon S3 backups.	August 1, 2023
Tag on restore feature now available in China Regions	Tags that are part of a backup can now be copied when you create a restore job in China (Beijing) or China (Ningxia) Regions. For more information, see Copy tags during a restore.	July 17, 2023

Change	Description	Date
AWS Backup now supports Amazon S3 in additional Regions	AWS Backup support for Amazon S3 is now available in Europe (Spain), Europe (Zurich), Asia Pacific (Hyderabad), and Asia Pacific (Melbourne) Regions. For more information, see Feature availability by Region.	July 6, 2023
Cross-account copy expands to additional Regions	AWS Backup now supports cross-account backup copy of most resources in the following Regions: Asia Pacific (Jakarta), Middle East (Bahrain), Asia Pacific (Hong Kong), Africa (Cape Town), Europe (Milan), Asia Pacific (Osaka), Middle East (UAE) , Europe (Spain), Europe (Zurich), Asia Pacific (Hyderabad), and Asia Pacific (Melbourne). For more information, see Feature availability by Region	July 5, 2023
Backup Audit Manager available in GovCloud Regions	AWS Backup has expanded AWS Backup Audit Manager into AWS GovCloud (US-East) and AWS GovCloud (US-West) . For more information, see Feature availability by Region	June 29, 2023

Change	Description	Date
Cross-account management now available in GovCloud Regions	AWS Backup now supports cross-account management of resources in AWS GovCloud (US-East) and AWS GovCloud (US-West). For more information, see Managing AWS Backup resources across multiple AWS accounts.	June 29, 2023
Support for cross-Region copies of Amazon Aurora in additional Regions	AWS Backup now supports cross-Region backup copies for Aurora clusters into and from the following Regions: Asia Pacific (Jakarta), Middle East (Bahrain), Asia Pacific (Hong Kong), Africa (Cape Town), Europe (Milan), Middle East (UAE) , Europe (Spain), Europe (Zurich), Asia Pacific (Hyderabad), and Asia Pacific (Melbourne).	June 5, 2023
Copy tags when restoring	Tags that are part of a backup can now be copied when you create a restore job. For more information, see Copy tags during a restore.	May 22, 2023

Change	Description	Date
AWS Backup integrates with AWS User Notifications	You can now choose to receive notifications related to backup, copy, and restore events through the AWS User Notifications console. For more information, see Getting started with AWS User Notifications.	May 10, 2023
Cross-Region backups available in four new Regions	AWS Backup now supports cross-Region backup in Middle East (UAE) Region, Europe (Spain) Region, Europe (Zurich) Region, and Asia Pacific (Hyderabad) Region.	April 28, 2023
Expanded cross-Region AWS Backup copy support	Cross-Region backups of Amazon EFS, VMware, and DynamoDB resources can now be conducted within the following Regions: Asia Pacific (Jakarta), Middle East (Bahrain), Asia Pacific (Hong Kong), Africa (Cape Town), and Europe (Milan).	April 28, 2023
Amazon S3 backup and restore in South America (São Paulo) Region	AWS Backup support for Amazon S3 (Amazon Simple Storage Service) is now available in South America (São Paulo) Region. For more information, see Amazon S3 backups.	April 20, 2023

Change	Description	Date
AWS Backup expands to Asia Pacific (Melbourne) Region	AWS Backup is now available in Asia Pacific (Melbourne) Region. For more information, see Feature availability by AWS Region.	April 20, 2023
Expanded Regional support for Amazon S3	AWS Backup support for Amazon S3 (Amazon Simple Storage Service) is now available in AWS GovCloud (US-East) and AWS GovCloud (US-West) Regions For more information, see Amazon S3 backups.	April 19, 2023
Backup and restore SAP HANA databases on Amazon EC2 instances	AWS Backup now offers the ability to backup and restore SAP HANA databases running on Amazon EC2 instances in most Regions. For more information, see SAP HANA databases on Amazon EC2 instances backup.	April 17, 2023

