AMS Advanced Concepts and Procedures

AMS Advanced User Guide

Version April 12, 2024

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

AMS Advanced User Guide: AMS Advanced Concepts and Procedures

Copyright © 2024 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 Managed Services? 1
About this guide 3
Operations plans 3
Accelerate operations plan 4
Advanced operations plan 4
Getting started 4
Key terms 5
Service description 11
AMS features 11
What we do, what we do not do 16
AMS responsibility matrix (RACI) 17
AMS environment basic components 33
AMS account limits 36
AMS service level objectives (SLOs) 38
Supported AWS services 40
Supported configurations 42
Unsupported operating systems 45
AMS Advanced interfaces 46
VPC endpoints 47
AMS protected namespaces 48
AMS reserved prefixes 49
AMS maintenance window 50
AMS information resources 50
AMS compliance 51
AMS Supported Compliance Standards 52
Shared Responsibility 55
AMS Amazon Machine Images (AMIs) 55
Security enhanced AMIs 59
How integration between AD FS and AMS works 59
AMS Managed Active Directory 61
Application deployments 64
Service management 65
Account governance 65
Service commencement 66

Customer relationship management (CRM) ... 66
CRM Process ... 67
CRM meetings ... 68
CRM Meeting Arrangements ... 69
CRM monthly reports ... 70
Cost optimization ... 71
Cost optimization framework ... 71
Cost optimization responsibility matrix ... 73
Planned event management (PEM) ... 75
AMS PEM criteria ... 75
Types of PEM ... 76
The AMS PEM process ... 76
PEM FAQs ... 76
Service hours ... 77
Getting help ... 78
Network architecture ... 79
MALZ network architecture ... 79
About multi-account landing zone network architecture ... 79
Choosing single MALZ or multiple MALZs ... 82
Multi-Account Landing Zone accounts ... 87
SALZ network architecture ... 132
AMS Single-account landing zone shared services ... 134
Setting up AMS ... 136
AMS default settings ... 137
DNS resolution defaults (MALZ) ... 137
EC2 IAM instance profile ... 138
Alerts from baseline monitoring in AMS ... 144
Log retention and rotation defaults ... 160
Using the AMS consoles ... 161
Using the AMS API and CLI ... 163
AMS API HTTP endpoints for REST calls ... 163
Installing or upgrading the AMS CLI ... 163
Using the AMS API in CLI, Ruby, Python, and Java ... 165
AMS bring your own EPS ... 174
Using BYOEPS ... 176
Receiving AMS notifications ... 177

AMS AMI notifications with SNS ... 179
Service notifications ... 183
RFC state change notifications ... 183
Setting up private and public DNS ... 187
AMS egress traffic management ... 191
Deploying IAM resources ... 192
Automated IAM Provisioning ... 193
Setting permissions with IAM roles and profiles ... 215
Requesting a new IAM user role or instance profile ... 215
Restrict permissions with IAM role policy statements ... 215
Restrict permissions with Amazon EC2 IAM instance profiles ... 218
AD FS claim rule and SAML settings ... 219
ADFS claim rule configurations ... 219
Web console ... 220
API and CLI access with SAML ... 220
Restrict with network ACL ... 223
AMS on Outposts ... 224
AWS Outposts installation and operational management ... 224
Provisioning AMS managed resources on AWS Outposts ... 227
Limitations of AMS on AWS Outposts ... 227
AMS on AWS Outposts compliance ... 228
AMS on AWS Outposts FAQs ... 229
Using tags ... 231
AMS infrastructure automatic tagging ... 233
AMS recommended tags ... 235
Tag bulk update notes ... 239
Resource Scheduler ... 241
Deploying Resource Scheduler ... 242
Customizing Resource Scheduler ... 242
Using Resource Scheduler ... 243
AMS Resource Scheduler cost estimator ... 243
AMS Resource Scheduler best practices ... 244
AWS Systems Manager in AMS Advanced ... 246
Available AMS Advanced SSM documents ... 247
AMS Advanced SSM document versions ... 247
Systems Manager pricing ... 248

Offboard AMS accounts ... 248
Offboard from single-account landing zone accounts ... 248
Offboard from multi-account landing zone accounts ... 252
Change management modes ... 273
Modes overview ... 274
Types of modes and accounts in AMS ... 274
AMS modes and applications or workloads ... 279
Real world use cases for AMS modes ... 287
RFC mode ... 291
Getting started with RFC mode ... 291
What are RFCs? ... 291
How RFCs work ... 293
What are change types? ... 338
Troubleshooting RFC errors ... 350
Direct Change mode ... 361
Getting Started with Direct Change mode ... 362
Security and compliance ... 364
Change management in Direct Change mode ... 369
Creating stacks using Direct Change mode ... 372
Direct Change Mode use cases ... 375
Developer mode ... 376
Getting started with Developer mode ... 377
Security and compliance ... 379
Change management ... 381
Provisioning infrastructure ... 386
Detective controls ... 387
Logging, monitoring, and event management ... 387
Incident management ... 387
Patch management ... 387
Continuity management ... 388
Security and access management ... 388
Self-Service Provisioning mode ... 388
Getting started with SSP mode ... 389
Amazon API Gateway ... 390
Alexa for Business ... 391
Amazon AppStream 2.0 ... 392

Amazon Athena ... 394
Amazon CloudSearch ... 395
Amazon CloudWatch Synthetics ... 396
Amazon Cognito ... 397
Amazon Comprehend ... 399
Amazon Connect ... 399
Amazon DevOps Guru ... 401
Amazon DocumentDB (with MongoDB compatibility) ... 402
Amazon DynamoDB ... 403
Amazon Elastic Container Registry ... 404
EC2 Image Builder ... 405
Amazon ECS on AWS Fargate ... 407
Amazon EKS on AWS Fargate ... 409
Amazon EMR ... 412
Amazon EventBridge ... 415
Amazon Forecast ... 416
Amazon FSx ... 418
Amazon FSx OpenZFS ... 420
Amazon FSx NetApp ONTAP ... 421
Amazon Inspector Classic ... 422
Amazon Managed Service for Apache Flink ... 423
Amazon Managed Service for Prometheus ... 424
Amazon Data Firehose ... 425
Amazon Kinesis Data Streams ... 426
Amazon Kinesis Video Streams ... 427
Amazon Lex ... 428
Amazon MQ ... 429
Amazon Managed Streaming for Apache Kafka ... 430
Amazon Personalize ... 431
Amazon QuickSight ... 433
Amazon Rekognition ... 435
Amazon SageMaker ... 436
Amazon Simple Email Service ... 439
Amazon Simple Workflow Service ... 440
Amazon Textract ... 441
Amazon Transcribe ... 441

Amazon WorkDocs ... 442
Amazon WorkSpaces ... 443
AMS Code services ... 445
AWS Amplify ... 448
AWS AppSync ... 449
AWS App Mesh ... 450
AWS Audit Manager ... 451
AWS Batch ... 452
AWS Certificate Manager ... 453
AWS Private Certificate Authority ... 454
AWS CloudEndure ... 457
AWS CloudHSM ... 457
AWS CodeBuild ... 459
AWS CodeCommit ... 460
AWS CodeDeploy ... 461
AWS CodePipeline ... 462
AWS Compute Optimizer ... 464
AWS DataSync ... 465
AWS Device Farm ... 467
AWS Elastic Disaster Recovery ... 467
AWS Elemental MediaConvert ... 469
AWS Elemental MediaLive ... 469
AWS Elemental MediaPackage ... 470
AWS Elemental MediaStore ... 471
AWS Elemental MediaTailor ... 472
AWS Global Accelerator ... 473
AWS Glue ... 473
AWS Lake Formation ... 474
AWS Lambda ... 475
AWS License Manager ... 476
AWS Migration Hub ... 477
AWS Outposts ... 478
AWS Resilience Hub ... 479
AWS Secrets Manager ... 480
AWS Security Hub ... 483
AWS Service Catalog AppRegistry ... 483

AWS Shield ... 484
AWS Snowball ... 486
AWS Step Functions ... 487
AWS Systems Manager Parameter Store ... 488
AWS Systems Manager Automation ... 489
AWS Transfer Family ... 492
AWS Transit Gateway ... 493
AWS WAF ... 494
AWS Well-Architected Tool ... 495
AWS X-Ray ... 496
VM Import/Export ... 497
Customer Managed mode ... 498
Getting started with Customer Managed mode ... 499
AMS and AWS Service Catalog ... 499
Getting started with Service Catalog ... 499
Service Catalog in AMS before you begin ... 500
Finding the data you need (SKMS) ... 504
What Is service knowledge management? ... 504
Find VPC IDs ... 505
Find subnet IDs ... 507
Find AMI IDs ... 509
Find security group (SG) IDs ... 511
Find IAM entities ... 511
Find stack IDs ... 512
Find instance IDs or IP addresses ... 514
Find ARNs ... 516
Find resources by ARN ... 518
Find account settings ... 519
Find FQDNs ... 521
Find availability zones (AZs) ... 521
Find SNS topics ... 522
Find backup settings ... 523
Access management ... 524
What is Access Management? ... 524
Why and when we access your account ... 525
How and when to use root ... 530

AMS Advanced console and Amazon EC2 access ... 531
Accessing the AWS Management console and the AMS console ... 532
Temporary AMS console access ... 533
Accessing instances using bastions ... 534
DNS friendly bastion names ... 535
Saving costs on Single-account landing zone (SALZ) bastions ... 537
Using bastion IP addresses ... 537
Instance access examples ... 538
Team, or role, based access control in an AMS account ... 550
Automated EC2 instance configuration ... 551
Prerequisites for automated instance configuration ... 551
SSM Agent automatic installation ... 552
Automated changes ... 553
Automatically update code on Linux instances ... 553
Automatically update PBIS on Linux instances ... 554
Automatically update the minimum version of SSM and CloudWatch agents ... 555
CloudWatch configuration files, update details ... 556
Automatically configured logs ... 556
Monitoring and event management ... 558
What is monitoring? ... 559
What does the AMS monitoring system monitor? ... 560
Single-Account Landing Zone proactive monitoring of Active Directory Trust ... 560
How monitoring works ... 561
Alert notification ... 562
Tag-based alert notification ... 562
Viewing the monitoring configuration for an account ... 563
Changing the monitoring configuration for an account ... 563
Using OpsCenter ... 564
Alert notifications ... 564
Receiving alerts ... 565
Tag-based alert notifications ... 565
AMS automatic remediation of alerts ... 566
Creating additional CloudWatch alarms ... 569
Creating custom CloudWatch metrics and alarms ... 570
Using CloudWatch Application Insights for .Net and SQL server ... 571
AMS Event Router ... 572

Amazon EventBridge Managed Rules deployed by AMS ... 573
Creating Managed Rules for AMS ... 573
Editing Managed Rules for AMS ... 573
Deleting Managed Rules for AMS ... 573
Log management ... 575
What is log management? ... 575
How AMS logging works ... 575
Accessing your logs ... 576
AMS aggregated service logs ... 576
AMS shared services logs ... 589
Amazon Elastic Compute Cloud (Amazon EC2) - system level logs ... 591
Integrating with Splunk ... 593
Customizing your log configuration ... 594
Altering CloudWatch log retention ... 594
Enabling logging for supported services ... 595
Security management ... 596
Data protection in AMS ... 596
Amazon Macie ... 597
GuardDuty ... 598
AWS Certificate Manager (ACM) certificate ... 601
Data encryption in AMS ... 601
Identity and access management ... 603
Multi-Account Landing Zone (MALZ) IAM safeguards ... 603
Authenticating with identities ... 604
Security Incident Response ... 632
Security event logging and monitoring ... 655
Endpoint Security (EPS) ... 655
Malware mitigation process ... 659
Amazon Inspector security ... 661
AMS incident response ... 662
Compliance validation ... 663
Multi-Account Landing Zone viewing the compliance status of your AWS Config Rules ... 664
AMS multi-account landing zone service control policy restrictions ... 666
Resilience ... 666
Infrastructure security ... 666
Security best practices ... 699

AMS multi-account landing zone EPS non-default settings ... 699
AMS Guardrails ... 699
MALZ Service control policies ... 700
Continuity management ... 701
What is continuity management? ... 701
How continuity management works ... 701
Backup plans ... 702
Backup vaults ... 705
Backup change types ... 706
Monitoring and reporting for backups ... 707
Disaster recovery response ... 707
Disaster recovery planning ... 708
Multi-site or highly available (HA) ... 709
Warm standby ... 710
Pilot light ... 712
Backup and restore ... 714
Patch management ... 721
AMS Patch Orchestrator: a tag-based patching model ... 722
Using Patch Orchestrator ... 723
Patch Orchestrator prerequisites ... 727
Patch windows ... 728
Patch notifications ... 729
Patch baselines ... 732
Patch Orchestrator reserved tags ... 732
On-demand patching ... 732
AMS standard patching ... 733
Supported operating systems ... 734
Supported patches ... 734
Patching and infrastructure design ... 737
How AMS standard patching works ... 737
AMS standard patching failures ... 743
Actions you can take in AMS standard patching ... 743
AMS standard patching FAQs ... 746
Patching service commitments ... 749
Standard patching ... 749
Critical patching ... 751

Reports and options 754
On-request reporting 754
AMS Patch reporting 754
AMS Backup reporting 763
Billing Charges Details report 767
Self-service reporting 768
Daily Patch reports 770
Monthly billing report 778
Daily backup report 781
Weekly Incident report 785
Data retention policy 787
Offboarding from SSR 788
Get support 789
Incident management 789
What is incident management? 790
Incident management service commitments 792
Incident management examples 794
Service request management 803
How service request management works 804
Service request management examples 804
Billing questions 812
Operations On Demand 814
Requesting AMS Operations On Demand 821
Making changes to Operations on Demand offerings 822
Document history 823
AWS Glossary 850

What is AWS Managed Services?

Welcome to AWS Managed Services (AMS), infrastructure operations management for Amazon Web Services (AWS). AMS is an enterprise service that provides ongoing management of your AWS infrastructure.

This user guide is intended for IT and application developer professionals. A basic understanding of IT functionality, networking, and application deployment terms and practices is assumed.

AMS implements best practices and maintains your infrastructure to reduce your operational overhead and risk. AMS provides full-lifecycle services to provision, run, and support your infrastructure, and automates common activities such as change requests, monitoring, patch management, security, and backup services. AMS enforces your corporate and security infrastructure policies, and enables you to develop solutions and applications using your preferred development approach.

To better understand AMS architecture, see these diagrams.

Topics

- About this AMS user guide
- AMS operations plans
- Getting started with AWS Managed Services
- AMS key terms
- Service description
- AMS information resources
- AMS compliance
- AMS Amazon Machine Images (AMIs)
- How integration between AD FS and AMS works
- AMS Managed Active Directory
- AMS application deployments

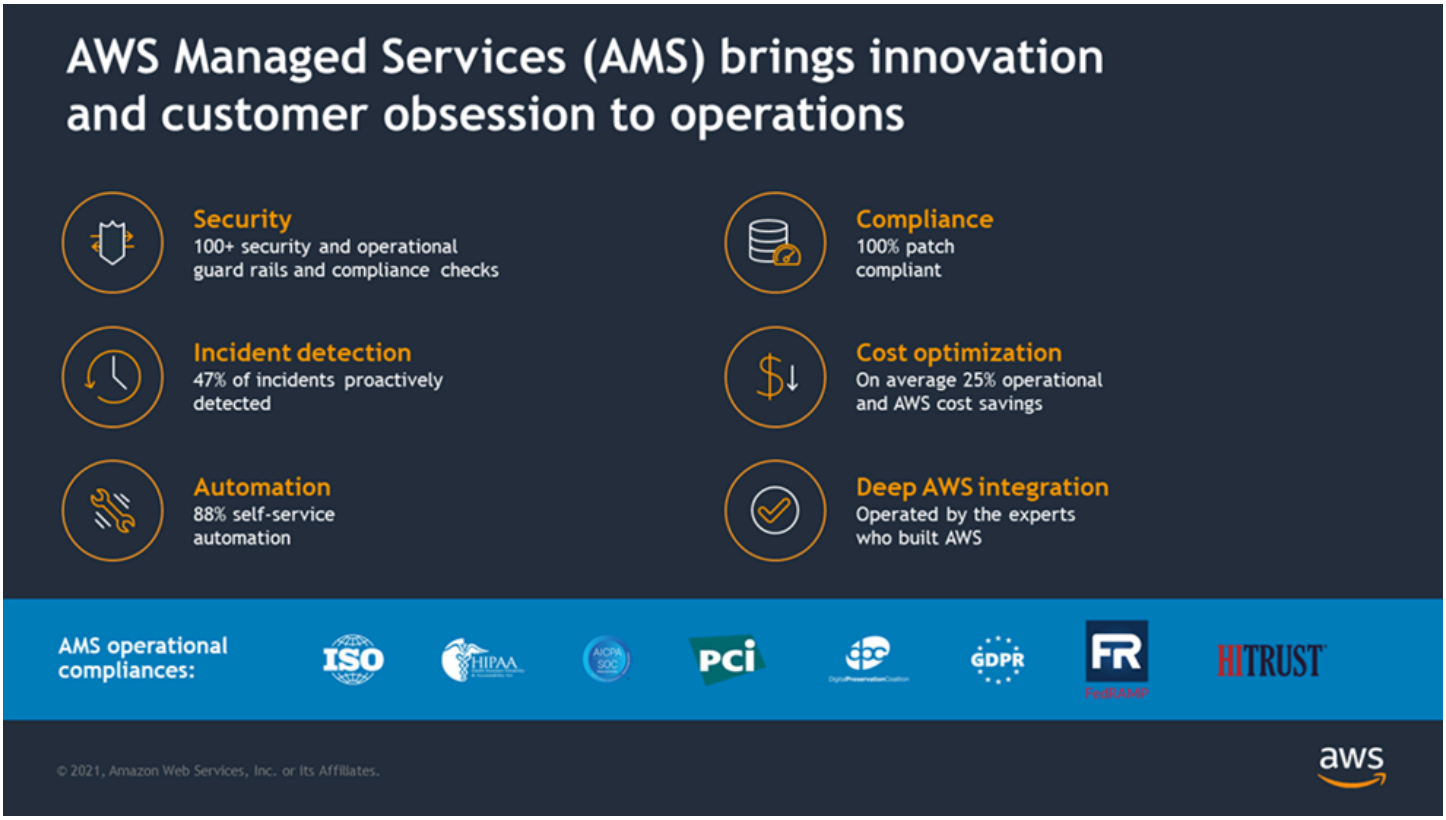

> **Note**
>
> New AWS Regions are added frequently. For the most recent AMS-supported AWS Regions, and the most recent AMS-supported operating systems, see Supported configurations. To learn more about AWS Regions, see Managing AWS Regions.

AMS seeks to continuously improve our services based on your feedback. We use several mechanisms to enable your self-service, to automate repetitive tasks, and to implement new AWS services and features as they are released. You can submit an AMS service request at any time to suggest new features or feature improvements.

AMS business hours are 24 hours a day, 7 days a week, 365 days a year.

AMS follows a set of practices for IT service management (ITSM) that focuses on aligning IT services with the needs of your business.

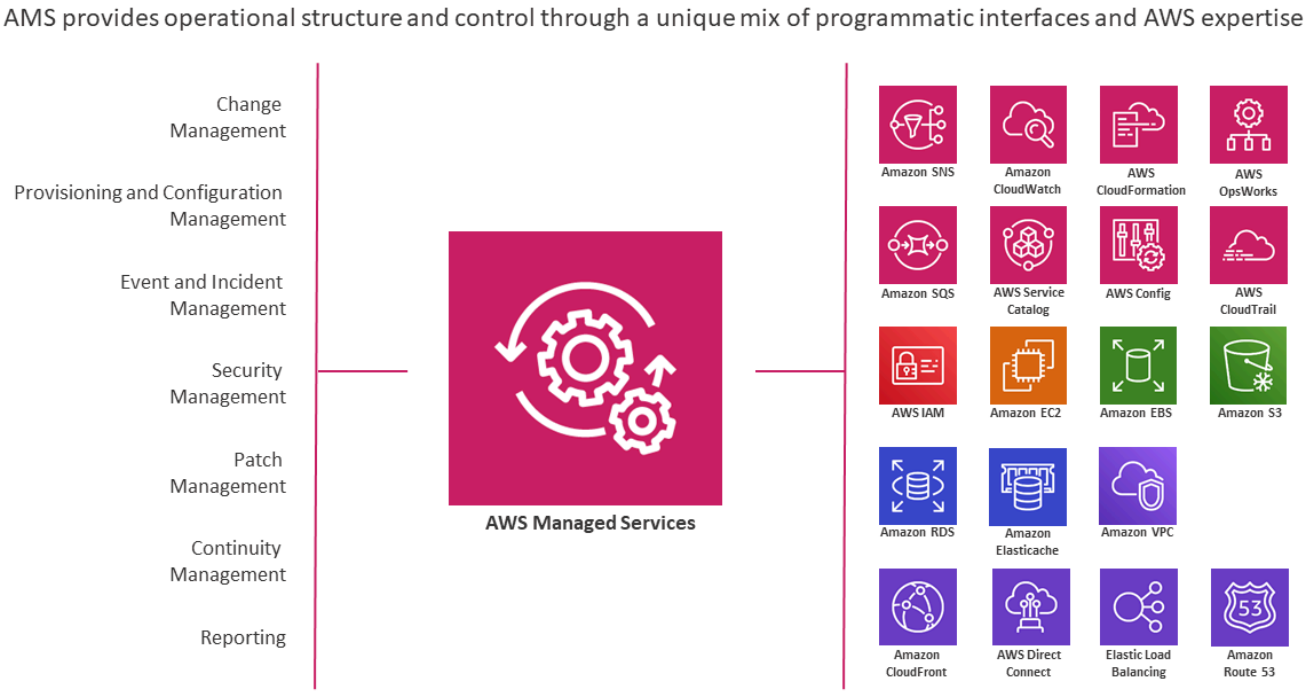

About this AMS user guide

This user guide is intended for AMS Advanced customers with either a multi-account or single-account landing zone. For more details about the AMS landing zone offerings, see the AMS Key Terms; also see Multi-Account Landing Zone architecture and Single-Account Landing Zone architecture.

AMS operations plans

AWS Managed Services (AMS) is available with two operations plans: AMS Accelerate and AMS Advanced. An operations plan offers a specific set of features and has differing levels of service, technical capabilities, requirements, price, and restrictions. Our operations plans give you the flexibility to select the right-sized operational capabilities for each of your AWS workloads. This section outlines the capabilities and differences, as well as the responsibilities, features, and benefits associated with each plan, so that you can understand which operations plan is best for your accounts.

For a detailed feature comparison of the two operations plans, see AWS Managed Services Features.

AMS Accelerate operations plan

AMS Accelerate is the AMS operations plan that helps you operate the day-to-day infrastructure management of your new or existing AWS environment. AMS Accelerate provides operational services, such as monitoring, incident management, and security. AMS Accelerate also offers an optional patch add-on for Amazon EC2-based workloads that require regular patching.

With AMS Accelerate, you decide which AWS accounts you want AMS Accelerate to operate, the AWS Regions you want AMS Accelerate to operate in, the add-ons you require, and the service-level agreements (SLAs) you need. For more details, see Using the AMS Accelerate operations plan and Service Description.

AMS Advanced operations plan

AMS Advanced provides full-lifecycle services to provision, run, and support your infrastructure. In addition to the operational services provided by AMS Accelerate, AMS Advanced also includes additional services, such as landing zone management, infrastructure changes and provisioning, access management, and endpoint security.

AMS Advanced deploys a landing zone to which you migrate your AWS workloads and receive AMS operational services. Our managed multi-account landing zones are pre-configured with the infrastructure to facilitate authentication, security, networking, and logging.

AMS Advanced also includes a change and access management system that protects your workloads by preventing unauthorized access or the implementation of risky changes to your AWS infrastructure. Customers need to create a request for change (RFC) using our change management system to implement most changes in your AMS Advanced accounts. You create RFCs from a library of automated changes that are pre-vetted by our security and operations teams or request manual changes that are reviewed and implemented by our operations team if they are deemed both safe and supported by AMS Advanced.

AMS Advanced also offers different SLAs. For more information, see the AWS Managed Services AMS Advanced service description.

Getting started with AWS Managed Services

For details about getting started with the multi-account landing zone AMS service, see the AWS Managed Services Onboarding Introduction. The two onboarding guides provide descriptions of

the service and questions to consider to help you get started. Review the feature set AWS Managed Services Features, and current resources at AWS Managed Services Resources.

AMS key terms

- *AMS Advanced*: The services described in the "Service Description" section of the AMS Advanced Documentation. See Service Description.
- *AMS Advanced Accounts*: AWS accounts that at all times meet all requirements in the AMS Advanced Onboarding Requirements. For information on AMS Advanced benefits, case studies, and to contact a sales person, see AWS Managed Services.
- *AMS Accelerate Accounts*: AWS accounts that at all times meet all requirements in the AMS Accelerate Onboarding Requirements. See Getting Started with AMS Accelerate.
- *AWS Managed Services*: AMS and or AMS Accelerate.
- *AWS Managed Services Accounts*: The AMS accounts and or AMS Accelerate accounts.
- *Critical Recommendation*: A recommendation issued by AWS through a service request informing you that your action is required to protect against potential risks or disruptions to your resources or the AWS services. If you decide not to follow a Critical Recommendation by the specified date, you are solely responsible for any harm resulting from your decision.
- *Customer-Requested Configuration*: Any software, services or other configurations that are not identified in:
  - Accelerate: Supported Configurations or AMS Accelerate; Service Description.
  - AMS Advanced: Supported Configurations or AMS Advanced; Service Description.
- *Incident Communication*: AMS communicates an Incident to you or you request an Incident with AMS via an Incident created in Support Center for AMS Accelerate and in the AMS Console for AMS. The AMS Accelerate Console provides a summary of Incidents and Service Requests on the Dashboard and links to Support Center for details.
- *Managed Environment*: The AMS Advanced accounts and or the AMS Accelerate accounts operated by AMS.

  For AMS Advanced, these include multi-account landing zone (MALZ) and single-account landing zone (SALZ) accounts.
- *Billing start date*: The next business day after AWS receives the Customer's information requested in the AWS Managed Services Onboarding Email. The AWS Managed Services Onboarding Email refers to the email form sent by AWS to the Customer to collect the information necessary in order to activate AWS Managed Services on the Customer's accounts.

For accounts subsequently enrolled by the Customer, the billing start date will be the next business day after AWS Managed Services sends an AWS Managed Services Activation Notification for the enrolled account. An AWS Managed Services Activation Notification occurs when:

1. Customer grants access to a compatible AWS account and hands it over to AWS Managed Services.
2. AWS Managed Services designs and builds the AWS Managed Services Account.

- *Service Termination*: The Customer may terminate the AWS Managed Services for all AWS Managed Services Accounts, or for a specified AWS Managed Services Account for any reason by providing AWS at least 30 days notice through a service request. On the Service Termination Date, either:
  1. AWS will hand over the controls of all AWS Managed Services Accounts or the specified AWS Managed Services Accounts, as applicable, to customer, or
  2. The parties will remove the AWS Identity and Access Management roles that give AWS access from all AWS Managed Services Accounts or the specified AWS Managed Services Accounts, as applicable.
- *Service Termination Date*: The last day of the calendar month following the end of the 30 days requisite termination notice period; provided that, if the end of the requisite termination notice period falls after the 20th day of the calendar month, the Service Termination Date will be the last day of the following calendar month. Examples:
  - If a customer issued termination notice on April 12, 30 days notice ends on May 12. The Service Termination Date is May 31.
  - If a customer issued termination notice on April 29, 30 days notice ends on May 29. The Service Termination Date is June 30.
- *Provision of AWS Managed Services*: AWS will make available to customer and customer may access and use AWS Managed Services for each AWS Managed Services account from the service commencement date.
- *Termination for specified AWS Managed Services accounts*: customer may terminate the AWS Managed Services for a specified AWS Managed Services account for any reason by providing AWS notice through a service request ("AMS Account Termination Request").

Incident management terms:

- *Event*: A change in your AMS environment.

- *Alert*: Whenever an event from a supported AWS service exceeds a threshold and triggers an alarm, an alert is created and notice is sent to your contacts list. Additionally, an incident is created in your Incident list.
- *Incident*: An unplanned interruption or performance degradation of your AMS environment or AWS Managed Services that results in an impact as reported by AWS Managed Services or you.
- *Problem*: A shared underlying root cause of one or more incidents.
- *Incident Resolution* or *Resolve an Incident*:
  - AMS has restored all unavailable AMS services or resources pertaining to that incident to an available state, or
  - AMS has determined that unavailable stacks or resources cannot be restored to an available state, or
  - AMS has initiated an infrastructure restore authorized by you.
- *Incident Response Time*: The difference in time between when you create an incident, and when AMS provides an initial response by way of the console, email, service center, or telephone.
- *Incident Resolution Time*: The difference in time between when either AMS or you creates an incident, and when the incident is resolved.
- *Incident Priority*: How incidents are prioritized by AMS, or by you, as either Low, Medium, or High.
  - *Low*: A non-critical problem with your AMS service.
  - *Medium*: An AWS service within your managed environment is available but is not performing as intended (per the applicable service description).
  - *High*: Either (1) the AMS Console, or one or more AMS APIs within your managed environment are unavailable; or (2) one or more AMS stacks or resources within your managed environment are unavailable and the unavailability prevents your application from performing its function.

  AMS may re-categorize incidents in accordance with the above guidelines.
- *Infrastructure Restore*: Re-deploying existing stacks, based on templates of impacted stacks, and initiating a data restore based on the last known restore point, unless otherwise specified by you, when incident resolution is not possible.

Infrastructure terms:

- *Managed production environment*: A customer account where the customer's production applications reside.

- *Managed non-production environment*: A customer account that only contains non-production applications, such as applications for development and testing.
- *AMS stack*: A group of one or more AWS resources that are managed by AMS as a single unit.
- *Immutable infrastructure*: An infrastructure maintenance model typical for Amazon EC2 Auto Scaling groups (ASGs) where updated infrastructure components, (in AWS, the AMI) are replaced for every deployment, rather than being updated in-place. The advantages to immutable infrastructure is that all components stay in a synchronous state since they are always generated from the same base. Immutability is independent of any tool or workflow for building the AMI.
- *Mutable infrastructure*: An infrastructure maintenance model typical for stacks that are not Amazon EC2 Auto Scaling groups and contain a single instance or just a few instances. This model most closely represents traditional, hardware-based, system deployment where a system is deployed at the beginning of its life cycle and then updates are layered onto that system over time. Any updates to the system are applied to the instances individually, and may incur system downtime (depending on the stack configuration) due to application or system restarts.
- *Security groups*: Virtual firewalls for your instance to control inbound and outbound traffic. Security groups act at the instance level, not the subnet level. Therefore, each instance in a subnet in your VPC could have a different set of security groups assigned to it.
- *Service Level Agreements (SLAs)*: Part of AMS contracts with you that define the level of expected service.
- SLA *Unavailable* and *Unavailability*:
  - An API request submitted by you that results in an error.
  - A Console request submitted by you that results in a 5xx HTTP response (the server is incapable of performing the request).
  - Any of the AWS service offerings that constitute stacks or resources in your AMS-managed infrastructure are in a state of "Service Disruption" as shown in the Service Health Dashboard.
  - Unavailability resulting directly or indirectly from an AMS exclusion is not considered in determining eligibility for service credits. Services are considered available unless they meet the criteria for being unavailable.
- *Service Level Objectives (SLOs)*: Part of AMS contracts with you that define specific service goals for AMS services.

Patching terms:

- *Mandatory patches*: Critical security updates to address issues that could compromise the security state of your environment or account. A "Critical Security update" is a security update rated as "Critical" by the vendor of an AMS-supported operating system.
- *Patches announced versus released*: Patches are generally announced and released on a schedule. Emergent patches are announced when the need for the patch has been discovered and, usually soon after, the patch is released.
- *Patch add-on*: Tag-based patching for AMS instances that leverages AWS Systems Manager (SSM) functionality so you can tag instances and have those instances patched using a baseline and a window that you configure.
- *Patch methods*:
  - *In-place patching*: Patching that is done by changing existing instances.
  - *AMI replacement patching*: Patching that is done by changing the AMI reference parameter of an existing EC2 Auto Scaling group launch configuration.
- *Patch provider* (OS vendors, third party): Patches are provided by the vendor or governing body of the application.
- *Patch Types*:
  - *Critical Security Update (CSU)*: A security update rated as "Critical" by the vendor of a supported operating system.
  - *Important Update (IU)*: A security update rated as "Important" or a non-security update rated as "Critical" by the vendor of a supported operating system.
  - *Other Update (OU)*: An update by the vendor of a supported operating system that is not a CSU or an IU.
- *Supported patches*: AMS supports operating system level patches. Upgrades are released by the vendor to fix security vulnerabilities or other bugs or to improve performance. For a list of currently supported OSs, see Support Configurations.

Security terms:

- *Detective Controls*: A library of AMS-created or enabled monitors that provide ongoing oversight of customer managed environments and workloads for configurations that do not align with security, operational, or customer controls, and take action by notifying owners, proactively modifying, or terminating resources.

Service Request terms:

- *Service request*: A request by you for an action that you want AMS to take on your behalf.
- *Alert notification*: A notice posted by AMS to your **Service requests** list page when an AMS alert is triggered. The contact configured for your account is also notified by the configured method (for example, email). If you have contact tags on your instances/resources, and have provided consent to your cloud service delivery manager (CSDM) for tag-based notifications, the contact information (key value) in the tag is also notified for automated AMS alerts.
- *Service notification*: A notice from AMS that is posted to your **Service request** list page.

Miscellaneous terms:

- *AWS Managed Services Interface*: For AMS: The AWS Managed Services Advanced Console, AMS CM API, and AWS Support API. For AMS Accelerate: The AWS Support Console and AWS Support API.
- *Customer satisfaction (CSAT)*: AMS CSAT is informed with deep analytics including Case Correspondence Ratings on every case or correspondence when given, quarterly surveys, and so forth.
- *DevOps*: DevOps is a development methodology that strongly advocates automation and monitoring at all steps. DevOps aims at shorter development cycles, increased deployment frequency, and more dependable releases by bringing together the traditionally-separate functions of development and operations over a foundation of automation. When developers can manage operations, and operations informs development, issues and problems are more quickly discovered and solved, and business objectives are more readily achieved.
- *ITIL*: Information Technology Infrastructure Library (called ITIL) is an ITSM framework designed to standardize the lifecycle of IT services. ITIL is arranged in five stages that cover the IT service lifecycle: service strategy, service design, service transition, service operation, and service improvement.
- *IT service management (ITSM)*: A set of practices that align IT services with the needs of your business.
- *Managed Monitoring Services (MMS)*: AMS operates its own monitoring system, Managed Monitoring Service (MMS), that consumes AWS Health events and aggregates Amazon CloudWatch data, and data from other AWS services, notifying AMS operators (online 24x7) of any alarms created through an Amazon Simple Notification Service (Amazon SNS) topic.
- *Namespace*: When you create IAM policies or work with Amazon Resource Names (ARNs), you identify an AWS service by using a namespace. You use namespaces when identifying actions and resources.

Service description

AWS Managed Services (AMS) is a service for managing operations of your AWS infrastructure. AMS provides routine infrastructure operations such as patch, continuity management, security management, and IT management processes such as incident, change and service request management. For a list of supported services, see Supported AWS services.

YouTube Video: How can AMS help me achieve operational excellence in the cloud?

Topics

- AMS features
- What we do, what we do not do
- AMS responsibility matrix (RACI)
- AMS environment basic components
- AMS account limits
- AMS service level objectives (SLOs)
- Supported AWS services
- Supported configurations
- Capabilities for unsupported operating systems in AMS
- AMS Advanced interfaces
- AMS VPC endpoints
- AMS protected namespaces
- AMS reserved prefixes
- AMS maintenance window

AMS features

AWS Managed Services (AMS) offers the following features for supported AWS services:

- **Logging, Monitoring, Guardrails, and Event Management**:

  AMS configures and monitors your managed environment for logging activity and defines alerts based on a variety of health checks. Alerts are investigated by AMS for applicable AWS services, and those that negatively impact your usage of those services result in the creation of incidents. AMS aggregates and stores all logs generated as a result of all operations in CloudWatch,

CloudTrail, and system logs in S3. Upon request, you can ask for additional alerts to be put in place. In addition to AMS' preventative controls, AMS deploys configuration guardrails and detective controls to provide ongoing protection for you from misconfigurations that could reduce the operational and security integrity of the managed accounts, to enforce your controls such as tagging and compliance. When a monitored control is detected an alarm is generated that results in notification, modification, or termination of resources based on pre-defined AMS defaults that can be modified by you.

- **Continuity management** (Backup and Restore):

  AMS provides backups of resources using standard, existing AWS Backup functionality on a scheduled interval determined by you. Restore actions from specific snapshots can be performed by AMS with your RFC. Data changes that occur between snapshot intervals are the responsibility of you to backup. You can submit an RFC for backup or snapshot requests outside of scheduled intervals. In the case of Availability Zone (AZ) unavailability in an AWS Region, with your permission, AMS restores the managed environment by recreating new stack(s) based on templates and available EBS snapshots of the impacted Stacks.

- **Security and access management**:

  AMS provides endpoint security (EPS) such as configuring anti-virus and anti-malware protection. You can also use your own EPS tool and processes and not use AMS for EPS using a feature called bring your own EPS (BYOEPS). AMS also configures default AWS security capabilities that are approved by you during onboarding, such as identity access management (IAM) roles and EC2 security groups, and uses standard AWS tools (e.g. SecurityHub, Macie, GuardDuty) to monitor and respond to security issues. You manage your users through an approved directory service provided by you. For a list of approved directory services, see Supported configurations.

  AMS includes endpoint security (EPS), which is inclusive of antivirus (AV), and anti-malware protection, malware and intrusion detection (Trend Micro). Security groups are defined per stack template and are modified at launch depending on the visibility of the application (public/private) security groups.

  Access to systems is requested through change management requests for change (RFCs). Access management provides access to distinct resources, such as Amazon EC2 instances, the AWS Management Console, and APIs. After establishing a one-way trust with an AMS Microsoft Active Directory deployment during onboarding and federating to AWS, you can use your existing corporate credentials for all interactions.

- **Patch management**:

AMS applies and installs updates to EC2 instances for supported operating systems (OSs) and software pre-installed with supported operating systems. For a list of supported operating systems, see Supported configurations.

AMS offers two models for patching:

- AMS standard patch for traditional account-based patching, and
- AMS Patch Orchestrator, for tag-based patching.

In AMS standard patch, a monthly maintenance window is chosen by you for AMS to perform most patching activities. AMS applies *critical security updates* outside of the selected maintenance window (with appropriate notifications) and *important updates* during the selected maintenance window. AMS additionally applies updates to infrastructure management tools during the selected maintenance window. You can exclude stacks from patch management or reject updates, if you want.

With AMS Patch Orchestrator, a default maintenance window per account, is defined by you for AMS to perform patching activities. You can schedule additional custom maintenance windows for AMS to patch a specific set of instances defined by you with tags. AMS applies all available updates, but you can filter or reject updates by creating a custom patch baseline. For both models, if you approve or reject an update provided under patch management but later change your mind, you are responsible for initiating the update via an RFC. AMS tracks the patch status of resources and highlights systems that aren't current in the monthly business review. Patch management is limited to stacks in the managed environment, including all AMS managed applications and supported AWS services with patching capabilities (for example, RDS). In order to support all types of infrastructure configurations when an update is released, AMS a) updates the EC2 instance and b) provides an updated AMS AMI for you to use. It is your responsibility to install, configure, patch, and monitor any additional applications not specifically covered above.

- **Change management**:

  AMS change management is the mechanism for you to control changes in your managed environment. AMS uses a combination of preventative and detective controls to facilitate this process and provides different level of control and associated risk depending on the AMS mode selected.

  All actions in your AMS environment are logged in AWS CloudTrail.

For more information about AMS Change Management and different modes, see AMS Change Management guide and AMS Modes.

- **Automated and self-service provisioning management**:

  You can provision AWS resources on AMS Advanced in several ways:

  - Submit provisioning and configuration Requests for Change (RFCs)
  - Deploy through AWS Service Catalog
  - Deploy through Direct Change mode
  - Deploy through Developer mode. Remember that the resources created through the Developer mode are not managed by AMS.
  - Configure AWS services directly using self-service provisioning for select AWS services (see Supported AWS services).

- **Incident management**:

  AMS proactively notifies you of incidents detected by AMS. AMS responds to both customer-submitted and AMS-generated incidents and resolves incidents based on the incident priority. Unless otherwise instructed by you, incidents that are determined by AMS to be a risk to the security of your managed environment, and incidents relating to the availability of AMS and other AWS services, are proactively actioned. AMS takes action on all other incidents once your authorization is received. Recurring incidents are addressed by the problem management process.

- **Problem management**:

  AMS performs trend analysis to identify and investigate problems and to identify the root cause. Problems are remediated either with a workaround or a permanent solution that prevents recurrence of similar future service impact. A post incident report (PIR) may be requested for any "High" incident, upon resolution. The PIR captures the root cause and preventative actions taken, including implementation of preventative measures.

- **Reporting**:

  AMS provides you with a monthly service report that summarizes key performance metrics of AMS, including an executive summary and insights, operational metrics, managed resources, AMS service level agreement (SLA) adherence, and financial metrics around spending, savings, and cost optimization. Reports are delivered by the AMS cloud service delivery manager (CSDM) assigned to you.

- **Service request management** :

  To request information about your managed environment, AMS, or AWS service offerings, submit service requests using the AMS console. You can submit a service request for "How to" questions about AWS services and features or to request additional AMS services.

- **Service Desk** :

  AMS staffs engineering operations with full-time Amazon employees to fulfill non-automated requests including incident management, service request management, and change management. The Service Desk operates 24 x 7 365 days a year.

- **Designated resources**:

  Each customer is assigned a Cloud Service Delivery Manager (CSDM) and a Cloud Architect (CA).

  - CSDMs can be contacted directly. They perform service reviews, and delivery reporting and insights through all phases of the implementation, migration and operational life cycle. CSDMs conduct monthly business reviews and detail items such as financial spend, cost-saving recommendations, service utilization, and risk reporting. They dive deep into operational performance statistics and provide recommendations of areas of improvements.
  - CAs can be contacted directly and provide technical expertise to help you optimize your use of the AWS cloud. Example CA activities include, selecting workloads for migration, assisting with the onboarding additional accounts and workloads, acting as the technical lead in operational activities such as game days, disaster recovery testing, problem management, and technical advice to get the most out of AMS and AWS. CAs drive technical discussions at all levels of your organization and assist with incident management, making trade-offs, establishing best practices, and technical risk mitigation.

- **Developer mode** :

  This feature enables you to iterate infrastructure designs and deployments quickly within AMS-configured accounts[1] by allowing direct access to AWS service APIs and the AWS console in addition to access to the AMS change management process. Resources provisioned or configured with developer mode permissions outside of the change management process are your responsibility to manage (See "Automated and Self-Service Provisioning Management"). Resources provisioned through the AMS change management process are supported like other change management-provisioned workloads on AMS.