Change	Description	Date
AWS Backup now available in Europe (Spain), Europe (Zurich), and Asia Pacific (Hyderabad) Regions	AWS Backup support has expanded to new Regions, including Europe (Spain), Europe (Zurich), and Asia Pacific (Hyderabad). Supported resources can be backed up and restored within these Regions. For more information, see Feature availability by AWS Region.	April 13, 2023
Updated AWS managed policy `AWSBackupAuditAccess`	Updated AWS managed policy AWSBackupAuditAccess. AWS Backup replaced the resource selection within the API `config:DescribeComplianceByConfigRule` with a wildcard resource. For more information see Policy updates for AWS Backup.	April 11, 2023
Hypervisors with Amazon CloudWatch Logs	AWS Backup gateway users can now integrate hypervisors with CloudWatch Logs to maintain logs. For more information, see Editing a hypervisor configuration and CloudWatch Logs.	March 29, 2023

Change	Description	Date
Expanded Regional support for Amazon S3	AWS Backup support for Amazon S3 is now available in Asia Pacific (Jakarta) and Middle East (UAE) Regions.	March 22, 2023
Virtual machine incremental backup improvement	VMware VM (virtual machine) backups that experience CBT (Changed Block Tracking) data issues now contain additional information to help remedy and troublesh oot. For more information, see Incremental VM backups and Troubleshoot your virtual machines.	March 15, 2023
AWS Backup support for multiple network adapters	AWS Backup gateway now supports configuring multiple network adapters For more information on configuring your network adapters, see Configure your gateway for multiple NICs in VMware in the *AWS Backup Developer Guide*.	March 8, 2023

Change	Description	Date
AWS Backup support for vSphere 8	AWS Backup now supports backup and restore of virtual machines which run on VMware vSphere 8. For more information on supported VMware options, see Supported VMs in the *AWS Backup Developer Guide*.	March 8, 2023
AWS Backup Audit Manager supports Amazon RDS Multi-AZ backups	Backup Audit Manager now offers support for Amazon Relational Database Service Multi-Availability Zone backups. For more information, see how to audit backups and create reports with AWS Backup Audit Manager.	February 1, 2023
AWS Backup offers increment al backup for Amazon Timestream tables	AWS Backup now offers expanded backup capabilit ies for Timestream backups. Backup plans can now take incremental backups to lower the time required to backup Timestream resources and lower storage costs. For more information, see Amazon Timestream backups.	January 23, 2023

Change	Description	Date
AWS Backup now available in Dubai	AWS Backup has expanded to the Middle East (UAE) Region. Supported resources can be backed up and restored within this Region.	January 17, 2023
Cross-Region copying available in additional Regions	AWS Backup now offers cross-Region backups in Asia Pacific (Jakarta) Region, Middle East (Bahrain) Region, Asia Pacific (Hong Kong) Region, Africa (Cape Town) Region, and Europe (Milan) Region for most resources. For more information, see Creating backup copies across AWS Regions.	December 21, 2022

Change	Description	Date
Backup Gateway Bandwidth Limits and Throttling	AWS Backup Gateway now allows limits on the upload throughput from gateways to AWS Backup to control the amount of network bandwidth used by the gateway. To support this feature, AWS Backup has created and updated managed policies, including `AWSBackupFullAccess` and `AWSBackupOperatorAccess`. For more information, see Backup Gateway bandwidth throttling.	December 15, 2022

Change	Description	Date
Backup Gateway VMware tag support	AWS Backup Gateway now supports VMware tags. Users have the additional flexibili ty to create AWS tags that match tags used for virtual machines. To support this feature, AWS Backup has created and updated managed policies, including `AWSBackup GatewayServiceRole PolicyForVirtualMa chineMetadataSync` , `AWSBackupFullAccess` , and `AWSBackupOperatorA ccess` . For more information, see VMware tags.	December 15, 2022
AWS Backup support for Amazon Timestream	AWS Backup now supports backing up and restoring Amazon Timestream tables. For more information, see Amazon Timestream backup.	December 13, 2022
AWS Backup offers Legal Hold	AWS Backup introduces a new tool to help protect recovery points through a legal hold. For more information see Legal hold.	November 27, 2022