- **AWS support**:

AMS customers can choose the level of AWS Support they require to complement their AMS Operations plan. Accounts enrolled in AMS can be subscribed to either Business Support or Enterprise Support. To learn about the differences in Support Plans, see AWS Support Plans.

- **Customer-managed account**:

  This feature enables you to request AWS accounts within the same managed environment but the ongoing operations of workloads and AWS resources within those accounts are your responsibility. AMS provisions customer-managed accounts, but once the accounts are created, no other AMS features or services are provided to those accounts. AWS will not enroll customer-managed accounts in enterprise-level premium support. It will be your responsibility to enroll customer-managed accounts in AWS support at the support rate you choose.

- **Firewall management**:

  AMS provides an optional managed firewall solution for Supported Firewall Services, which enables internet-bound egress traffic filtering for networks in your managed environment. This excludes public-facing services that do not use the AWS network infrastructure and whose traffic goes directly to the internet. The solution combines industry-leading firewall technology with AMS infrastructure management capabilities to deploy, monitor, manage, scale, and restore the firewall infrastructure.

When you onboard AMS, you receive a complete list of your AMS network infrastructure. To get an updated list of services running in support of your AMS infrastructure at any time, file a service request with specifics about the information you want. To request a change to your network design, create a service request describing the changes you want to make—for example, adding a VPC or requesting a security group rule change.

What we do, what we do not do

AMS gives you a standardized approach to deploying AWS infrastructure and provides the necessary ongoing operational management. For a full description of roles, responsibilities, and supported services, see Service Description.

> **Note**
>
> To request that AMS provide an additional AWS service, file a service request. For more information, see Making Service Requests.

- **What we do**:

  After you complete onboarding, the AMS environment is available to receive requests for change (RFCs), incidents, and service requests. Your interaction with the AMS service revolves around the lifecycle of an application stack. New stacks are ordered from a preconfigured list of templates, launched into specific virtual private cloud (VPC) subnets, modified during their operational life through requests for change (RFCs), and monitored for events and incidents 24/7.

  Active application stacks are monitored and maintained by AMS, including patching, and require no further action for the life of the stack unless a change is required or the stack is decommissioned. Incidents detected by AMS that affect the health and function of the stack generate a notification and may or may not need your action to resolve or verify. How-to questions and other inquiries can be made by submitting a service request.

  Additionally, AMS allows you to enable compatible AWS services that are not managed by AMS. For information about AWS-AMS compatible services, see Self-service provisioning mode.

- **What we DON'T do**:

  While AMS simplifies application deployment by providing a number of manual and automated options, you're responsible for the development, testing, updating, and management of your application. AMS provides troubleshooting assistance for infrastructure issues that impact applications, but AMS can't access or validate your application configurations.

AMS responsibility matrix (RACI)

> **Note**
>
> In order to fulfill its obligations in a timely manner, AWS Managed Services (AMS) may require inputs from you for deciding an appropriate course of action. AMS will contact the designated customer contact for all such clarifications and inputs. AMS will expect a response to such queries within 24 business hours. In case there is no reply within 24 business hours, AMS may choose an action on your behalf.

The AMS responsible, accountable, consulted, and informed, or RACI, matrix assigns primary responsibility either to the customer or AMS for a variety of activities.

AMS manages your AWS infrastructure. The following table provides an overview of the responsibilities of customer and AMS for activities in the lifecycle of an application running within an AMS managed environment.

AMS is not responsible for any of the following activities for Customer Managed accounts or the infrastructure running within them; therefore this RACI is not applicable.

- **R** stands for responsible party that does the work to achieve the task.
- **C** stands for consulted; a party whose opinions are sought, typically as subject matter experts; and with whom there is bilateral communication.
- **I** stands for informed; a party which is informed on progress, often only on completion of the task or deliverable.
- **Self-service Provisioning** refers to resources that are provisioned by the customer with self-service through the AWS API or Console, including Developer Mode and Self-Service Provisioned Services.

  To provide self-service provisioning capabilities, AMS has created elevated IAM roles with permission boundaries to limit unintended changes from direct AWS service access. Roles do not prevent all changes and you are responsible to adhere to your internal controls, compliance, and to validate that all AWS services being used meet the required certifications. We call this the Self-Service Provisioning mode. For details on AWS compliance requirements, see AWS Compliance.

  For resources that you provision through self-service, AMS provides incident management, detective controls and guardrails, reporting, designated resources (Cloud Service Delivery Manager and Cloud Architect), Security & access, and technical support through service requests. Additionally, where applicable, you assume responsibility for continuity management, patch management, infrastructure monitoring, and change management for resources provisioned or configured outside of the AMS change management system.

Activity	Customer	AWS Managed Services (AMS)
Application lifecycle		
Application development	R	I

Activity	Customer	AWS Managed Services (AMS)
Application infrastructure requirements analysis and design	R	C
Design and optimization for non-standard AMS stacks	R	C
Design and optimization of AMS standard stack	I	R
Application deployment	R	C
AWS Infrastructure deployment	C	R
Application monitoring	R	I
Application testing/optimization	R	I
AWS infrastructure optimization guidance	I	R
AWS infrastructure monitoring	I	R
Troubleshoot and resolve application issues	R	C
Troubleshoot and resolve AWS network issues	C	R
Troubleshoot and resolve operating system and infrastructure issues	C	R
Self-Service Provisioning	*R*	*C*
Application and ITSM Integration		
Application integration with AWS Service Offerings	R	C
ITSM integration with the AWS Managed Services Interface	R	C

Activity	Customer	AWS Managed Services (AMS)
Networking		
Managed Environment VPC and VPC set-up and configuration	C	R
Allocate private address space for VPCs (e.g. /16)	R	C
Configure & Operate non-AWS Managed Services, Customer managed Firewalls/Proxy/ Bastions/HOSTs	R	C
Configure & Operate AWS Security Groups/ NAT/Customer Bastions/NACL inside the Managed Environment	I	R
Networking (e.g. DirectConnect) configuration and implementation within customer network	R	C
Networking configuration and implementation within the Managed Environment	C	R
Managed environment configuration		
Define default Auto Scaling settings for baseline Stack templates	I	R
Recommend RI optimization	C	R
Purchase RI and PIOP capacity	R	C
Remove capacity when capacity is over provisioned (when supported by customer application)	C	R
Create/update AWS customer specific information for AWS Managed Services	C	R

Activity	Customer	AWS Managed Services (AMS)
S3 configuration	C	R
Self-service provisioning	R	C
Glacier configuration	C	R
Define archival policy	R	C
Archival policy configuration	C	R
Selecting customer maintenance window	R	I
AWS RDS Management		
Monitor source/replica/RO replication health	I	R
Identify RCA of source failover	I	R
Automated snapshot (backup) configuration	C	R
Self-service provisioning	R	C
Coordinate and schedule DB engine patch management	C	R
Self-service provisioning	R	C
Recommend DB storage and PIOP capacity	C	R
Self-service provisioning	R	C
Recommend instance sizing for running databases	C	R
Self-service provisioning	R	C

Activity	Customer	AWS Managed Services (AMS)
Recommend RI optimization for Managed Environment	C	R
Self-service provisioning	R	C
RDS performance monitoring (CloudWatch)	C	R
Self-service provisioning	R	C
RDS event subscription configuration (SNS)	C	R
Self-service provisioning	R	C
RDS security group configuration	C	R
Self-service provisioning	R	C
RDS engine parameter/option configuration	R	C
DB table design	R	I
DB indexing	R	I
DB log analysis	R	I
AMS Change Management		
Creating customer RFCs (e.g. access to resources creating/updating/deleting managed stacks, deploying/updating applications, changes to configuration of AWS Service Offerings)	R	I
Approving Customer RFCs	I	R

Activity	Customer	AWS Managed Services (AMS)
Creating AWS Managed Services RFCs (e.g. access to resources, creating resources on customer's behalf, applying updates to OS as part of Patch Management)	I	R
Approving non-automated RFCs	R	I
Submitting request for new Change Types	R	C
Creating new Change Types	I	R
Maintenance of application change calendar	R	C
Notice of upcoming Maintenance Window	I	R
AWS Service Catalog		
Create portfolios and products	R	I
Distribute products to end users	R	I
Create tags and tag option library	R	C
Sharing portfolios and products with end users	R	I
Revise / update portfolios and products	R	I
Create and assign constraints to portfolios and products	R	C
Associate Service Actions to products	R	C
Update provisioned resources with new version of product	R	I
Provisioning		

Activity	Customer	AWS Managed Services (AMS)
Customer specific additions to AWS Managed Services baseline AMI	R	C
Configure additional approved Change Types used to provision Stack templates	C	R
Launch managed Stacks and associated AWS resources submitted through AMS change management process or AWS Service Catalog. *Self-service provisioning*	I	R
	R	I
Install/Update custom and 3rd party applicati ons on Instances provisioned through AMS change management process or AWS Service Catalog.	R	I
Provisioning - Stack Architecture		
Providing OS licenses (including usage fees for the applicable AWS services – e.g. EC2 and RDS) *Self-Service Provisioning*	I	R
	R	*I*
Define baseline infrastructure templates (Stacks) for application deployment through AMS change management system. *Self-Service Provisioning*	I	R
	R	*I*
Creating baseline approved AMIs[8]	I	R
Evaluate customer application inventory and determine fit with available infrastructure templates (Stacks)	R	C

Activity	Customer	AWS Managed Services (AMS)
Define unique Stacks that are in addition to the baseline template offerings	R	C
Logging, Monitoring and Event Management		
Recording AWS infrastructure change logs	I	R
Recording all application change logs	R	C
Installation and configuration of agents and scripts for patching, security, monitoring, etc. of AWS infrastructure provisioned through the AMS change management process. *Self-Service Provisioning*	I *R*	R *C*
Define customer specific monitoring and incident requirements	R	C
Configuring alerts for Managed Environment	I	R
Monitoring all AMS configured alerts *Self-Service Provisioning*	I *R*	R *C*
Investigating infrastructure Alerts for Incident notification *Self-Service Provisioning*	I *R*	R *C*
Investigating application alarms	R	C
Incident Management		
Proactively notify Incidents on AWS infrastru cture based on monitoring *Self-Service Provisioning*	I *R*	R *C*

Activity	Customer	AWS Managed Services (AMS)
Handle application performance issues and outages	R	I
Categorize Incident priority	I	R
Provide Incident response	I	R
Provide Incident resolution / infrastructure restore **Note** SLAs do not apply to instance-based resources provisioned outside AMS change management, including those provisioned using self-service provisioning and developer mode.	C	R
Problem Management		
Identify Problems in Managed Environment	C	R
Perform RCA for Problems in Managed Environment	C	R
Remediation of Problems in Managed Environment	C	R
Identify and remediate application problems	R	I
Security Management		

Activity	Customer	AWS Managed Services (AMS)
Customer infrastructure security and/or establishing baseline for security compliance process as determined and agreed to during customer onboarding. *Self-Service Provisioning*	C *R*	R *C*
Maintaining valid licenses for Managed EPS	R	C
Configure Managed EPS	I	R
Self-Service Provisioning	*R*	*C*
Update Managed EPS	I	R
Self-Service Provisioning	*R*	*C*
Monitoring malware on instances provisioned through the AMS CM process. *Self-Service Provisioning*	I *R*	R *C*
Maintaining and updating virus signatures.	I	R
Self-Service Provisioning	*R*	*C*
Remediating instances infected with malware.	C	R
Self-Service Provisioning	*R*	*C*
Security event management	C	R
Security - Access Management		
Manage the lifecycle of users, and their permissions for local directory services, which are used to access AWS Managed Services	R	I

Activity	Customer	AWS Managed Services (AMS)
Operate federated authentication system(s) for customer access to AWS console/APIs	R	C
Accept and maintain Active Directory (AD) trust from AWS Managed Services AD to customer managed AD	R	C
During onboarding, create cross-account IAM Admin roles within each managed account	R	C
Secure the AWS root credential for each account	I	R
Define IAM resources for Managed Environme nt	C	R
Manage privileged credentials for OS access for AMS engineers	I	R
Manage privileged credentials for OS access provided to customer by AMS	R	I
Patch Management[9]		
Monitor for applicable updates to supported OS and software preinstalled with supported OS for EC2 instances. *Self-Service Provisioning*	I *R*	R *C*
Notify customer of upcoming updates (*applies to AMS Standard Patch only*)	I	R
Exclude certain updates and/or certain Stacks from patching activities	R	I

Activity	Customer	AWS Managed Services (AMS)
Define default and custom maintenance windows schedules and other parameters (e.g. maintenance window duration) to apply patches (*applies to AMS Patch Orchestrator only*)	R	I
Define custom Patch Baselines to filter and exclude specific patches (*applies to AMS Patch Orchestrator only*)	R	I
Tag instances to associate them with custom maintenance windows and Patch Baselines (*applies to AMS Patch Orchestrator only*)	R	I
Track the patch status of resources and highlight systems that aren't current in the monthly business review.	C	R
Patch the Windows operating system, and Microsoft packages installed on the operating system which are governed by Windows Update	I	R
Self-Service Provisioning	*R*	-
Patch installed applications, software, or application dependencies not managed by Windows Update	R	I
Self-service provisioning	*R*	-

Activity	Customer	AWS Managed Services (AMS)
Patch the Linux operating system and any package that is enabled for management by the operating system's native package manager (for example Yum, Apt, Zypper)	I	R
Self-service provisioning	*R*	-
Patch installed applications, software, or application dependencies not managed by the Linux operating system's native package manager	R	I
Self-service provisioning	*R*	-
Continuity Management		
Specify backup schedules	R	I
Execute backups per schedule.	I	R
Self-Service Provisioning	*R*	*C*
Validate backups	R	I
Request backup restoration activities	R	I
Execute backup restoration activities.	I	R
Self-Service Provisioning	*R*	*C*
Restore affected Stacks and VPCs.	I	R
Self-Service Provisioning	*R*	*C*
Restore affected custom/3rd party application	R	C
Reporting		

Activity	Customer	AWS Managed Services (AMS)
Prepare and deliver monthly service report	I	R
AMS on AWS Outposts	R	I
Configure and retrieve API audit history on demand (CloudTrail).	I	R
Self-service provisioning	R	I
Provide access to incident history through AWS Managed Services Interface	I	R
Provide access to change history through AWS Managed Services Interface.	I	R
Self-service provisioning	N/A	N/A
Service Request Management		
Request information using service requests	R	I
Reply to service requests	I	R
Managed Firewall		
Request the deployment of AMS-Managed Firewall	R	I
Design and optimization of AMS-Managed Firewall architecture	I	R
Deployment of AWS Infrastructure and AMS-Managed Firewall appliance	I	R
Providing Firewall licenses (including usage fees for the applicable AWS services – e.g. EC2)	R	I

Activity	Customer	AWS Managed Services (AMS)
Define default domain allow-list	I	R
Request to add, modify, and delete custom allow-lists and security policies	R	I
Configuring alerts for AMS-Managed Firewall	I	R
Monitoring all AMS-Managed Firewall configured alerts	I	R
Execute Backups of firewall configuration	I	R
Request backup restoration activities	R	I
Update provisioned resources with new version of product	I	R
Recording AMS-Managed Firewall logs	I	R
Forward logs from AMS-Managed Firewall to CloudWatch	I	R
Request configuration changes in the AMS-Managed Firewall	R	I
Approve configuration changes in the AMS-Managed Firewall	I	R
Execute configuration changes in the AMS-Managed Firewall	I	R

[8]AMS provides AMIs for Amazon EC2 only

[9]AMS is responsible for End of Life OSes only when the customer signs an extended support agreement with OS vendor

AMS environment basic components

Multi-Account Landing Zone

This is an estimate of the components, and potential costs, of the infrastructure in the core accounts. This does not include other costs such as bandwidth, CloudWatch detailed monitoring, logging, alarms, Route53, Amazon S3, Simple Notification Service (Amazon SNS), snapshots, or reserved Amazon EC2 instances.

You pay for the components required by the AMS-Managed AWS landing zone infrastructure. Estimates place the cost of a plain AMS multi-account landing zone environment at $2,450 per month and $50 for a plain application account.

For information about pricing, see AWS pricing.

Basic Environment Components

Component	Est. Cost	Description
Management account	$60	An AWS Organizations Management account; creates and financially manages member accounts. It contains the AWS Landing Zone (ALZ) framework, account configuration stack sets, and AWS Organization service control policies (SCPs). • Directory Service: $35 • CloudTrail: $7 • CloudWatch: $6 • Others: $12
Shared Services Account	$2000	Contains infrastructure and resources required for access management (i.e., Active Directory), end-point security management (Trend Micro), and your bastions (SSH/RDP); estimate is $2400 a month. This estimate does not include the cost of the Trend Micro licenses. • EC2: $800 (with the minimum number of Bastions) • RDS: $300 (EPS) • VPC (endpoints): $400

Component	Est. Cost	Description
		• Directory Service: $300 • CloudWatch: $100 • GuardDuty : $15 • Secrets Manager: $10 • Data Transfer: $10 • Config: $10 • Others: $45
Networking Account	$350	The central hub for network routing between AMS accounts, your on-premise network, and egress traffic to the Internet. Additionally, contains public DMZ bastions (the entry point for AMS engineers to access hosts in your AMS environment). Price may increase depending on traffic traversing the Transit Gateway and Direct Connect. • EC2: $250 (Bastions) • VPC: $80 • Others: $20
Log Archive Account	$20	An S3 bucket with copies of AWS CloudTrail and AWS Config log files from each of your AMS environment accounts. Costs increase as more logs are collected. • S3: $10 • CloudWatch: $5 • Others: $5

Component	Est. Cost	Description
Security Account	$20	The central hub for security related operations, and the main point for funneling notifications and alerts to AMS control plane services. Additionally, houses the Amazon Guard Duty management account. Costs increase as more events are analyzed using Amazon GuardDuty. • CloudWatch: $15 • Others: $5

Single-Account Landing Zone

The following table lists the components of an example AMS-managed infrastructure.

Basic Environment Components, Last Updated 2020/07/09

Name	Instance Type	OS	# of Components
mc-eps-dsm	m5.large	Linux	2
mc-management	m5.large	Windows	2
mc-bastion-dmz-ssh	m5.large	Linux	2
mc-bastion-customer-rdp	m5.large	Windows	2
mc-eps-relay	m5.large	Linux	2
directory services	N/A	N/A	
additional components	N/A	N/A	

For information about pricing, see AWS Pricing.

AMS account limits

There are three distinct types of limits to consider within AMS multi-account landing zone: AMS API limits, AMS resource limits, and AWS limits.

There are two distinct types of limits to consider within AMS single-account landing zone: AMS API limits, and AWS limits.

AMS account API limits

This section describes the account level limits after which AWS Managed Services (AMS) throttles the AMS SKMS API service. This means, if you call any of the listed APIs more than 10 times in a second, one of the calls is "throttled" (you receive a `ThrottleException`). Under rare situations, an external or downstream dependency might throttle the AMS API and then AMS may throttle your API calls at a possibly lower rate.

> **Note**
>
> For information on the AMS SKMS API, download the reference through the **Reports** tab of the AWS Artifact console.

For each AMS SKMS API listed, the operation is throttled after 10 TPS (transactions per second):

- `GetStack`
- `GetSubnet`
- `GetVpc`
- `ListAmis`
- `ListStackSummaries`
- `ListSubnetSummaries`
- `ListVpcSummaries`

AMS multi-account landing zone account resource limits

Account resource limits relate to AMS multi-account landing zone application accounts and VPCs and subnets.

Application account resource limits

There is a soft limit of 50 application accounts per organization. If you have a use case for more than 50 application accounts, contact your cloud service delivery manager (CSDM) to relay your requirements.

VPCs and subnets resource limits

There is a soft limit of 10 VPCs per application account within the pre-defined AWS Region for the organization.

Each VPC may have 1 to 10 private subnet tiers spanned across 2 to 3 availability zones. Additionally, each VPC may have 0 to 5 public subnet tiers spanned across 2 to 3 availability zones. If you have requirements beyond these limits, inform your CSDM or Cloud Architect to review your use case.

AMS multi-account landing zone application to account ratio

One account per application is supported in AMS multi-account landing zone; however, each Application account has a small cost, and you are charged for the number of connections to the Transit Gateway per hour, and the amount of traffic that flows through AWS Transit Gateway. So, the more segregated applications are into accounts or VPCs, the higher the costs.

To reduce costs and still ensure an appropriate segregation of duties, AMS recommends that you 1) group applications by teams with tightly coupled business processes, and 2) do not mix applications that are in different stages (prod vs. non-prod) or managed by different teams. In this way, you will have fewer accounts, access management and the segregation of duties will be easier, and traffic cost could be mitigated.

For example: An enterprise has in production a Trading application and a Portfolio Management application, both applications are managed by the Investments IT team and exchange a lot of traffic with each other. In this scenario the company can benefit from grouping both applications in the same account and same the VPC, because the Investments IT team won't have to request access to multiple Application accounts and the company will save on traffic costs. In this case, the company should create another account for the same applications in development stage and provide access to the development team.

In another scenario, the enterprise has in production a Payroll application and an Accounting application, managed by the Human Resources IT and Accounting IT teams respectively. Although

the Payroll application has to exchange information with the Accounting application, we recommend segregating both applications in different accounts, one per team, and establishing a connection between both application's VPCs using the Networking account. In this way, the company will prevent HR IT team request changes affecting the accounting application infrastructure, of which they would have no knowledge.

Tips on how to group accounts into organizational units (OUs). An OU is logical grouping mechanism that enables you to categorize (group) accounts and apply policies and configurations to based on those groups. The recommended approach for creating OUs is to base them on policies that need to be applied to a specific group of accounts, not on the internal hierarchy of teams within your reporting structure. An OU is not equivalent to an Active Directory's OU, and attempting to replicate the AD OU structure in AWS Orgs is discouraged and results in a difficult to maintain and/or operate structure.

AWS account limits

AWS account limits apply to your AWS Managed Services (AMS) accounts. The easiest method to determine default and current limits for AWS services is by leveraging AWS Service Quotas. AMS recommends right-sizing individual service limits to the appropriate size to run the service(s) in the account. Limits act like guard-rails to protect your accounts for security and cost runaways. If you would like to raise a specific limit, submit a service request with AMS, and AMS Operations will raise the limit on your behalf. For example, the default limit (or quota) for RDS instances is 40; if your workload requires 50 RDS instances, raise a service request for AMS Operations to raise the limit to your needed value.

AMS service level objectives (SLOs)

The following table describes the goals of the AWS Managed Services (AMS) service. Service Level Agreements (SLAs) for other aspects of the AMS service, including incident management, are covered in the SLA document shared with you when you subscribed to AMS. For more information, speak to your CSDM.

AMS Service Level Objectives

Feature	Performance Indicator (PI)	Plus (Business Days, M-F 8AM to 6PM local time)	Premium (Calendar Days, 24 x 7)
Change management	Time taken to schedule or reject automated RFCs	<=30 min	<=30 min
	Time of initiation of scheduled RFCs compared to scheduled execution time	<=1 min	<=1 min
	Time taken to approve/reject non-automated RFCs, available in CT catalog	<=48 hours	<=24 hours
	Time taken to approve/reject non-automated RFCs not available in CT catalog	<=5 days	<=5 days
Problem management	Time taken to complete root cause analysis (RCA)	<=10 days	<=10 days
Service request management	Response time for first and every subsequent reply	<=8 hours	<=4 hours

Supported AWS services

AWS Managed Services (AMS) provides operational management support services for the following AWS services. Each AWS service is distinct, and as a result AMS's level of operational management, support varies depending on the nature and characteristics of the underlying AWS service. Specific AWS services are grouped based on the complexity and scope of the operational management support service provided by AMS.

Note

The three groups, A, B, and C, indicate pricing as a percentage of total monthly spend per account for the AMS service, based on support plan (Plus or Premium), for AMS customers before March 16, 2021. AMS customers onboarded after March 16, 2021 should submit a service request for additional pricing information. Group A indicates no additional charge. Group B indicates an additional charge of 12% (Plus) or 18% (Premium). Group C indicates an additional charge of 25% (Plus) or 42% (Premium).

One star (*) indicates services that are deployed within an AMS managed environment by a customer using the AWS Console and APIs. See 'Automated and self-service provisioning management' in AMS features for additional details on customer responsibilities when provisioning and configuring services in this manner.

Two stars (**) indicate that Amazon EC2 on AWS Outposts will be billed as a Group B service; all other resources hosted on AWS Outposts will be billed at their standard rate.

Supported AWS services

Group A	Group B	Group C
Amazon Alexa for Business*	Amazon API Gateway*	Amazon Aurora
Amazon Managed Streaming for Apache Kafka*	Amazon AppStream*	Amazon CloudWatch
Amazon CloudFront	Amazon Athena*	Amazon Elastic Block Store (EBS)
Amazon Elastic File System	Amazon CloudSearch*	Amazon Elastic Compute Cloud**
Amazon Glacier	Amazon Cognito*	Amazon Elastic Load Balancing (classic, application, and network; not gateway)
Amazon Simple Storage Service	Amazon Comprehend*	Amazon ElastiCache
AWS Amplify*	Amazon Connect*	
	Amazon Document DB (with MongoDB compatibility)*	
	Amazon DynamoDB*	

Group A	Group B	Group C
AWS AppMesh*	Amazon EC2 Container Registry (ECR)*	Amazon OpenSearch Service
AWS Auto Scaling	Amazon Elastic Container Service (ECS) on AWS Fargate*	Amazon GuardDuty
AWS Backup	Amazon Elastic Kubernetes Service (EKS) on Fargate*	Amazon Macie
AWS CloudFormation	Amazon Elemental MediaConvert*	Amazon Redshift
AWS Compute Optimizer	Amazon Elemental MediaPackage*	Amazon Relational Database Service
AWS Global Accelerator*	Amazon Elemental MediaStore*	Amazon Route 53
AWS Identity and Access Management	Amazon Elemental MediaTailor*	Amazon Simple Email Service
AWS License Manager*	Amazon Elastic MapReduce*	Amazon Simple Notification Service
AWS Management Console	AmazonEventBridge*	Amazon Simple Queue Service
AWS Marketplace	Amazon Forecast*	Amazon Virtual Private Cloud (VPC)
AWS Lake Formation*	Amazon FSx*	AWS CloudTrail
AWS Well Architected Tool*	Amazon Inspector*	AWS Config
VM Import/ Export*	Amazon Kinesis Analytics*	AWS Database Migration Service
	Amazon Kinesis Data Stream*	AWS Data Transfer
	Amazon Kinesis Firehose*	AWS Direct Connect
	Amazon Kinesis Video Streams*	AWS Directory Service
	Amazon Lex*	AWS Key Management Service
	Amazon MQ*	AWS Systems Manager (SSM)
	Amazon Personalize**	
	Amazon QuickSight*	
	Amazon Rekognition*	
	Amazon SageMaker*	
	Amazon SimpleDB*	
	Amazon Simple Workflow*	
	Amazon Textract*	
	Amazon Transcribe*	
	Amazon Translate*	
	Amazon WorkDocs*	

Group A	Group B	Group C
	Amazon WorkSpaces* AWS AppSync* AWS Audit Manager* AWS Batch* AWS Certificate Manager* AWS CloudEndure* AWS CloudHSM* AWS CodeBuild* AWS CodeCommit* AWS CodeDeploy* AWS CodePipeline* AWS DataSync* AWS Elemental MediaLive* AWS Glue* AWS Lambda* AWS MigrationHub* AWS Outposts** AWS Secrets Manager* AWS Security Hub* AWS Service Catalog AWS Transfer for SFTP* AWS Shield* AWS Snowball* AWS Step Functions* AWS Transit Gateway* AWS WAF* AWS X-Ray*	

If you request AWS Managed Services to provide services for any software or service that is not expressly identified as supported below, any AWS Managed Services provided for such customer requested configurations will be treated as a "Beta Service" under the Service Terms.

Supported configurations

These are the configurations AWS Managed Services (AMS) supports:

- Language: AMS is available in English.
- Firewall Services: Palo Alto VM-Series Next-Generation Firewall

- Security software: Deep Security from Trend Micro (Required). AWS Marketplace: Trend Micro Deep Security
- Approved directory services: Microsoft Active Directory (AD)
- Supported AWS services.
- Supported AWS Regions:

  AMS operates in a subset of all AWS Regions; however, the AMS API/CLI runs out of the "USA East (N. Virginia)" Region only. If you run either the AMS change management API (`amscm`) or the AMS service knowledge management API (`amsskms`), in a non-USA East Region, you must add `--region us-east-1` to the command.

  - US East (Virginia)
  - US West (N. California)
  - US West (Oregon)
  - US East (Ohio)
  - Canada (Central)
  - South America (São Paulo)
  - EU (Ireland)
  - EU (Frankfurt)
  - EU (London)
  - EU West (Paris)
  - Asia Pacific (Mumbai)
  - Asia Pacific (Seoul)
  - Asia Pacific (Singapore)
  - Asia Pacific (Sydney)
  - Asia Pacific (Tokyo)

  New Regions are added frequently. To learn more about AWS Regions and availability zones, see Regions and Zones.
- Amazon machine images (AMIs): AMS provides security enhanced images (AMIs) based on CIS Level 1 benchmark for a subset of AMS's supported operating systems. To find out which operating systems have a security enhanced image available, see the AMS Security User Guide, which is available through AWS Artifact -> Reports tab filtered for AWS Managed Services. To

access AWS Artifact, can contact your CSDM for instructions or go to Getting Started with AWS Artifact.

- Supported operating system architecture: x86-64

  ARM64 architecture is not supported.

- Supported operating systems:
  - Amazon Linux 2023
  - Amazon Linux 2 (**expected AMS support end date June 30, 2025**)
  - CentOS 7.x
  - Oracle Linux 7: minor versions 7.5 and above. Oracle Linux 8: minor versions up to 8.3.
  - Red Hat Enterprise Linux (RHEL) 8.x, 7.x
  - SUSE Linux Enterprise Server 15 SP5 and SAP specific versions, SUSE Linux Enterprise Server 12 SP5 and SAP specific versions.
  - Microsoft Windows Server 2022, 2019, 2016
- Supported End of Support (EOS) operating systems:

> **Note**
>
> End of Support (EOS) operating systems are outside of the general support period of the operating system manufacturer and have increased security risk. EOS operating systems are considered supported configurations only if AMS-required agents support the operating system and...
>
> 1. you have extended support with the operating system vendor that allows you to receive updates, or
> 2. any instances using an EOS operating system follow the security controls as specified by AMS in the Advanced User Guide, or
> 3. you comply with any other compensating security controls required by AMS.
>
> In the event AMS is no longer able to support an EOS operating system, AMS issues a Critical Recommendation to upgrade the operating system.
> AMS-required agents may include but are not limited to: AWS Systems Manager, Amazon CloudWatch, Endpoint Security (EPS) agent, and Active Directory (AD) Bridge (Linux only).

  - Amazon Linux (**expected AMS support end date December 31, 2023**)

- Ubuntu Linux 18.04
- SUSE Linux Enterprise Server 15 SP3 and SP4
- Microsoft Windows Server 2012/2012 R2

Capabilities for unsupported operating systems in AMS

An *unsupported* operating system is any operating system not listed in the Supported configurations. AMS considers instances with unsupported operating systems to be "Customer-Requested Configurations" that are subject to the AWS Betas and Previews service terms.

The following limited set of AMS capabilities are available to instances with unsupported operating systems:

Capability	Notes
Incident management	AMS provides incident response.
Service request management	AMS responds to service requests.
Requests for change (RFCs)	AMS evaluates RFCs for execution. Unsupported operating systems may impact the ability to execute RFCs.
Monitoring	AMS monitors and responds to Amazon EC2 system status checks and instance status checks. System status checks include: loss of network connectivity, loss of system power, software issues on the physical host, and hardware issues on the physical host that impact network reachability. Instance status checks include: incorrect networking or startup configuration, exhausted memory, corrupted file system, and incompatible kernel.
Security management	AMS monitors and responds to Amazon EC2 GuardDuty findings.

Capability	Notes
Backup management	AMS provides Continuity management in AMS Advanced for EC2 using AMS-customized AWS Backup plans and vaults.

AMS Advanced interfaces

- *AMS Consoles*: AMS has a console for each of the operations plans: AMS Advanced and AMS Accelerate. Each are available through the AWS Management Console, once you have an account with that operations plan.

  You use the AMS Accelerate console to view summaries all your current incidents and service requests, and resource security status including compliance and real-time threat detection, and to quickly access configuration panels.

  You use the AMS console to create RFCs, report and respond to incidents, make service requests, and find information on existing VPCs and stacks. When in doubt of what to do, or when you need help with AMS or your managed resources, create a service request by using this interface.
- *AWS Management Console*: Many AWS consoles can be useful for viewing AMS information, for example:
  - *Amazon EC2 console*: Use to view instance information including bastion IP addresses, Amazon EC2 Auto Scaling groups, and load balancers.
  - *Multi-Account Landing Zone AWS Config Rules compliance*: You can view compliance status across your accounts and identify non-compliant resources.
  - *AWS CloudFormation console*: Use to view stack information including stack IDs (you can find Amazon RDS stacks and Amazon RDS instance IDs here, and event information).
  - *Amazon RDS console*: Use to view event information such as a post made to a WordPress app on a site in your account. Note you must have the Amazon RDS instance ID.

  Depending on the mode of your login role, you have different level of access to the AWS Management Console. For more information on modes, see AMS modes.
- *AMS change management API* – Read/Write: Use the change management API (CM API) to request additions and specific changes to your managed infrastructure including resource monitoring, log, backup, and patch configurations. Also, use this API to request access to

resources, delete resources, create AMIs, and create IAM instance profiles. You can access the CM API through the AMS CLI and SDKs.

- *AMS SKMS API* – Read-Only: Use this API to list managed resources and get information needed for reporting or preparing requests for change.
- *AWS Support API*: Use the standard AWS Support API to programmatically create and respond to incidents and service requests. To learn more, see Getting Started with AWS Support.
- *AWS APIs* – Read Only: Your main IT administrator can use the AWS APIs to see all resources under management, view CloudTrail logs, billing information, and many other read functions.

AMS VPC endpoints

A VPC endpoint lets you privately connect your VPC to AWS services without requiring an Internet gateway. Instances in your VPC do not require public IP addresses to communicate with resources in the service.

Endpoints are virtual devices. They are horizontally scaled, redundant, and highly available VPC components that allow communication between instances in your VPC and services without imposing availability risks or bandwidth constraints on your network traffic. To learn more, see VPC Endpoints.

There are two types of VPC endpoints: interface endpoints and gateway endpoints.

- Gateway endpoints: The VPC in the account has an Amazon S3 Gateway endpoint enabled by default.
- Interface endpoints: Instances in your AMS environment can talk to supported services without leaving the Amazon network. This is optional for **single-account landing zone** and it is not enabled in the account by default; submit a service request to AMS operations to get this enabled. However, for **multi-account landing zone**, interface endpoints are enabled by default in the Shared Services account.

  List of interface endpoints supported by AMS:

  - AWS CloudFormation
  - AWS CloudTrail
  - AWS Config
  - Amazon EC2 API
  - AWS Key Management Service

- Amazon CloudWatch
- Amazon CloudWatch Events
- Amazon CloudWatch Logs
- AWS Secrets Manager
- Amazon SNS
- AWS Systems Manager
- AWS Security Token Service

AMS protected namespaces

The list of protected namespaces for AWS Managed Services (AMS). When you work with AWS resources, prevent conflict with AMS by not using these namespaces. For details on other AWS service namespaces, see Amazon Resource Names (ARNs) and AWS Service Namespaces.

- `ams-*` (this is the preferred naming standard for new resources)
- `AWSManagedServices*` (this is the preferred naming standard for resources where CamelCase is appropriate)
- `/ams/*` (this is the preferred naming standard for path-based resources)
- `ams*` and `AMS*` and `Ams*`
- `sentinel*` and `Sentinel*`
- `Managed_Services*`
- `NewAMS*`
- `AWS_*` and `aws*`
- `VPC_*`
- `CloudTrail*` and `Cloudtrail*`
- `*/aws_reserved/*`
- `INGEST*`
- `EPSDB*`
- `MMS*`
- `TemplateId*`
- `StackSet-ams*`
- `StackSet-AWS-Landing-Zone`

- IAMPolicy*
- customer-mc-*
- Root*
- LandingZone*
- StateMachine*
- codedeploy_service_role
- managementhost
- sentinel.int.
- eps
- UnhealthyInServiceBastion
- ms-

AMS reserved prefixes

AMS resource attributes must comply with certain patterns; for example, IAM instance profile names, BackupVault names, tag names, and so forth, must not start with AMS reserved prefixes. Those reserved prefixes are:

```
*/aws_reserved/*
ams-*
/ams/*
ams*
AMS*
Ams*
aws*
AWS*
AWS_*
AWSManagedServices*
codedeploy_service_role
CloudTrail*
Cloudtrail*
customer-mc-*
eps
EPSDB*
IAMPolicy*
INGEST*
LandingZone*
Managed_Services*
```

```
managementhost
mc*
MC*
Mc*
MMS*
ms-
NewAMS*
Root*
sentinel*
Sentinel*
sentinel.int.
StackSet-ams*
StackSet-AWS-Landing-Zone
StateMachine*
TemplateId*
VPC_*
UnhealthyInServiceBastion
```

AMS maintenance window

The AWS Managed Services Maintenance Window (or Maintenance Window) performs maintenance activities for AWS Managed Services (AMS) and recurs the second Thursday of every month from 3 PM to 4 PM Pacific Time. AMS may change the maintenance window with 48 hours notice. This is for AWS Managed Services (AMS); to perform maintenance activities for managed infrastructures, such as deploying new AMS AMIs.

Your maintenance window is when AMS will apply patching and you determine your maintenance window at onboarding. You can also agree to the proposed patching window provided in your patching service notification, or suggest a different window.

For guidance on creating a maintenance window, see Maintenance Window.

AMS information resources

AMS provides several information resources to help you succeed.

- **AMS Accelerate User Guide**: Helps you understand the components and features that AMS Accelerate provides and how to use them. Look here for AMS Accelerate background information and details on default settings, finding resources, and how-to examples. HTML index, PDF

- **AMS Advanced User Guide**: Helps you understand the components and features that AMS Advanced provides and how to use them. Look here for AMS Advanced background information and details on default settings, finding resources, and how-to examples. HTML index, PDF
- **AMS Advanced Application Guide**: Describes the steps for deploying applications to AWS Managed Services infrastructure. Look here for information on application deployment and maintenance methodologies and considerations. HTML index, PDF.
- **AMS Advanced Onboarding Guide**: Describes the initial steps for creating the basic AWS Managed Services multi-account, or single-account, landing zone infrastructure in an AMS account. Look here for information on AMS account basics, validation, and questions to prepare you for onboarding to AMS. HTML index, PDF.
- **AMS Advanced Change Type Reference**: Provides reference material on the current change types that AWS Managed Services provides, including change type schemas and example walkthroughs for each change type and tips. Includes general information about change types Helps you understand all aspects of requests for change (RFCs) and AMS change types (CTs). Look here for specifics on change types, including links to relevant information. HTML index, PDF.
- **AMS CM (change management) API Reference**: Describes the AWS Managed Services CM API, which provides operations for creating and monitoring change requests and provides information about your resources that are managed by Managed Services. HTML index, PDF.
- **AMS Security Guides**: Describe proprietary AMS security information.

  Private; available on the AMS **Reports** tab in the AWS Artifact Console.
- **AMS Developer's Resources**: Access to the AMS CLI and SDK, for both amscm and amsskms. See https://console.aws.amazon.com/managedservices/.
- **AMS YouTube Videos**: Key customer operations explained in video. See AWS Managed Services YouTube Instructional Videos.
- **AMS Blog posts**: Specialty information on AWS Managed Services. See AWS Blogs.

AMS compliance

AMS has undergone auditing for the following standards and is eligible for use as part of solutions for which you must obtain compliance certification.

AMS Supported Compliance Standards

AMS supports AWS compliance standards. To learn more about AWS compliance programs, see AWS Compliance.

These are the current compliance standards supported by AMS.

FedRAMP: The US Federal Government is dedicated to delivering its services to the American people in the most innovative, secure, and cost-efficient fashion. Cloud computing plays a key part in how the federal government can achieve operational efficiencies and innovate on demand to advance their mission across the nation. That is why many federal agencies today are using AWS cloud services to process, store, and transmit federal government data.

For more information, see FedRAMP.

HIPAA: AWS has expanded its Health Insurance Portability and Accountability Act (HIPAA) compliance program to include AMS as a HIPAA Eligible Service. If you have a Business Associate Agreement (BAA) with AWS, you can use AMS to help build your HIPAA-compliant applications.

See HIPAA-focused whitepaper to learn how to leverage AMS for the processing and storage of health information. For more information, see HIPAA Compliance.

HITRUST: The Health Information Trust Alliance Common Security Framework (HITRUST CSF) leverages nationally and internationally accepted standards and regulations such as GDPR, ISO, NIST, PCI, and HIPAA to create a comprehensive set of baseline security and privacy controls.

For more information, see HITRUST CSF.

	ISO 27001: ISO/IEC 27001:2013 is a security management standard that specifies security management best practices and comprehensive security controls following the ISO/IEC 27002 best practice guidance. The basis of this certification is the development and implementation of a rigorous security program, which includes the development and implementation of an Information Security Management System (ISMS) which defines how AWS perpetually manages security in a holistic, comprehensive manner. For more information, see ISO/IEC 27001:2013.
	ISO 27017: ISO/IEC 27017:2015 provides guidance on the information security aspects of cloud computing, recommend ing the implementation of cloud-specific information security controls that supplement the guidance of the ISO/IEC 27002 and ISO/IEC 27001 standards. This code of practice provides additional information security controls implementation guidance specific to cloud service providers. For more information, see ISO/IEC 27017:2015 Compliance.
	ISO 27018: ISO/IEC 27018:2019 is a code of practice that focuses on protection of personal data in the cloud. It is based on ISO/IEC information security standard 27002 and provides implementation guidance on ISO/IEC 27002 controls applicabl e to public cloud Personally Identifiable Information (PII). It also provides a set of additional controls and associate d guidance intended to address public cloud PII protection requirements not addressed by the existing ISO/IEC 27002 control set. For more information, see ISO/IEC 27018:2019 Compliance.

ISO 9001: ISO 9001:2015 outlines a process-oriented approach to documenting and reviewing the structure, responsibilities, and procedures required to achieve effective quality management within an organization. Specific sections of the standard contain information on topics such as:

- Requirements for a quality management system, including documentation of a quality manual, document control, and determining process interactions
- Responsibilities of management
- Management of resources, including human resources and an organization's work environment
- Service development, including the steps from design to delivery
- Customer satisfaction
- Measurement, analysis, and improvement of the QMS through activities like internal audits and corrective and preventive actions

For more information, see ISO 9001:2015 Compliance.

PCI: AMS has an Attestation of Compliance for Payment Card Industry (PCI) Data Security Standard (DSS) version 3.2 at Service Provider Level 1. Customers who use AWS products and services to store, process, or transmit cardholder data can use AMS as they manage their own PCI DSS compliance certification.

For more information about PCI DSS, including how to request a copy of the AWS PCI Compliance Package, see PCI DSS Level 1. Importantly, you must configure fine-grained password policies in AMS to be consistent with PCI DSS version 3.2 standards. For details on which policies must be enforced, see Enable PCI Compliance for Your AWS Microsoft AD Directory.

SOC: AMS System & Organization Control (SOC) Reports are independent, third-party examination reports that demonstrate how AMS achieves key compliance controls and objectives. The purpose of these reports is to help you and your auditors understand the AMS controls established to support operations and compliance. There are three types of AMS SOC reports:

- AWS SOC 1 Report - Download with AWS Artifact
- AWS SOC 2: Security, Availability, & Confidentiality Report - Download with AWS Artifact
- AWS SOC 3: Security, Availability, & Confidentiality Report

For more information, see SOC Compliance.

Shared Responsibility

Security, including PCI compliance, is a shared responsibility. It is important to understand that AMS compliance status does not automatically apply to applications that you run in the AWS Cloud. You need to ensure that your use of AWS services complies with the standards. For more details on how AMS works together with customers across specific activities, see the AMS AMS responsibility matrix (RACI).

AMS Amazon Machine Images (AMIs)

AMS produces updated Amazon Machine Images (AMIs) every month for AMS supported operating systems. In addition, AMS also produces security enhanced images (AMIs) based on CIS Level 1 benchmark for a subset of AMS's supported operating systems. To find out which operating systems have a security enhanced image available, see the AMS Security User Guide, which is available through AWS Artifact -> Reports page (find the **Reports** option in the left navigation pane) filtered for AWS Managed Services. To access AWS Artifact, can contact your CSDM for instructions or go to Getting Started with AWS Artifact.

To receive alerts when new AMS AMIs are released, you can subscribe to an Amazon Simple Notification Service (Amazon SNS) notification topic called "AMS AMI". For details, see AMS AMI notifications with SNS.

The AMS AMI naming convention is: `customer-ams-<operating system>-<release date> - <version>`. (for example, `customer-ams-rhel6-2018.11-3`)

Only use AMS AMIs that start with `customer`.

AMS recommends always using the most recent AMI. You can find the most recent AMIs by either:

- Looking in the AMS console, on the **AMIs** page.
- Viewing the latest AMS AMI CSV file, available from your CSDM or through this ZIP file: AMS 02.2024 AMI contents and CSV file in a ZIP.

  For past AMI ZIP files, see the Doc History.
- Running this AMS SKMS command (AMS SKMS SDK required):

```
aws amsskms list-amis --vpc-id VPC_ID --query "Amis.sort_by(@,&Name)[?
 starts_with(Name,'customer')].[Name,AmiId,CreationTime]" --output table
```

AMS AMI content added to base AWS AMIs, by operating system (OS)

- Linux AMIs:
  - AWS CLI Tools
  - NTP
  - Trend Micro Endpoint Protection Service Agent
  - Code Deploy
  - PBIS / Beyond Trust AD Bridge
  - SSM Agent
  - Yum Upgrade for critical patches
  - AMS custom scripts / management software (controlling boot, AD join, monitoring, security, and logging)
- Windows Server AMIs:
  - Microsoft .NET Framework 4.5
  - PowerShell 5.1
  - AWS Tools for Windows PowerShell
  - AMS PowerShell Modules controlling boot, AD join, monitoring, security, and logging
  - Trend Micro Endpoint Protection Service Agent

- SSM Agent
- CloudWatch Agent
- EC2Config service (through Windows Server 2012 R2)
- EC2Launch (Windows Server 2016 and Windows Server 2019)
- EC2LaunchV2 (Windows Server 2022 and later)

Linux-based AMIs:

- Amazon Linux 2023 (Latest Minor Release)
- Amazon Linux 2 (Latest Minor Release)
- Amazon Linux (Latest 2018.03 Release)
- Red Hat Enterprise 7 (Latest Minor Release)
- Red Hat Enterprise 8 (Latest Minor Release)
- SUSE Linux Enterprise Server 12 SP5
- SUSE Linux Enterprise Server 15 SP3
- SUSE Linux Enterprise Server 15 SP4
- Ubuntu Linux 18.04
- Ubuntu Linux 20.04
- Ubuntu Linux 22.04
- CentOS 7 (Latest Minor Release)

> **Note**
>
> To use the CentOS AMIs, you must opt in to the no cost Cent OS license from the AWS Marketplace. To do this, go to AWS Marketplace and follow the instructions for opting in.

- Amazon Linux: For product overview, pricing information, usage information, and support information, see Amazon Linux AMI (HVM / 64-bit) and Amazon Linux 2.

  For more information, see Amazon Linux 2 FAQs.
- RedHat Enterprise Linux (RHEL): For product overview, pricing information, usage information, and support information, see Red Hat Enterprise Linux (RHEL) 7 (HVM).

- Ubuntu Linux 18.04: For product overview, pricing information, usage information, and support information, see Ubuntu 18.04 LTS - Bionic.
- CentOS: To use the CentOS AMIs, you must opt in to the no cost Cent OS license from the AWS Marketplace. To do this, go to AWS Marketplace and follow the instructions for opting, or re-opting, in. You do not incur software charges for using this product, but you are responsible for other AWS charges, including EC2 usage.
- SUSE Linux Enterprise Server for SAP applications 15:
  - Run the following steps once per account:
    1. Navigate to the **AWS Marketplace**.
    2. Search for the respective SUSE 15 SAP product. We currently support:
       - SUSE Linux Enterprise Server for SAP Applications 15
       - SUSE Linux Enterprise Server for SAP Applications 15 SP1
    3. Click **Continue to subscribe**.
    4. Click **Accept terms**.
  - Complete the following steps **every time** you need to launch a new **SUSE Linux Enterprise Server for SAP Applications 15** instance:
    1. Note the AMI ID for the subscribed **SUSE Linux Enterprise Server for SAP Applications 15** AMI.
    2. Create a manual (Management | Other | Other | Create) RFC with the following wording; replace *`AMI ID`* with the AWS Marketplace AMI ID you have subscribed to.

Windows-based AMIs:

Microsoft Windows Server (2012, 2012 R2, 2016, 2019 and 2022), based on latest Windows AMIs, (for example: Windows_Server-2012-R2_RTM-English-64Bit-Base-*).

Additional information: For product overview, pricing, usage, and support, see Microsoft Windows Server 2012 RTM.

For examples of creating AMIs, see Create AMI.

Offboarding AMS AMIs:

AMS does not unshare any AMIs from you during offboarding to avoid impact for any of your depedencies. If you want to remove AMS AMIs from your account, you can use the `cancel-`

`image-launch-permission` API to hide specific AMIs. For example, you can use the script below to hide all of the AMS AMIs that were shared with your account earlier:

```
for ami in $(aws ec2 describe-images --executable-users self --owners 027415890775 --
query 'Images[].ImageId' --output text) ;
    do
    aws ec2 cancel-image-launch-permission --image-id $ami ;
    done
```

You must have the AWS CLI v2 installed for the script to execute without any errors. For AWS CLI installation steps, see Installing or updating the latest version of the AWS CLI. For details on the `cancel-image-launch-permission` command, see `cancel-image-launch-permission`.

Security enhanced AMIs

AMS provides security enhanced images (AMIs) based on CIS Level 1 benchmark for a subset of AMS's supported operating systems. To find out which operating systems have a security enhanced image available, see the AMS Security User Guide, which is available through AWS Artifact -> Reports page (find the **Reports** page link in the left navigation pane) filtered for AWS Managed Services. To access AWS Artifact, can contact your CSDM for instructions or go to Getting Started with AWS Artifact.

How integration between AD FS and AMS works

A one-way trust between your on-premises network and the AMS domain is the default means for access to stacks and VPCs. When a VPC and stack are created, access is granted via pre-configured Active Directory security groups. In addition, access to the AWS Management Console can be configured using Active Directory Federation Service (AD FS), or any federation software that supports SAML, for a single sign-on (SSO) to the AWS Management Console.

> **Note**
>
> AMS can federate to many federation services, Ping, Okta, and so on. You aren't limited to AD FS; we provide here an example of one federation technology available to you.

Information here is duplicated from this blog post: Enabling Federation to AWS Using Windows Active Directory, AD FS, and SAML 2.0.

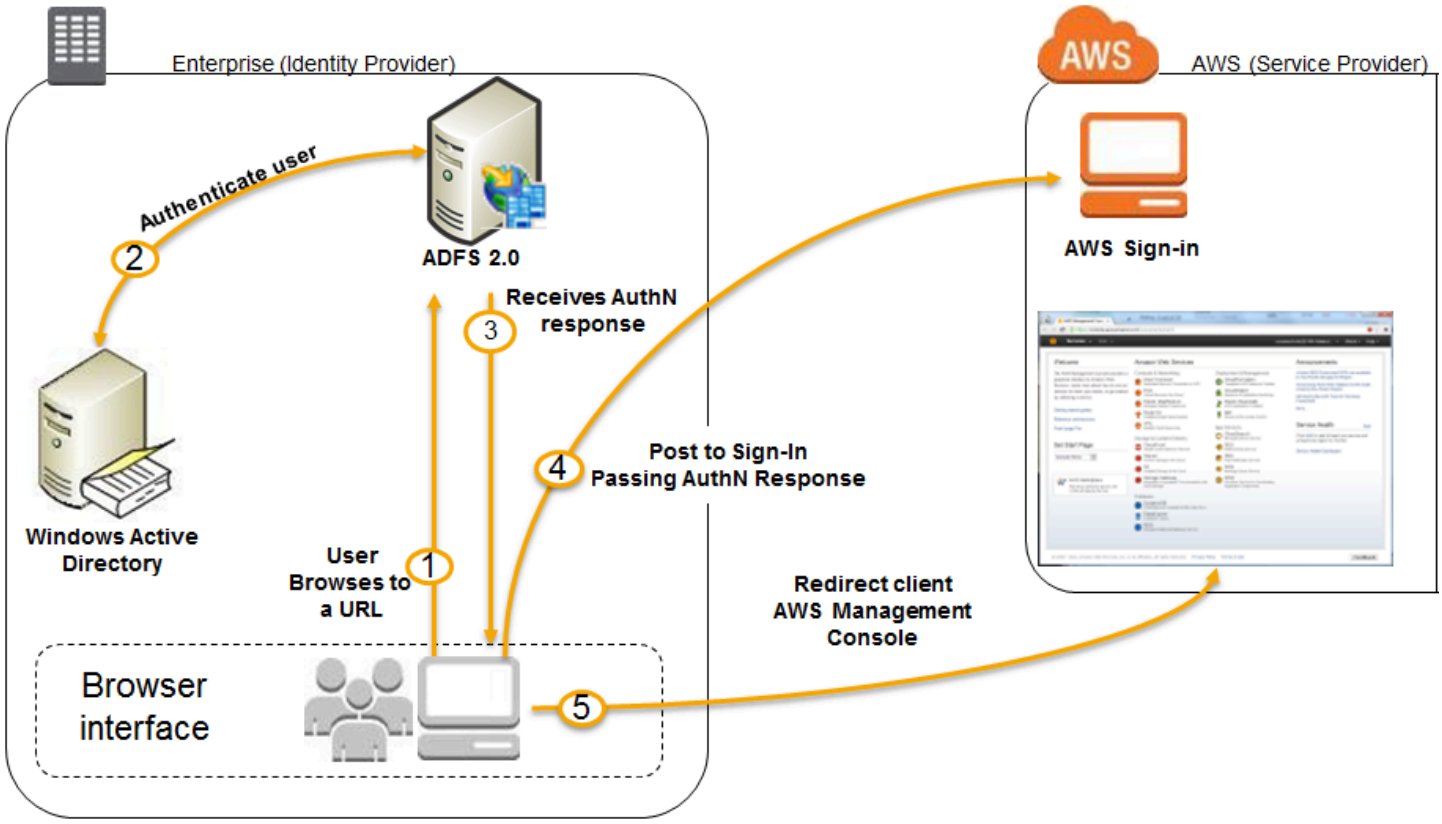

1. The flow is initiated when a user (let's call him Bob) browses to the AD FS sample site (https://Fully.Qualified.Domain.Name.Here/adfs/ls/IdpInitiatedSignOn.aspx) inside his domain. When you install AD FS, you get a new virtual directory named **adfs** for your default website, which includes this page.
2. The sign-on page authenticates Bob against AD. Depending on the browser Bob is using, he might be prompted for his AD username and password.
3. Bob's browser receives a SAML assertion in the form of an authentication response from AD FS.
4. Bob's browser posts the SAML assertion to the AWS sign-in endpoint for SAML (https://signin.aws.amazon.com/saml). Behind the scenes, sign-in uses the AssumeRoleWithSAML API to request temporary security credentials and then constructs a sign-in URL for the AWS Management Console.
5. Bob's browser receives the sign-in URL and is redirected to the console.

From Bob's perspective, the process happens transparently. He starts at an internal website and ends up at the AWS Management Console, without ever having to supply any AWS credentials.

Note

More information on configuring federation to the AMS console is provided in:

- **Multi-Account Landing Zone**: Configuring Federation to the AMS Console
- **Single-Account Landing Zone**: Configuring Federation to the AMS Console

Additionally, see Appendix: AD FS claim rule and SAML settings. For information about using AWS Microsoft AD to support your Active Directory–aware applications, in the AWS Cloud, that are subject to compliance requirements, see Manage Microsoft AD Compliance.

AMS Managed Active Directory

AMS is now offering a new service called Managed Active Directory (aka Managed AD) that allows AMS to take care of your Active Directory (AD) infrastructure operations, while keeping you in control of your Active Directory administration.

AMS support for Managed AD is similar to AMS support for the Amazon Relational Database Service (Amazon RDS). In both cases, AWS (including AMS) supports the creation and management of the infrastructure running the service, while you perform access control and all administration functions. This model has the following advantages:

- Limits security risks: AWS and AMS don't need administrative privileges to your domain.
- Direct integrations: You can use your current authorization model and integrate it with AD without needing to interface with AMS.

Notes:

- Neither AMS nor you will have access to your Managed AD domain controllers, so no software can be installed on the domain controllers. This is important because third-party solutions that require software to be installed on domain controllers is not allowed.

  Access works like this:

  - AWS Directory Service team: Has access to domain controllers.
  - AMS: Has access to Directory Service APIs to perform certain actions on the domain. These actions include taking AD snapshots, changing AD schema, and others actions.

  - You: Have access to the domain (AD) for creating users, groups, and so on.
- We recommend that you perform a proof of concept on Managed AD before migrating your corporate AD, because not all functionality from a traditional AD environment is available in a Managed AD environment.
- AMS will not manage or provide guidance on your AD management. For example, AMS will not provide guidance on Organizational Unit structure, group policy structure, AD user naming conventions, and so forth.

It works like this:

1. AMS onboards a new AWS account for you, separate from and in addition to your AMS account, and provisions an Active Directory (AD) environment through AWS Directory Service (see also What Is AWS Directory Service?).

   The following is the information a systems integrator would need to gather from you in order for AMS to on board Managed AD:

   - Account information
     - Account ID of the AWS account that was created for your AMS-Managed AD: AWS account number
     - Region to onboard your Managed AD to: AWS Region
   - Managed Active Directory information:
     - Microsoft AD Edition: Standard/Enterprise. AWS Microsoft AD (Standard Edition) includes 1 GB of directory object storage. This capacity can support up to 5,000 users or 30,000 directory objects, including users, groups, and computers. AWS Microsoft AD (Enterprise Edition) includes 17 GB of directory object storage, which can support up to 100,000 users or 500,000 objects.

       For more information, see AWS Directory Service FAQs.
     - Domain FQDN: The FQDN for your AMS Managed AD domain.
     - Domain NetBIOS name: The NetBIOS name for your AMS Managed AD domain.
     - Account numbers of AMS-standard accounts you would like Managed AD integration to (AMS configures a one way trust from the AMS-standard account's AD to the Managed AD)
     - Are Active Directory Schema modifications required and if so, what modifications?
     - By default, two domain controllers are provisioned. Do you require more? If so, how many do you require and for what reason?

- Networking for Managed Active Directory information:
  - Managed AD VPC CIDR for domain controllers (a CIDR in your private subnet range for the Managed AD domain controllers):
    - Subnet CIDR 1 for domain controllers: [your CIDR, needs to be part of AMS Managed AD VPC CIDR]
    - Subnet CIDR 2 for domain controllers: [your CIDR, needs to be part of AMS Managed AD VPC CIDR]

    For example:

    - Managed AD VPC CIDR: 192.168.0.0/16
    - CIDR 1 for domain controllers: 192.168.1.0/24
    - CIDR 2 for domain controllers: 192.168.2.0/24

    To avoid IP address conflicts, be sure that the Managed AD VPC CIDR you specify does not conflict with any other private subnet CIDR you are using in your corporate network.
  - VPN Technology (optional): [Direct Connect/Direct Connect and VPN]
    - Your gateway's BGP Autonomous System Number (ASN): [Customer-provided ASN]
    - The Internet-routable IP address for your gateway's outside interface, the address must be static: [Customer Provided IP Address]
    - Whether or not your VPN connection requires static routes: [yes/no]

2. AMS provides you with the Admin account password for the AD environment and asks you to reset the password so AMS engineers can no longer access your AD environment.
3. To reset the Admin account password, connect to your Active Directory environment using Active Directory Users and Computers (ADUC). ADUC and other Remote Server Administration Tools (RSAT) should be installed and run on Administrative hosts provisioned by you on non-AMS infrastructure. Microsoft has best practices for securing such administrative hosts. For information, see Implementing Secure Administrative Hosts. You manage your Active Directory environment using these Administrative hosts.
4. In daily operations, AMS manages the AWS account up to the AWS Directory Service side of things; for example, VPC configuration, AD backups, AD trust creation and deletion, and so forth. You use, and manage, your AD environment; for example, user creation, group creation, group policy creation, and so forth.

For the most recent RACI table, see the "Roles and Responsibilities" section in the See Service description.

AMS application deployments

AMS Application Developer's guide provides detailed descriptions and walkthroughs for the following deployments:

- The AMS workload ingest CT allows you and an AMS cloud migration partner to easily move your existing workloads into an AMS-managed VPC. Using AMS workload ingest, you can create an AMS AMI by submitting an RFC with the Deployment | Ingestion | Stack from migration partner migrated instance | Create CT (ct-257p9zjk14ija). You must have an instance migrated from your on-premises to AWS by a migration partner, as well as a target AMS VPC and subnet, into which the instance will be ingested.

  For details, see the *AMS Application Developer's* guide at Workload Ingest.
- The AWS CloudFormation ingest change type (ct-36cn2avfrrj9v) feature allows you to easily use an existing CloudFormation template to deploy custom stacks in an AMS-managed VPC.

  For details, see the *AMS Application Developer's* guide at CloudFormation Template Ingest.
- You can import your on-premises database into a new database to your AMS-managed Amazon S3 bucket or Amazon RDS instance. You do this using a Deployment | Advanced stack components | Database Migration Service (DMS) change types, including Create replication instance (ct-27apldkhqr0ol), Create replication subnet group (ct-2q5azjd8p1ag5), Create replication task (ct-1d2fml15b9eth), Create source endpoint (ct-0attesnjqy2cx) or Create source endpoint (S3) (ct-2oxl37nphsrjz), and Create target endpoint (ct-3gf8dolbo8x9p) or Create target endpoint (S3) (ct-05muqzievnxk5).

  For details, see the *AMS Application Developer's* guide at Database Migration Service.
- You can import your on-premises MS SQL database into a new database on your AMS-managed RDS SQL instance. You do this using a variety of AMS change types, and the Amazon RDS API, plus AWS consoles.

  For details, see the AMS Application Guide at Database (DB) Import to MS SQL RDS.

Service management in AWS Managed Services

Topics

- Account governance in AWS Managed Services
- Service commencement in AWS Managed Services
- Customer relationship management (CRM)
- Cost optimization in AWS Managed Services
- Planned event management (PEM)
- Service hours in AWS Managed Services
- Getting help in AWS Managed Services

How the AMS service works for you.

Account governance in AWS Managed Services

This section covers AMS account governance.

You are designated a cloud service delivery manager (CSDM) who provides advisory assistance across AMS, and has a detailed understanding of your use case and technology architecture for the managed environment. CSDMs work with account managers, technical account managers, AWS Managed Services cloud architects (CAs), and AWS solution architects (SAs), as applicable, to help launch new projects and give best-practices recommendations throughout the software development and operations processes. The CSDM is the primary point of contact for AMS. Key responsibilities of your CSDM are:

- Organize and lead monthly service review meetings with customers.
- Provide details on security, software updates for environment and opportunities for optimization.
- Champion your requirements including feature requests for AMS.
- Respond to and resolve billing and service reporting requests.
- Provide insights for financial and capacity optimization recommendations.

Service commencement in AWS Managed Services

Service Commencement: The *Service Commencement Date* for an AWS Managed Services account is the first day of the first calendar month after which AWS notifies you that the activities set out in the Onboarding Requirements for that AWS Managed Services account have been completed; provided that if AWS makes such notification after the 20th day of a calendar month, the Service Commencement Date is the first day of the second calendar month following the date of such notification.

Service Commencement

- **R** stands for responsible party that does the work to achieve the task.
- **I** stands for informed; a party which is informed on progress, often only on completion of the task or deliverable.

Service commencement

Step #	Step title	Description	Custome	AMS
1.	Customer AWS account handover	Customer creates a new AWS account and hands it over to AWS Managed Services	R	I
2.	AWS Managed Services Account - design	Finalize design of AWS Managed Services Account	I	R
3.	AWS Managed Services Account - build	An AWS Managed Services account is built per the design in Step 2	I	R

Customer relationship management (CRM)

AWS Managed Services (AMS) provides a customer relationship management (CRM) process to ensure that a well-defined relationship is established and maintained with you. The foundation of this relationship is based on AMS's insight into your business requirements. The CRM process facilitates accurate and comprehensive understanding of:

- Your business needs and how to fill those needs
- Your capabilities and constraints
- AMS and your different responsibilities and obligations

The CRM process allows AMS to use consistent methods to deliver services to you and provide governance for your relationship with AMS. The CRM process includes:

- Identifying your key stakeholders
- Establishing a governance team
- Conducting and documenting service review meetings with you
- Providing a formal service complaint procedure with an escalation procedure
- Implementing and monitoring your satisfaction and feedback process
- Managing your contract

CRM Process

The CRM process includes these activities:

- Identifying and understanding your business processes and needs. Your agreement with AMS identifies your stakeholders.
- Defining the services to be provided to meet your needs and requirements.
- Meeting with you in the service review meetings to discuss any changes in the AMS service scope, SLA, contract, and your business needs. Interim meetings may be held with you to discuss performance, achievements, issues, and action plans.
- Monitoring your satisfaction by using our customer satisfaction survey and feedback given at meetings.
- Reporting performance on monthly internally-measured performance reports.
- Reviewing the service with you to determine opportunities for improvements. This includes frequent communication with you regarding the level and quality of the AMS service provided.

CRM meetings

AMS cloud service delivery managers (CSDMs) conduct meetings with you regularly to discuss service tracks (operations, security, and product innovations) and executive tracks (SLA reports, satisfaction measures, and changes in your business needs).

Meeting	Purpose	Mode	Participants
Weekly status review (optional)	Outstanding issues or incidents, patching, security events, problem records 12-week operational trend (+/- 6) Application operator concerns Weekend schedule	On-site customer location/ Telecom/Chime	AMS: CSDM and cloud architect (CA) Customer assigned team members (ex: Cloud/ Infrastructure, Application Support, Architecture teams, etc.)
Monthly business review	Review service level performance (reports, analysis, and trends) Financial analysis Product roadmap CSAT	On-site customer location/ Telecom/Chime	AMS: CSDM, cloud architect (CA), AMS account team, AMS technical product manager (TPM) (optional), AMS OPS manager (optional) You: Application Operator representative

Meeting	Purpose	Mode	Participants
Quarterly business review	Scorecard and service level agreement (SLA) performance and trends (6 months) Upcoming 3/6/9/12 months plans/ migrations Risk and risk mitigations Key improvement initiatives Product roadmap items Future direction aligned opportuni ties Financials Cost savings initiatives Business optimization	On-site customer location	AMS: CSDM, cloud architect , AMS account team, AMS service director, AMS operation manager You: Applicati on operator representative, service represent ative, service director

CRM Meeting Arrangements

The AMS CSDM is responsible for documenting the meeting, including:

- Creating the agenda, including action items, issues, and list of attendees.
- Creating the list of action items reviewed at each meeting to ensure items are completed and resolved on schedule.
- Distributing meeting minutes and the action item list to meeting attendees by email within one business day after the meeting.
- Storing meeting minutes in the appropriate document repository.

In absence of the CSDM, the AMS representative leading the meeting creates and distributes minutes.

Note

Your CSDM works with you to establish your account governance.

CRM monthly reports

Your AMS CSDM prepares and sends out monthly service performance presentations. The presentations include information on the following:

- Report date
- Summary and Insights:
  - Key Call Outs: total and active stack count, stack patching status, account onboarding status (during onboarding only), customer-specific issues summaries
  - Performance: Stats on incident resolution, alerts, patching, requests for change (RFCs), service requests, and console and API availability
  - Issues, challenges, concerns, and risks: Customer-specific issues status
  - Upcoming items: Customer-specific onboarding or incident resolution plans
- Managed Resources: Graphs and pie charts of stacks
- AMS Metrics: Monitoring and event metrics, incident metrics, AMS SLA adherence metrics, service request metrics, change management metrics, storage metrics, continuity metrics, Trusted Advisor metrics, and cost summaries (presented several ways). Feature requests. Contact information.

Note

In addition to the described information, your CSDM also informs you of any material change in scope or terms, including use of subcontractors by AMS for operational activities. AMS generates reports about patching and backup that your CSDM includes in your monthly report. As part of the report generating system, AMS adds some infrastructure to your account that is not accessible to you:

- An S3 Bucket, with the raw data reported
- An Athena instance, with query definitions to query the data
- A Glue Crawler to read the raw data from the S3 bucket

Cost optimization in AWS Managed Services

AWS Managed Services provides a detailed cost utilization and savings reports every month to you during your monthly business reviews (MBRs).

AMS follows a standard set of processes and mechanisms to identify cost saving avenues in your managed accounts and assist you to plan and roll-out the changes to optimize your AWS spend.

> **Note**
>
> AMS is developing a video to help with cost optimization. The first step is providing you with a PDF and an Excel spreadsheet of cost optimization best practices. To access these resources, open the Quick guide to cost optimization ZIP file.

Cost optimization framework

AMS follows a three-staged approach with you to optimize your AWS costs:

1. Identify cost optimization avenues in your managed environment
2. Present a cost optimization plan to you
3. Assist in achieving cost optimization in a measurable way

Identify cost optimization avenues in the managed environment

AMS utilizes AWS native tools like Cost explorer, and Trusted Advisor while leveraging over 20 cost savings patterns across architecture optimization, EC2 instance, and AWS account-focused optimizations to build tailored cost savings recommendations for you.

Some of the optimization recommendations include the following.

Architectural optimization recommendations:

- **Optimal S3 storage class use**: Amazon S3 offers a range of storage classes to meet various workload requirements based on data access, resiliency, and cost. S3 Intelligent-Tiering and S3 storage class analysis based on the workload needs allow you to manage the S3 costs efficiently.
- **Using caching architectures**: Leveraging cache instances, where applicable, can help you replace some database instances, while simultaneously meeting your IOPS requirements.

- **EBS upgrade savings**: Migrating your EBS volumes from gp2 to gp3 provides a cost savings of up to 20% and you can take advantage of predictable 3,000 IOPS baseline performance and 125 MiB/s, regardless of volume size.
- **Using elasticity**: The auto-scaling capabilities that AWS provides allow effective resource utilization and avenues for cost optimization. Reviewing and updating the instance scaling policies regularly based on need, further provides cost savings.

EC2 instance-focused recommendations

- **Instance rightsizing**: Recommendations focused on sizing the instances and optimal configurations based on the usage. Recommendations also include utilizing Amazon EC2 Auto Scaling feature and replacing EC2 instances where applicable with AWS Lambda or static web content on Amazon S3, etc.
- **Instance scheduling**: Using AMS Resource Scheduler to automatically start and stop instances based on a time schedule helps contain costs, especially for non-production instances that are not utilized during non-business hours.
- **Subscribing to Savings plans**: Savings plan is the easiest way to save on AWS usage. The EC2 Instance Savings Plans offer up to 72% savings compared to On-Demand pricing on your Amazon EC2 instances usage. The Amazon SageMaker Savings Plans offer up to 64% savings on your Amazon SageMaker services usage. AMS provides appropriate recommendations on Savings plans based on your AWS resource usage.
- **Reserved instance (RI) usage and consumption guidance**: Amazon EC2 Reserved Instances (RI) provide a significant discount (up to 75%) compared to On-Demand pricing and provide a capacity reservation when used in a specific availability zone.
- **Spot instance usage**: Fault tolerant workloads can utilize Spot instances and reduce prices up to 90%.
- **Idle instance termination**: Identifying and reporting instances that are idle or have low utilization that can be terminated.

Account-focused recommendations

- **Account cleanup**: At an account level, AMS also identifies un-utilized EBS volumes, duplicate CloudTrail trails, empty accounts with unused resources, and so forth, and provides recommendations for clean-up.

- **SLA recommendations**: Further, AMS regularly reviews your Plus and Premium accounts and recommends choosing the right SLA level for the accounts.
- **AMS automation optimization**: AMS continuously optmizes AMS automation and infrastructure used to provide AMS services.

Present to customers and assist in planning

AMS conducts monthly business reviews (MBRs) with the key customer stakeholders and present the cost saving avenues, mechanisms and recommendations identified along with potential cost savings. We further work with you to plan the changes needed.

Assist in recommendation implementation and measure the cost impact

AMS assists in achieving and measuring cost impacts and optimization changes.

You assess the application impact, risk and success criteria of the recommended changes, and raise the appropriate requests for change (RFCs) through the AMS console. AMS collaborates with you and implements the changes related to cost optimization in your managed accounts. AMS measures the cost impact and include the savings realised in the monthly business reviews (MBRs).

Cost optimization responsibility matrix

Responsibilities in AMS cost optimization.

Cost optimization RACI

Activity	Customer	AMS
Compiling cost saving recommendations and preparing the report	I	R
Presenting cost	C	R

Activity	Customer	AMS
savings report		
Planning changes associated with cost savings	R	C
Assessing the change impact and risk	R	C
Raising RFCs for implementing the changes	R	C
Reviewing the RFCs and implementing the changes	C	R

Activity	Customer	AMS
Testing the application and validating the change implementation	R	C
Measuring the cost impact post change and presenting to customer	I	R

Planned event management (PEM)

AWS Managed Services (AMS) planned event management (PEM) is an AMS service offering. PEM is used to engage, coordinate, and assist during customer events and projects using AMS services. PEM assists in coordinating a set of related RFCs that align with the agreed scope and timeline of the PEM event or project.

AMS PEM criteria

A planned event is a scope-bound and time-bound project. AMS uses the details that you provide (including plan and scope, expected outcomes, and changes that AMS operations are expected to perform) to effectively support you during PEM activity. Your Cloud Architects (CAs) then review and assess the PEM activity for completeness, technical implementation, and operations engagement. After CA review, AMS operations reviews the plans and coordinates with your cloud service delivery manager (CSDM) for operations team engagement.

Types of PEM

There are two types of PEM with differing activities:

- **Planned Event Management (PEM) workflow**: Subtypes: gamedays, disaster recovery (DR) events, and security events, application roll-outs, and new application deployments.
- **Environments and Workload Migrations (EWM) workflow**: Subtypes: on-boarding and migration workloads.

The workflows help you engage with AMS for planned work and give AMS visibility based on the PEM type. For additional details on your requirements, contact your CSDM.

The AMS PEM process

The PEM process consists of the following phases:

- **PEM initiation:** You work with your CSDM and Technical Delivery Managers (TDMs) to define your objective for the planned event and determine what is needed from AMS Operations. AMS CAs review the technical aspects of the PEM plan. The CAs work with AMS Security and Operations on compliance, execution optimization and automation, and to define pre-PEM execution tasks and deliverables. Then, your CSDM creates the PEM ticket and provides AMS with the project information and technical details. AMS requires a lead time of 14 calendar days to allow the AMS Operations team time to plan, provide technical review, and assign resources.
- **PEM review:** The AMS Operations team reviews the PEM request and works with your CSDM to verify that the information in the PEM plan is correct and complete.
- **PEM acceptance:** AMS reviews the provided information and communicates to the CSDM what the level of support will be during the PEM activity. If the PEM contains complete information and your CSDM agrees with the scope of work, then the PEM is approved.
- **Readiness and execution:** AMS makes sure that tasks needed before the PEM begins are completed and facilitates internal and customer communications. AMS also ensures execution of the PEM plan and provides status and progress reporting.

PEM FAQs

How do I engage AMS with a RFC/Service Request (SR) during a PEM event?

- Use the PEM ID shared by your CSDM in the RFC/SR subject line in the format *PEM/EWM*-s*ID*.

- You can also create a Service Request (SR) to discuss your use cases or for questions about your planned event. If you use an SR, then the PEM doesn't have to be valid.

What validations are performed when a PEM-related RFC or SR is submitted?

- Verification that the Account ID is listed on the PEM.
- Verification that the PEM status is approved and active between the provided start and end dates.

Are there SLAs or SLOs for PEM requests?

- PEMs are not associated with SLAs or SLOs.
- SLAs/SLOs for PEM-related work items (RFC/Service Request) are defined by AMS SLOs. For more information, see AMS service level objectives (SLOs).

Service hours in AWS Managed Services

Feature	AMS Advanced
	Premium Tier
Service request	24/7
Incident management (P2-P3)	24/7
Backup and recovery	24/7
Patch management	24/7
Monitoring and alerting	24/7
Automated request for change (RFC)	24/7
Non-automated request for change (RFC)	24/7
Cloud service delivery manager (CSDM)	Monday to Friday: 08:00–17:00, local business hours

Getting help in AWS Managed Services

AMS supports you with Incident Management, Service Request Management, and Change Management 24 hours a day, 7 days a week, 365 days a year (in accordance with the AMS Service Level Agreement applied to the account).

To report an AWS or AMS service performance issue that impacts your managed environment, use the AMS console and submit an incident report. For details, see Reporting an incident. For general information about AMS incident management, see Incident response.

To ask for information or advice, or to request additional services from AMS, use the AMS console and submit a service request. For details, Creating a Service Request. For general information about AMS service requests, see Service Request Management.

Network architecture

AWS Managed Services (AMS) offers two network architectures:

- Multi-account landing zone (MALZ): provides common services - such as access, end point security, networking - from shared accounts for workloads that are deployed in separate member accounts.
- Single-account landing zone (SALZ): provides self contained accounts where common services such as access, end point security, networking are deployed in the same account as the workload. It is recommended for workloads that require a high level of isolation as it incurs higher AWS costs.

MALZ network architecture

About multi-account landing zone network architecture

Before you start the onboarding process to AWS Managed Services (AMS) Multi-account landing zone (MALZ), it is important to understand the baseline architecture, or landing zone, that AMS creates on your behalf, its components, and functions.

AMS multi-account landing zone is a multi-account architecture, pre-configured with the infrastructure to facilitate authentication, security, networking, and logging.

> **Note**
>
> For estimates of costs, see AMS multi-account landing zone environment basic components.

Topics

- Service region
- Organizational units
- Service control policies and AWS Organization

The following diagram outlines at a high level the account structure and how infrastructure is segregated into each of the accounts:

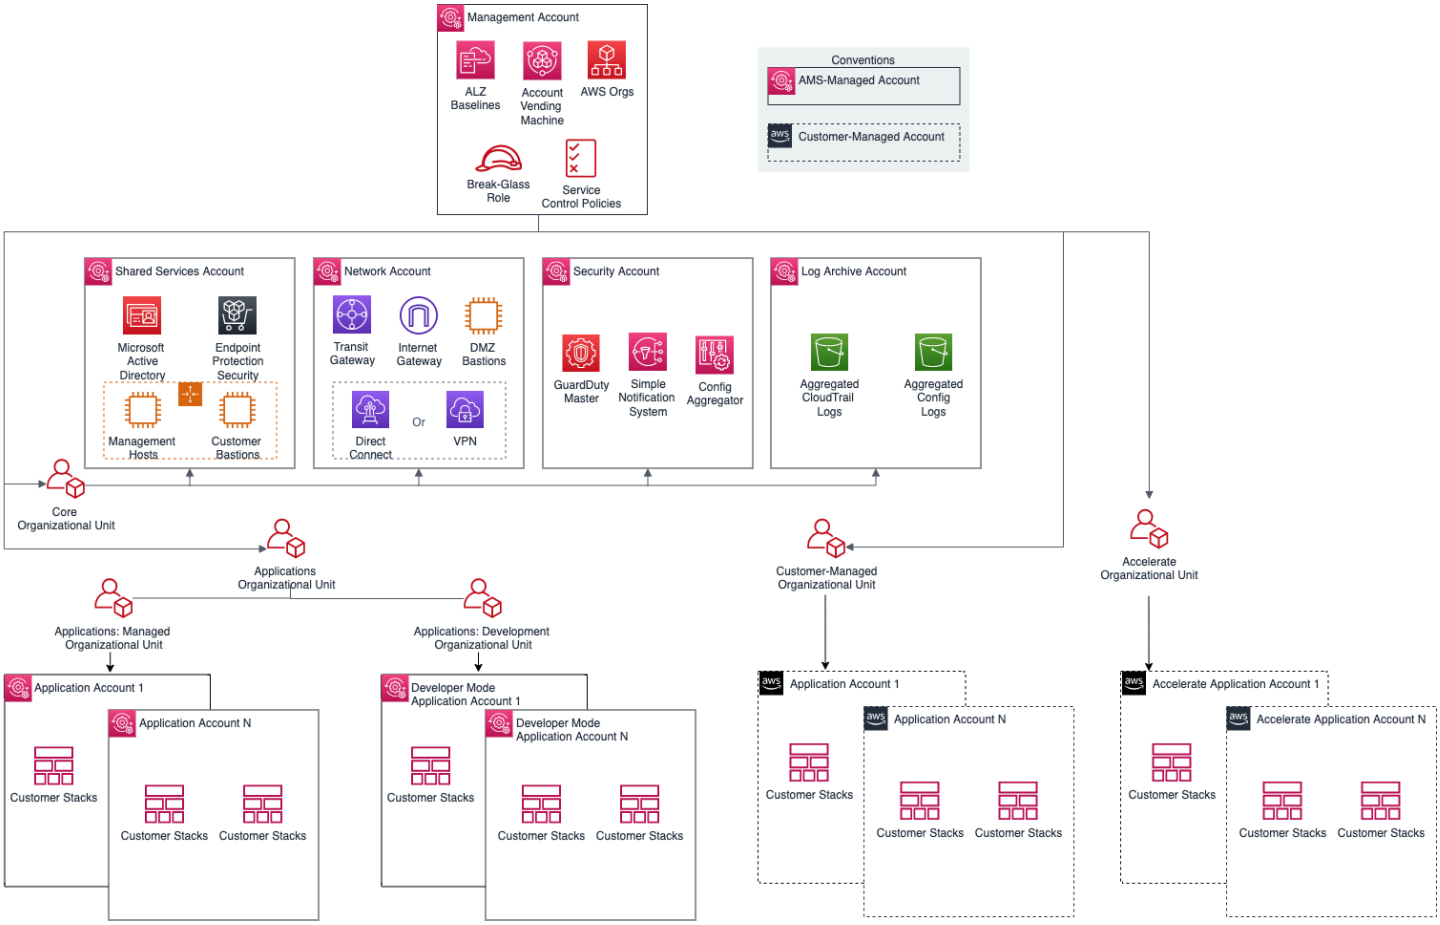

Service region

All resources within an AMS multi-account landing zone are deployed within a single AWS Region of your choice, due to current cross region limitation with Active Directory and Transit Gateway.

Organizational units

A typical AMS multi-account landing zone consists of four top-level organizational units (OUs):

- The core Organizational unit (OU) (used to group accounts together to administer as a single unit)
- The applications OU
- The customer-managed OU
- The accelerate OU

AMS-managed multi-account landing zone also enables you to create custom OUs for grouping and organizing AWS Accounts and to associate custom SCPs with them; for examples on doing this,

see Management account | Create Custom OUs and Management account | Create Custom SCP (review required), respectively. AMS provides four existing OUs under which new OUs and accounts can be requested: accelerate, applications > managed, applications > development, and customer-managed.

- accelerate OU:

  This is a top-level OU in AMS multi-account landing zone (MALZ). Accounts under this OU are provisioned by AMS with an RFC (Deployment | Managed landing zone | Management account | Create Accelerate account, change type ID: ct-2p93tyd5angmi). In these accelerate application accounts, you can benefit from accelerate operational services such as monitoring and alerting, incident management, security management, and backup management. For more details, see AMS Accelerate accounts.

- applications > managed OU:

  In this sub organizational unit of the Application OU, accounts are fully managed by AMS including all operational tasks. The operational tasks include service request management, incident management, security management, continuity management, patch management, cost optimization, monitoring and event management. These tasks are carried out for your infrastructure's management. Multiple child OUs can be created as needed, until a maximum limit of nested OUs is reached for AWS organizations. For details, see Quotas for AWS Organizations.

- applications > development OU:

  Under this sub-OU of the application OU in AMS-managed landing zone, accounts are Developer mode accounts that provide you with elevated permissions to provision and update AWS resources outside of the AMS change management process. This OU also supports the creation of new children OU as needed.

- customer-managed OU:

  This is a top-level OU in AMS multi-account landing zone. Accounts under this OU are provisioned by AMS with an RFC. In these accounts, the operations of workloads and AWS resources are your responsibility. This OU also supports the creation of new children OU as needed.

As a best practice, we recommend that accounts under these OUs and custom-requested sub-OUs be grouped based on their functionalities and policies.

Service control policies and AWS Organization

AWS provides service control policies (SCPs) for permissions management in an AWS Organization. SCPs are used to define additional guardrails for what actions users can perform in which OUs. By default, AMS provides a set of SCPs deployed in management accounts which provide protections at different default OU levels. For SCP restrictions, please contact your CSDM.

You can also create custom SCPs and attach them to specific OUs. They can be requested from your Management account using change type ct-33ste5yc7hprs. AMS then reviews the custom SCPs requested before applying them to the target OUs. For examples, see Management account | Create Custom OUs and Management account | Create Custom SCP (Review Required).

Choosing single MALZ or multiple MALZs

The following table provides some high level considerations on deciding between a single multi-account landing zone (MALZ) vs multiple multi-account landing zones (for example, two multi-account landing zones - Prod and non-Prod). In general, the choice depends upon individual needs, legal requirements, and operating practices.

Single multi-account landing zone vs. multiple multi-account landing zones

Entity	Single AMS landing zone	Multiple (two or more) landing zones
Base cost	Lower, optimized at approximately $3,000 per month.	Higher, an additional cost of approximately $3,000 per environme nt.
Billing	Single bill, due to single Billing/ Management account.	Separate bill for each multi-account landing zone. Currently AWS Org does not support multi-Management accounts with a single bill.
Portability of existing reserved instances (RIs)	Low. AWS RIs are currently not convertible across multiple billing accounts. You would repurpose existing RIs for multi-account landing zone.	Lower. You would repurpose and distribute RIs across all multi-account landing zone.