Change	Description	Date
AWS Backup Audit Manager cross-Region and cross-acc ount reporting	AWS Backup Audit Manager brings additional functiona lity to compliance and job reports. Users can generate reports incorporating multiple Regions and multiple accounts. For more information, see Working with audit reports.	November 27, 2022
AWS Backup supports Amazon Redshift	AWS Backup now offers support to backup Amazon Redshift clusters and to restore Amazon Redshift clusters and tables. For more information, see Amazon Redshift backups.	November 27, 2022
AWS Backup offers support to backup AWS CloudFormation application stacks	AWS Backup provides the capability to backup CloudFormation and restore applications containing multiple resources by backing up a stack and restoring the resources within it. For more information, see Application stack backups.	November 27, 2022

Change	Description	Date
AWS Backup offers delegated administrator accounts and backup policy delegation	AWS Backup accounts enrolled in AWS Organizations can designate member accounts as delegated administrator accounts. For more information, see Managing multiple accounts with AWS Organizations.	November 27,2022
Public Preview of SAP HANA on Amazon EC2 Instances backup and restore	AWS Backup and AWS Backint are offering an integrated public preview of functionality to backup and restore SAP HANA databases on EC2 instances. For more information, see our Public Preview of SAP HANA on Amazon EC2 instances. To support this preview, AWS Backup has provided policy updates and new AWS Managed Policies for these features.	November 20, 2022

Change	Description	Date
Restore VMware to Amazon EC2 instances	AWS Backup now offers the ability to restore virtual machines to Amazon EC2 instances, in addition to the ability to restore machines to EBS, VMware, VMware Cloud on AWS, and VMware Cloud on AWS Outposts. For more information, see documentation on how to Use the AWS Backup console to restore virtual machine recovery points.	November 9, 2022
Expanded AWS Backup Vault Lock functionality	AWS Backup Vault Lock can be now created in governanc e mode for additional IAM protections or in compliance mode to ensure immutability. Learn more at AWS Backup Vault Lock.	October 4, 2022
AWS Backup Audit Manager now available in Africa (Cape Town) Region and Europe (Milan) Region	AWS Backup Audit Manager has expanded to Africa (Cape Town) Region and Europe (Milan) Region. For more information on Backup Audit Manager, see Audit backups and create reports with AWS Backup Audit Manager.	September 14, 2022

Change	Description	Date
AWS Backup brings Amazon CloudWatch metrics to Backup console dashboard	AWS Backup enhances its Backup console dashboard to display integrated Amazon CloudWatch metrics for backup and restore jobs for additional monitoring capability and flexibility.	September 8, 2022
Support for additional Amazon EBS encryption flexibility during restore	AWS Backup now offers additional choices of encryption during restoration of Amazon EBS snapshots.	September 1, 2022
AWS Backup supports Amazon S3 cross-account and cross-Region backup copying	AWS Backup now offers cross-Region and cross-account backup copying for Amazon S3 backups. For more information see Amazon S3 backups.	July 28, 2022
AWS Backup Audit Manager offers additional control support for FSx for ONTAP	AWS Backup Audit Manager now offers additional controls to support monitoring and auditing FSx for ONTAP volumes, including Backup resources are included in at least one backup plan and Last recovery point created. For more information, see AWS Backup Audit Manager controls and remediation.	July 22, 2022

Change	Description	Date
AWS Backup adds support to backup and restore Amazon RDS Multi-AZ clusters for PostgreSQL and MySQL clusters	AWS Backup has added a Multi-Availability Zone cluster backup and restore option with one primary and two readable standby database instances. To learn more, see Amazon RDS Multi-AZ backups.	July 20, 2022
AWS Backup Audit Manager adds new control for recovery point creation	AWS Backup Audit Manager offers a new audit control for increased compliance support. `Last recovery point created` is an optional additional control to ensure recovery points are created within specified time frames. To learn more, see Last recovery point created control.	June 29, 2022
Added AWS Backup Gateway endpoint sample	AWS Backup Gateway provided a sample endpoint to assist users with connectin g to VPNs (Virtual Private Networks).	June 14, 2022