Entity	Single AMS landing zone	Multiple (two or more) landing zones
Product tiering discounts	High. See Volume discounts.	Low. See Volume discounts.
Initial setup overhead (on project/ migration timelines)	Low. Active Directory, networking and single sign-on (SSO) integrations once only.	High. You would perform Active Directory, network integration, and SSO integrations for every landing zone. This could cause potential delays to any migration project.
Common services configura bility	Low efforts. You configure common/ shared services like DNS, backup, monitoring, logging etc.	High efforts. Additional planning is required to address where the common infrastructure or services will be sitting. Traffic traversing across multiple transit gateways (TGWs) in each landing zone, could lead to extra cost.
Scalability	Medium. AMS has a current practical limit of 150 accounts per multi-acc ount landing zone. Multiple teams or vendors running applications in same account could have access to stacks owned by different teams. This limitation can be mitigated by controlling access to application-specific stacks at the ServiceNow layer (by integrating the AMS ServiceNow Connector application and making use of tags). Ask AMS technical delivery managers (TDMs) or cloud architects (CAs) how to implement this.	High. Ability to leverage multiple multi-account landing zone to distribute the accounts while achieving an account or application level of segregation. Managing large numbers of accounts could lead to operational or cost overhead.

Entity	Single AMS landing zone	Multiple (two or more) landing zones
Operational Risk	(Depends) Low. Operational integration and readiness once only. Less chance of process drifts.	(Depends) Low. Multiple integration and operational activities. Drift in multiple landing zones over the period could lead to operational risks.
Multi AWS Region	Single AWS Region. AMS multi-account landing zone is restricted to a single AWS Region. To span multiple AWS Regions, use multiple multi-account landing zone.	Multi AWS Region. With multiple multi-account landing zones, you can have each MALZ deployed in one region and interconnect them using transit gateway (TGW) peering.
Account migration or portability	Yes. Moving accounts from one OU to another within the same AWS Organization is possible.	No. AMS doesn't support migration of an account across landing zones; that is, across AWS Organizations. Workloads can reach across landing zones with transit gateway (TGW) or VPC peering.
Change management	Medium. Making destructive changes to common components like TGW, Active Directory (AD), or outbound (egress) can impact all workloads in a multi-account landing zone. However, changes to AMS-managed components are tested internally and are pushed in rolling updates.	Low. Making destructive changes to common components like TGW, AD, or outbound (egress) can impact only the workloads in that specific multi-account landing zone.
Data and access controls	(Depends) Low control if you'd like to connect to different on-premise ADs and networks for Prod vs Non-Prod workloads. SAML federation, TGW domains, and security groups (SGs) can help implement required controls too.	(Depends) High control if you'd like to connect to different on-premise ADs and networks for Prod vs Non-Prod workloads. Use separate landing zones for strict compliance requirements.

Entity	Single AMS landing zone	Multiple (two or more) landing zones
Compliance and Security	(Depends) Low if there are strict compliance needs to completely segregate material vs non-material workloads. AMS standard preventative and detective controls in place.	(Depends) High as multiple multi-account landing zone could help achieve strict compliance requirements by completely segregating material vs non-material workloads. AMS standard preventative and detective controls in place.

Recommendation: Without strict Compliance or multi-Region need, starting with single AMS multi-account landing zone would strike a good balance among cost, security, operational excellence, and migration complexity. You can always setup additional landing zone, if any account or business constraints are encountered.

Single multi-account landing zone vs. Multiple multi-account landing zone FAQs

Some commonly asked questions when choosing to set up a single multi-account landing zone or multiple multi-account landing zones:

Q1: Can I start with a single multi-account landing zone and move to multiple multi-account landing zone, if any account limits or business constraints are encountered?

A: Yes. You can choose to set up another multi-account landing zone at any given time:

- A new billing payer account will be required to be setup (currently AWS doesn’t support multi-payer accounts in a single AWS org).
- Multi-Account Landing Zone base build takes up to 2 weeks lead time once the multi-account landing zone questionnaire is filled out.
- Every multi-account landing zone means an addition of ~3K USD / month running cost.
- N/W, AD, DNS, and SSO integration will be required to establish for new MALZ.
- Any Reserved Instances (RIs), Cost Saving plans will be needed to be setup for the new multi-account landing zone (RIs are not transferrable).
- AMS multi-account landing zone doesn't support migration of an account across multi-account landing zone accounts; for example, across AWS Orgs. However, to move applications from one account to another is possible using standard migration methods.

Q2: What is AMS approach to MALZ updates/changes to underlying/shared infrastructure and quantify the risk to customers? Provide details on what assurances are wrapped into the process. How do Customers get comfortable that MALZ updates/changes will not impact customers? Is there any measures Customer need to take to prevent disruption?

A: AMS follows a strict change methodology using internal tools that enables us to define, review, schedule and execute changes to customers' environments.

The process to release updates enforces code reviews, integration testing, deployment in gamma and beta environments, and additional baking time and testing in beta and gamma environments before releasing to customers environments. All releases include rollback procedures and are closely monitored by the releases team and the team who created and requested the change. The scope of the releases are confined to stacks owned and provisioned by AMS. On average, we execute at least one release per week.

In addition:

- AMS SLA are applicable. As per AMS service description any incident raised post shared infra maintenance activity would adhere to entitled SLA for resolution or credits.
- No special preventive measures are required by Customers to prevent disruption to common infrastructure. Customers have Read-Only permissions at AWS Org or Core OU accounts, so customers can’t make any destructive changes to the MALZ core env. All customer’s requests to Core infrastructure requires AMS review and approval.
- Customers can test certain Org level changes like SCPs/Roles at individual non-prod account levels before propagating changes at App OU level. It is on the AMS roadmap to allow multiple APP OUs (Q2 2020), which would further alleviate risk in making some of the ORG level changes. MALZ team has already released separate OU for “Build Mode” accounts, to ensure clear segregation of customer ownership and separate controls.
- Most of these are changes that allow AMS to operate the workload in effective and efficient manner and does not necessarily impact customers workload. Where AMS believes a shared infra change can have an impact to customers’ workload they are then aligned with customers’ change window.

High level recommendation, start with multiple multi-account landing zones if:

- If it helps you achieve any specific compliance.
- If you need to use Multi-Region.

- If you have different on-prem ADs and Networks for Prod/Material vs Non-Prod/Non-Material workloads, to clearly segregate b/w the workloads.

Multi-Account Landing Zone accounts

Topics

- Management account
- Networking account
- Shared Services account
- Log Archive account
- Security account
- Application account types
- AMS Tools account (migrating workloads)

Management account

The management account is your initial AWS account when you begin onboarding with AMS. It utilizes AWS Organizations as a management account (also known as the payer account that pays the charges of all the member accounts), which gives the account the ability to create and financially manage member accounts. It contains the AWS landing zone (ALZ) framework, account configuration stack sets, AWS Organization service control policies (SCPs), etc.

For more information on using a management account, see Best practices for the management account.

The following diagram provides a high-level overview of the resources contained in the management account.

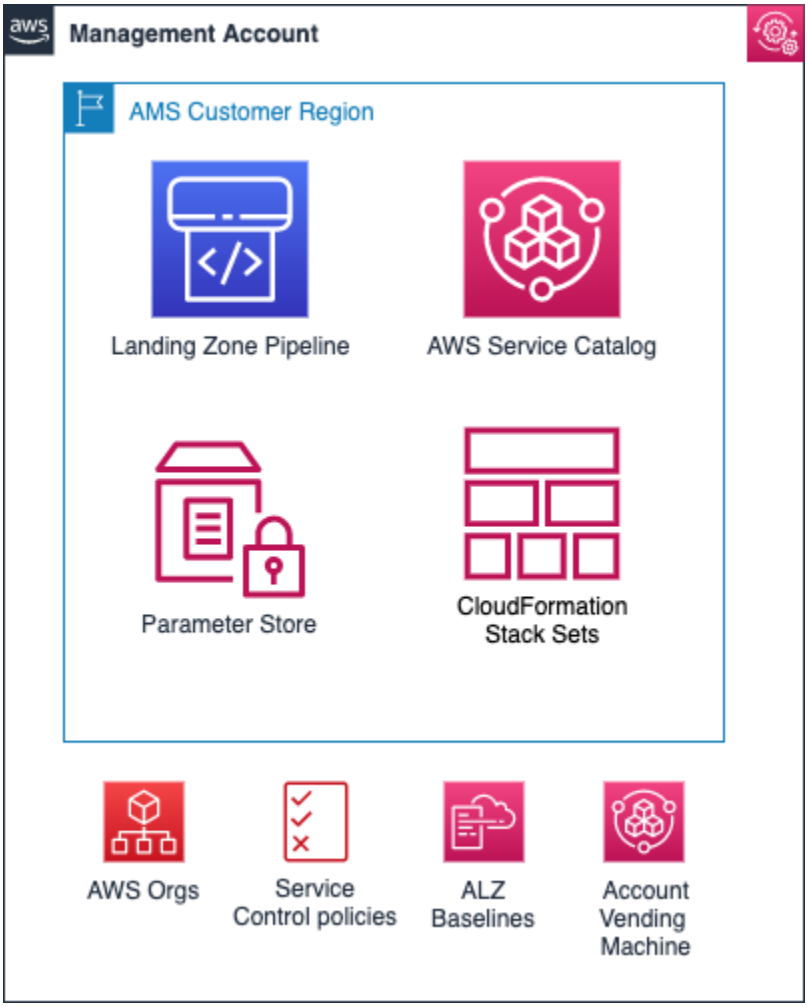

Resources in the management account

Other than the above standard services, no additional AWS resources are created in the management account during onboarding. The following inputs are required during onboarding to AMS:

- *Management account ID*: AWS Account ID that is created initially by you.
- *Core Accounts emails*: Provide the emails to be associated with each of the core accounts: Networking, Shared Services, Logging, and Security account.
- *Service Region*: Provide the AWS region to which all resources of your AMS landing zone will be deployed.

Networking account

The Networking account serves as the central hub for network routing between AMS multi-account landing zone accounts, your on-premises network, and egress traffic out to the Internet. In addition, this account contains public DMZ bastions that are the entry point for AMS engineers to access hosts in the AMS environment. For details, see the following high-level diagram of the networking account below.

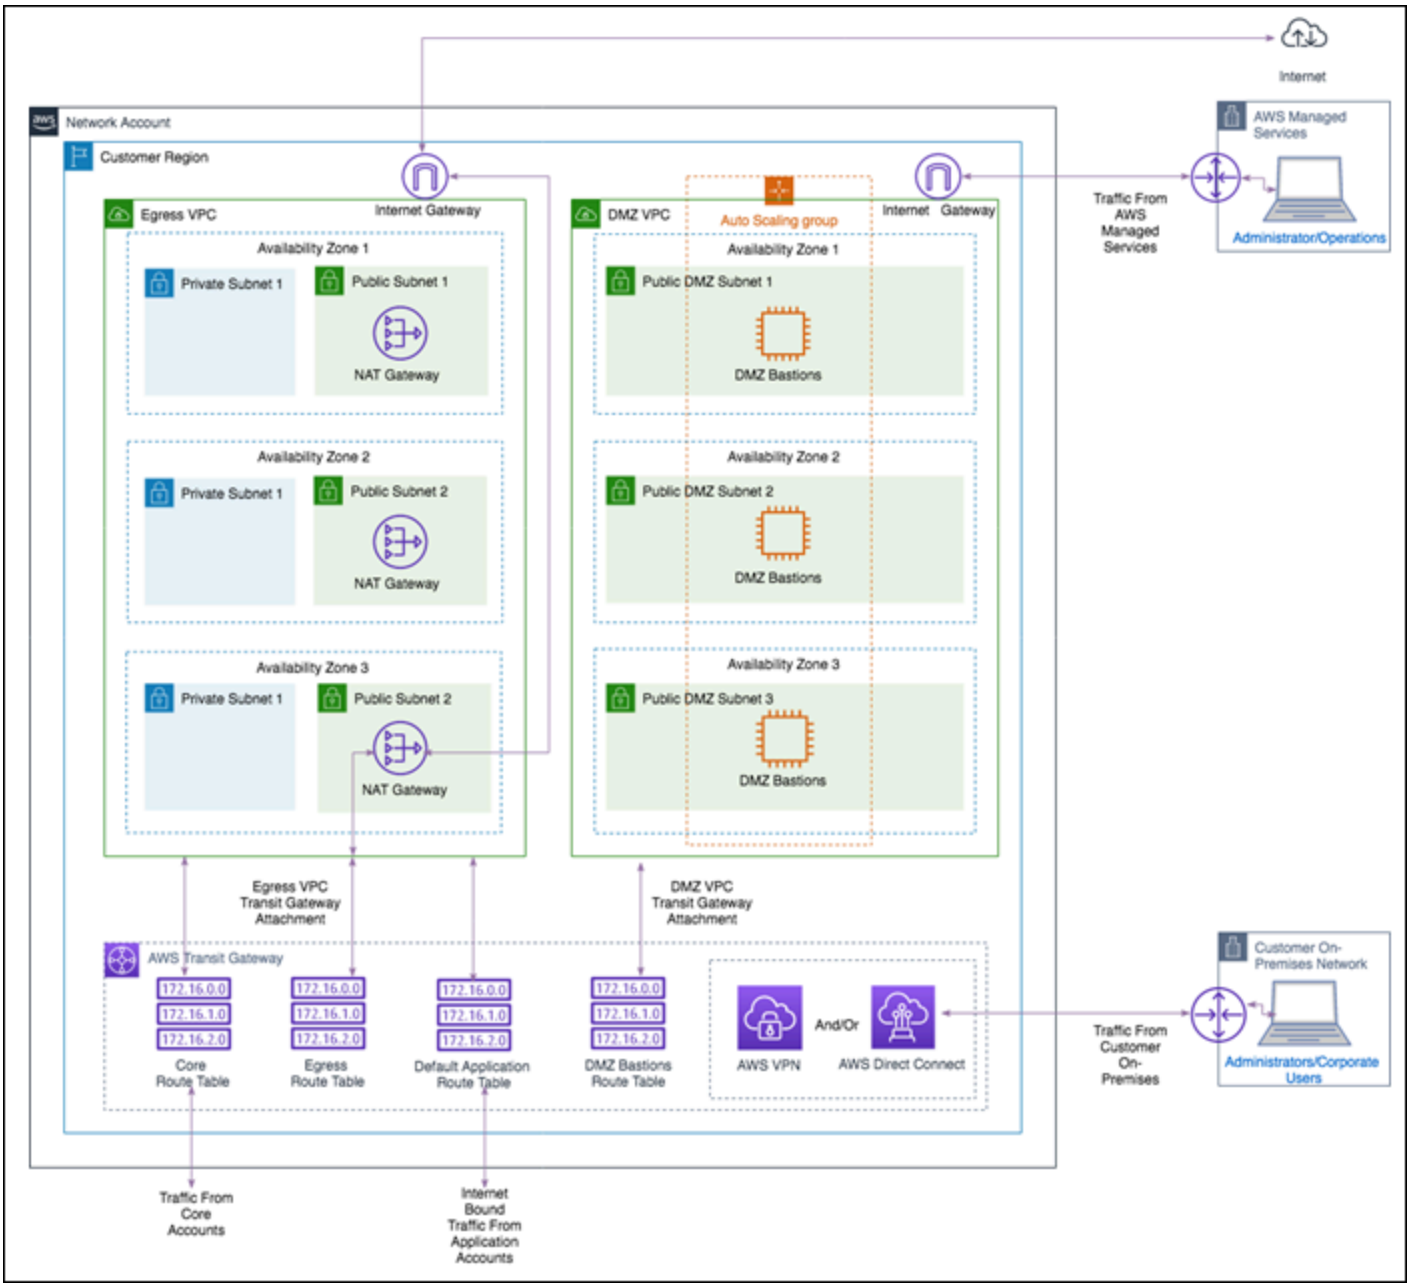

Networking account architecture

The following diagram depicts the AMS multi-account landing zone environment, showcasing network traffic flows across account, and is an example of a highly-available setup.

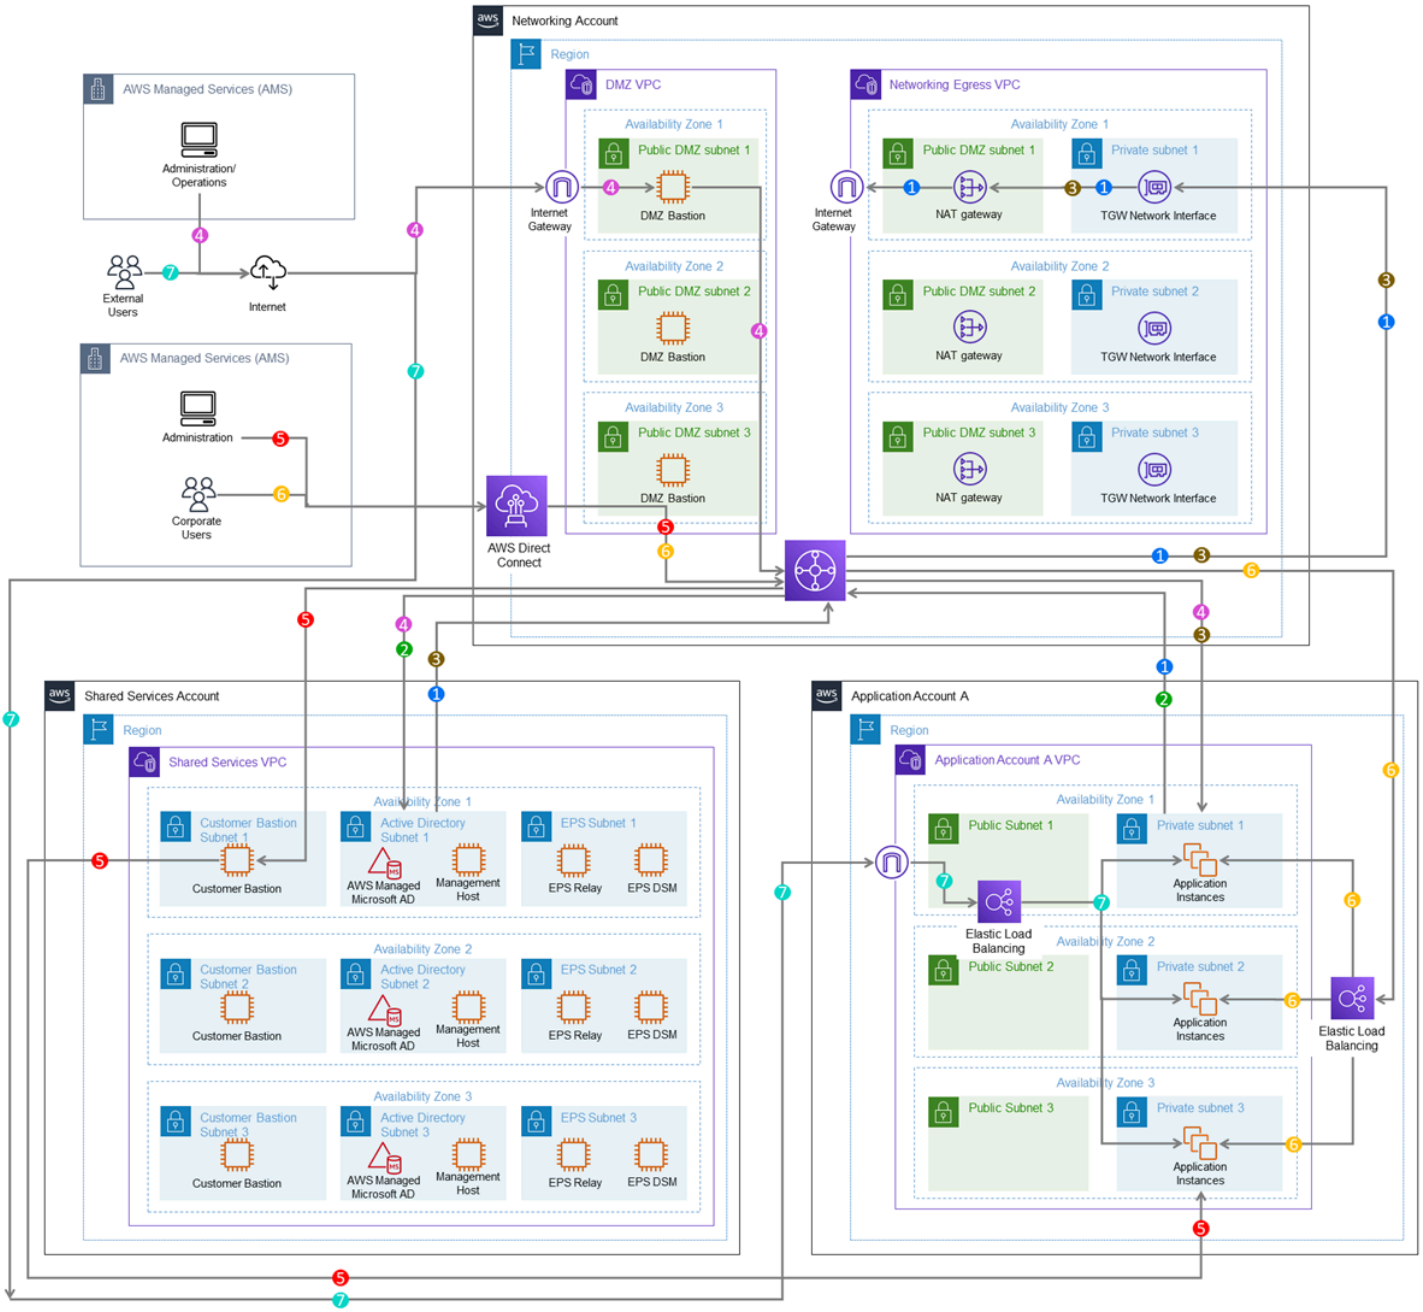

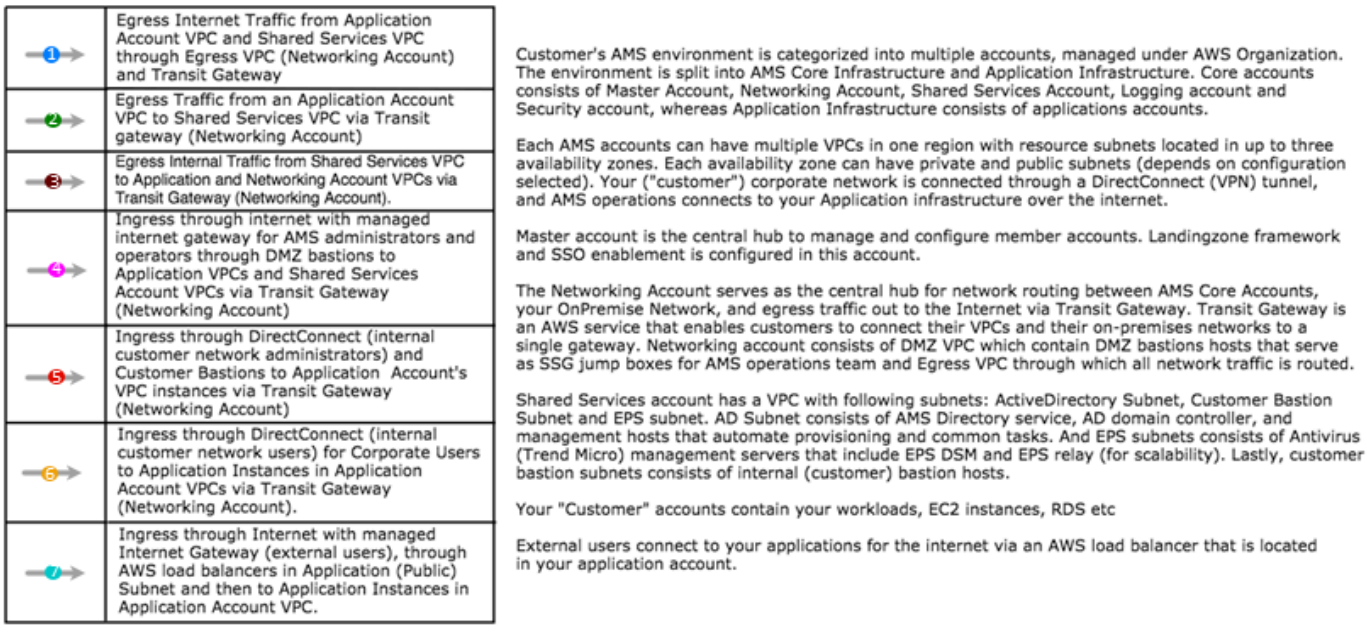

AMS configures all aspects of networking for you based on our standard templates and your selected options provided during onboarding. A standard AWS network design is applied to your AWS account, and a VPC is created for you and connected to AMS by either VPN or Direct Connect. For more information about Direct Connect, see AWS Direct Connect. Standard VPCs include the DMZ, shared services, and an application subnet. During the onboarding process, additional VPCs might be requested and created to match your needs (for example, customer divisions, partners). After onboarding, you are provided with a network diagram: an environment document that explains how your network has been set up.

> **Note**
>
> For information about default service limits and constraints for all active services, see the AWS Service Limits documentation.

Our network design is built around the Amazon "Principle of Least Privilege". In order to accomplish this, we route all traffic, ingress and egress, through a DMZ, except traffic coming from a trusted network. The only trusted network is the one configured between your on-premises environment and the VPC through the use of a VPN and/or an AWS Direct Connect (DX). Access is granted through the use of bastion instances, thereby preventing direct access to any production resources. All of your applications and resources reside inside private subnets that are reachable through public load balancers. Public egress traffic flows through the NAT Gateways in the egress

VPC (in the Networking account) to the Internet Gateway and then to the Internet. Alternatively, the traffic can flow over your VPN or Direct Connect to your on-premises environment.

Private network connectivity to AMS Multi-account landing zone environment

AWS offers private connectivity via either virtual private network (VPN) connectivity, or dedicated lines with AWS Direct Connect. Private connectivity in your multi-account environment, is set up using one of the methods described next:

- Centralized Edge connectivity using Transit Gateway
- Connecting Direct Connect (DX) and/or VPN to account virtual private clouds (VPCs)

Centralized edge connectivity using transit gateway

AWS Transit Gateway is a service that enables you to connect your VPCs and your on-premises networks to a single gateway. Transit gateway (TGW) can be used to consolidate your existing edge connectivity and route it through a single ingress/egress point. Transit gateway is created in the networking account of your AMS multi-account environment. For more details about transit gateway, see AWS Transit Gateway.

AWS Direct Connect (DX) gateway is used to connect your DX connection over a transit virtual interface to the VPCs or VPNs that are attached to your transit gateway. You associate a Direct Connect gateway with the transit gateway. Then, create a transit virtual interface for your AWS Direct Connect connection to the Direct Connect gateway. For information on DX virtual interfaces, see AWS Direct Connect Virtual Interfaces.

This configuration offers the following benefits. You can:

- Manage a single connection for multiple VPCs or VPNs that are in the same AWS Region.
- Advertise prefixes from on-premises to AWS, and from AWS to on-premises.

> **Note**
>
> For information about using a DX with AWS services, see the Resiliency Toolkit section Classic. For more information, see Transit Gateway associations.

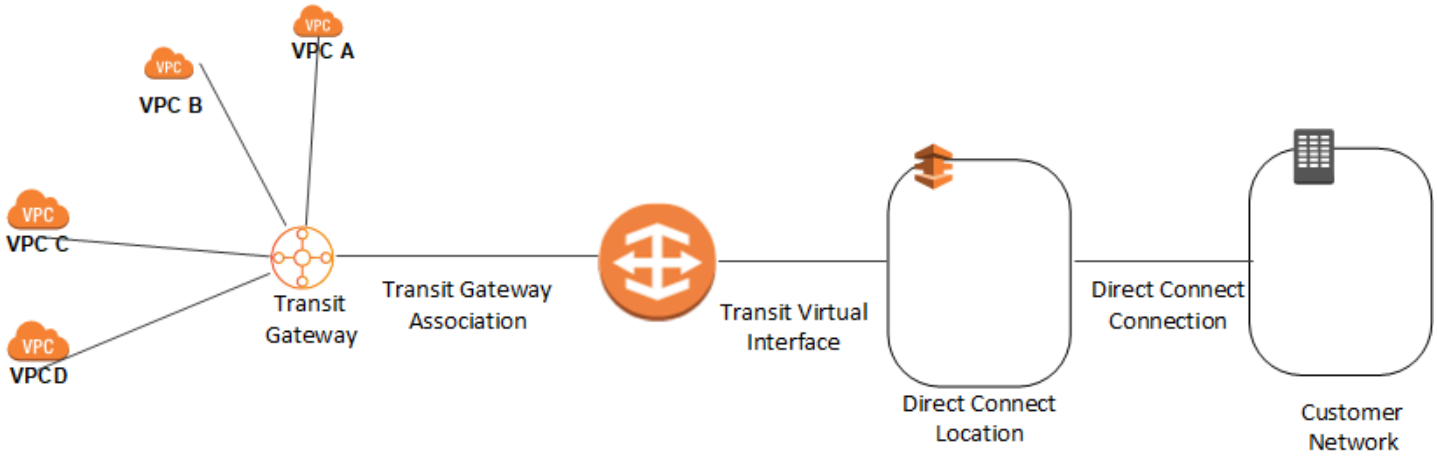

To increase the resiliency of your connectivity, we recommend that you attach at least two transit virtual interfaces from different AWS Direct Connect locations to the Direct Connect gateway. For more information, see the AWS Direct Connect resiliency recommendation.

Connecting DX or VPN to account VPCs

With this option, the VPCs in your AMS multi-account landing zone environments are directly connected to Direct Connect or VPN. The traffic directly flows from the VPCs to Direct Connect or VPN without traversing through the transit gateway.

Resources in the networking account

As shown in the networking account diagram, the following components are created in the account and require your input.

The Networking account contains two VPCs: **Egress VPC** and **DMZ VPC** also known as the **Perimeter** VPC.

AWS Network Manager

AWS Network Manager is a service that enables you to visualize your transit gateway (TGW) networks at no additional cost to AMS. It provides centralized network monitoring on both AWS resources and on on-premises networks, a single global view of their private network in a topology diagram and in a geographical map, and utilization metrics, such as bytes in/out, packets in/out, packets dropped, and alerts for changes in the topology, routing, and up/down connection status. For information, see AWS Network Manager.

Use one of the following roles to access this resource:

- AWSManagedServicesCaseRole

- AWSManagedServicesReadOnlyRole
- AWSManagedServicesChangeManagementRole

Egress VPC

The Egress VPC is primarily used for egress traffic to the Internet and is composed of public/private subnets in up to three availability zones (AZs). Network address translation (NAT) gateways are provisioned in the public subnets, and transit gateway (TGW) VPC attachments are created in the private subnets. Egress, or outbound, internet traffic from all networks enter through the private subnet via TGW, where it is then routed to a NAT via VPC route tables.

For your VPCs that contain public-facing applications in a public subnet, traffic originating from the internet is contained within that VPC. Return traffic is not routed to the TGW or Egress VPC, but routed back through the internet gateway (IGW) in the VPC.

> **Note**
>
> Networking VPC CIDR range: When you create a VPC, you must specify a range of IPv4 addresses for the VPC in the form of a Classless Inter-Domain Routing (CIDR) block; for example, 10.0.16.0/24. This is the primary CIDR block for your VPC.
> The AMS multi-account landing zone team recommends the range of 24 (with more IP address) to provide some buffer in case other resources/appliances, are deployed in the future.

Managed Palo Alto egress firewall

AMS provides a Managed Palo Alto egress firewall solution, which enables internet-bound outbound traffic filtering for all networks in the Multi-Account Landing Zone environment (excluding public facing services). This solution combines industry-leading firewall technology (Palo Alto VM-300) with AMS' infrastructure management capabilities to deploy, monitor, manage, scale, and restore infrastructure within compliant operating environments. Third parties, including Palo Alto Networks, do not have access to the firewalls; they are managed solely by AMS engineers.

Traffic control

The managed outbound firewall solution manages a domain allow-list composed of AMS-required domains for services such as backup and patch, as well as your defined domains. When outbound

internet traffic is routed to the firewall, a session is opened, traffic is evaluated, and if it matches an allowed domain, the traffic is forwarded to the destination.

Architecture

The managed egress firewall solution follows a high-availability model, where two to three firewalls are deployed depending on number of availability zones (AZs). The solution utilizes part of the IP space from the default egress VPC, but also provisions a VPC extension (/24) for additional resources required for managing the firewalls.

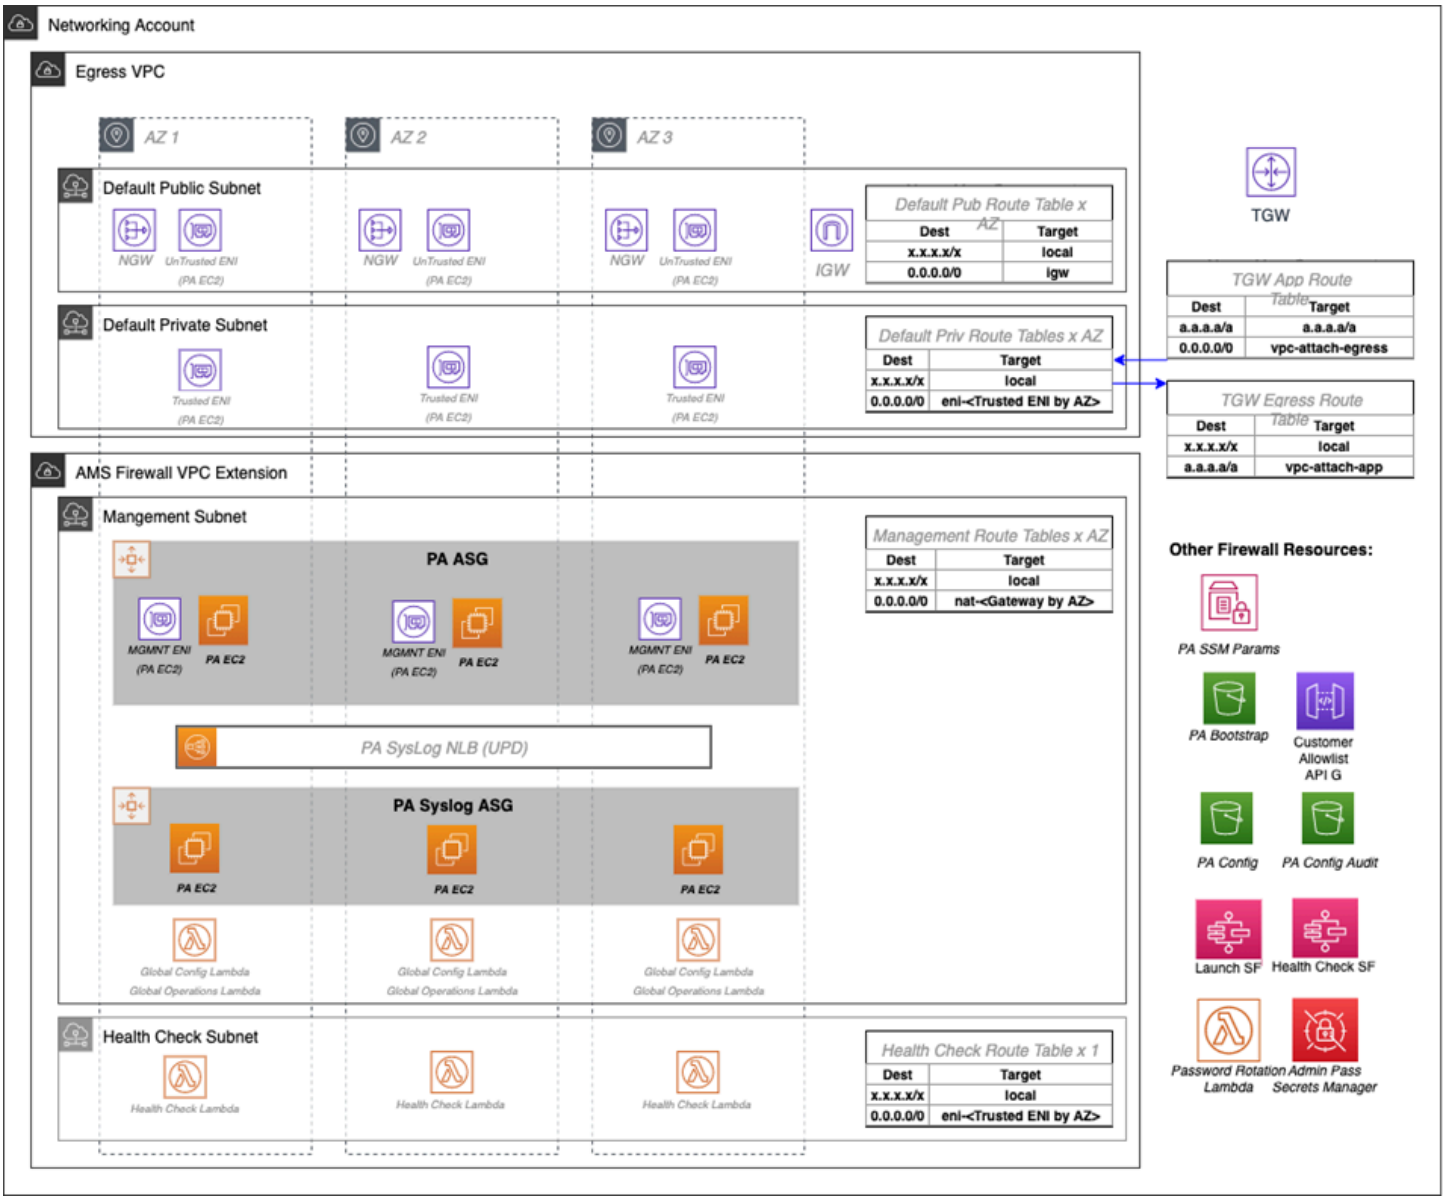

Network flow

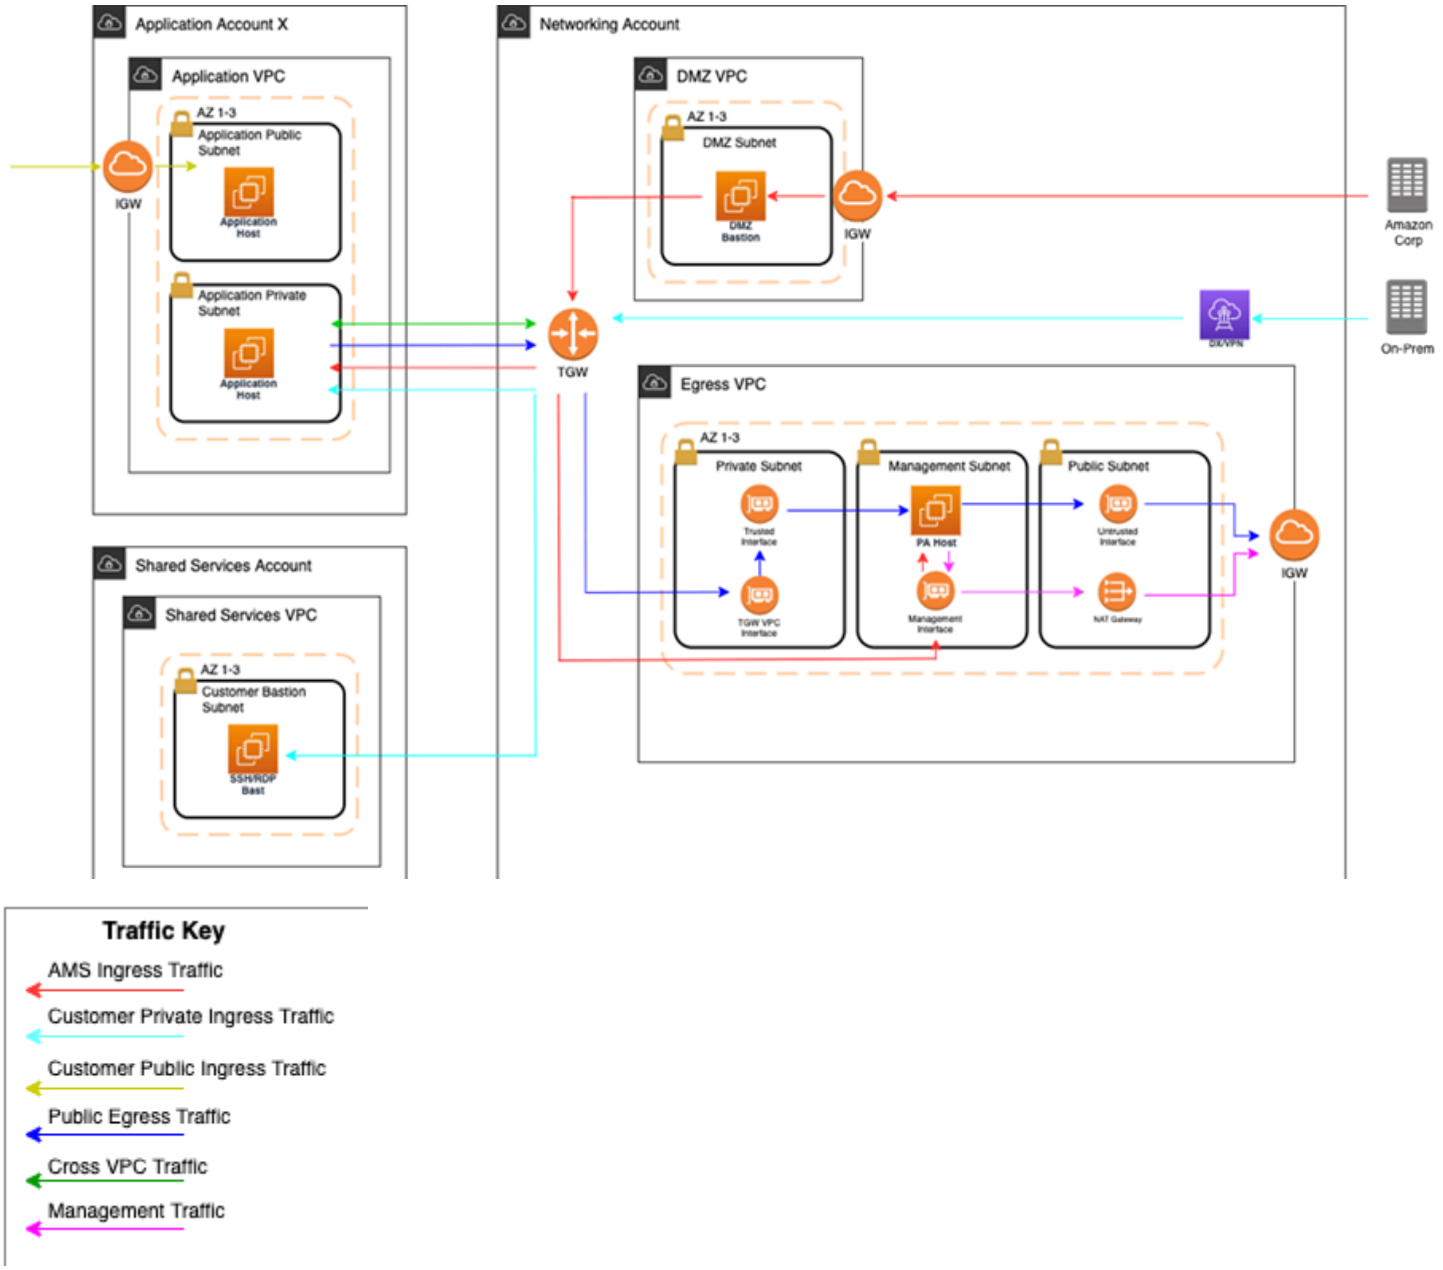

At a high level, public egress traffic routing remains the same, except for how traffic is routed to the internet from the egress VPC:

1. Egress traffic destined for the internet is sent to the Transit Gateway (TGW) through VPC route table
2. TGW routes traffic to the egress VPC via the TGW route table
3. VPC routes traffic to the internet via the private subnet route tables

a. In the default Multi-Account Landing Zone environment, internet traffic is sent directly to a network address translation (NAT) gateway. The managed firewall solution reconfigures the private subnet route tables to point the default route (0.0.0.0/0) to a firewall interface instead.

The firewalls themselves contain three interfaces:

1. Trusted interface: Private interface for receiving traffic to be processed.
2. Untrusted interface: Public interface to send traffic to the internet. Because the firewalls perform NAT, external servers accept requests from these public IP addresses.
3. Management interface: Private interface for firewall API, updates, console, and so on.

Throughout all the routing, traffic is maintained within the same availability zone (AZ) to reduce cross-AZ traffic. Traffic only crosses AZs when a failover occurs.

Allow-list modification

After onboarding, a default allow-list named `ams-allowlist` is created, containing AMS-required public endpoints as well as public endpoints for patching Windows and Linux hosts. Once operating, you can create RFC's in the AMS console under the Management | Managed Firewall | Outbound (Palo Alto) category to create or delete allow-lists, or modify the domains. Be aware that `ams-allowlist` cannot be modified. The RFC's are handled with full automation (they are not manual).

Custom security policy

Security policies determine whether to block or allow a session based on traffic attributes, such as the source and destination security zone, the source and destination IP address, and the service. Custom security policies are supported with fully automated RFCs. CTs to create or delete security policy can be found under Management | Managed Firewall | Outbound (Palo Alto) category, and the CT to edit an existing security policy can be found under Deployment | Managed Firewall | Outbound (Palo Alto) category. You'll be able to create new security policies, modify security policies, or delete security policies.

> **Note**
>
> The default security policy `ams-allowlist` cannot be modified

CloudWatch PA egress dashboards

Two dashboards can be found in CloudWatch to provide an aggregated view of Palo Alto (PA). the **AMS-MF-PA-Egress-Config-Dashboard** provides a PA config overview, links to allow-lists, and a list of all security policies including their attributes. The **AMS-MF-PA-Egress-Dashboard** can be customized to filter traffic logs. For example, to create a dashboard for a security policy, you can create an RFC with a filter like:

```
fields @timestamp, @message
| filter @logStream like /pa-traffic-logs/
| filter @message like /<Security Policy Name>/
| parse @message
 "*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*,*
 as x1, @x2, @x3, @x4, @type, @x6, @x7, @source_ip, @destination_ip, @source_nat_ip,
 @dest_nat_ip, @rule, @x13, @x14, @application, @x16, @from_zone, @to_zone,
 @x19, @x20, @x21, @x22, @session_id, @x24, @source_port, @destination_port,
 @source_nat_port, @destination_nat_port, @x29, @protocol, @action, @bytes,
 @bytes_sent, @bytes_recieved, @packets, @x36, @x37, @category, @x39, @x40, @x41,
 @source_country, @destination_country, @x44, @packets_sent, @packets_recieved,
 @session_end_reason, @x48, @x49, @x50
| display @timestamp, @rule, @action, @session_end_reason, @protocol, @source_ip,
 @destination_ip, @source_port, @destination_port, @session_id, @from_zone,
 @to_zone, @category, @bytes_sent, @bytes_recieved, @packets_sent, @packets_recieved,
 @source_country, @destination_country
```

Failover model

The firewalls solution includes two-three Palo Alto (PA) hosts (one per AZ). Healthy check canaries run on a constant schedule to evaluate the health of the hosts. If a host is identified as unhealthy, AMS is notified and the traffic for that AZ is automatically shifted to a healthy host in a different AZ via route table change. Since the health check workflow is running constantly, if the host becomes healthy again due to transient issues or manual remediation, then traffic is shifted back to the correct AZ with the healthy host.

Scaling

AMS monitors the firewall for throughput and scaling limits. When throughput limits exceed lower watermark thresholds (CPU/Networking), AMS receives an alert. A low watermaker threshold indicates that resources are approaching saturation, reaching a point where AMS will evaluate the metrics over time and reach out to suggest scaling solutions.

Backup and Restore

Backups are created during initial launch, after any configuration changes, and on a regular interval. Initial launch backups are created on a per host basis, but configuration change and regular interval backups are performed across all firewall hosts when the backup workflow is invoked. AMS engineers can create additional backups outside of those windows or provide backup details if requested.

AMS engineers can perform restoration of configuration backups if required. If a restoration is required, it will occur across all hosts to keep configuration between hosts in sync.

Restoration also can occur when a host requires a complete recycle of an instance. An automatic restoration of the latest backup occurs when a new EC2 instance is provisioned. In general, hosts are not recycled regularly, and are reserved for severe failures or required AMI swaps. Host recycles are initiated manually, and you are notified before a recycle occurs.

Other than the firewall configuration backups, your specific allow-list rules are backed up separately. A backup is automatically created when your defined allow-list rules are modified. Restoration of the allow-list backup can be performed by an AMS engineer, if required.

Updates

AMS Managed Firewall Solution requires various updates over time to add improvements to the system, additional features, or updates to the firewall operating system (OS) or software.

Most changes will not affect the running environment such as updating automation infrastructure, but other changes such as firewall instance rotation or OS update may cause disruption. When a potential service disruption due to updates is evaluated, AMS will coordinate with you to accommodate maintenance windows.

Operator access

AMS operators use their ActiveDirectory credentials to log into the Palo Alto device to perform operations (e.g., patching, responding to an event, etc.). The solution retains standard AMS Operator authentication and configuration change logs to track actions performed on the Palo Alto Hosts.

Default logs

By default, the logs generated by the firewall reside in local storage for each firewall. Overtime, local logs will be deleted based on storage utilization. The AMS solution provides real-time

shipment of logs off of the machines to CloudWatch logs; for more information, see CloudWatch Logs integration.

AMS engineers still have the ability to query and export logs directly off the machines if required. In addition, logs can be shipped to a customer-owned Panorama; for more information, see Panorama integration.

The Logs collected by the solution are the following:

RFC Status Codes

Log Type	Description
Traffic	Displays an entry for the start and end of each session. Each entry includes the date and time, source and destination zones, addresses and ports, application name, security rule name applied to the flow, rule action (allow, deny, or drop), ingress and egress interface, number of bytes, and session end reason. The Type column indicates whether the entry is for the start or end of the session, or whether the session was denied or dropped. A "drop" indicates that the security rule that blocked the traffic specified "any" application, while a "deny" indicates the rule identified a specific application. If traffic is dropped before the application is identified, such as when a rule drops all traffic for a specific service, the application is shown as "not-applicable".
Threat	Displays an entry for each security alarm generated by the firewall. Each entry includes the date and time, a threat name or URL, the source and destination zones, addresses, and ports, the application name, and the alarm action (allow or block) and severity. The Type column indicates the type of threat, such as "virus" or "spyware;" the Name column is the threat description or URL; and the Category column is the threat category (such as "keylogger") or URL category.

Log Type	Description
URL Filtering	Displays logs for URL filters, which control access to websites and whether users can submit credentials to websites.
Configuration	Displays an entry for each configuration change. Each entry includes the date and time, the administrator user name, the IP address from where the change was made, the type of client (web interface or CLI), the type of command run, whether the command succeeded or failed, the configuration path, and the values before and after the change.
System	Displays an entry for each system event. Each entry includes the date and time, the event severity, and an event description.
Alarms	The alarms log records detailed information on alarms that are generated by the system. The information in this log is also reported in Alarms. Refer to "Define Alarm Settings".
Authentication	Displays information about authentication events that occur when end users try to access network resources for which access is controlled by Authentication policy rules. Users can use this information to help troubleshoot access issues and to adjust user Authentication policy as needed. In conjunction with correlation objects, users can also use Authentication logs to identify suspicious activity on the users network, such as brute force attacks. Optionally, users can configure Authentication rules to Log Authentic ation Timeouts. These timeouts relate to the period of time when a user needs authenticate for a resource only once but can access it repeatedly. Seeing information about the timeouts helps users decide if and how to adjust them.

Log Type	Description
Unified	Displays the latest Traffic, Threat, URL Filtering, WildFire Submissio ns, and Data Filtering log entries in a single view. The collective log view enables users to investigate and filter these different types of logs together (instead of searching each log set separately). Or, users can choose which log types to display: click the arrow to the left of the filter field and select traffic, threat, url, data, and/or wildfire to display only the selected log types.

Event management

AMS continually monitors the capacity, health status, and availability of the firewall. Metrics generated from the firewall, as well as AWS/AMS generated metrics, are used to create alarms that are received by AMS operations engineers, who will investigate and resolve the issue. The current alarms cover the following cases:

Event Alarms:

- Firewall Dataplane CPU Utilization
  - CPU Utilization - Dataplane CPU (Processing traffic)
- Firewall Dataplane Packet Utilization is above 80%
  - Packet utilization - Dataplane (Processing traffic)
- Firewall Dataplane Session Utilization
- Firewall Dataplane Session Active
- Aggregate Firewall CPU Utilization
  - CPU Utilization across all CPUs
- Failover By AZ
  - Alarms when a fail over occurs in an AZ
- Unhealthy Syslog Host
  - Syslog host fails health check

Management Alarms:

- Health Check Monitor Failure Alarm

- When health check workflow fails unexpectedly
- This is for the workflow itself, not if a firewall health check fails

- Password Rotation Failure Alarm
  - When password rotation fails
  - API/Service user password is rotated every 90 days

Metrics

All metrics are captured and stored in CloudWatch in the Networking account. These can be viewed by gaining console access to the Networking account and navigating to the CloudWatch console. Individual metrics can be viewed under the metrics tab or a single-pane dashboard view of select metrics and aggregated metrics can be viewed by navigating to the Dashboard tab, and selecting **AMS-MF-PA-Egress-Dashboard**.

Custom Metrics:

- Health Check
  - Namespace: AMS/MF/PA/Egress
    - PARouteTableConnectionsByAZ
    - PAUnhealthyByInstance
    - PAUnhealthyAggregatedByAZ
    - PAHealthCheckLockState
- Firewall Generated
  - Namespace: AMS/MF/PA/Egress/<instance-id>
    - DataPlaneCPUUtilizationPct
    - DataPlanePacketBuffferUtilization
    - panGPGatewayUtilizationPct
    - panSessionActive
    - panSessionUtilization

CloudWatch Logs integration

CloudWatch Logs integration forwards logs from the firewalls into CloudWatch Logs, which mitigates the risk of losing logs due to local storage utilization. Logs are populated in real-time as the firewalls generate them, and can be viewed on-demand through the console or API.

Complex queries can be built for log analysis or exported to CSV using CloudWatch Insights. In addition, the custom AMS Managed Firewall CloudWatch dashboard will also show a quick view of specific traffic log queries and a graph visualization of traffic and policy hits over time. Utilizing CloudWatch logs also enables native integration to other AWS services such as a AWS Kinesis.

> **Note**
>
> PA logs cannot be directly forwarded to an existing on-prem or 3rd party Syslog collector. AMS Managed Firewall solution provides real-time shipment of logs off of the PA machines to AWS CloudWatch Logs. You can use CloudWatch Logs Insight feature to run ad-hoc queries. In addition, logs can be shipped to your Palo Alto's Panorama management solution. CloudWatch logs can also be forwarded to other destinations using CloudWatch Subscription Filters. Learn more about Panorama in the following section. To learn more about Splunk, see Integrating with Splunk.

Panorama integration

AMS Managed Firewall can, optionally, be integrated with your existing Panorama. This allows you to view firewall configurations from Panorama or forward logs from the firewall to the Panorama. Panorama integration with AMS Managed Firewall is read only, and configuration changes to the firewalls from Panorama are not allowed. Panorama is completely managed and configured by you, AMS will only be responsible for configuring the firewalls to communicate with it.

Licensing

The price of the AMS Managed Firewall depends on the type of license used, hourly or bring your own license (BYOL), and the instance size in which the appliance runs. You are required to order the instances size and the licenses of the Palo Alto firewall you prefer through AWS Marketplace.

- Marketplace Licenses: Accept the terms and conditions of the VM-Series Next-Generation Firewall Bundle 1 from the networking account in MALZ.

- BYOL Licenses: Accept the terms and conditions of the VM-Series Next-Generation Firewall (BYOL) from the networking account in MALZ and share the "BYOL auth code" obtained after purchasing the license to AMS.

Limitations

At this time, AMS supports VM-300 series or VM-500 series firewall. Configurations can be found here: VM-Series Models on AWS EC2 Instances,

> **Note**
>
> The AMS solution runs in Active-Active mode as each PA instance in its AZ handles egress traffic for their respected AZ. So, with two AZs, each PA instance handles egress traffic up to 5 Gbps and effectively provides overall 10 Gbps throughput across two AZs. The same is true for all limits in each AZ. Should the AMS health check fail, we shift traffic from the AZ with the bad PA to another AZ, and during the instance replacement, capacity is reduced to the remaining AZs limits.
> AMS does not currently support other Palo Alto bundles available on AWS Marketplace; for example, you cannot ask for the "VM-Series Next-Generation Firewall Bundle 2". Note that the AMS Managed Firewall solution using Palo Alto currently provides only an egress traffic filtering offering, so using advanced VM-Series bundles would not provide any additional features or benefits.

Onboarding requirements

- You must review and accept the Terms and Conditions of the VM-Series Next-Generation Firewall from Palo Alto in AWS Marketplace.
- You must confirm the instance size you want to use based on your expected workload.
- You must provide a /24 CIDR Block that does not conflict with networks in your Multi-Account Landing Zone environment or On-Prem. It must be of same class as the Egress VPC (the Solution provisions a /24 VPC extension to the Egress VPC).

Pricing

AMS Managed Firewall base infrastructure costs are divided in three main drivers: the EC2 instance that hosts the Palo Alto firewall, the software license Palo Alto VM-Series licenses, and CloudWatch Integrations.

The following pricing is based on the VM-300 series firewall.

- EC2 Instances: The Palo Alto firewall runs in a high-availability model of 2-3 EC2 instances, where instance is based on expected workloads. Cost for the instance depends on the region and number of AZs
  - Ex. us-east-1, m5.xlarge, 3AZs
    - $0.192 * 24 * 30 * 3 = $414.72
  - https://aws.amazon.com/ec2/pricing/on-demand/
- Palo Alto Licenses: The software license cost of a Palo Alto VM-300 next-generation firewall depends on the number of AZ as well as instance type.
  - Ex. us-east-1, m5.xlarge, 3AZs
    - $0.87 * 24 * 30 * 3 = $1879.20
    - https://aws.amazon.com/marketplace/pp/B083M7JPKB?ref_=srh_res_product_title#pdp-pricing
- CloudWatch Logs Integration: CloudWatch logs integration utilizes SysLog servers (EC2 - t3.medium), NLB, and CloudWatch Logs. The cost of the servers is based on region and number of AZs, and the cost of the NLB/CloudWatch logs varies based on traffic utilization.
  - Ex. us-east-1, t3.medium, 3AZ
    - $0.0416 * 24 * 30 * 3 = $89.86
  - https://aws.amazon.com/ec2/pricing/on-demand/
  - https://aws.amazon.com/cloudwatch/pricing/

Perimeter (DMZ) VPC

The Perimeter, or DMZ, VPC contains the necessary resources for AMS Operations engineers to access AMS networks. It contains public subnets across 2-3 AZs, with SSH Bastions hosts in an Auto Scaling group (ASG) for AMS Operations engineers to log into or tunnel through. The security groups attached to the DMZ bastions contain port 22 inbound rules from **Amazon Corp Networks**.

DMZ VPC CIDR range: When you create a VPC, you must specify a range of IPv4 addresses for the VPC in the form of a Classless Inter-Domain Routing (CIDR) block; for example, 10.0.16.0/24. This is the primary CIDR block for your VPC.

> **Note**
>
> The AMS team recommends the range of 24 (with more IP address) to provide some buffer in case other resources, such as a firewall, are deployed in the future.

AWS Transit Gateway

AWS Transit Gateway (TGW) is a service that enables you to connect your Amazon Virtual Private Clouds (VPCs) and your on-premises networks to a single gateway. Transit gateway is the networking backbone that handles the routing between AMS account networks and external networks. For information about Transit Gateway, see AWS Transit Gateway.

Provide the following input to create this resource:

- *Transit Gateway ASN number**: Provide the private Autonomous System Number (ASN) for your transit gateway. This should be the ASN for the AWS side of a Border Gateway Protocol (BGP) session. The range is 64512 to 65534 for 16-bit ASNs.

Shared Services account

The Shared Services account serves as the central hub for most AMS data plane services. The account contains infrastructure and resources required for access management (AD), end-point security management (Trend Micro), and it contains the customer bastions (SSH/RDP). A high-level overview of the resources contained within Shared Services Account is shown in the following graphic.

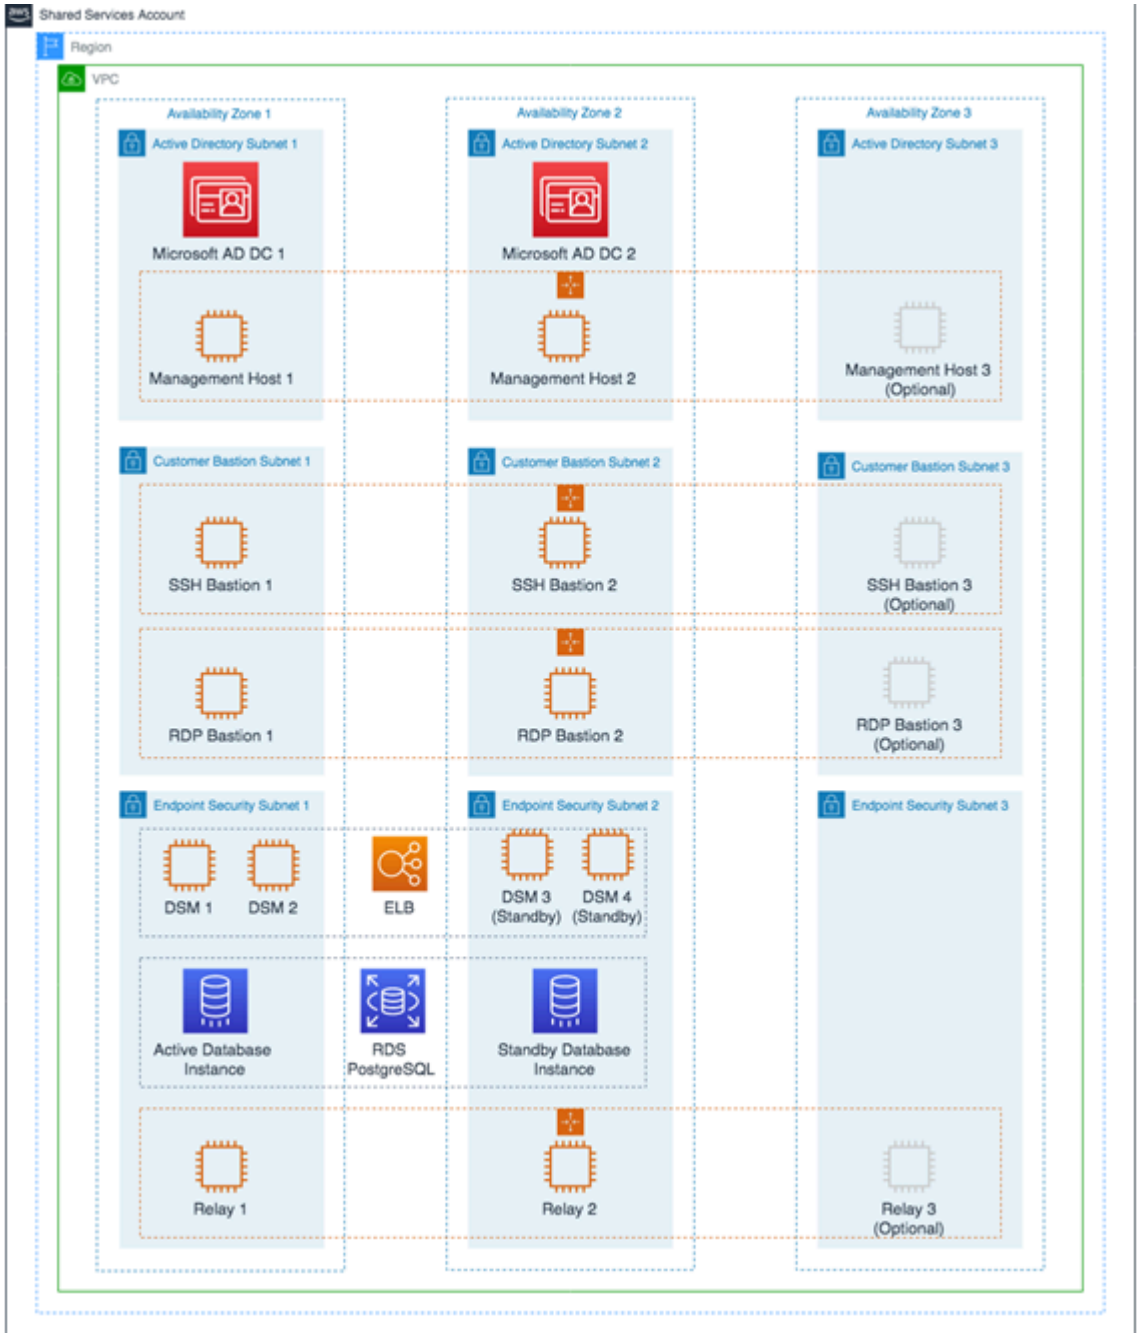

The Shared Services VPC is composed of the AD subnet, the EPS subnet, and the customer bastions subnet in the three availability zones (AZs). The resources created in the Shared Services VPC are listed below and require your input.

- *Shared Services VPC CIDR range:* When you create a VPC, you must specify a range of IPv4 addresses for the VPC in the form of a Classless Inter-Domain Routing (CIDR) block; for example, 10.0.1.0/24. This is the primary CIDR block for your VPC.

> **Note**
>
> The AMS team recommends the range of /23.

- *Active Directory Details*: Microsoft Active Directory (AD) is utilized for user/resource management, authentication/authorization, and DNS, across all of your AMS multi-account landing zone accounts. AMS AD is also configured with a one-way trust to your Active Directory for trust-based authentication. The following input is required to create the AD:
  - Domain Fully Qualified Domain Name (FQDN): The fully qualified domain name for the AWS Managed Microsoft AD directory. The domain should not be an existing domain or child domain of an existing domain in your network.
  - Domain NetBIOS Name: If you don't specify a NetBIOS name, AMS defaults the name to the first part of your directory DNS. For example, corp for the directory DNS corp.example.com.

- *Trend Micro – endpoint protection security (EPS)*: Trend Micro endpoint protection (EPS) is the primary component within AMS for operating system security. The system is comprised of Deep Security Manager (DSM), EC2 instances, relay EC2 instances, and an agent present within all data plane and customer EC2 instances.

  You must assume the `EPSMarketplaceSubscriptionRole` in the Shared Services account, and subscribe to either the Trend Micro Deep Security (BYOL) AMI, or the Trend Micro Deep Security (Marketplace).

  The following default inputs are required to create EPS (if you want to change from the defaults):
  - Relay Instance Type: Default Value - m5.large
  - DSM Instance Type: Default Value - m5.xlarge
  - DB Instance Size: Default Value - 200 GB
  - RDS Instance Type: Default Value - db.m5.large

- *Customer bastions*: You are provided with SSH or RDP bastions (or both) in the Shared Services Account, to access other hosts in your AMS environment. In order to access the AMS network as a user (SSH/RDP), you must use "customer" Bastions as the entry point. The network path originates from the on-premise network, goes through DX/VPN to the transit gateway (TGW),

and then is routed to the Shared Services VPC. Once you are able to access the bastion, you can jump to other hosts in the AMS environment, provided that the access request has been granted.

- The following inputs are required for SSH bastions.
  - SSH Bastion Desired Instance Capacity: Default Value - 2.
  - SSH Bastion Maximum Instances: Default Value - 4.
  - SSH Bastion Minimum Instances: Default Value -2.
  - SSH Bastion Instance Type: Default Value - m5.large (can be changed to save costs; for example a t3.medium).
  - SSH Bastion Ingress CIDRs: IP address ranges from which users in your network access SSH Bastions.
- The following inputs are required for Windows RDP bastions.
  - RDP Bastion Instance Type: Default Value - t3.medium.
  - RDP Bastion Desired Minimum Sessions: Default Value - 2.
  - RDP Maximum Sessions: Default Value -10.
  - RDP Bastion Configuration Type: You can choose one of the below configuration
    - SecureStandard = A user receives one bastion and only one user can connect to the bastion.
    - SecureHA = A user receives two bastions in two different AZ's to connect to and only one user can connect to the bastion.
    - SharedStandard = A user receives one bastion to connect to and two users can connect to the same bastion at once.
    - SharedHA = A user receives two bastions in two different AZ's to connect to and two users can connect to the same bastion at once.
  - Customer RDP Ingress CIDRs: IP address ranges from which users in your network will access RDP Bastions.

Updates to shared services: Multi-Account Landing Zone

AMS applies data plane releases to managed accounts on a monthly basis, without prior notice.

AMS uses the core OU to provide shared services such as access, networking, EPS, log storage, alert aggregation in your Multi-Account Landing Zone. AMS is responsible for addressing vulnerabilities, patching, and deployments of these shared services. AMS regularly updates the resources used for

providing these shared services so that users have access to latest features, and security updates. The updates typically happen on a monthly basis. Resources that are part of these updates are:

- Accounts that are part of the core OU.

  The management account, shared services account, network account, security account, and log archive account have resources for RDP and SSH bastions, proxies, management hosts, and endpoint security (EPS), that are typically updated every month. AMS uses immutable EC2 deployments as part of the shared services infrastructure.
- New AMS AMIs incorporating the latest updates.

> **Note**
>
> AMS operators utilize an internal alarm suppression change type (CT) when executing data plane changes and the RFC for that CT appears in your RFC list. This is because, as the data plane release is deployed, various infrastructure may be shut down, rebooted, taken offline, or there may be CPU spikes or other effects of the deployment that trigger alarms that, during the data plane deployment, are extraneous. Once the deployment is complete, all infrastructure is verified to be running properly and alarms are re-enabled.

Log Archive account

The Log Archive account serves as the central hub for archiving logs across your AMS multi-account landing zone environment. There is an S3 bucket in the account that contains copies of AWS CloudTrail and AWS Config log files from each of the AMS multi-account landing zone environment accounts. You could use this account for your Centralised Logging solution with AWS Firehose, or Splunk, and so forth. AMS access to this account is limited to a few users; restricted to auditors and security teams for compliance and forensic investigations related to account activity.

Security account

The Security account is the central hub for housing security related operations and the main point for funneling notifications and alerts to the AMS control plane services. In addition, the Security account houses the Amazon Guard Duty management account and the AWS Config aggregator.

Application account types

Application accounts are AWS accounts within the AMS-managed landing zone architecture that you use to host your workloads. AMS offers three types of Application accounts:

- AMS-managed application accounts
- AMS Accelerate accounts
- Customer Managed application accounts

Application accounts are grouped in different OUs in AWS Organizations depending on the Application account type:

- Root OU:
  1. Applications OU
     - Managed OU: AMS-managed accounts
     - Development OU: AMS-managed accounts with Developer mode enabled
  2. Accelerate OU: AMS Accelerate Application accounts
  3. Customer-managed OU: Customer-managed Application accounts

Application accounts are provisioned through an RFC submitted from the Management account:

- Create Application Account With VPC ct-1zdasmc2ewzrs
- Create Accelerate Account ct-2p93tyd5angmi
- Create Customer-Managed Application Account ct-3pwbixz27n3tn

AMS-managed application accounts

Application accounts that are fully managed by AMS are referred to as AMS-managed application accounts, where some or all operational tasks, like service request management, incident management, security management, continuity management (backup), patch management, cost-optimization, or monitoring and event management of infrastructure, are performed by AMS.

The amount of tasks performed by AMS depends on the Change Management mode that you select. AMS-managed accounts support different modes for change management:

- RFC mode
- Direct Change mode in AMS
- AMS and AWS Service Catalog
- AMS Advanced Developer mode
- Self-Service Provisioning mode

For more information about change management and different modes, see Change management modes.

There are some AWS services that you can use in your AMS-managed account without AMS management. The list of these AWS services and how to add them into your AMS account are described in the Self-service provisioning section.

AMS Accelerate accounts

AMS Accelerate is the AMS operations plan that can operate AWS infrastructure supporting workloads. You can benefit from AMS Accelerate operational services such as monitoring and alerting, incident management, security management, and backup management, without going through a new migration, experiencing downtime, or changing how you use AWS. AMS Accelerate also offers an optional patch add-on for EC2 based workloads that require regular patching.

With AMS Accelerate you have the freedom to use, configure, and deploy all AWS services natively, or with your preferred tools. You will use your preferred access and change mechanisms while AMS consistently applies proven practices that help scale your team, optimize costs, increase security and efficiency, and improve resiliency.

> **Note**
>
> AMS Accelerate accounts in AMS Advanced do not have AMS change management (RFCs) or the AMS Advanced console. Instead, they have the AMS Accelerate console and functionality.

Accelerate accounts can only be provisioned from your AMS multi-account landing zone Management account. Accelerate offers different operational capabilities. To learn more see the Accelerate service description.

- You will continue to enjoy some of the features from the multi-account landing zone (MALZ) core accounts such as centralized logging, single billing, Config Aggregator in the security account and SCPs.
- AMS Accelerate does not provide some AMS Advanced services like EPS, Access management, Change management and provisioning. We recommend you follow the next steps to gain access and configure the transit gateway (TGW).

For more details about Accelerate, see What is Accelerate.

Creating your Accelerate account

To create an Accelerate account, follow the steps outlined here Create an Accelerate account.

Accessing your Accelerate account

After you provision an Accelerate account in your multi-account landing zone (MALZ) account, a role with Administrative access permissions, `AccelerateDefaultAdminRole`, is in the account for you to assume.

To access the new Accelerate account:

1. Log into the IAM console for the management account with the `CustomerDefaultAssumeRole` role.

2. In the IAM console, on the navigation bar, choose your username.
3. Choose **Switch Role**. If this is the first time choosing this option, a page appears with more information. After reading it, choose **Switch Role**. If you clear your browser cookies, this page can appear again.
4. On the **Switch Role** page, type the Accelerate account ID and the name of the role to assume: `AccelerateDefaultAdminRole`.

Now that you have access, you can create new IAM Roles to continue to access your environment. If you would like to leverage SAML Federation for your Accelerate account, see Enabling SAML 2.0 federated users to access the AWS Management Console.

Connecting your Accelerate account with Transit Gateway

AMS does not manage the network setup of an Accelerate account. You have the option of managing your own network using AWS APIs (see Networking Solutions) or connecting to the MALZ network managed by AMS, using the existing Transit Gateway (TGW) deployed in AMS MALZ.

> **Note**
>
> You can only have a VPC attached to the TGW if the Accelerate account is in the same AWS Region. For more information see Transit gateways.

To add your Accelerate account to Transit Gateway, request a new route using the Deployment | Managed landing zone | Networking account | Add static route (ct-3r2ckznmt0a59) change type, include this information:

- **Blackhole**: True to indicate that the route's target isn't available. Do this when the traffic for the static route is to be dropped by the Transit Gateway. False to route the traffic to the specified TGW attachment ID. Default value is false.
- **DestinationCidrBlock**: The IPV4 CIDR range used for destination matches. Routing decisions are based on the most specific match. Example: 10.0.2.0/24.
- **TransitGatewayAttachmentId**: The TGW Attachment ID that will serve as the route table target. If **Blackhole** is false, this parameter is required, otherwise leave this parameter blank. Example: tgw-attach-04eb40d1e14ec7272.
- **TransitGatewayRouteTableId**: The ID of the TGW route table. Example: tgw-rtb-06ddc751c0c0c881c.

Create routes in the TGW route tables to connect to this VPC:

1. By default this VPC will not be able to communicate with any of the other VPCs in your MALZ network.
2. Decide with your solutions architect what VPCs you want this Accelerate VPC to communicate with.
3. Submit a Deployment | Managed landing zone | Networking account | Add static route (ct-3r2ckznmt0a59) change type, include this information:
   - **Blackhole**: True to indicate that the route's target isn't available. Do this when the traffic for the static route is to be dropped by the Transit Gateway. False to route the traffic to the specified TGW attachment ID. Default value is false.
   - **DestinationCidrBlock**: The IPV4 CIDR range used for destination matches. Routing decisions are based on the most specific match. Example: 10.0.2.0/24.
   - **TransitGatewayAttachmentId**: The TGW Attachment ID that will serve as the route table target. If **Blackhole** is false, this parameter is required, otherwise leave this parameter blank. Example: tgw-attach-04eb40d1e14ec7272.
   - **TransitGatewayRouteTableId**: The ID of the TGW route table. Example: tgw-rtb-06ddc751c0c0c881c.

Connecting a new Accelerate account VPC to the AMS Multi-Account Landing Zone network (creating a TGW VPC attachment):

1. In your multi-account landing zone Networking account, open the Amazon VPC console.
2. On the navigation pane, choose **Transit Gateways**. Record the TGW ID of the transit gateway you see.
3. In your Accelerate account, open the Amazon VPC console.
4. In the navigation pane, choose **Transit Gateway Attachments** > **Create Transit Gateway Attachment**. Make these choices:
   - For the **Transit Gateway ID**, choose the transit gateway ID you recorded in Step 2.
   - For **Attachment type**, choose **VPC**.
   - Under **VPC Attachment**, optionally type a name for **Attachment name tag**.
   - Choose whether to enable **DNS Support** and **IPv6 Support**.
   - For **VPC ID**, choose the VPC to attach to the transit gateway. This VPC must have at least one subnet associated with it.

- For **Subnet IDs**, select one subnet for each Availability Zone to be used by the transit gateway to route traffic. You must select at least one subnet. You can select only one subnet per Availability Zone.

5. Choose **Create attachment**. Record the ID of the newly created TGW Attachment.

Associating the TGW attachment to a route table:

1. Decide which TGW route table you want to associate the VPC with. We recommend creating a new application route table for Accelerate account VPCs using Deployment | Managed landing zone | Networking account | Create transit gateway route table (ct-3dscwaeyi6cup) change type.
2. Submit a Management | Managed landing zone | Networking account | Associate TGW attachment (ct-3nmhh0qr338q6) RFC on the Networking account to associate the VPC or TGW attachment to the route table you select.

Create routes in the TGW route tables to connect to this VPC:

1. By default, this VPC will not be able to communicate with any of the other VPCs in your multi-account landing zone network.
2. Decide with your solutions architect what VPCs you want this Accelerate account VPC to communicate with.
3. Submit a Deployment | Managed landing zone | Networking account | Add static route (ct-3r2ckznmt0a59) RFC against the networking account to create the TGW routes you need.

Configuring your VPC Route tables to point at the AMS multi-account landing zone transit gateway:

1. Decide with your solutions architect what traffic you want to send to the AMS Multi-Account Landing Zone transit gateway.
2. Submit a Deployment | Managed landing zone | Networking account | Add static route (ct-3r2ckznmt0a59) RFC against the networking account to create the TGW routes you need.

Customer Managed application accounts

You can create accounts that AMS doesn't manage in the standard way. Those accounts are called Customer Managed accounts and they give you full control to self-operate the infrastructure within the accounts while enjoying the benefits of the centralized architecture managed by AMS.

Customer Managed accounts do not have access to the AMS console or any of the services we provide (patch, backup, and so on).

Customer Managed accounts can only be provisioned from your AMS multi-account landing zone management account.

Different AMS modes work with Application accounts differently; to learn more about the modes, see AWS Managed Services modes.

To create your Customer Managed application account, see Management account | Create Customer-Managed Application Account.

To delete a Customer Managed application account, use Management account | Offboard Application Account. (The Confirm Offboarding CT does not apply to Customer Managed application accounts.)

Accessing your Customer Managed account

After you provision a Customer Managed account (CMA) in multi-account landing zone, (MALZ) an Admin role, `CustomerDefaultAdminRole`, is in the account for you to assume, through SAML federation, to configure the account.

To access the CMA:

1. Log into the IAM console for the management account with the **CustomerDefaultAssumeRole** role.
2. In the IAM console, on the navigation bar, choose your username.
3. Choose **Switch Role**. If this is the first time choosing this option, a page appears with more information. After reading it, choose **Switch Role**. If you clear your browser cookies, this page can appear again.
4. On the **Switch Role** page, type the Customer Managed account ID and the name of the role to assume: **CustomerDefaultAdminRole**.

Now that you have access, you can create new IAM Roles to continue to access your environment. If you would like to leverage SAML Federation for your CMA Account, see Enabling SAML 2.0 federated users to access the AWS Management Console.

Connecting your CMA with Transit Gateway

AMS does not manage the network setup of Customer Managed accounts (CMAs). You have the option of managing your own network using AWS APIs (see Networking Solutions) or connecting to the multi-account landing zone network managed by AMS, using the existing Transit Gateway (TGW) deployed in AMS MALZ.

> **Note**
>
> You can only have a VPC attached to the TGW if the CMA is in the same AWS Region. For more information see Transit gateways.

To add your CMA to Transit Gateway, request a new route with the Networking account | Add static route (ct-3r2ckznmt0a59) change type and include this information:

- **Blackhole**: True to indicate that the route's target isn't available. Do this when the traffic for the static route is to be dropped by the Transit Gateway. False to route the traffic to the specified TGW attachment ID. Default value is false.
- **DestinationCidrBlock**: The IPV4 CIDR range used for destination matches. Routing decisions are based on the most specific match. Example: `10.0.2.0/24`.
- **TransitGatewayAttachmentId**: The TGW Attachment ID that will serve as route table target. If **Blackhole** is false, this parameter is required, otherwise leave this parameter blank. Example: `tgw-attach-04eb40d1e14ec7272`.
- **TransitGatewayRouteTableId**: The ID of the TGW route table. Example: `tgw-rtb-06ddc751c0c0c881c`.

Connecting a new customer-managed VPC to the AMS Multi-Account Landing Zone network (creating a TGW VPC attachment):

1. In your multi-account landing zone Networking account, open the Amazon VPC console.
2. In the navigation pane, choose **Transit Gateways**. Record the TGW ID of the transit gateway you see.

3. In your Customer Managed account, open the Amazon VPC console.
4. In the navigation pane, choose **Transit Gateway Attachments** > **Create Transit Gateway Attachment**. Make these choices:
    a. For the **Transit Gateway ID**, choose the transit gateway ID you recorded in Step 2.
    b. For **Attachment type**, choose **VPC**.
    c. Under **VPC Attachment**, optionally type a name for **Attachment name tag**.
    d. Choose whether to enable **DNS Support** and **IPv6 Support**.
    e. For **VPC ID**, choose the VPC to attach to the transit gateway. This VPC must have at least one subnet associated with it.
    f. For **Subnet IDs**, select one subnet for each Availability Zone to be used by the transit gateway to route traffic. You must select at least one subnet. You can select only one subnet per Availability Zone.
5. Choose **Create attachment**. Record the ID of the newly created TGW Attachment.

Associating the TGW attachment to a route table:

Decide which TGW route table you want to associate the VPC with. We recommend creating a new application route table for Customer Managed VPCs by submitting a Deployment | Managed landing zone | Networking account | Create transit gateway route table (ct-3dscwaeyi6cup) RFC. To associate the VPC or TGW attachment to the route table you select, submit a Deployment | Managed landing zone | Networking account | Associate TGW attachment (ct-3nmhh0qr338q6) RFC on the Networking account.

Create routes in the TGW route tables to connect to this VPC:

1. By default, this VPC will not be able to communicate with any of the other VPCs in your Multi-Account Landing Zone network.
2. Decide with your solutions architect what VPCs you want this customer-managed VPC to communicate with. Submit a Deployment | Managed landing zone | Networking account | Add static route (ct-3r2ckznmt0a59) RFC against the networking account to create the TGW routes you need.

> **Note**
>
> This CT (ct-3r2ckznmt0a59) does not allow adding static routes to core route table EgressRouteDomain; if your CMA needs to allow egress traffic, submit a Management | Other | Other (MOO) RFC with ct-0xdawir96cy7k.

Configuring your VPC Route tables to point at the AMS Multi-Account Landing Zone transit gateway:

Decide with your solutions architect what traffic you want to send to the AMS Multi-Account Landing Zone transit gateway. Update your VPC route tables to send traffic to TGW attachment created earlier

Getting operational help with your Customer Managed accounts

AMS can help you operate the workloads you deployed in your Customer Managed accounts by on-boarding the account into AMS Accelerate. With AMS Accelerate you can benefit from operational services such as monitoring and alerting, incident management, security management, and backup management, without going through a new migration, experiencing downtime, or changing how you use AWS. AMS Accelerate also offers an optional patch add-on for EC2-based workloads that require regular patching. With AMS Accelerate you continue using, configuring, and deploying all AWS services natively, or with your preferred tools; as you do with AMS Advanced Customer Managed accounts. You use your preferred access and change mechanisms while AMS applies proven practices that help scale your team, optimize costs, increase security and efficiency, and improve resiliency. To learn more see the Accelerate service description.

To onboard your Customer Managed account into Accelerate, contact your CSDM and follow the steps from Getting Started with AMS Accelerate.

> **Note**
>
> AMS Accelerate accounts in AMS Advanced do not have AMS change management (requests for change or RFCs) or the AMS Advanced console. Instead, they have the AMS Accelerate console and functionality.

AMS Tools account (migrating workloads)

Your Multi-Account Landing Zone tools account (with VPC) helps accelerate migration efforts, increases your security position, reduces cost and complexity, and standardizes your usage pattern.

A tools account provides the following:

- A well-defined boundary for access to replication instances for system integrators outside of your production workloads.
- Enables you to create an isolated chamber to check a workload for malware, or unknown network routes, before placing it into an account with other workloads.
- As a defined account setup, it provides faster time to onboard and get set up for migrating workloads.
- Isolated network routes to secure traffic from on-premise -> CloudEndure -> Tools account -> AMS ingested image. Once an image has been ingested, you can share the image to the destination account via an AMS Management | Advanced stack components | AMI | Share (ct-1eiczxw8ihc18) RFC.