Change	Description	Date
AWS Backup now offers Amazon VPC endpoints for VMware	AWS Backup now supports Amazon VPC endpoints for VMware, enabling you to use a virtual private network between your VMware environments and AWS using AWS PrivateLink. For more information, see Creating a gateway and AWS Backup and AWS PrivateLink.	June 1, 2022
AWS Backup Audit Manager offers additional control support for Amazon S3	Backup Audit Manager now offers support for the compliance control **Backup resources protected by backup plan** for S3 resource types. For more information, see AWS Backup Audit Manager controls and remediation.	May 25, 2022
AWS Backup Audit Manager offers additional control support for Storage Gateway	Backup Audit Manager now offers support for the compliance control **Backup resources protected by backup plan** for Storage Gateway resource types. For more information, see AWS Backup Audit Manager controls and remediation.	May 25, 2022

Change	Description	Date
Support for Amazon FSx for OpenZFS	AWS Backup now offers added management of data protection for backing up and restoring to FSx for OpenZFS file systems.	May 18, 2022
AWS Backup Audit Manager support for VMware	AWS Backup now provides support for virtual machines in the Backup Audit Manager controls and remediation. For more information, see AWS Backup Audit Manager controls and remediation.	May 11, 2022
Amazon FSx now supported in Asia Pacific (Osaka) Region	AWS Backup now offers backing up Amazon FSx in, and cross-Region copies to and from, the Asia Pacific (Osaka) Region.	April 26, 2022
Support for Amazon FSx for Lustre Persistent_2	AWS Backup now offers general availability of support for Amazon FSx for Lustre, which supports higher levels of throughput per storage unit as compared to Persisten t_1 file systems.	April 5, 2022
VMware Enhancements	AWS Backup now offers restoring to Amazon EBS Volume, disk level restore, and support for VMware on AWS Outposts. For more information, see Restoring a virtual machine.	March 31, 2022

Change	Description	Date
AWS Backup Availability for Asia Pacific (Jakarta)	AWS Backup is now available to customers in the Asia Pacific (Jakarta) Region.	March 17, 2022
New Controls for AWS Backup Audit Manager	AWS Backup Audit Manager introduces three new audit controls: Cross-Region copy, Cross-account copy, and Backup Vault Lock. For more information, see AWS Backup Audit Manager controls and remediation.	March 17, 2022
Support for AWS PrivateLink	With AWS PrivateLink for AWS Backup, you can connect directly to AWS Backup using an interface endpoint in your VPC instead of connectin g over the public internet. Interface endpoints are directly accessible from applications that are on premises or in a different AWS Region. For more informati on, see AWS Backup and AWS PrivateLink.	February 28, 2022

Change	Description	Date
Support for Amazon Simple Storage Service (Amazon S3)	General availability of AWS Backup for Amazon S3 in all AWS Regions is available except for China (Beijing) Region, China (Ningxia) Region, AWS GovCloud (US-West), and AWS GovCloud (US-East) Regions. For more information, see Working with Amazon S3 data.	February 14, 2022
Support for Advanced DynamoDB backup in AWS China Regions	Advanced DynamoDB backup is now available in China (Beijing) Region and China (Ningxia) Region. For more information, see Advanced DynamoDB backup.	January 18, 2022
Public preview of support for Amazon S3	AWS Backup offers a public preview of Amazon S3 backups. For more informati on, see Working with Amazon S3 data.	November 30, 2021
Support for VMware virtual machines (VMs)	You can now use AWS Backup to automatically back up VMware VMs. For more information, see Virtual machine backups.	November 30, 2021