High level architecture diagram:

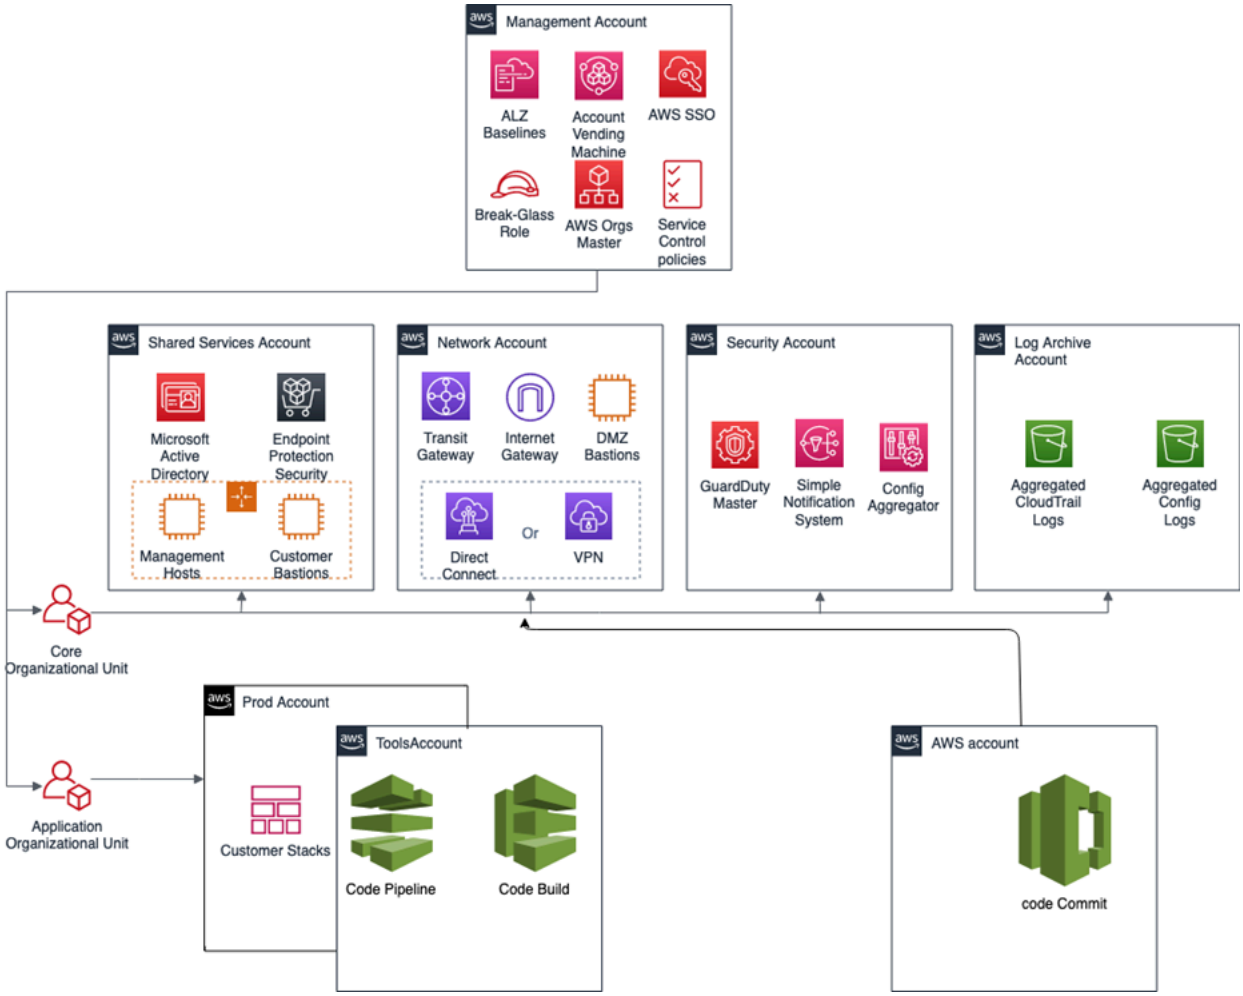

Use the Deployment | Managed landing zone | Management account | Create tools account (with VPC) change type (ct-2j7q1hgf26x5c), to quickly deploy a tools account and instantiate a Workload Ingestion process within a Multi-Account Landing Zone environment. See Management account, Tools account: Creating (with VPC).

Note

We recommend having two availability zones (AZs), since this is a migration hub.
By default, AMS creates the following two security groups (SGs) in every account. Confirm the that the two SGs are present, and, if not, open a new Management | Other | Other | Create CT (ct-1e1xtak34nx76) to request them:

- SentinelDefaultSecurityGroupPrivateOnlyEgressAll
- InitialGarden-SentinelDefaultSecurityGroupPrivateOnly

Ensure that CloudEndure replication instances are created in the private subnet where there are routes back to on-premise. You can confirm that by ensuring that the route tables for the private subnet has a default route back to TGW. However, performing a CloudEndure machine cut over should go into the "isolated" private subnet where there is no route back to on-premise, only Internet outbound traffic is allowed. It is critical to ensure cutover occurs in the isolated subnet to avoid potential issues to the on-premise resources.

Prerequisites:

1. Either **Plus** or **Premium** support level.
2. The application account IDs for the KMS key where the AMIs are deployed.
3. The tools account, created as described previously.

AWS Application Migration Service (AWS MGN)

AWS Application Migration Service (AWS MGN) can be used in your MALZ Tools account through the `AWSManagedServicesMigrationRole` IAM role that is created automatically during Tools account provisioning. You can use AWS MGN to migrate applications and databases that run on supported versions of Windows and Linux operating systems.

For the most up-to-date information on AWS Region support, see the AWS Regional Services List.

If your preferred AWS Region is not currently supported by AWS MGN, or the operating system on which your applications run is not currently supported by AWS MGN, consider using the CloudEndure Migration in your Tools account instead.

Requesting AWS MGN Initialization

AWS MGN must be initialized by AMS before first use. To request this for a new Tools account, submit a Management | Other | Other RFC from the Tools account with these details:

```
RFC Subject=Please initialize AWS MGN in this account
RFC Comment=Please click 'Get started' on the MGN welcome page here:

    https://console.aws.amazon.com/mgn/home?region=MALZ_PRIMARY_REGION#/welcome using
 all default values
```

```
    to 'Create template' and complete the initialization process.
```

Once AMS successfully completes the RFC and initializes AWS MGN in your Tools account, you can use `AWSManagedServicesMigrationRole` to edit the default template for your requirements.

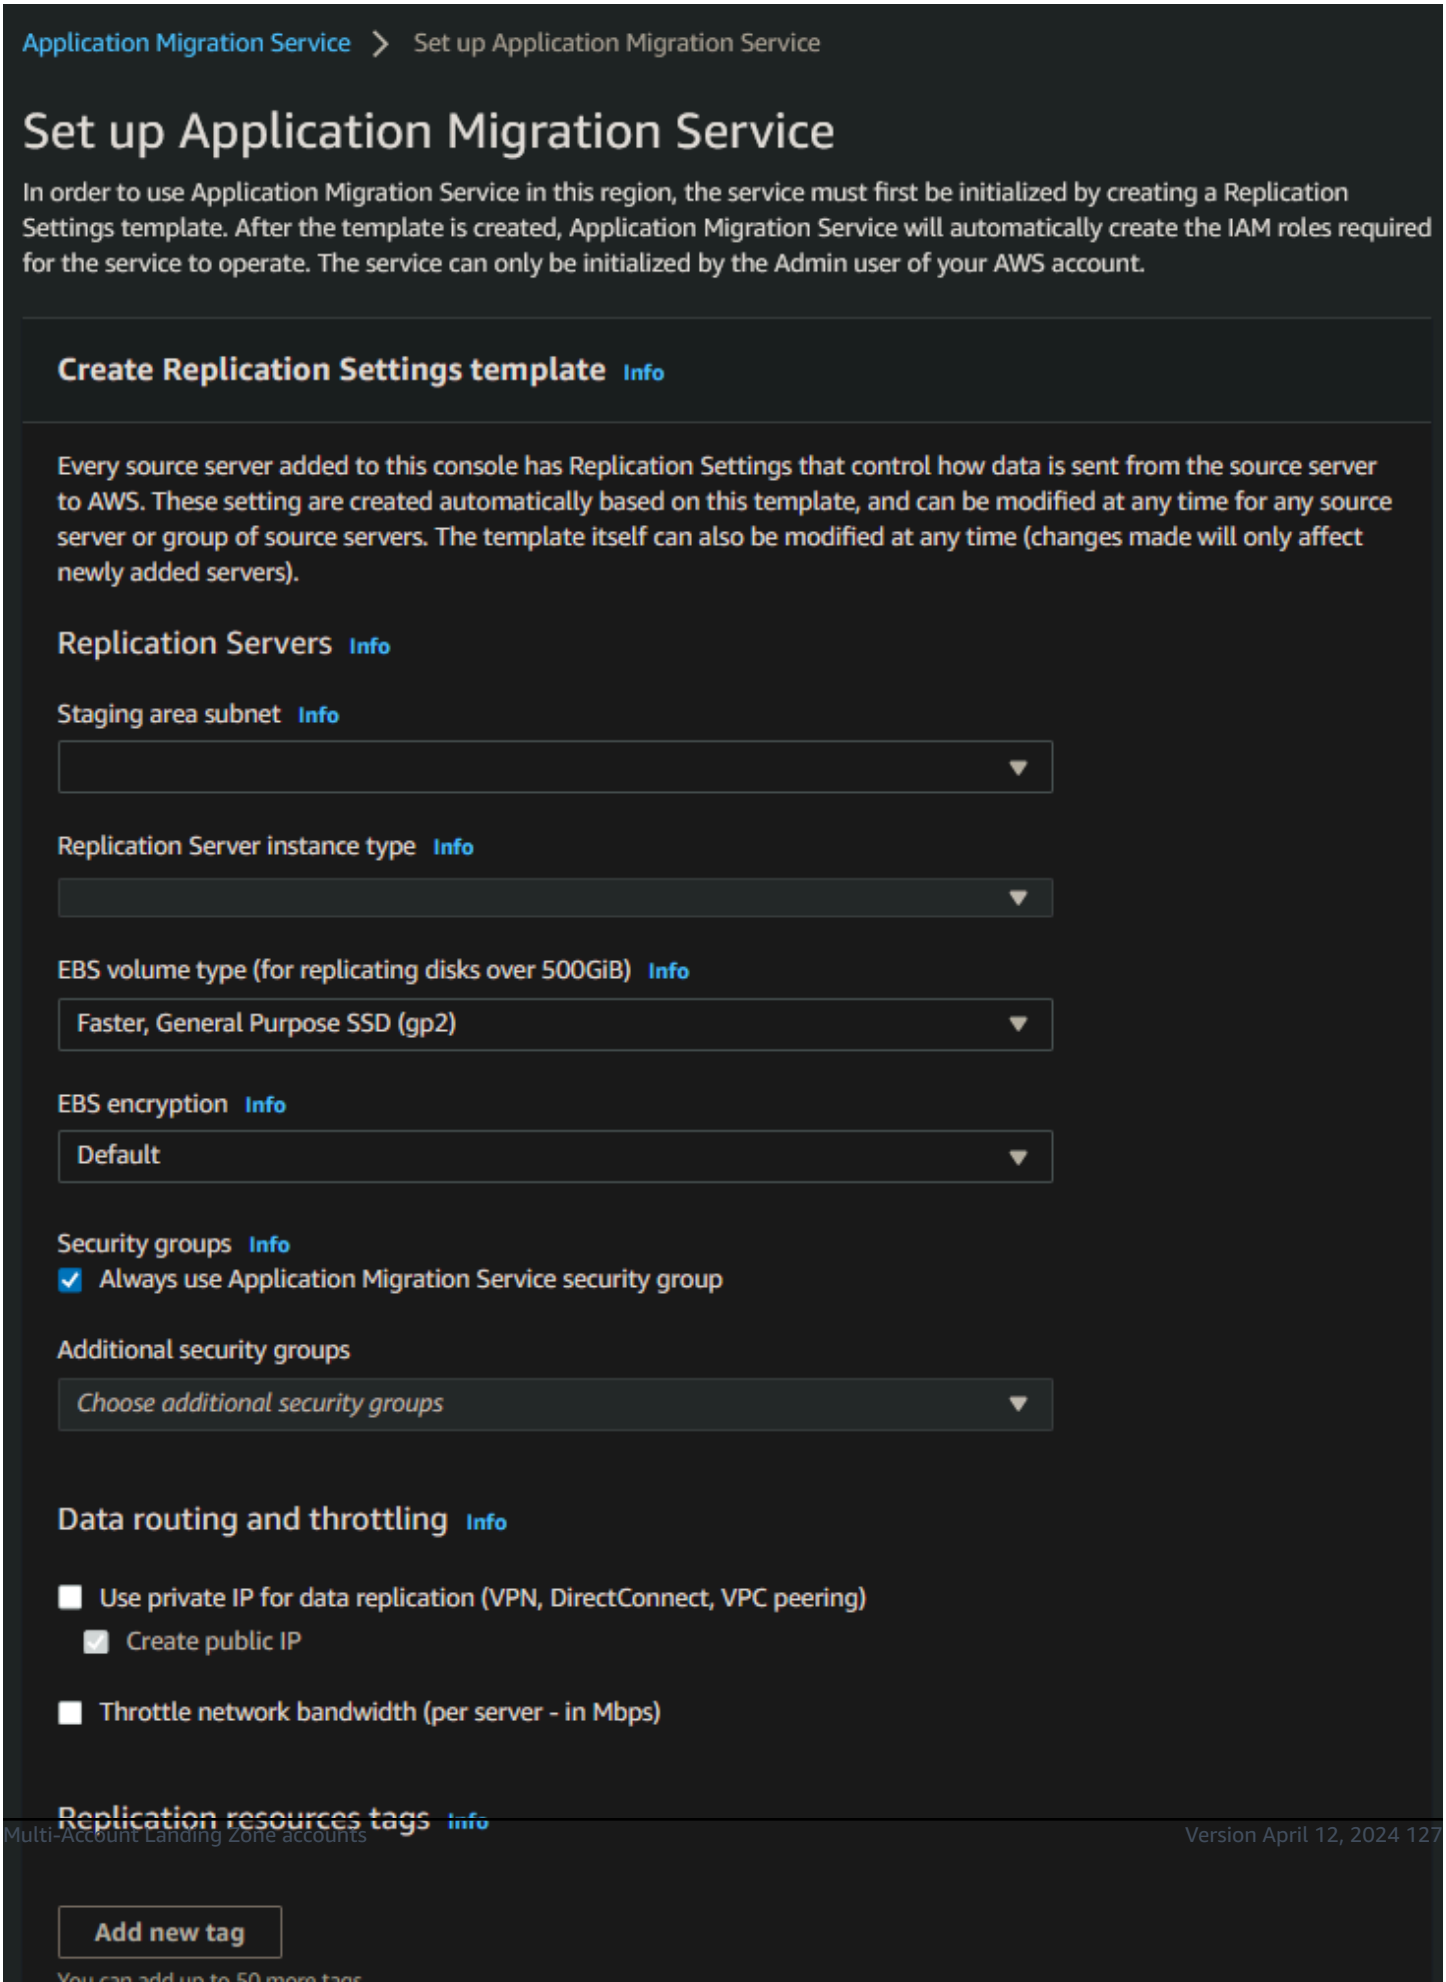
Application Migration Service > Set up Application Migration Service
Set up Application Migration Service
In order to use Application Migration Service in this region, the service must first be initialized by creating a Replication Settings template. After the template is created, Application Migration Service will automatically create the IAM roles required for the service to operate. The service can only be initialized by the Admin user of your AWS account.
Create Replication Settings template Info
Every source server added to this console has Replication Settings that control how data is sent from the source server to AWS. These setting are created automatically based on this template, and can be modified at any time for any source server or group of source servers. The template itself can also be modified at any time (changes made will only affect newly added servers).
Replication Servers Info
Staging area subnet Info
Replication Server instance type Info
EBS volume type (for replicating disks over 500GiB) Info
Faster, General Purpose SSD (gp2)
EBS encryption Info
Default
Security groups Info
Always use Application Migration Service security group
Additional security groups
Choose additional security groups
Data routing and throttling Info
Use private IP for data replication (VPN, DirectConnect, VPC peering)
Create public IP
Throttle network bandwidth (per server - in Mbps)
Replication resources tags Info
Add new tag
You can add up to 50 more tags

Enable access to the new AMS Tools account

Once the tools account is created, AMS provides you with an account ID. Your next step is to configure access to the new account. Follow these steps.

1. Update the appropriate Active Directory groups to the appropriate account IDs.

   New AMS-created accounts are provisioned with the ReadOnly role policy as well as a role to allow users to file RFCs.

   The Tools account also has an additional IAM role and user available:

   - IAM role: `AWSManagedServicesMigrationRole`
   - IAM user: `customer_cloud_endure_user`

2. Request policies and roles to allow service integration team members to set up the next level of tools.

   Navigate to the AMS console and file the following RFCs:

   a. Create KMS key. Use either Create KMS Key (auto) or Create KMS Key (review required).

      As you use KMS to encrypt ingested resources, using a single KMS key that is shared with the rest of the Multi-Account Landing Zone application accounts, provides security for ingested images where they can be decrypted in the destination account.

   b. Share the KMS key.

      Use the Management | Other | Other | Create (ct-1e1xtak34nx76) change type to request that the new KMS key be shared with your application accounts where ingested AMIs will reside.

Example graphic of a final account setup:

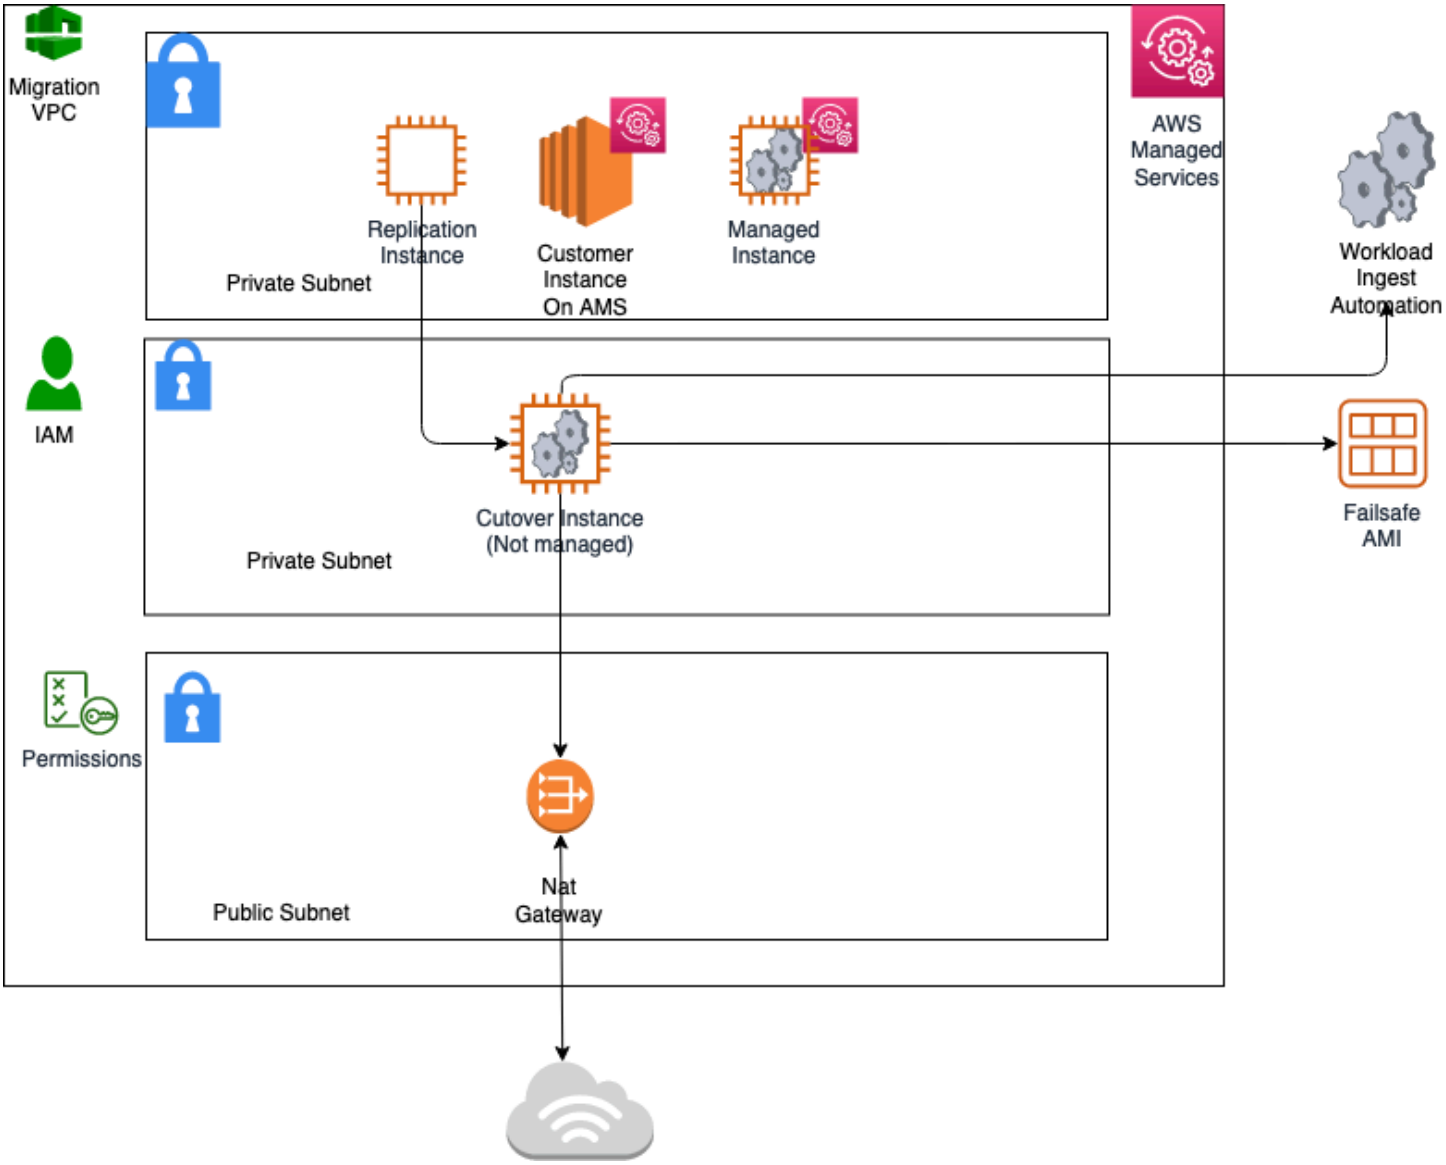

Example AMS pre-approved IAM CloudEndure policy

To see an AMS pre-approved IAM CloudEndure policy: Unpack the WIGS Cloud Endure Landing Zone Example file and open the `customer_cloud_endure_policy.json`.

Testing AMS Tools account connectivity and end-to-end setup

1. Start with configuring CloudEndure and installing the CloudEndure agent on a server that will replicate to AMS.
2. Create a project in CloudEndure.
3. Enter the AWS credentials shared when you performed the prerequisites, though secrets manager.

4. In **Replication settings**:
   a. Select both AMS "Sentinel" security groups (Private Only and EgressAll) for the **Choose the Security Groups to apply to the Replication Servers** option.
   b. Define cutover options for the machines (instances). For information, see Step 5. Cut over
   c. **Subnet**: Private subnet.
5. **Security Group**:
   a. Select both AMS "Sentinel" security groups (Private Only and EgressAll).
   b. Cutover instances have to communicate to the AMS-managed Active Directory (MAD) and to AWS public endpoints:
      i. **Elastic IP**: None
      ii. **Public IP**: no
      iii. **IAM role**: customer-mc-ec2-instance-profile
   c. Set tags as per your internal tagging convention.
6. Install the CloudEndure agent on the machine and look for the replication instance to come up in your AMS account in the EC2 console.

The AMS ingestion process:

AMS Ingestion Process

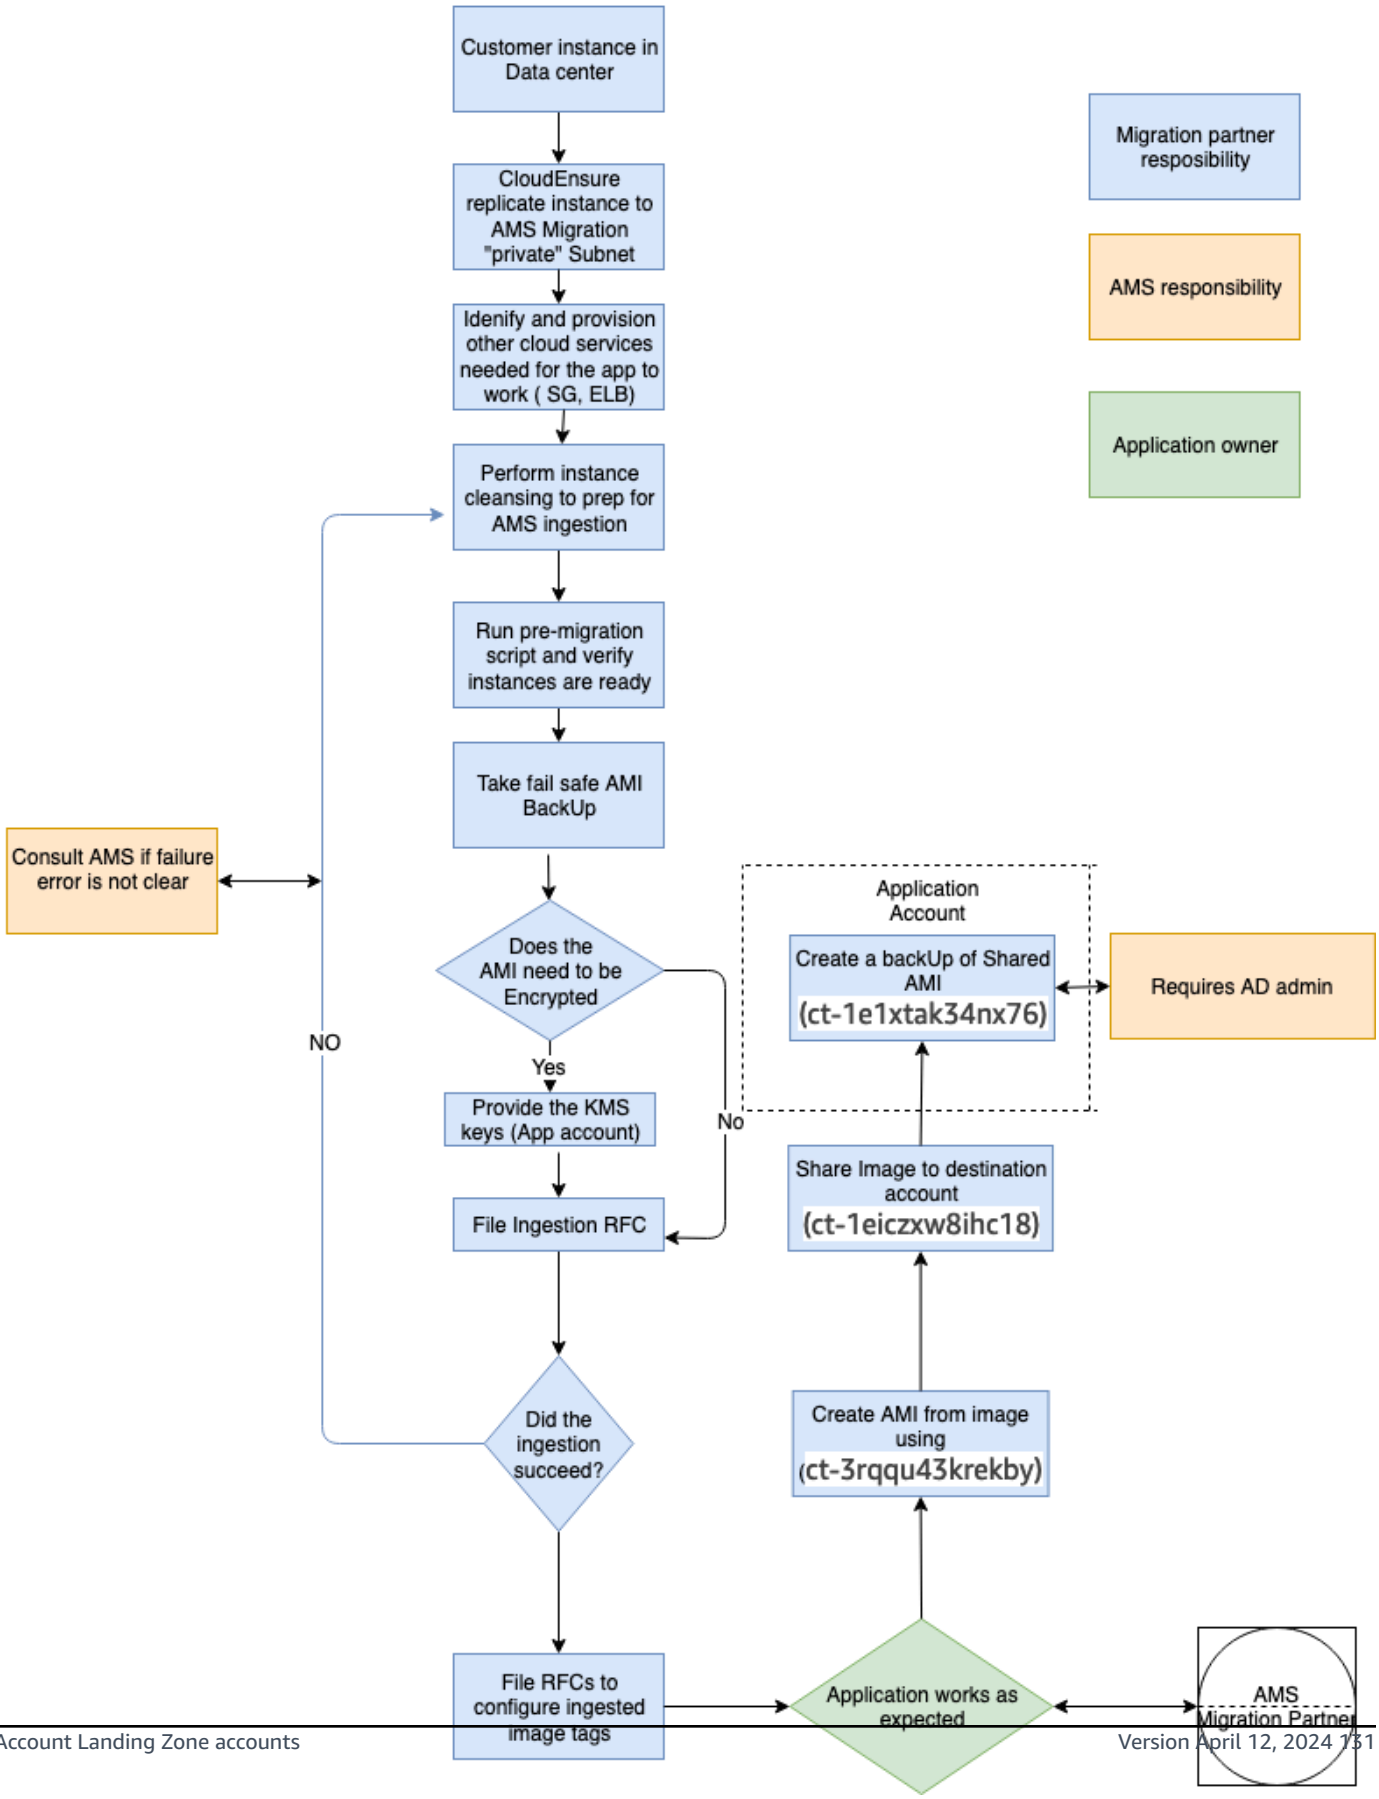

AMS Tools account hygiene

You'll want to clean up after you are done in the account have shared the AMI and no longer have a need for the replicated instances:

- Post instance WIGs ingestion:
  - Cutover instance: At a minimum, stop or terminate this instance, after the work has been completed, via the AWS console
  - Pre-Ingestion AMI backups: Remove once the instance has been ingested and the on-premise instance terminated
  - AMS-ingested instances: Turn off the stack or terminate once the AMI has been shared
  - AMS-ingested AMIs: Delete once sharing with the destination account is completed
- End of migration clean up: Document the resources deployed through Developer mode to ensure clean-up happens on regular basis, for example:
  - Security groups
  - Resources created via Cloud-formation
  - Network ACK
  - Subnet
  - VPC
  - Route Table
  - Roles
  - Users and accounts

Migration at scale - Migration Factory

See Introducing AWS CloudEndure Migration Factory Solution.

SALZ network architecture

The following diagram depicts the AWS Managed Services (AMS) single-account landing zone (SALZ) VPC network layout and is an example of the highly available setup.

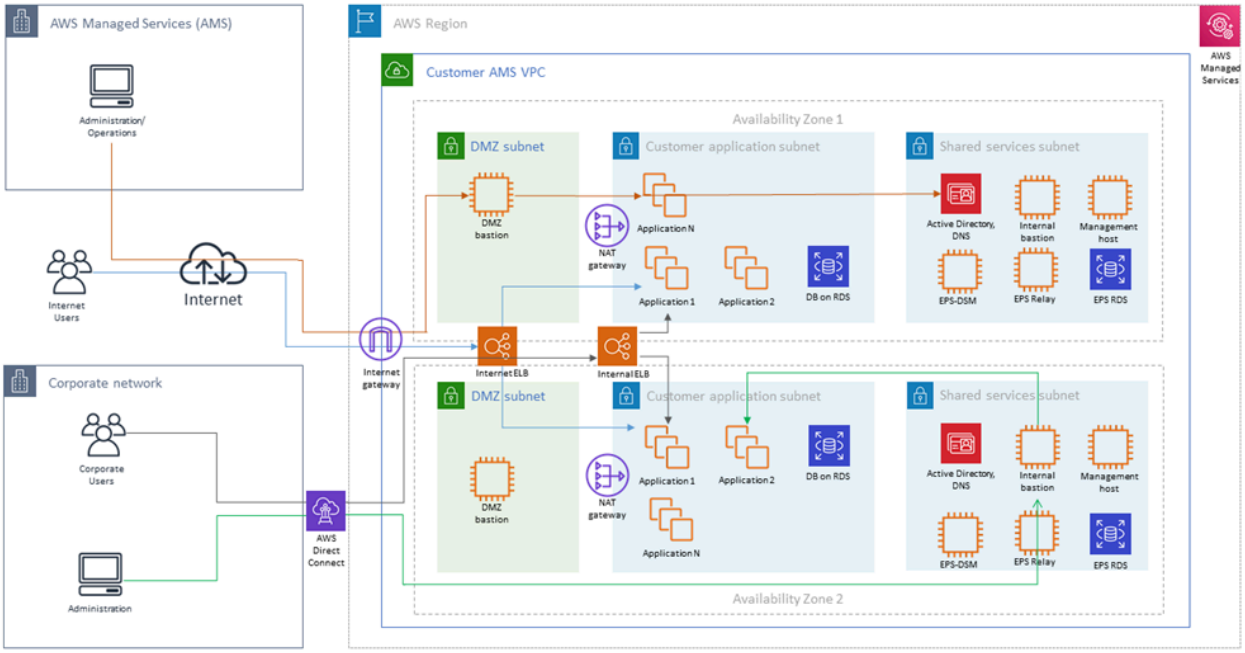

— 1 →	Ingress through DirectConnect (internal customer network users) and Internet with managed Internet Gateway (external users), through AWS load balancers to customers subnet applications. **Note** that traffic for external users goes through load balancers in DMZ (Public) Subnet, while traffic for internal users goes through load balancers in Application (Private) Subnet
— 2 →	Ingress through Internet with managed Internet Gateway for AMS administrators and operators through DMZ bastions to customer and shared services subnets
— 3 →	Ingress through DirectConnect (internal customer network administrators) and internal bastions to customer subnets

Each AMS account has a VPC in one region with resource subnets located in two availability zones. Each availability zone has three subnets: DMZ, Customer, and Shared Services. Your ("customer") corporate network is connected through a DirectConnect (VPN) tunnel, and AMS Operations connects to your managed VPC over the Internet.

Shared services subnets contain AMS Directory Services with one AD Domain Controller per shared services subnet, and AMS Management Hosts that automate provisioning and common tasks, Antivirus (TrendMicro) management servers that include EPS DSM and EPS relay (for scalability), and internal (customer) bastion hosts.

DMZ subnets contain Internet load balancers, your DMZ instances, and DMZ bastion hosts that serve as SSH jump boxes for the AMS Operations team. DMZ bastions, as well as other AMS infrastructure in the Shared services subnet, have two nodes for high availability.

Your "customer" subnets contain your workloads, EC2 instances, RDS, etc.

External users connect to your applications for the Internet via an AWS Load Balancer that is located in your DMZ.

AMS configures all aspects of networking for you based on our standard templates and your selected options provided during onboarding. A standard AWS network design is applied to your AWS account, and a virtual private cloud (VPC) is created for you and connected to AMS by either VPN or Direct Connect. Learn more about Direct Connect at AWS Direct Connect. Standard VPCs include the DMZ, shared services, and an application subnet. During the onboarding process, additional VPCs might be requested and created to match your needs (for example, customer divisions, partners). After onboarding, you're provided with a network diagram. an environment document that explains how your network has been set up.

> **Note**
>
> To learn about default service limits and constraints for all active services, see the AWS Service Limits documentation.

Our network design is built around the Amazon "Principle of Least Privilege". In order to accomplish this, we route all traffic, inbound and outbound, through gateways, except traffic coming from a trusted network. The only trusted network is the one configured between your on-premises environment and the VPC through the use of a VPN and/or an AWS Direct Connect (DX). Access is granted through the use of bastion instances, thereby preventing direct access to any production resources. All of your applications and resources reside inside private subnets that are reachable through public load balancers. Public egress traffic flows through our forward proxies to the Internet Gateway and then to the Internet. Alternatively, the traffic can flow over your VPN or Direct Connect to your on-premises environment.

AMS Single-account landing zone shared services

Shared services subnets contain AMS Directory Services, the Management Host that automates provisioning and common tasks, antivirus (TrendMicro) management server, and internal bastion hosts:

- AMS Directory Services = AD Domain Controller

  Creates an Active Directory in AMS accounts, creates the AMS domain, joins managed stacks to the domain on launch.
- Management hosts = AMS Management Host (automate provisioning and common tasks)

Act as an API endpoint to modify AWS Directory Service, interact with AWS Directory Service domain controllers.

- Security services: Antivirus (TrendMicro) management server = EPS DSM + EPS Relay

  Leverages Trend Micro™ Deep Security software (DSM), operates in a client-server model and has a back-end database, includes Deep Security managers, agents, and relays.
- Internal bastion hosts = Customer bastions

  Special purpose servers designed to be the primary access point from the Internet and act as a proxy to your other Amazon EC2 instances.

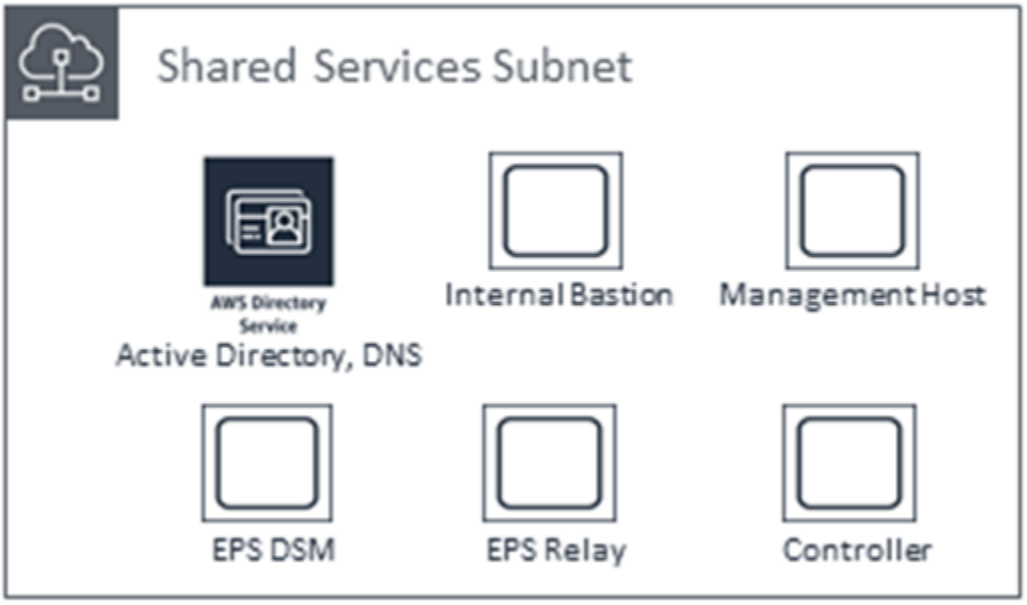

Setting up AMS

Topics

- AMS default settings
- Using the AMS consoles
- Using the AMS API and CLI
- AMS bring your own EPS
- Receiving AMS notifications
- Setting up private and public DNS
- AMS egress traffic management
- Deploying IAM resources in AMS Advanced
- Setting permissions in AMS with IAM roles and profiles
- AD FS claim rule and SAML settings
- Restrict with network ACL
- AMS on Outposts
- Using tags in AMS
- AWS Managed Services Resource Scheduler
- AWS Systems Manager in AMS Advanced
- Offboard AMS accounts

Some AMS setup tasks might be completed at onboarding.

For a full description of roles and responsibilities, including the AMS Supported AWS services, see AMS responsibility matrix (RACI).

> **Note**
>
> To request that AMS provide an additional AWS service, file a service request. For information about how to make this request, see Service request management.

AMS default settings

Topics

- DNS resolution defaults (MALZ)
- EC2 IAM instance profile
- Alerts from baseline monitoring in AMS
- Log retention and rotation defaults

Your AWS Managed Services (AMS) network is configured in a standardized manner with defaults for most services.

This section describes the default settings that AMS uses for access, monitoring, and logging, management.

For an example of multi-account landing zone or single-account landing zone infrastructure costs, see AMS environment basic components.

DNS resolution defaults (MALZ)

AWS Managed Services (AMS) multi-account landing zone: In AWS environments, domain name system (DNS) resolution between Route 53 Resolver and DNS resolvers in a VPC can be integrated by configuring Resolver forwarding rules. Before these rules can be used for forwarding DNS queries, inbound and outbound resolver endpoints need to be set up to which these queries can be forwarded.

By default, DNS queries within application account VPCs in multi-account settings in AMS are forwarded to the conditional forwarders of the AWS Directory Service for Microsoft Active Directory (also known as Managed AD) domain present in the shared services account. AMS optionally enables you to make use of the AmazonProvidedDNS; for example, AmazonProvidedDNS to forward DNS queries to. This helps you utilize VPC endpoints that today only support Amazon-provided DNS through Amazon Route 53. Correspondingly, Resolver Rules are also automatically set up for common VPC endpoints that are deployed by default in the shared services account. For more information on these common VPC endpoints, see AMS VPC endpoints.

To configure Dynamic Host Configuration Protocol (DHCP) Option Sets in all of your application account VPCs to use Amazon-provided DNS for VPC endpoints, and have Route53 Resolver rules

pointing to the common VPC endpoints in your shared services accounts (with an optional Resolver Rule for on-premises domain), create a Management | Other | Other | Create request for change (RFC) specifying the shared services account, and requesting enablement of the application account VPC local DNS and Route 53 Resolver rules for VPC endpoints.

EC2 IAM instance profile

An instance profile is a container for an IAM role that you can use to pass role information to an EC2 instance when the instance starts.

MALZ

Currently there are two AMS default instance profiles, `customer-mc-ec2-instance-profile` and `customer-mc-ec2-instance-profile-s3`, these instance profiles provide the permissions described in the following table.

Policy descriptions

Profile	Policies
`customer-mc-ec2-instance-profile`	`AmazonSSMManagedInstanceCore` : Allows Ec2 instances to use the SSM agent.
	`AMSInstanceProfileLoggingPolicy` : Allows Ec2 instances to push logs to S3 and CloudWatch.
	`AMSInstanceProfileManagementPolicy` : Allows Ec2 instances to perform booting actions, like joining Active Directory.
	`AMSInstanceProfileMonitoringPolicy` : Allows Ec2 instances to report findings to AMS monitoring services.
	`AMSInstanceProfilePatchPolicy` : Allows Ec2 instances to receive patches.
`customer-mc-ec2-instance-profile-s3`	`AMSInstanceProfileBYOEPSPolicy` : Allows Ec2 instances to use AMS bring your own EPS.

Profile	Policies
	`AMSInstanceProfileLoggingPolicy` : Allows Ec2 instances to push logs to S3 and CloudWatch.
	`AMSInstanceProfileManagementPolicy` : Allows Ec2 instances to perform booting actions, like joining Active Directory.
	`AMSInstanceProfileMonitoringPolicy` : Allows Ec2 instances to report findings to AMS monitoring services.
	`AMSInstanceProfilePatchPolicy` : Allows Ec2 instances to receive patches.
	`AMSInstanceProfileS3WritePolicy` : Allows Ec2 instances to read/write to customer S3 buckets.

SALZ

Currently there is one AMS default instance profile, `customer-mc-ec2-instance-profile`, this instance profile provides the permissions described in the following table. The profile grants permissions to the applications. running on the instance, not to users logging into the instance.

Policies often include multiple statements, where each statement grants permissions to a different set of resources or grants permissions under a specific condition.

CW = CloudWatch. ARN = Amazon Resource Name. * = wildcard (any).

EC2 default IAM instance profile permissions

CW = CloudWatch. ARN = Amazon Resource Name. * = wildcard (any).			
Policy statement	**Effect**	**Actions**	**Description and resource (ARN)**
Amazon Elastic Compute Cloud (Amazon EC2)			

CW = CloudWatch. ARN = Amazon Resource Name. * = wildcard (any).			
Policy statement	**Effect**	**Actions**	**Description and resource (ARN)**
EC2 Message Actions	Allow	AcknowledgeMessage, DeleteMessage, FailMessage, GetEndpoint, GetMessages, SendReply	Allows EC2 Systems Manager messaging actions in your account.
Ec2 Describe	Allow	* (All)	Allows the console to display configuration details of an EC2 in your account.
Iam Get Role ID	Allow	GetRole	Allows EC2 to get your IAM ID from `aws:iam::*:role/customer-*` and `aws:iam::*:role/customer_*` .
Instance To Upload Log Events	Allow	Create Log Group	Allows logs to be created in: `aws:logs:*:*:log-group:i-*`
		Create Log Stream	Allows logs to be streamed to: `aws:logs:*:*:log-group:i-*`

CW = CloudWatch. ARN = Amazon Resource Name. * = wildcard (any).			
Policy statement	**Effect**	**Actions**	**Description and resource (ARN)**
CW For MMS	Allow	DescribeAlarms, PutMetricAlarm, PutMetricData	Allows CloudWatch to retrieve alarms in your account. Allows CW to create or update an alarm and associate it with the specified metric. Allows CW to publish metric data points to your account.
Ec2 Tags	Allow	CreateTags, DescribeTags,	Allows tags to be added, overwritten, and described on the specified instances in your account.
Explicitly Deny CW Logs	Deny	DescribeLogStreams, FilterLogEvents, GetLogEvents	Disallows listing, filtering, or getting the log streams for: `aws:logs:*:*:log-group:/mc/*`
Amazon EC2 Simple Systems Manager (SSM)			
SSM Actions	Allow	DescribeAssociation, GetDocument, ListAssociations, UpdateAssociationStatus, UpdateInstanceInformation	Allows a variety of SSM functions in your account.

CW = CloudWatch. ARN = Amazon Resource Name. * = wildcard (any).			
Policy statement	**Effect**	**Actions**	**Description and resource (ARN)**
SSM Access In S3	Allow	GetObject, PutObject, AbortMultipartUpload, ListMultipartUploadPorts, ListBucketMultipartUploads	Allows the SSM on the EC2 to get and update objects in, and to abort a multi-part object upload to, and list ports and buckets available for, multi-part uploads in `aws:s3:::mc-*-internal-*/aws/ssm*` .
Amazon EC2 Simple Storage Service (S3)			
Get Object In S3	Allow	Get List	Allows EC2 applications to retrieve and list objects in S3 buckets in your account.
Customer Encrypted Log S3 Access	Allow	PutObject	Allows EC2 applications to update objects in `aws:s3:::mc-*-logs-*/encrypted/app/*`
Patch Data Put Object S3	Allow	PutObject	Allows EC2 applications to upload patching data to your S3 buckets at `aws:s3:::awsms-a*-patch-data-*`
Uploading Own Logs To S3	Allow	PutObject	Allows EC2 applications to upload custom logs to: `aws:s3:::mc-a*-logs-*/aws/instances/*/${aws:userid}/*`

CW = CloudWatch. ARN = Amazon Resource Name. * = wildcard (any).			
Policy statement	**Effect**	**Actions**	**Description and resource (ARN)**
Explicitly Deny MC Namespace S3 Logs	Deny	GetObject* Put*	Disallows EC2 applications getting or putting any objects from or to: `aws:s3:::mc-*-logs-*/encrypted/mc*` , `aws:s3:::mc-*-logs-*/mc/*`, `aws:s3:::mc-a*-logs-*-audit/*`
Explicitly Deny S3 Delete	Deny	* (all)	Disallows EC2 applications taking any action on objects in: `aws:s3:::mc-a*-logs-*/*` , `aws:s3:::mc-a*-internal-*/*` ,
Explicitly Deny S3 CFN Bucket	Deny	Delete*	Disallows EC2 applications deleting any objects from: `aws:s3:::cf-templates-*`
Explicitly Deny List Bucket S3	Deny	ListBucket	Disallows you listing any encrypted, audit log, or reserved (mc) objects from: `aws:s3:::mc-*-logs-*`

If you're unfamiliar with Amazon IAM policies, see Overview of IAM Policies for important information.

Note

Policies often include multiple statements, where each statement grants permissions to a different set of resources or grants permissions under a specific condition.

Alerts from baseline monitoring in AMS

Learn about AMS monitoring defaults; for more information, see Monitoring and event management in AMS.

The following table shows what is monitored, and the default alerting thresholds. You can change the alerting thresholds with a Management | Other | Other | Update (ct-0xdawir96cy7k) RFC after determining what changes you want and subscribing to the relevant CloudWatch Amazon SNS topic. For information about creating and subscribing to topics, see Subscribe to a Topic. For general information, see Amazon SNS FAQs. To be notified directly when alarms cross their threshold, in addition to AMS's standard alerting process, follow these instructions about how to overwrite alarm configurations, Receiving alerts generated by AMS.

Amazon CloudWatch provides extended retention of metrics. For more information, see CloudWatch Limits.

Note

AMS calibrates its baseline monitoring on a periodic basis. New accounts are always onboarded with the latest baseline monitoring and the table describes the baseline monitoring for an account that is newly onboarded. AMS updates the baseline monitoring in existing accounts on a periodic basis and you may experience a time lag before the updates are in place. For more information, see Viewing the monitoring configuration for an AMS account.

Note

The EC2 instance alert `Non-root volume usage` is **DISABLED** by default. If you require alert generation based on this alarm, then you must enable it using the RFC Change Type ct-0erkoad6uyvvg

Alerts from baseline monitoring

Resource	Security alert	Alert name and trigger condition	Notes
For starred (*) alerts, AMS proactively assesses impact and remediates when possible; if remediation is not possible, AMS creates an incident. Where automation fails to remediate the issue, AMS informs you of the incident case and an AMS engineer is engaged. In addition, these alerts can be sent directly to your email (if you have opted in to the Direct-Customer-Alerts SNS topic).			
Application Load Balancer (ALB) instance	No	RejectedConnectionCount sum > 0 for 1 min, 5 consecutive times.	CloudWatch alarm if the number of connections that were rejected because the load balancer reached its maximum.
Application Load Balancer (ALB) target	No	TargetConnectionErrorCount sum > 0 for 1 min, 5 consecutive times.	CloudWatch alarm if number of connections were unsuccess fully established between the load balancer and the registered instances.
Aurora instance	No	CPUUtilization > 85% for 5 mins, 2 consecutive times.	CloudWatch alarm.
AWS Backup	Yes	DeleteRecoveryPoint An unexpected IAM role principal or IAM user principal has deleted an AWS Backup recovery point.	CloudWatch event. Emitted when a backup recovery point is deleted.
AWS Outposts	Yes	AMSOutpostsInstanceFamilyCapacityAvailability InstanceFamilyCapacityAvailability	CloudWatch alarm on instance family capacity availability of the AWS Outposts resource.

Resource	Security alert	Alert name and trigger condition	Notes
		= 80% for 5 minutes, 12 consecutive times.	
		AMSOutpostsInstanceTypeCapacityAvailability TypeCapacityAvailability = 80% for 5 minutes, 12 consecutive times.	CloudWatch alarm on instance type capacity availability of the AWS Outposts resource.
		AMSOutpostsConnectedStatusConnectedStatus < 1 for 5 minutes, 1 consecutive time.	CloudWatch alarm on AWS Outposts service link connection, less than 1 count is impaired.
		AMSOutpostsCapacityExceptionCapacityExceptions 0 for 5 minutes, 1 consecutive time.	CloudWatch alarm on insufficient capacity errors for instance launches for AWS Outpostss resource .
EC2 instance - all OSs	No	CPUUtilization* >= 95% for 5 mins, 6 consecutive times.	CloudWatch alarm. High CPU utilization is an indicator of a change in application state such as dead locks, infinite loops, malicious attacks, and other anomalies.
		StatusCheckFailed > 0 for 5 minutes, 3 consecutive times.	CloudWatch alarm.

Resource	Security alert	Alert name and trigger condition	Notes
		Root Volume Usage >= 95% for 5 mins, 6 consecutive times.	
		Non-root Volume Usage > 85% for 5 mins, 2 consecutive times. **Disabled by default; for details, see Additional Information.**	
		Memory Free* MemoryFree < 5% for 5 minutes, 6 consecutive times.	
	Yes	EPS Malware Malware found on instance.	CloudWatch event.
Amazon EC2 instance - Linux	No	Root Volume Inode Usage Average >= 95% for 5 mins, 6 consecutive times.	CloudWatch alarm. Applied to Linux instances only.
		Swap Free* Memory Swap < 5% for 5 minutes, 6 consecutive times.	

Resource	Security alert	Alert name and trigger condition	Notes
ElastiCache Cluster	No	CurrConnections = 65000	This alarm notifies AMS of the maximum connection limit of an ElastiCache Host. CloudWatch Alarm. If you would like to update this threshold, contact AMS support.

Resource	Security alert	Alert name and trigger condition	Notes
ElastiCache Node	No	CPUUtilization Average > predefined value for 15 mins, 2 consecutive times.	CloudWatch alarm. Default is 90. If Redis, use one the following values based on instance type: • cache.t1.micro: 90% • cache.m1.small: 90% • cache.m1.medium: 90% • cache.m1.large: 45% • cache.m1.xlarge: 22.5% • cache.m2.xlarge: 45% • cache.m2.4xlarge: 11.25% • cache.c1.xlarge: 11.25% • cache.t2.micro: 90% • cache.t2.small: 90% • cache.t2.medium: 45% • cache.m3.medium: 90% • cache.m3.large: 45% • cache.m3.xlarge: 22.5% • cache.m3.2xlarge: 11.25% • cache.r3.large: 45% • cache.r3.xlarge: 22.5% • cache.r3.2xlarge: 11.25% • cache.r3.4xlarge: 5.625% • cache.r3.8xlarge: 2.8125%
ElastiCache Node - memcached	No	SwapUsage maximum > 50,000,000 bytes for 5 mins, 5 consecutive times.	CloudWatch alarm. Applied to memcached only.

Resource	Security alert	Alert name and trigger condition	Notes
OpenSearch cluster	No	ClusterStatus.red maximum is >= 1 for 1 minute, 1 consecutive time. *AMS takes pro-active actions to reduce operational impact, when this alert is triggered.*	CloudWatch alarm. At least one primary shard and its replicas are not allocated to a node. To learn more, see Red Cluster Status.
OpenSearch domain	No	KMSKeyError >= 1 for 1 minute, 1 consecutive time.	CloudWatch alarm. The KMS encryption key that is used to encrypt data at rest in your domain is disabled. Re-enable it to restore normal operations. To learn more, see Encryption of Data at Rest for OpenSearch Service Service.
		ClusterStatus.yellow maximum is >= 1 for 1 minute, 1 consecutive time *AMS takes pro-active actions to reduce operational impact, when this alert is triggered.*	At least one replica shard is not allocated to a node. To learn more, see Yellow Cluster Status.
		FreeStorageSpace minimum is <= 20480 for 1 minute, 1 consecutive time *AMS takes pro-active actions to reduce operational impact, when this alert is triggered.*	A node in your cluster is down to 20 GiB of free storage space. To learn more, see Lack of Available Storage Space.

Resource	Security alert	Alert name and trigger condition	Notes
		ClusterIndexWritesBlocked >= 1 for 5 minutes, 1 consecutive time *AMS takes pro-active actions to reduce operational impact, when this alert is triggered.*	The cluster is blocking write requests. To learn more, see ClusterBlockException.
		Nodes minimum is < x for 1 day, 1 consecutive time *AMS takes pro-active actions to reduce operational impact, when this alert is triggered.*	x is the number of nodes in your cluster. This alarm indicates that at least one node in your cluster has been unreachable for one day. To learn more, see Failed Cluster Nodes.
		CPUUtilization average is >= 80% for 15 minutes, 3 consecutive times *AMS takes pro-active actions to reduce operational impact, when this alert is triggered.*	100% CPU utilization is common, but sustained high averages are problematic. Consider using larger instance types or adding instances.

Resource	Security alert	Alert name and trigger condition	Notes
		JVMMemoryPressure maximum is >= 80% for 5 minutes, 3 consecutive times *AMS takes pro-active actions to reduce operational impact, when this alert is triggered.*	The cluster could encounter out of memory errors if usage increases. Consider scaling vertically. Amazon ES uses half of an instance's RAM for the Java heap, up to a heap size of 32 GiB. You can scale instances vertically up to 64 GiB of RAM, at which point you can scale horizontally by adding instances .
		MasterCPUUtilization average is >= 50% for 15 minutes, 3 consecutive times *AMS takes pro-active actions to reduce operational impact, when this alert is triggered.*	Consider using larger instance types for your dedicated master nodes. Because of their role in cluster stability and blue/green deployments, dedicated master nodes should have lower average CPU usage than data nodes.
		MasterJVMMemoryPressure maximum is >= 80% for 15 minutes, 1 consecutive time *AMS takes pro-active actions to reduce operational impact, when this alert is triggered.*	Consider using larger instance types for your dedicated master nodes. Because of their role in cluster stability and blue/green deployments, dedicated master nodes should have lower average CPU usage than data nodes.

Resource	Security alert	Alert name and trigger condition	Notes
OpenSearch instance	No	AutomatedSnapshotFailure maximum is >= 1 for 1 minute, 1 consecutive time.	CloudWatch alarm. An automated snapshot failed. This failure is often the result of a red cluster health status. See Red Cluster Status.
Elastic Load Balancing instance	No	SurgeQueueLength > 100 for 1 minute, 15 consecutive times.	CloudWatch alarm if an excess number of requests are pending routing.
		HTTPCode_ELB_5XX_Count sum > 0 for 5 min, 3 consecutive times.	CloudWatch alarm on excess number of HTTP 5XX response codes that originate from the load balancer.
		SpilloverCount > 1 for 1 minute, 15 consecutive times.	CloudWatch alarm if an excess number of requests that were rejected because the surge queue is full.
GuardDuty service	Yes	Not applicable; all findings (threat purposes) are monitored. Each finding corresponds to an alert. Changes in the GuardDuty findings. These changes include newly generated findings or subsequent occurrences of existing findings.	List of supported GuardDuty finding types are on GuardDuty Active Finding Types.

Resource	Security alert	Alert name and trigger condition	Notes
Health	Varies	AWS Personal Health Dashboard.	Notifications sent when there are changes in the status of AWS Personal Health Dashboard (AWS Health) events. Service event example: Scheduled EC2 instance store retirement. These Health events are not monitored: • AWS_VPN_SINGLE_TUNNEL_NOTIFICATION • AWS_ELASTICACHE_UPDATE_COMPLETED • AWS_STORAGEGATEWAY_SOFTWARE_UPDATE_AVAILABLE • AWS_BILLING_NOTIFICATION • AWS_VPN_REDUNDANCY_LOSS
AWS Managed Microsoft AD	No	Active Directory Status AWS Managed Microsoft AD instance sends an active status event.	Service event. Emitted when the directory is operating normally after an event.

Resource	Security alert	Alert name and trigger condition	Notes
		Impaired Directory Status AWS Managed Microsoft AD instance sends an impaired directory status event.	Service event. Emitted when the directory is running in a degraded state. One or more issues have been detected, and not all directory operations may be working at full operational capacity.
		Inoperable Directory Status AWS Managed Microsoft AD instance sends an inoperable status event.	Service event. Emitted when the directory is not functiona l. All directory endpoints have reported issues.
		Deleting Directory Status AWS Managed Microsoft AD instance sends a deleting directory status event.	Service event. Emitted when the directory is currently being deleted.
		Failed Directory Status AWS Managed Microsoft AD instance sends a failed status event.	Service event. Emitted when the directory could not be created.
		RestoreFailed Directory Status AWS Managed Microsoft AD instance sends a restore failed directory status event.	Service event. Emitted when restoring the directory from a snapshot failed.
Amazon RDS instance	No	Low Storage alert triggers when the allocated storage for the DB instance has been exhausted.	RDS-EVENT-0007, see details at Using Amazon RDS event notification.

Resource	Security alert	Alert name and trigger condition	Notes
		DB instance fail The DB instance has failed due to an incompatible configuration or an underlying storage issue. Begin a point-in-time-restore for the DB instance.	Service event. RDS-EVENT-0031, Amazon RDS Event Categories and Event Messages.
		Failover not attempted Amazon RDS is not attempting a requested failover because a failover recently occurred on the DB instance.	Service event. RDS-EVENT-0034, Amazon RDS Event Categories and Event Messages.
		DB instance invalid parameters For example, MySQL could not start because a memory-related parameter is set too high for this instance class, so the customer action would be to modify the memory parameter and reboot the DB instance.	Service event. RDS-EVENT-0035, Amazon RDS Event Categories and Event Messages.
		Invalid subnet IDs DB instance The DB instance is in an incompatible network. Some of the specified subnet IDs are invalid or do not exist.	Service event. RDS-EVENT-0036, Amazon RDS Event Categories and Event Messages.

Resource	Security alert	Alert name and trigger condition	Notes
		DB instance read replica error An error has occurred in the read replication process. For more information, see the event message. For information on troubleshooting Read Replica errors, see Troubleshooting a MySQL Read Replica Problem.	Service event. RDS-EVENT-0045, Amazon RDS Event Categories and Event Messages.
		DB instance read replication ended Replication on the Read Replica was ended.	Service event. RDS-EVENT-0057, Amazon RDS Event Categories and Event Messages.
		Error create statspack user account Error while creating Statspack user account PERFSTAT. Drop the account before adding the Statspack option.	Service event. RDS-EVENT-0058, Amazon RDS Event Categories and Event Messages.
		DB instance recovery start The SQL Server DB instance is re-establishing its mirror. Performance will be degraded until the mirror is reestablished. A database was found with non-FULL recovery model. The recovery model was changed back to FULL and mirroring recovery was started. (<dbname>: <recovery model found>[,...]).	Service event. RDS-EVENT-0066, Amazon RDS Event Categories and Event Messages.

Resource	Security alert	Alert name and trigger condition	Notes
		A failover for the DB cluster has failed.	RDS-EVENT-0069, see details at Amazon RDS Event Categories and Event Messages.
		Invalid permissions recovery S3 bucket The IAM role that you use to access your Amazon S3 bucket for SQL Server native backup and restore is configured incorrectly. For more information, see Setting Up for Native Backup and Restore.	Service event. RDS-EVENT -0081, Amazon RDS Event Categories and Event Messages.
		Aurora was unable to copy backup data from an Amazon S3 bucket.	RDS-EVENT-0082, see details at Amazon RDS Event Categories and Event Messages.
		Low storage alert when the DB instance has consumed more than 90% of its allocated storage	RDS-EVENT-0089, see details at Amazon RDS Event Categories and Event Messages.
		Notification service when scaling failed for the Aurora Serverless DB cluster.	RDS-EVENT-0143, see details at Amazon RDS Event Categories and Event Messages.
		The DB instance is in an invalid state. No actions are necessary. Autoscaling will retry later.	RDS-EVENT-0219, see details at Amazon RDS Event Categories and Event Messages.
		The DB instance has reached the storage-full threshold, and the database has been shut down.	RDS-EVENT-0221, see details at Amazon RDS Event Categories and Event Messages.

Resource	Security alert	Alert name and trigger condition	Notes
		This event indicates the RDS instance storage autoscaling is unable to scale, there could be multiple reasons for why the autoscaling failed.	RDS-EVENT-0223, see details at Amazon RDS Event Categories and Event Messages.
		Storage autoscaling has triggered a pending scale storage task that would reach the maximum storage threshold.	RDS-EVENT-0224, see details at Amazon RDS Event Categories and Event Messages.
		The DB instance has a storage type that's currently unavailable in the Availability Zone. Autoscaling will retry later.	RDS-EVENT-0237, see details at Amazon RDS Event Categories and Event Messages.
		RDS couldn't provision capacity for the proxy because there aren't enough IP addresses available in your subnets.	RDS-EVENT-0243, see details at Amazon RDS Event Categories and Event Messages.
		The storage for your AWS account has exceeded the allowed storage quota.	RDS-EVENT-0254, see details at Amazon RDS Event Categories and Event Messages.
		CPUUtilization Average CPU utilization > 75% for 15 mins, 2 consecutive times.	CloudWatch alarm.
		DiskQueueDepth Sum is > 75 for 1 mins, 15 consecutive times.	

Resource	Security alert	Alert name and trigger condition	Notes
		FreeStorageSpace Average < 1,073,741,824 bytes for 5 mins, 2 consecutive times.	
		SwapUsage Average >= 104,857,600 bytes for 5 mins, 2 consecutive times.	
Amazon Redshift cluster	No	RedshiftClusterStatus The health of the cluster when not in maintenance mode < 1 for 5 min.	1 represents a healthy cluster.
Amazon Macie	Yes	Newly generated alerts and updates to existing alerts. Macie finds any changes in the findings. These changes include newly generated findings or subsequent occurrences of existing findings.	Amazon Macie alert. For a list of supported Macie alert types, see Analyzing Amazon Macie Findings. Note that Macie is not enabled for all accounts.

AMS takes pro-active actions (scaling the cluster) when this alert is triggered.

For information on remediation efforts, see AMS automatic remediation of alerts.

Watch Andrew's video to learn more (7:03)

Log retention and rotation defaults

This section describes AMS log management defaults; for more information, see Log Management.

- Rotation = Log turnover inside the instances
- Retention = Period of time we keep the logs in Amazon CloudWatch Logs and Amazon Simple Storage Service (S3)

The logs are retained in CloudWatch Logs as needed (you can configure this), and in S3. They don't expire or get deleted and are subject to service durability. For detailed S3 durability information, see Data protection in Amazon S3.

You can request a change to log retention for all logs, except AWS CloudTrail logs, which are kept indefinitely for audit and security reasons.

Log rotation is configured inside the instances. By default, operating system and security logs rotate hourly if they reach over 100MB, this is done to ensure that you don't run short on disk in the instances.

The log agent inside the instances uploads the log online to CloudWatch Logs, from there the logs are archived to S3.

The logs are stored in CloudWatch Logs and S3 in the raw format they are generated, there is no pre-processing.

Using the AMS consoles

The AMS consoles in the AWS Management Console are available for you to interact with AMS and operate your AMS Advanced-managed and AMS Accelerate resources. The AMS consoles generally behave like any AWS console; however, because AMS is a private organization, only accounts enabled for AMS can access the console. Once AMS is enabled in your account, you can access the console by searching for "Managed Services" in the unified search bar.

> **Note**
>
> Depending on your account role, you access the AMS Advanced console or the AMS Accelerate console.

When using the AMS consoles, be aware of the following caveats:

- The AMS console is account specific. So, if you are in a "Test" account for your organization, you won't be able to see resources in the "Prod" account for that organization. Likewise, you must have an AMS Advanced role to access the AMS Advanced console.
- The AMS consoles apply an IAM policy when you authenticate that determines which console you can access and what you can do there. Your administrator may apply additional polices to the default AMS policy to restrict what you can see and do in the console.

The AMS Advanced console has these features:

- Opening page: The opening page has information boxes and links to facilitate your access to your existing RFCs, incidents, service request, and reports.
- Feature pages, links in the left-hand navigation pane:
  - **Dashboard**: Provides an overview of the current status of your account including:
    - **Requests for change**: See how many RFCs are **Awaiting your response**, and jump to the RFC list page with that filter active. See how many RFCs are **Awaiting your approval**, and jump to the RFC list page with that filter active. See how many RFCs are **Open**, and jump to the RFC list page with that filter active. Open the list page for RFCs by clicking the **View all** link.
    - **Incidents**: See how many incident cases are **Awaiting your response**, and jump to the incident list page with that filter active. See and how many are **Open**, and jump to the incident list page with that filter active. Open the incident list page by clicking the **View all** link.
    - **Service requests**: See how many service requests are **Awaiting your response**, and jump to the service request list page with that filter active. See and how many are **Open**, and jump to the service request list page with that filter active. Open the service request list page by clicking the **View all** link.
    - **Recently updated RFCs**: Date, link to the RFC details, and status
    - **Recently created incidents and service requests**: Date, link to the case details, and type (incident or service request)
  - **RFCs**: Opens a list of the existing RFCs for the account
  - **Incidents**: Opens a list of the open incidents for the account
  - **Service requests**: Opens a list of the open service requests for the account
  - **Reports**: Opens the Reports page and the default reports, **Daily Backup** and **Daily Patch** and **Monthly Billling**
  - **Resources**:
    - **VPCs**: Opens a list of the existing VPCs for the account
    - **Stacks**: Opens a list of existing stacks for the account
    - **AMIs**: Opens a list of available AMS AMIs
- **Feature spotlight**: Information on the latest updates to the console
- **Developer's Resources**: A page of downloadable files, including the AMS Advanced change management SDK and more

- **Documentation**: The AWS Managed Services documentation landing page

Using the AMS API and CLI

The AWS Managed Services (AMS) API is similar to the APIs for other AWS services. You can read about the AMS API in the AMS API Reference.

AMS API HTTP endpoints for REST calls

Besides the various SDKs, AMS provides a CLI; you can also invoke REST API calls against the AMS endpoint.

There are two AMS APIs (the endpoint for both resides in us-east-1):

- Change Management: Use this API to request access to or changes to your infrastructure, including creating and updating RFCs, deploying new instances, updating and deleting instances, getting information on CTs, and creating AMIs. The HTTP endpoint is:

  ```
  https://amscm.us-east-1.amazonaws.com
  ```

- SKMS: Use this API to get information about your infrastructure, including VPCs, stacks, subnets, and AMIs. The HTTP endpoint is:

  ```
  https://amsskms.us-east-1.amazonaws.com
  ```

Installing or upgrading the AMS CLI

The AMS CLI is an easy way to interact with the AMS API and is used in the examples in this section. For usage conventions for the AWS CLI and AMS CLI, see Using the AWS command Line Interface.

For information on installing SAML, see AD FS claim rule and SAML settings.

To install or upgrade the AMS CLI, follow these instructions:

> **Note**
>
> You must have administrator credentials for this procedure.

The AWS CLI is a prerequisite for using the AWS Managed Services (AMS) CLIs (Change Management and SKMS).

1. To install the AWS CLI, see Installing the AWS Command Line Interface, and follow the appropriate instructions. Note that at the bottom of that page there are instructions for using different installers, Linux, MS Windows, macOS, Virtual Environment, Bundled Installer (Linux, macOS, or Unix).

   After the installation, run `aws help` to verify the installation.

2. Once the AWS CLI is installed, to install or upgrade the AMS CLI, download either the AMS **AMS CLI** or **AMS SDK** distributables zip file and unzip. You can access the AMS CLI distributables through the **Developer's Resources** link in the left nav of the AMS console.

3. The README file provides instructions for any install.

   Open either:

   - CLI zip: Provides the AMS CLI only.
   - SDK zip: Provides all of the AMS APIs and the AMS CLI.

   For **Windows**, run the appropriate installer (only 32 or 64 bits systems):

   - 32 Bits: **ManagedCloudAPI_x86.msi**
   - 64 Bits: **ManagedCloudAPI_x64.msi**

   For **Mac/Linux**, run the file named: **AWSManagedServices_InstallCLI.sh** by running this command: `sh AWSManagedServices_InstallCLI.sh`. Note that the **amscm** and **amsskms** directories and their contents must be in the same directory as the **AWSManagedServices_InstallCLI.sh** file.

4. If your corporate credentials are used through federation with AWS (the AMS default configuration) you must install a credential management tool that can access your federation service. For example, you can use this AWS Security Blog How to Implement Federated API and CLI Access Using SAML 2.0 and AD FS for help configuring your credential management tooling.

5. After the installation, run `aws amscm help` and `aws amsskms help` to see commands and options.

> **Note**
>
> The AMS CLI must be installed for these commands to work. To install the AMS API or CLI, go to the AMS console **Developers Resources** page. For reference material on the AMS CM API or AMS SKMS API, see the AMS Information Resources section in the User Guide. You may need to add a `--profile` option for authentication; for example, `aws amsskms` *`ams-cli-command`* `--profile SAML`. You may also need to add the `--region` option as all AMS commands run out of us-east-1; for example `aws amscm` *`ams-cli-command`* `--region=us-east-1`.

Using the AMS API in CLI, Ruby, Python, and Java

The following is a list of code snippets for the AMS API `ListChangeTypeClassificationSummaries` operation, in all available languages.

For the Python, Ruby, and Java SDKs, see Tools for Amazon Web Services and scroll down to the SDKs section. Each SDK installer contains a README with additional code snippets.

AMS API to CLI example

After you have installed the AMS CLI (requires the AWS CLI; see Installing or upgrading the AMS CLI), you can run any AMS API operation by reforming the call first specifying which AMS API, `aws amscm` or `aws amsskms`, and then giving the action with hyphens replacing camel case. Finally, provide credentials, such as SAML.

To learn more, see Using the AWS Command Line Interface.

Example:

- API: `'ChangeTypeClassificationSummaries[].[Category,Subcategory,Item,Operation,ChangeTypeId]'`
- CLI: `amscm list-change-type-classification-summaries --query "ChangeTypeClassificationSummaries[*].[Category,Subcategory,Item,Operation,ChangeTypeId]" --output table`

Note

If you authenticate with SAML, add `aws --profile saml` to the beginning of the command. For example,

```
aws --profile saml amscm list-change-type-classification-
summaries --query "ChangeTypeClassificationSummaries[*].
[Category,Subcategory,Item,Operation,ChangeTypeId]" --output table
```

AMS API to Python example

In order to use the AMS API with Python, install the AMS CLI and install boto3. Follow these steps:

1. Install the AMS CLI. See Installing or upgrading the AMS CLI.
2. Install boto3, the AWS SDK for Python. For more information, see this blog post Now Available – AWS SDK For Python (Boto3).

   ```
   import boto3
   ```
3. Get the AMS Change Management client:

   ```
   cm = boto3.client('amscm')
   ```
4. Get the AMS CTs:

   ```
   cts = cm.list_change_type_classification_summaries()

   print(cts)
   ```

Python examples

The following are some examples for using Python in AMS, to create EC2 instances, and/or use Lambda.

Python example to create an EC2

This example shows how you can use the amscm RESTFul API from within Python code to file and perform RFC processes.

1. Install the AMS CLI somewhere you have access to; you need the files it supplies.

2. Call Python libraries and create the EC2 instance:

```
import boto3
import json
import time

# Create the amscm client
cm = boto3.client('amscm')

# Define the execution parameters for EC2 Create
AMSExecParams = {
    "Description": "EC2-Create",
    "VpcId": "VPC_ID",
    "Name": "My-EC2",
    "TimeoutInMinutes": 60,
    "Parameters": {
        "InstanceAmiId": "INSTANCE_ID",
        "InstanceSubnetId": "SUBNET_ID"
    }
}

# Create the AMS RFC
cts = cm.create_rfc(
    ChangeTypeId="ct-14027q0sjyt1h",
    ChangeTypeVersion="3.0",
    Title="Python Code RFC Create",
    ExecutionParameters=json.dumps(AMSExecParams)
)

# Extract the RFC ID from the response
NewRfcID = cts['RfcId']

# Submit the RFC
RFC_Submit_Return=cm.submit_rfc(RfcId=NewRfcID)

# Check the RFC status every 30 seconds
RFC_Status = cm.get_rfc(RfcId=NewRfcID)
RFC_Status_Code = RFC_Status['Rfc']['Status']['Name']

while RFC_Status_Code != "Success":
    if RFC_Status_Code == "PendingApproval":
        print(RFC_Status_Code)
```

```
        time.sleep(30)
    elif RFC_Status_Code == "InProgress":
        print(RFC_Status_Code)
        time.sleep(30)
    elif RFC_Status_Code == "Failure":
        print(RFC_Status_Code)
        break
    else:
        print(RFC_Status_Code)

    RFC_Status = cm.get_rfc(RfcId=NewRfcID)
    RFC_Status_Code = RFC_Status['Rfc']['Status']['Name']
```

Python example with Lambda

This example shows how to bundle the AMS models with your code so you can use it with Lambda, or EC2; places you won't, or can't, install `amscli`.

> **Note**
>
> AMS does not provide an importable AMS-specific Python SDK. The `amscli` install script installs the AMS service data models in the CLI's normal path. For CLI usage and system Python usage, that is fine, because both `awscli` and `boto3` read their service models from the same default locations (`~/.aws/models`). However, when you want to use AMS services via boto3 in Lambda (or any other non-local runtime), it breaks, because you no longer have the data models. The following is a method to fix this by packaging the data models with the function.

There are simple steps that you can take to run your AMS-integrated Python code in Lambda or another runtime like EC2, Fargate, etc. The following workflow shows the steps necessary for AMS-integrated Lambda functions.

By adding the data models to the code's deployment package and updating the SDK search path, you can simulate an SDK experience.

Important

This example and all of the non-python commands shown were tested on a Mac computer.

Example Workflow:

1. Install the `amscli`. This creates a folder at `~/.aws/models` on your computer (Mac).
2. Copy the models to a local directory: `cp ~/.aws/models ./models`.
3. Include the models into your code's deployment package.
4. Update your function code to add the new models to the SDK path. Note that this code must run before boto3 or botocore are imported!

```
# Force Python to search local directory for boto3 data models
import os
os.environ['AWS_DATA_PATH'] = './models'

import boto3
import botocore
```

Note

Because the example models are in a directory named `models`, we add `./models` to AWS_DATA_PATH. If the directory was named `/ams/boto3models`, we would add the following code:

```
import os.environ['AWS_DATA_PATH'] = './ams/boto3models'

import boto3
import botocore
```

Your code should successfully find the AMS models. As a more specific example re: packaging, here's the Lambda specific workflow.

Example AMS Lambda Workflow:

These steps apply the preceding generic example to creating an AWS Lambda function.

1. Install the amscli. This creates a folder at `~/.aws/models` on your computer (Mac).
2. Copy the models to a local directory:

```
cp ~/.aws/models ./models
```

3. Add the models to your function's deployment zip file:

```
zip -r9 function.zip ./models
```

Important

Update your function code to add the new models to the SDK path. Note that this code must run before boto3 or botocore are imported!

```
# Force Python to search local directory for boto3 data models
import os
os.environ['AWS_DATA_PATH'] = './models'

import boto3
import botocore
```

Note

Because the example models are in a directory named `models`, We add `./models` to AWS_DATA_PATH. If the directory was named `/ams/boto3models`, we would add the following code:

```
import os
os.environ['AWS_DATA_PATH'] = './ams/boto3models'

import boto3
import botocore
```

Now, deploy your function:

1. Add your function code to the deployment zip file (if you haven't done so already):

```
zip -g function.zip lambda-amscm-test.py
```

2. Create or update your function with the zip file you created (console or CLI):

```
aws lambda update-function-code --function-name lambda-amscm-test --zip-file fileb://
function.zip --region us-east-1
```

Your AMS-integrated Python Lambda should now work.

> **Note**
>
> Your function must have IAM permissions for `amscm` or you get a permissions error.

Sample Lambda function code to test amscm (contents of lambda-amscm-test.py):

```
import json

# Force lambda to search local directory for boto3 data models
import os
os.environ['AWS_DATA_PATH'] = './models'

import boto3
import botocore


def lambda_handler(event, context):
    use_session = boto3.session.Session(region_name="us-east-1")
    try:
        cm = use_session.client("amscm")
        cts = cm.list_change_type_categories()
        print(cts)
    except botocore.exceptions.UnknownServiceError:
        print("amscm not found")

    return {
        'statusCode': 200,
        'body': json.dumps('Hello from Lambda!')
    }
```

Test outputs (success):

Function Response:

```
{
  "statusCode": 200,
  "body": "\"Hello from Lambda!\""
}

Request ID:
"1cea13c0-ed46-43b1-b102-a8ea28529c27"
```

Function Logs:

```
START RequestId: 1cea13c0-ed46-43b1-b102-a8ea28529c27 Version: $LATEST
{'ChangeTypeCategories': ['Deployment', 'Internal Infrastructure Management',
 'Management'], 'ResponseMetadata': {'RequestId': 'e27276a0-e081-408d-
bcc2-10cf0aa19ece', 'HTTPStatusCode': 200, 'HTTPHeaders': {'x-amzn-requestid':
 'e27276a0-e081-408d-bcc2-10cf0aa19ece', 'content-type': 'application/x-amz-json-1.1',
 'content-length': '89', 'date': 'Sun, 10 May 2020 23:21:19 GMT'}, 'RetryAttempts': 0}}
END RequestId: 1cea13c0-ed46-43b1-b102-a8ea28529c27
```

AMS API to Ruby example

In order to use the AMS API with Ruby, install the AWS Ruby SDK and AMS CLI. Follow these steps:

1. Install the AMS CLI. See Installing or upgrading the AMS CLI.
2. Install the AWS Ruby SDK. See Tools for Amazon Web Services.
3. Configure Ruby with these commands:

   ```
   require 'aws-sdk'

   config = {

    region: 'us-east-1',

    credentials: Aws::Credentials.new('ACCESS_KEY','SECRET_KEY')}
   ```
4. Get the AMS CTs:

   ```
   ams_cm = Aws::amscm::Client.new(config)
   ```

```
cts = ams_cm.list_change_type_classification_summaries

print(cts)
```

AMS API to Java example

In order to use the AMS API with Java, install the AWS Java SDK and AMS CLI. Follow these steps:

1. Install the AMS CLI. See Installing or upgrading the AMS CLI.
2. Install the AWS Java SDK. See Tools for Amazon Web Services.
3. Configure Java with these commands:

```
import com.amazonaws.auth.BasicAWSCredentials;

import com.amazonaws.services.amscm.model.AWSManagedServicesCMClient;

import
com.amazonaws.services.amscm.model.ListChangeTypeClassificationSummariesReque

import
com.amazonaws.services.amscm.model.ListChangeTypeClassificationSummariesResul

public static void getChangeTypeClassificationSummaries() {
```

4. Set the credentials. We recommend that you do not hardcode this.

```
final BasicAWSCredentials awsCredsCm =

 new BasicAWSCredentials("ACCESS_KEY", "SECRET_KEY");
```

5. Create the AMS Change Management client:

```
final AWSManagedServicesCMClient cmClient =

 new AWSManagedServicesCMClient(awsCredsCm);
```

6. Get the AMS CTs:

```
final ListChangeTypeClassificationSummariesRequest listCtsRequest = new
ListChangeTypeClassification SummariesRequest();

final ListChangeTypeClassificationSummariesResult listCtsResult =
```

```
cmClient.listChangeTypeClassificationSummaries(listCtsRequest);

System.out.println("List of CTs");

listCtsResult.getChangeTypeClassificationSummaries().stream()

.map(x -> x.getCategory() + "/" + x.getSubcategory() + "/" +
x.getItem() + "/" + x.getOperation())

.forEach(System.out::println);

}
```

AMS bring your own EPS

You can use the AMS "bring your own end point security" (BYOEPS) feature to replace the default Trend Micro Deep Security agent with your own end point security solution, or Trend Micro license.

If you already have cost effective licenses for products other than Trend Micro Deep Security, or a team that provides your EPS, or if you want to use a specific EPS tool, use BYOEPS in your instances.

BYOEPS works at an account level and your instances in the account either use BYOEPS or the default, AMS-managed EPS:

- **multi-account landing zone (MALZ)**: You designate application accounts where you use BYOEPS or managed EPS.
- **single-account landing zone (SALZ)**: Your AMS accounts use BYOEPS or managed EPS.

If you use BYOEPS, your AWS bill reduces by the cost for Trend Micro Deep Security; however, you still incur a cost for EPS as the AMS-managed EPS is still required for protecting AMS-created and maintained EC2 instances required for access management (bastions, and management hosts). To calculate the total cost impact, you need to account for the the cost of licenses for your new tool, and the cost of managing EPS at the service levels you need.

The use of BYOEPS changes the AMS roles and responsibilities for security management:

- **R** stands for responsible party that does the work to achieve the task.

- **C** stands for consulted; a party whose opinions are sought, typically as subject matter experts; and with whom there is bilateral communication.
- **I** stands for informed; a party which is informed on progress, often only on completion of the task or deliverable.

Security management	Customer	AWS Managed Services
Maintaining valid licenses of Managed EPS for EC2 instances of AMS Shared Services	R	C
Configure Managed EPS for EC2 instances of AMS Shared Services	I	R
Update Managed EPS for EC2 instances of AMS Shared Services	I	R
Monitoring malware on EC2 instances of AMS Shared Services	I	R
Maintaining and updating virus signatures for EC2 instances of AMS Shared Services	I	R
Remediating instances infected with malware for EC2 instances of AMS Shared Services	C	R

When you use BYOEPS, you lose one of the security controls offered by AMS but have security management provided using tools such as AWS GuardDuty, AWS Macie; and process controls such

as reviews of IAM configuration to ensure the security of your AWS account. AMS compliance certifications and attestations are not affected if you use BYOEPS. However, many security framework and certifications have requirements for protection from malware and malicious code. To ensure account security and compliance, evaluate and ensure that your planned controls meet the security requirements for compliance certifications needed for your workload.