Change	Description	Date
Support for advanced DynamoDB backup	You can now use AWS Backup to perfrom the following features for all new DynamoDB table backups you create: cold storage tiering, cost allocation tagging, cross-Region copy, cross-account copy, indpendent encryption, and copying tags from source DynamoDB tables. For more information, see Advanced DynamoDB backup in the *Amazon DynamoDB Developer Guide* and Using AWS Backup with DynamoDB.	November 23, 2021
Support for AWS Backup resource assignment enhancement in AWS China Regions	AWS Backup resource assignment enhancement is now available in China (Beijing) Region and China (Ningxia) Region. For more information, see Assigning resources to a backup plan.	November 16, 2021

Change	Description	Date
Launch of AWS Backup resource assignment enhancement	Backup resource assignment enhancement gives you additional, fine-grained controls and new streamlined processes to deploy backup plans that protect hundreds of thousands of AWS resources. Use this feature to increase your speed, flexibility, and precision when protecting data using AWS Backup. For more information, see Assigning resources to a backup plan.	November 10, 2021
Support for Amazon Neptune	You can now use AWS Backup to back up Amazon Neptune clusters. To learn more, see What is AWS Backup?	November 5, 2021
Support for Amazon DocumentDB	You can now use AWS Backup to back up Amazon DocumentDB clusters. To learn more, see What is AWS Backup?	November 5, 2021
Support for AWS Backup Vault Lock in AWS China Regions	AWS Backup Vault Lock is now available in China (Beijing) Region and China (Ningxia) Region. For more information, see AWS Backup Vault Lock.	November 3, 2021

Change	Description	Date
Launch of AWS Backup Vault Lock	With AWS Backup Vault Lock, you can prevent deletion of backups stored in an AWS Backup backup vault. For more information, see AWS Backup Vault Lock.	October 7, 2021
Launch of AWS Backup Audit Manager compliance reports	With compliance reports, you can generate daily reports on the compliance of your backup activity and resources against the controls you defined in your AWS Backup Audit Manager frameworks. For more information, see Compliance report templates.	October 5, 2021
CloudFormation support for AWS Backup Audit Manager	With CloudFormation, you can now deploy AWS Backup Audit Manager framework s, controls, and report plans in a safe, repeatable manner at scale. For more informati on, see Backup audit and reports with AWS Backup Audit Manager.	October 4, 2021

Change	Description	Date
Launch of AWS Backup Audit Manager	With AWS Backup Audit Manager, you can now define controls for your backup activity and resources, and identify the activites and resources that do not comply with your controls. You can also use AWS Backup Audit Manager to generate daily and on-demand reports that serve as evidence of compliance with your defined controls over time. For more information, see Backup audit and reports with AWS Backup Audit Manager.	August 24, 2021
Support for new asynchron ous recovery point operations	AWS Backup now assumes a service-linked role to manage your backup lifecycle rules in the event that you modified or deleted your original IAM role. For more information, see Deleting backups.	August 23, 2021

Change	Description	Date
Support for Amazon EBS multi-volume, crash-consistent backup	Now, when you use AWS Backup to protect your Amazon EC2 instances, AWS Backup takes multi-volume, crash-consistent backups of all the Amazon EBS volumes attached to each Amazon EC2 instance by default. For more information, see Creating Amazon EBS multi-volume, crash-consistent backup.	June 14, 2021
Support for Amazon FSx in additional AWS Regions	You can now use AWS Backup to protect your Amazon FSx file systems in the following Regions: AWS GovCloud (US), Europe (Milan) Region, Africa (Cape Town) Region, and Middle East (Bahrain) Region. For more information, see AWS Backup endpoints and quotas in the *AWS General Reference*.	April 15, 2021

Change	Description	Date
Support for Amazon FSx cross-Region and cross-account backups	You can now use AWS Backup to copy Amazon FSx backups across AWS Regions and accounts. For more information, see Creating a Backup Copy. If you use customer managed policies, you should add the new permission `fsx:CopyBackup` to prevent existing backup jobs from failing. For that permission, see the last statement in the Amazon FSx backup policy in the Customer managed policies.	April 12, 2021
Support for cost allocation tags for Amazon EFS backups	You can now use cost allocation tags to track costs for your Amazon EFS backups on a detailed level, and view and filter those tags using AWS Cost Explorer. For more information, see Using Cost Allocation Tags.	April 7, 2021
FedRAMP High Authorization	AWS Backup is now authorized to support FedRAMP High workloads. For more information, see AWS Services in Scope by Compliance Program.	March 25, 2021

Change	Description	Date
New AWS Region	AWS Backup is now available in the Asia Pacific (Osaka) Region. In this Region, AWS Backup currently does not support Storage Gateway, Amazon FSx, and cross-acc ount backup in this Region. For more information, see AWS Backup endpoints and quotas in the *AWS General Reference*.	March 25, 2021
Support for recovery point batch operations	You can now use the AWS Backup console to automate batch operations to clean up recovery points in your backup vaults. For more information, see Deleting backups.	March 23, 2021
Support for restores to the Amazon EFS One Zone storage class	You can now restore your Amazon EFS backups to the Amazon EFS One Zone storage class. For more information, see Restoring an Amazon EFS file system.	March 12, 2021

Change	Description	Date
Support for Amazon Relational Database Service point-in-time restore and continuous backup	You can now use AWS Backup to automate Amazon RDS continuous backups and perform point-in-time restore (PITR), in addition to orchestrating your snapshot backups. For more informati on, see Restoring to a specified time using point-in-time recovery.	March 10, 2021
Support for Amazon CloudWatch	You can now use CloudWatc h to monitor AWS Backup metrics. For more informati on, see Monitoring Events and Metrics with Amazon CloudWatch and Amazon EventBridge.	February 3, 2021
Support for Amazon EventBridge	You can now use EventBrid ge to monitor AWS Backup events. For more informati on, see Monitoring Events and Metrics with Amazon CloudWatch and Amazon EventBridge.	February 3, 2021
Support for cross-account backups	You can now use AWS Backup to back up your resources across multiple AWS accounts. For more information, see Creating backup copies across AWS accounts.	November 18, 2020

Change	Description	Date
Support for backing up and restoring Amazon FSx file systems	You can now use AWS Backup to back up Amazon FSx file systems. For more informati on, see Working with Amazon FSx file systems.	November 9, 2020
New AWS Regions	AWS Backup is now available in the Africa (Cape Town) and Europe (Milan) AWS Regions. For more information, see AWS Backup endpoints and quotas in the *AWS General Reference*.	October 21, 2020
Support for VSS-Enabled Windows backup	You can now back up and restore VSS (Volume Shadow Copy Service)-enabled Windows applications running on Amazon EC2 instances . For more information, see Creating Windows VSS backups.	September 22, 2020
Support for Amazon EFS automatic backup	You can now use AWS Backup to automatically back up Amazon EFS file systems. For more information, see Getting started 4: Create Amazon EFS automatic backups.	July 16, 2020

Change	Description	Date
New AWS Region	AWS Backup is now available in the AWS GovCloud (US) Region. For more information, see AWS Backup endpoints and quotas in the *AWS General Reference*.	June 24, 2020
Support for managing backups across multiple AWS accounts	You can now manage backups across multiple AWS accounts by using AWS Organizations. For more information, see How Cross-Account Management Works.	June 24, 2020
Support for Amazon Aurora added to AWS Backup	You can now configure AWS Backup to back up resources for Amazon Aurora. For information, see Overview of Backing Up and Restoring an Aurora DB Cluster in the *Amazon Aurora User Guide*.	June 10, 2020
Support for configuring services to work with AWS Backup	You can now configure AWS Backup to back up resources for specific AWS services. For more information, see Opt in to managing services with AWS Backup.	May 20, 2020

Change	Description	Date
Support for backing up Amazon EC2 instances and also adds support for cross-Region backup	You can now back up entire Amazon EC2 instances and also copy resources across AWS Regions. For more information, see Creating backup copies across AWS Regions.	January 13, 2020
New guide	AWS launches AWS Backup and the *AWS Backup Developer Guide*.	January 15, 2019