Using BYOEPS

After ensuring that you want to use your own EPS solution, you are ready to request and begin using AMS BYOEPS.

Pre-requisites:

1. If you use an EC2 instance profile that is in addition to the default EC2 instance profile, `customer-mc-ec2-instance-profile`, allow the `ssm:GetParameter` action for the `/ams/end-point-security` resource, to your EC2 instance profile.
2. Update EC2 instance launch automations or processes using custom or AMS AMIs to use AMS AMIs released after December 2020.

Enable BYOEPS:

The use of BYOEPS changes the AMS responsibilities for Security Management. Consult your security and cloud platform team before enabling BYOEPS.

Request use of BYOEPS by submitting a "MOO" update RFC (Management | Other | Other | Update) with ct-0xdawir96cy7k, with the following details:

```
Please enable BYOEPS for this account/these accounts
Account IDs: IDs for the accounts for BYOEPS..
```

Accounts with EC2 instances using AMS-managed EPS:

If the accounts that you want BYOEPS for are using AMS-managed EPS, you need to work with AMS to uninstall the Trend Micro agents from those EC2 instances, and update the AMS code (i.e. boot scripts) on those instances. It is best to do this as part of a maintenance window as it may require a reboot. After AMS receives the MOO RFC, your cloud service delivery manager (CSDM) contacts you to decide on a maintenance window to perform this activity, and create a migration plan. A few things to consider as you plan:

- How many instances do you need to migrate in total? Divide the total number of instances into smaller, incremental batches.
- How would you divide the instances in batches? Options could be by resource groups, creating a list that can be shared with Operations etc.
- How much time would each batch take? How much total time is required? Consider that you might want to install your preferred EPS tooling in the same maintenance window. How much time would this take?

You will share this information with the AMS Operations engineers performing the migration. A spreadsheet, or some clear communication to your CSDM.

During the maintenance window, the following actions are performed on each instance that needs to be on-boarded to BYOEPS:

- Performed by AMS: Update AMS code (boot scripts, modules, and so forth.) to the latest. This is required because old AMS boot scripts do not have BYOEPS feature support and will re-install Trend Micro agent on every boot. Also, uninstall the Trend Micro Agent.
- Performed by You: Install, and configure your preferred EPS tooling.

Accounts without EC2 instances using AMS-managed EPS:

Accounts with new instance launches using the latest AMS AMIs can skip Trend Micro agent installation. Do not launch instances with AMIs older than December 2020 as they do not have the BYOEPS feature support. Update any automation using old AMIs to use the latest AMS AMIs with BYOEPS feature support. Once the feature is enabled, AMS confirms the action on the RFC.

Adding your agent on EC2 instances:

For help with deploying agents of tools such as Crowdstrike or Qualys, you can use AMS Patterns. Submit a service request for help.

Receiving AMS notifications

Communications between you and AMS occur for many reasons:

- An RFC created by AMS that requires your approval
- An AMS case created to investigate an RFC you created that has failed
- Events created by monitoring alerts

- Patching service notifications that inform you of upcoming patching
- Service requests and incident reports
- Monthly CRM reports
- Occasional important AWS announcements (your CSDM contacts you if any action on your part is required)

All of these notifications are sent to the default contact information (the root account email) that you provided AMS when you were onboarded. Because it's difficult to keep individual emails updated, we recommend that you use a group email that can be updated on your end. All notifications sent to you are also received by AMS operations and analyzed before making a response.

AMS notification service provides two additional ways to set up contacts for notifications:

- Tag your resources with contact tags (the tag Key Value being contact information) and provide the tag Key Name to your CSDM. Alarms on those resources will be sent to the contacts provided in the Key Value, in addition to the account contact created at onboarding. This is especially useful for application owners. For more information, see Tag-based alert notification.
- (Required at onboarding) Send to your CSDM named lists of contacts for non-resource based notifications. For example, you might have a list named "SecurityContacts" and another named "OperationsContacts", and so forth. AMS adds the list to the notification service, and alarms that apply to that list's context are sent to those contacts. This is especially useful for organizational matters.

This advanced alert routing feature is active for most of the essential CloudWatch alarms such as Amazon EC2 instance failure, Amazon Elastic Block Store (Amazon EBS) volume capacity utilization - Root usage, Amazon EBS NonRoot usage, High Memory utilization, High Swap usage, and High CPU utilization for Amazon EC2.

Additionally, when you file a service request, or incident report, you have the option of adding "CC Emails" (highly recommended) and those email addresses receive notifications about the service request or incident.

Important

While the CC email addresses provided in service requests and incident reports receive email notifications of communications, other notifications, such as patching notifications,

appear in your Service Request list (an email is also sent to the default contact), *without* explicit notification to you that you have a communication awaiting your attention. This is why we strongly recommend adding a CC email where you can, and setting up the default contact email as a group to which everyone using AMS is a member.

Additionally, you can request special notifications for new AMIs, for RFC state change, and for configuration changes in your AMS account. These optional notification services are discussed next.

AMS AMI notifications with SNS

AMS provides an AMI notification service. You can use it to subscribe to an Amazon Simple Notification Service (SNS) topic that notifies you when AMS AMI updates have been released. You can choose to receive notifications for only the AMS AMIs you use, or you can sign up to receive update notifications for all AMS AMIs. For more information on SNS topics, see What is Amazon Simple Notification Service?

Whenever AMIs are released, we send notifications to the subscribers of the corresponding topic; this section describes how to subscribe to the AMS AMI notifications.

Sample message


```
{
  "Type" : "Notification",
  "MessageId" : "example messageId",
  "TopicArn" : "arn:aws:sns:us-east-1:591688410472:customer-ams-windows2019",
  "Subject" : "New AMS AMIs are Now Available",
  "Message" : "{"v1": {"Message": "A new version of the AMS Amazon Machine Images has
 been released. You are now able to launch new EC2 stacks from these AMIs.
  Please use this time to update any dependencies such as CloudFormation or Autoscaling
 groups. Release Notes Windows - Contains latest Windows Patches:
  Microsoft Windows Server 2008 R2 Datacenter - (KB2819745, KB3018238, KB4507004,
 KB4507437) Microsoft Windows Server 2016 Datacenter Security Enhancedn - (KB4509091,
 KB4507459)
  Microsoft Windows Server 2016 Datacentern - (KB4509091, KB4507459) Microsoft Windows
 Server 2012 R2 Security Enhancedn - (KB3191564, KB3003057, KB3013172, KB3185319,
 KB4504418,
  KB4506996, KB4507463) Microsoft Windows Server 2012 R2 Standardn - (KB3003057,
 KB3013172, KB3185319, KB4504418, KB4506996, KB4507463) Linux - Contains latest Linux
 patches -
  All AMIs now force domainjoin-cli leave before domainjoin-cli join for better
 stability in the domain join process.", "images":
```

```
  {"images": {"image_name": "customer-ams-windows2019-2021.08-1", "image_id":
 "ami-05dfa45396fddaa5e"}}, "region": "us-east-1"}}",
  "Timestamp" : "2021-09-03T19:05:57.882Z",
  "SignatureVersion" : "1",
  "Signature" : "example sig",
  "SigningCertURL" : "example url",
  "UnsubscribeURL" : "example url"
}
```

Possible AMS AMI topics to subscribe to:

- **ALL**: Use `customer-ams-all-amis`. This topic subscription notifies you when any of the AMS AMIs are updated.
- **AMS AWS Linux AMIs**: For Amazon Linux, use `customer-ams-amazon1` and `customer-ams-amazon1-security-enhanced`. For Amazon Linux 2, use `customer-ams-amazon2` and `customer-ams-amazon2-security-enhanced`.
- **AMS SUSE Linux AMIs**: Use `customer-ams-sles12` or `customer-ams-sles15`.
- **AMS AWS RedHat AMIs**: Use `customer-ams-rhel8`, `customer-ams-rhel8-security-enhanced`, `customer-ams-rhel7`, `customer-ams-rhel7-security-enhanced`.
- **AMS AWS CentOs AMIs**: Use `customer-ams-centos7`, `customer-ams-centos7-security-enhanced`.
- **AMS Ubuntu AMIs**: Use `customer-ams-ubuntu18`.
- **AMS AWS Windows AMIs**: Use `customer-ams-windows2019`, `customer-ams-windows2019-security-enhanced`, `customer-ams-windows2016`, `customer-ams-windows2016-security-enhanced`, `customer-ams-windows2012`, `customer-ams-windows2012r2`, `customer-ams-windows2012r2-security-enhanced`, `customer-ams-windows2022`.

To subscribe to AMS new AMI notifications by using the Amazon SNS console:

1. Open the Amazon SNS console to the Dashboard.
2. In the upper-right corner, change to the AWS Region for the AMIs that you are subscribing to.
3. In the left-navigation pane, choose **Subscriptions**, and then choose **Create subscription**.
4. Provide the following information:

a. **Topic ARN**: `arn:aws:sns:{REGION}:287847593866:{AMS_AMI_NAME}` where REGION is the selected AWS Region (where the SNS notification was created) and AMS_AMI_NAME is the AMI that you want notifications about. Examples:

   - To subscribe to notifications of new AMS Amazon Linux AMIs in AWS Region us-east-1, use this **Topic ARN** = `arn:aws:sns:us-east-1:287847593866:customer-ams-amazon1`.
   - To subscribe to notifications of new AMS Window Server 2016 AMIs in AWS Region us-west-2, use this **Topic ARN** = `arn:aws:sns:us-west-2:287847593866:customer-ams-windows2016`

b. For **Protocol**, choose **Email**.

c. For **Endpoint**, enter an email address that you can use to receive the notifications. We recommend a distribution list rather than an individual's email.

5. Choose **Create subscription**.
6. When you receive a confirmation email with the subject line "AWS Notification - Subscription Confirmation," open the email and choose **Confirm subscription** to complete your subscription.

> **Note**
>
> You are not limited to email for the **Protocol**. For information on other acceptable protocols and how to use them, see subscribe.

To unsubscribe from AMS new AMI notifications by using the AWS SNS console:

1. Open the Amazon SNS console to the Dashboard.
2. In the navigation bar, change to the AWS Region of your choice. You must use the AWS Region in which you want to receive notifications for the corresponding AMIs.
3. In the navigation pane, choose **Subscriptions**, select the subscription, and then choose **Actions** -> **Delete subscriptions**.
4. When prompted for confirmation, choose **Delete**.

To subscribe to AMS New AMI notifications using the Deployment | Ingestion | Stack from CloudFormation Template | Create (ct-36cn2avfrrj9v):

1. To subscribe to the AmazonLinuxSubscription, create and save an execution parameters JSON file; this example names it CreateSubscribeAmiParams.json:

```
{
    "AWSTemplateFormatVersion": "2010-09-09",
    "Resources": {
        "AmazonLinuxSubscription":{
              "Type" : "AWS::SNS::Subscription",
              "Properties": {
                "TopicArn": "arn:aws:sns:{REGION}:287847593866:{AMS_AMI_NAME}",
                "Protocol": "email",
                "Endpoint": "username@yourdomain.com"
            }
        }
      }
}
```

2. Create and save the RFC parameters JSON file with the following content; this example names it CreateSubscribeAmiRfc.json file:

```
{
   "ChangeTypeId": "ct-36cn2avfrrj9v",
   "ChangeTypeVersion": "1.0",
   "Title": "cfn-ingest-subscribe-ami"
}
```

3. Create the RFC, specifying the CreateSubscribeAmiRfc file and the CreateSubscribeAmiParams file:

```
aws amscm create-rfc --cli-input-json file://CreateSubscribeAmiRfc.json  --
execution-parameters file://CreateSubscribeAmiParams.json
```

   You receive the ID of the new RFC in the response and can use it to submit and monitor the RFC. Until you submit it, the RFC remains in the editing state and does not start.

For examples of creating AMIs, see Create AMI.

For information on consuming AMIs programmatically, see EC2 stack: creating.

Service notifications

AMS sends outbound service requests, or service notifications, when you need to act on, or be aware of, something that might impact your account or resources, including:

- Infrastructure impact: AMS sends a service notification when there is an underlying AWS service impacting your infrastructure, and you need to take action before a certain date, or you may have an outage.
- EC2 Hardware issues: AMS sends service notifications out for EC2 hardware issues that require you to reboot an EC2 instance before a certain date, or letting you know that AMS will reboot the instance for you. This is an important notice because reboot can cause an outage and you must respond with an acceptable date, or create an RFC with ct-09qbhy7kvtxqw, to reboot the instance yourself. A service notification like this automatically closes in five days if you do not respond.

RFC state change notifications

AMS offers notifications for RFC state changes by email and CloudWatch Events:

- Emails by way of the AMS Console: There is an option on the second page of the Create RFC wizard, where you can add up to five email addresses to be notified when that RFC state changes.
- CloudWatch Events: You can configure different rules and targets for CloudWatch Events to receive notifications for every RFC state change.

Email notifications

You can add email addresses to receive RFC state changes to an RFC that you create in the AMS console, or by using the AMS API/CLI.

In the AMS console, use the **Email notifications** option, on the second page of the Create RFC wizard:

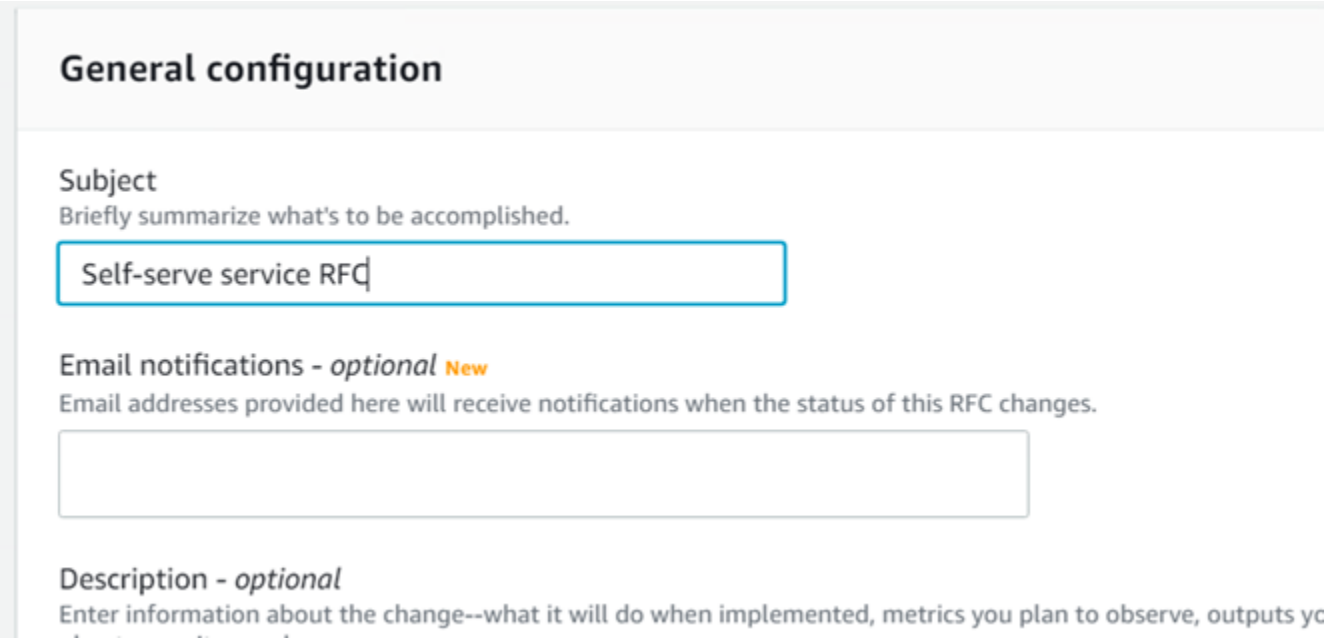

In the AMS API/CLI, add a line like this to the RFC parameters section of your RFC (do not add the line to the run parameters section):

```
--notification "{\"Email\": {\"EmailRecipients\" : [\"email@example.com\"]}}"
```

The behavior of the notifications varies depending the RFC scheduling type:

- Scheduled RFCs receive email notifications on : Submitted, Scheduled, InProgress, Completed, Rejected, Canceled, Auto-Rejected, or Auto-Canceled.
- ASAP RFCs receive email notification on: Submitted, InProgress, Completed, Rejected, Canceled, AutoRejected, or Auto-Canceled.

Note

- Email notifications are sent from this address: `no-reply@managedservices.amazonaws.com`.
- Special characters and URLs in your RFC title are redacted in the emails we send. This is a security measure.

CloudWatch Events notifications

AMS offers push notifications for the RFC State changes through CloudWatch Events. To get these notifications:

1. Create a topic and subscription where notifications will be sent. You can name the topic what you like; for information about doing this, see SNS Topic and Subscription: Creating.
2. Submit an RFC with the Management | Other | Other | Create change type and include the SNS topic and subscription in the request for RFC state change notices.

When you submit the Management | Other | Other RFC request for this feature, you can specify what RFC state changes you're interested in getting notified about and what change types, and set other filters. For example, you may want to request to be notified only when Admin Access change types are EventType = RfcSubmitted and EventType = RfcUpdated.

This is a template of CloudWatch event notifications that you can receive (with all possible values):

```
{
    "source ": "aws.managedservices",
    "detail-type": "AMS RFC State Change",
    "detail": {
        "ActionState": "null | AwsActionPending | AwsOperatorAssigned |
 CustomerActionPending | NotApplicable | NoActionPending",
        "ActualExecutionTimeRange": {
            "StartTime": "null | Actual Start Time",
            "EndTime": "null | Actual End Time"
        },
        "AutomationStatus": "Automated | Manual",
        "AwsAccountId": "AWS Account ID",
        "AwsApprovalStatus": "null | SubmissionPending | NotRequired | ApprovalPending
 | Rejected | Approved",
        "ChangeTypeId": "Change_Type_ID",
        "ChangeTypeVersion": "Change_Type_Version",
        "CreatedTime": "Created_Time",
        "CustomerApprovalStatus": "null | SubmissionPending | NotRequired |
 ApprovalPending | Rejected | Approved",
        "EventType": "RfcActionStateUpdated | RfcApproved | RfcAutoRejected |
 RfcCanceled | RfcCompleted | RfcCreated | RfcInProgress | RfcRejected | RfcSubmitted |
 RfcUpdated",
        "LastModifiedTime": "Last_Updated_Time",
        "LastSubmittedTime": "null | Last_Submitted_Time",
```

```
        "RequestedExecutionTimeRange": {
            "StartTime": "null | Expected_Start_Time",
            "EndTime": "null | Expected_End_Time"
        },
        "RfcId": "RFC_ID",
        "Status": "Editing | PendingApproval | Scheduled | Rejected | Canceled |
 ExecutionLock | InProgress | Success | Failure",
        "Title": "Title"
    }
}
```

The supported RFC state changes (EventType), as they appear in the actual CloudWatch Events notification are:

- RfcActionStateUpdated (no AMS console option): The RFC in one of the states, described later, changed.
- RfcApproved (no AMS console option): The RFC passed system and/or AMS operator validation and has been approved for completion.
- RfcAutoRejected (**Auto-Rejected**): The RFC failed system validation or AMS operator and has been rejected.
- RfcCanceled (**Canceled** or **Auto-Canceled**): The RFC was canceled by either the submitter or an AMS operator.
- RfcCompleted (**Completed**): The RFC run parameters have been completed, including UserData.
- RfcCreated (no AMS console option): The RFC was successfully created (the JSON and submitted parameters were valid).
- RfcInProgress (**InProgress**): The RFC run is still in progress.
- RfcRejected (**Rejected**): The RFC failed system or AMS operator validation has been rejected.
- RfcSubmitted (**Submitted**): The RFC has been submitted and is undergoing system validation.
- RfcUpdated (no AMS console option): The RFC has been manually updated by an AMS operator.

Additionally, you can send CloudWatch Events (CWE) notifications to any of the supported destinations and build your own systems on top of these automated notifications:

- Amazon EC2 instances
- AWS Lambda functions
- Streams in Amazon Kinesis Data Streams

- Delivery streams in Amazon Data Firehose
- Log groups in Amazon CloudWatch Logs
- Amazon ECS tasks
- Systems Manager Run Command
- Systems Manager Automation
- AWS Batch jobs
- Step Functions state machines
- Pipelines in CodePipeline
- CodeBuild projects
- Amazon Inspector assessment templates
- Amazon SNS topics
- Amazon SQS queues
- Built-in targets: EC2 CreateSnapshot API call, EC2 RebootInstances API call, EC2 StopInstances API call, and EC2 TerminateInstances API call.
- The default event bus of another AWS account

> **Note**
>
> We send CloudWatch Events notification for RFC state changes, on a best-effort basis.

Setting up private and public DNS

During onboarding, AMS sets up a private DNS service for communications between your managed resources and AMS.

You can use AMS Route 53 to manage the internal DNS names for your application resources (web servers, application servers, databases, and so forth) without exposing this information to the public Internet. This adds an additional layer of security, and also allows you to fail over from a primary resource to a secondary one (often called a "flip") by mapping the DNS name to a different IP address.

After you create private DNS resources using the Deployment | Advanced stack components | DNS (private) | Create (ct-0c38gftq56zj6) or Deployment | Advanced stack components | DNS

(public) | Create (ct-0vzsr2nyraedl), you can use the Management | Advanced stack components | DNS (private) | Update (ct-1d55pi44ff21u) and Management | Advanced stack components | DNS (public) | Update (ct-1hzofpphabs3i), CTs to configure additional, or update existing, record sets. For multi-account landing zone (MALZ) accounts, DNS resources created in the application account VPCs can be shared with the shared services account VPC to maintain centralized DNS using AMS AD.

MALZ

The following graphic illustrates a possible DNS configuration for Multi-Account Landing Zone AMS. It illustrates a hybrid DNS setup between AMS and a typical customer network. A Canonical Name Record (CNAME) in the customer network DNS server forwards to the AMS AD DNS in the shared services account with a conditional forward that has the CNAME of the AMS FQDN forwarded to the A record.

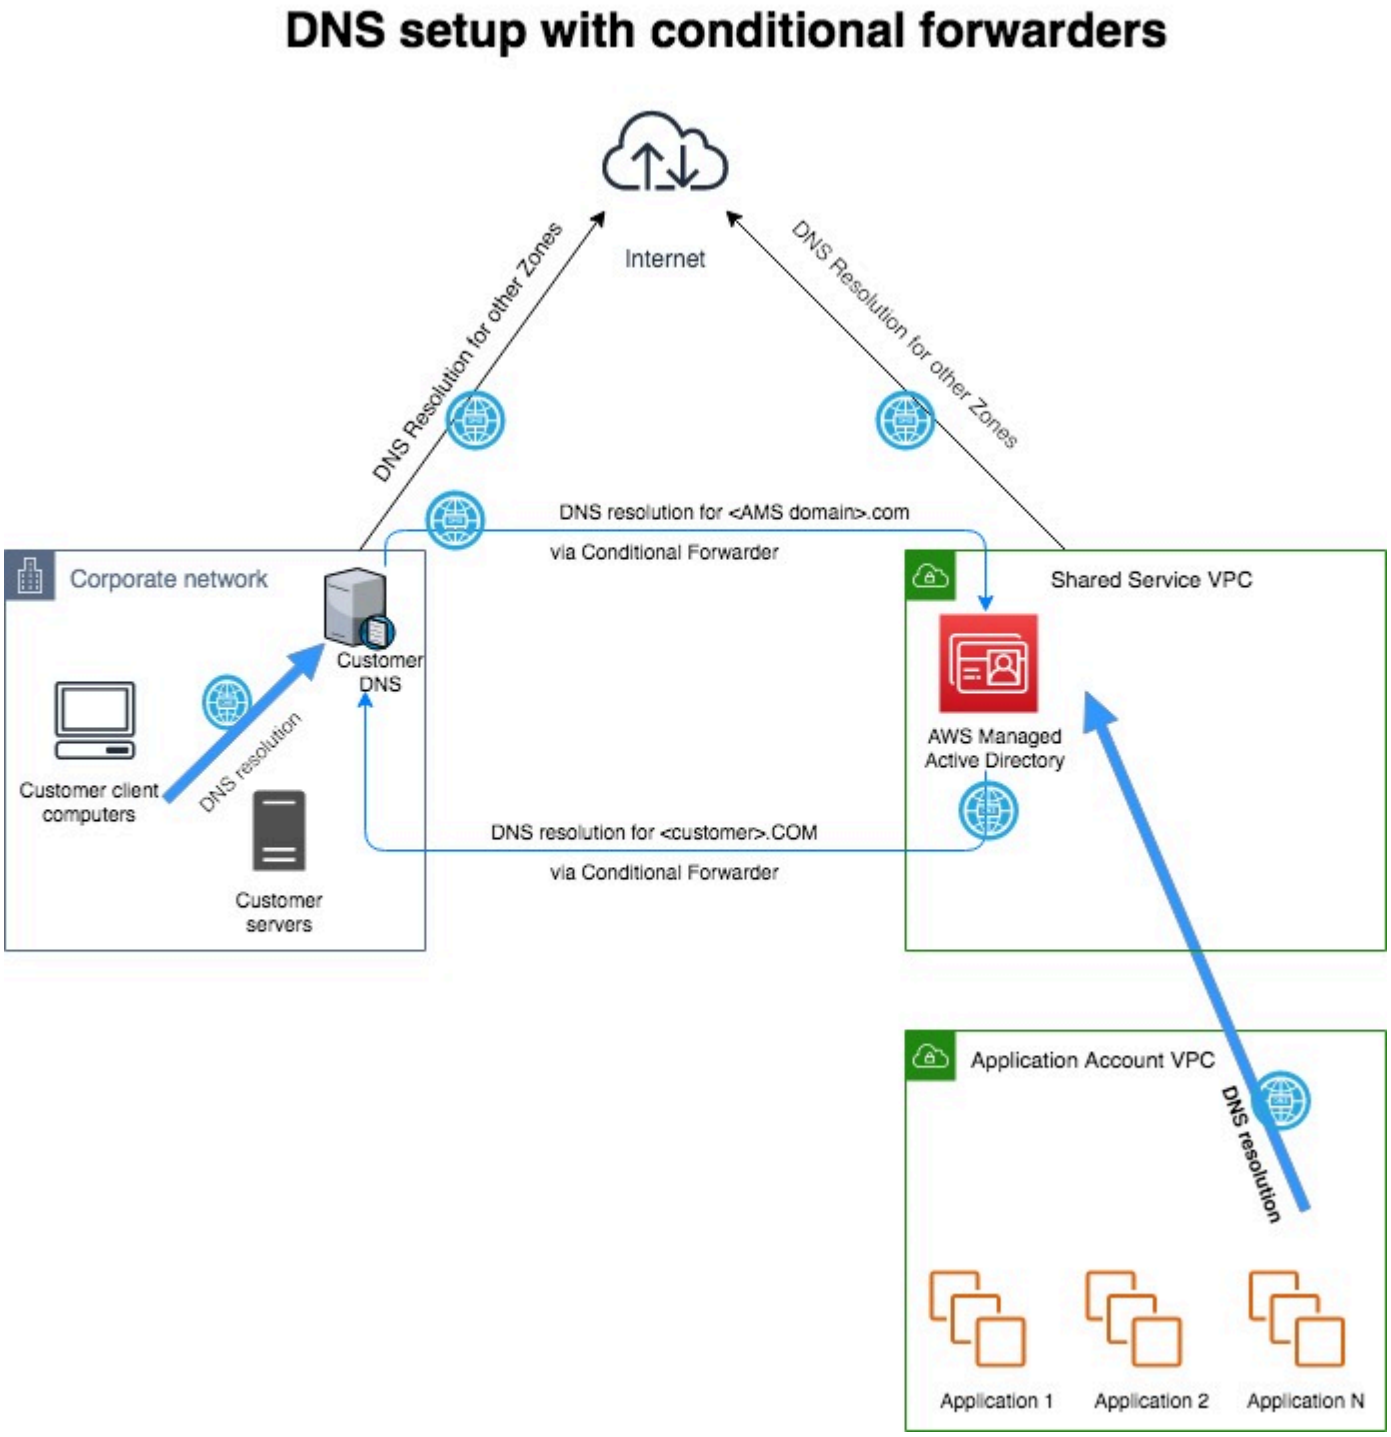

SALZ

The following graphic illustrates a possible DNS configuration for single-account landing zone (SALZ). It shows a hybrid DNS setup between AMS and a typical customer network. A CNAME in the customer network DNS server forwards to the AMS AD DNS with a conditional forward which has the CNAME of the AMS FQDN forwarded to the A record.

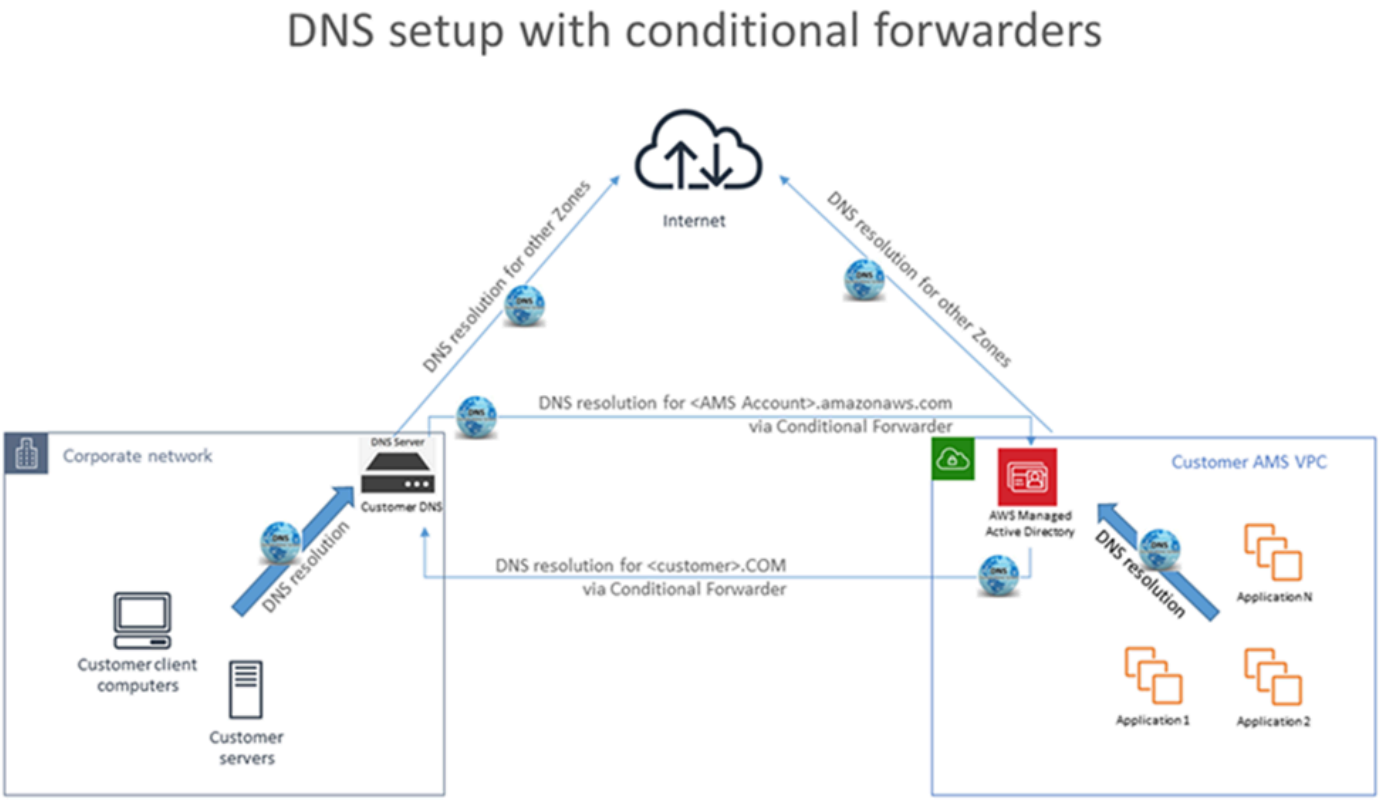

SALZ Route53 DNS

The following graphic illustrates a possible DNS configuration for single-account landing zone (SALZ). It shows a hybrid DNS setup between AMS and a typical customer network. A CNAME in the customer network DNS server forwards to the AMS AD DNS with a conditional forward which has the CNAME of the AMS FQDN forwarded to the A record. This also leverages Route53 for outbound network traffic so that any application in the account can have DNS Resolution in the account with the highest availability.

Route53 enabled Resolution paths:

- Instance attempting to resolve AMS MAD name --> VPC +2 (Route53/AmazonProvidedDNS) --> Conditional Forwarders evaluated --> Route53 MAD Conditional Forwarder rule matched --> Route53 Outbound resolver --> Managed AD DNS
- Instance attempting to resolve customer on-prem name --> VPC +2 (Route53/AmazonProvidedDNS) --> Conditional Forwarders evaluated --> Route53 On-prem Conditional Forwarder rule matched --> Route53 Outbound resolver --> Customer on-prem DNS
- Instance attempting to resolve Internet name --> VPC +2 (Route53/AmazonProvidedDNS) --> Conditional Forwarders evaluated --> No matching forwarder --> Internet DNS Service

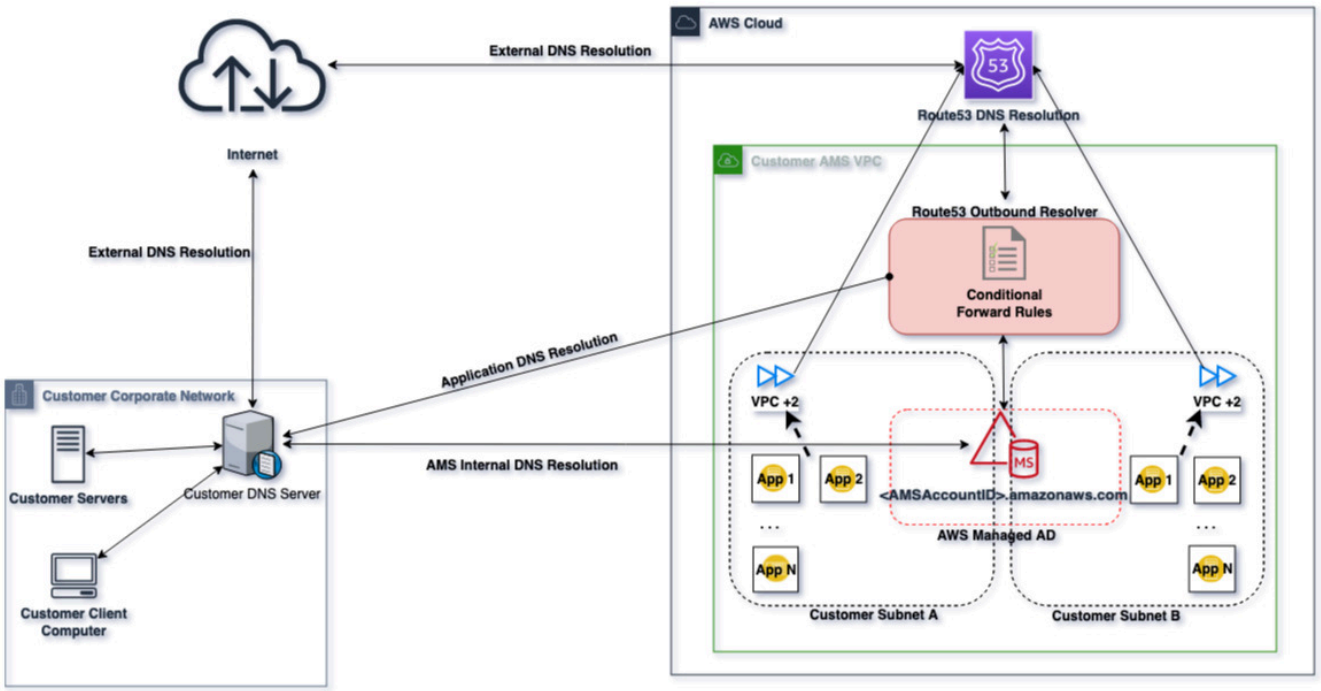

For more information, see Using DNS with Your VPC and Working with Private Hosted Zones.

AMS egress traffic management

By default, the route with a destination CIDR of 0.0.0.0/0 for AMS private and customer-applications subnets has a network address translation (NAT) gateway as the target. AMS services, TrendMicro and patching, are components that must have egress access to the Internet so that AMS is able to provide its service, and TrendMicro and operating systems can obtain updates.

AMS supports diverting the egress traffic to the internet through a customer-managed egress device as long as:

- It acts as an implicit (for example, transparent) proxy.

  and

- It allows AMS HTTP and HTTPS dependencies (listed in this section) in order to allow ongoing patching and maintenance of AMS managed infrastructure.

Some examples are:

- The transit gateway (TGW) has a default route pointing to the customer-managed, on-premises firewall over the AWS Direct Connect connection in the Multi-Account Landing Zone Networking account.
- The TGW has a default route pointing to an AWS endpoint in the Multi-Account Landing Zone egress VPC leveraging AWS PrivateLink, pointing to a customer-managed proxy in another AWS account.
- The TGW has a default route pointing to a customer-managed firewall in another AWS account, with site-to-site VPN connection as an attachment to the Multi-Account Landing Zone TGW.

AMS has identified the corresponding AMS HTTP and HTTPS dependencies, and develops and refines these dependencies on an ongoing basis. See Egress Management ZIP. Along with the JSON file, the ZIP contains a README.

Note

- This information isn't comprehensive--some required external sites aren't listed here.
- Do not use this list under a deny list or blocking strategy.
- This list is meant as a starting point for an egress filtering rule set, with the expectation that reporting tools will be used to determine precisely where the actual traffic diverges from the list.

To ask for information about filtering egress traffic, email your CSDM: ams-csdm@amazon.com.

Deploying IAM resources in AMS Advanced

AMS can deploy IAM resources in in your multi-account landing zone (MALZ) Application and single-account landing zone (SALZ) accounts in three different ways:

- Automated IAM Provisioning: This capability in AMS lets you submit create, update, or delete change types for IAM role or policy provisioning, without operator review, and with IAM and AMS validation checks run automatically.

  This capability must be explicitly enabled with the Management | Managed account | AMS Automated IAM Provisioning with read-write permissions | Enable (review required) change type (ct-1706xvvk6j9hf). To learn more, see Automated IAM Provisioning AMS. After AMS Automated

IAM Provisioning is enabled, you have access to Create, Update, and Delete change types to manage your IAM resources.

- Restricted IAM change type: This change type, Deployment | Advanced stack components | Identity and Access Management (IAM) | Create entity or policy (read permissions), (ct-19jq3ulr3g9zg), runs your resource deployment request without an operator review, but only resources with a subset of actions defined in the IAM console **ReadOnlyAccess** managed policy are allowed. AMS supports customization of this baseline (control) policy, subject to an AMS Security review. To learn more, see RFC security reviews and reach out to your cloud service delivery manager (CSDM) or cloud architect (CA). To see the baseline policy, download this ZIP file: baseline.zip.

  A number of validations are executed at the deployment time. These include IAM Policy linter, and checks that the IAM resource conforms to AMS technical standards (for example, only IAM roles trust policies between AMS accounts that belong to the same customer can be configured).
- Review required IAM change type: This change type, Deployment | Advanced stack components | Identity and Access Management (IAM) | Create entity or policy (review required) (ct-3dpd8mdd9jn1r), requires an AMS operator review, which can sometimes take a few days to complete if clarifications are needed.

> **Note**
>
> Whichever method is used, an IAM role is provisioned to the relevant account or accounts and, after the role is provisioned, you must onboard the role in your federation solution.

Automated IAM Provisioning AMS

AWS Managed Services (AMS) supports automated validation and provisioning for IAM resources including roles and policies using AMS Advanced requests for change (RFCs) and new change types (CTs). Previously, these requests went through a semi-automated process that sometimes resulted in long wait times. Now, you can use AMS Automated IAM Provisioning to provision IAM resources and get the results much more quickly.

How Automated IAM Provisioning in AMS works

Automated IAM Provisioning relies on automated run-time checks for IAM to validate changes to IAM resources. These automated checks, performed when the Create, Update, or Delete change

types are run, prevent IAM resources that are overly-permissive or have insecure patterns from being deployed into your account. This allows you to match the level of rigor in IAM reviews to the expertise of your team. We recommend that teams that are new to cloud services and need manual checks for all IAM resources changes use the existing review-required change type: Deployment | Advanced stack components | Identity and Access Management (IAM) | Create entity or policy (review required), (ct-3dpd8mdd9jn1r). Teams with AWS expertise and control of their environments can use Automated IAM Provisioning to speed up their deployments. You can use this feature to perform validation through automated run-time checks or to perform validation and provisioning of IAM resources after successful validation.

Important

AWS Managed Services has proactively implemented a list of validation runtime checks that prevent the creation of IAM resources or policies with certain permissions and conditions. For a description of these privileges and conditions, see Deploying IAM resources in AMS Advanced. The automated change types ct-1n9gfnog5x7fl, ct-1e0xmuy1diafq, and ct-17cj84y7632o6, allow users who are proficient in managing IAM resouces to provision IAM roles and policies that allow actions beyond Read Only privileges. In addition, you can use the roles created through the automated change types ct-1n9gfnog5x7fl, ct-1e0xmuy1diafq, and ct-17cj84y7632o6 to create the new resources. However, the resources can't follow the AMS naming standard and aren't part of the standard AMS stack. AMS provides the operational and security support of those specific resources on a best effort basis.

While both manual and automated processes aim to uphold our security standards, it's important to note that there are differences in the checks between the two. The automated provisioning allows for greater flexibility in creating and updating roles and policies; therefore, they are not the same. It's recommended that your organization carefully review the validation runtime checks listed in the AMS User Guide to ensure that they align with your organization's expectations and requirements.

Validation flow

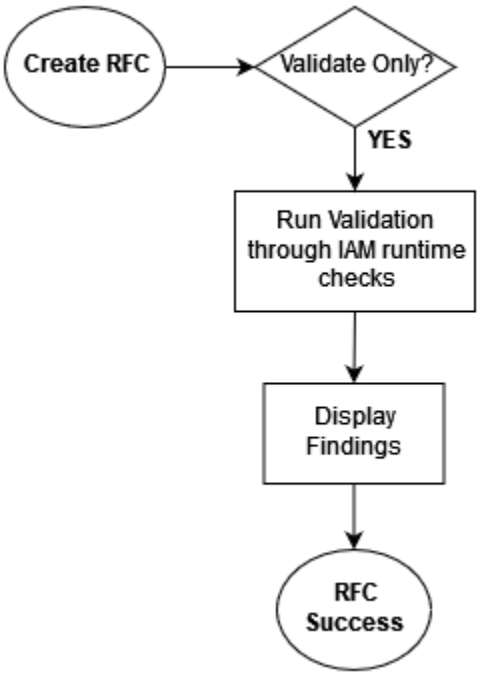

Validation and provisioning flow

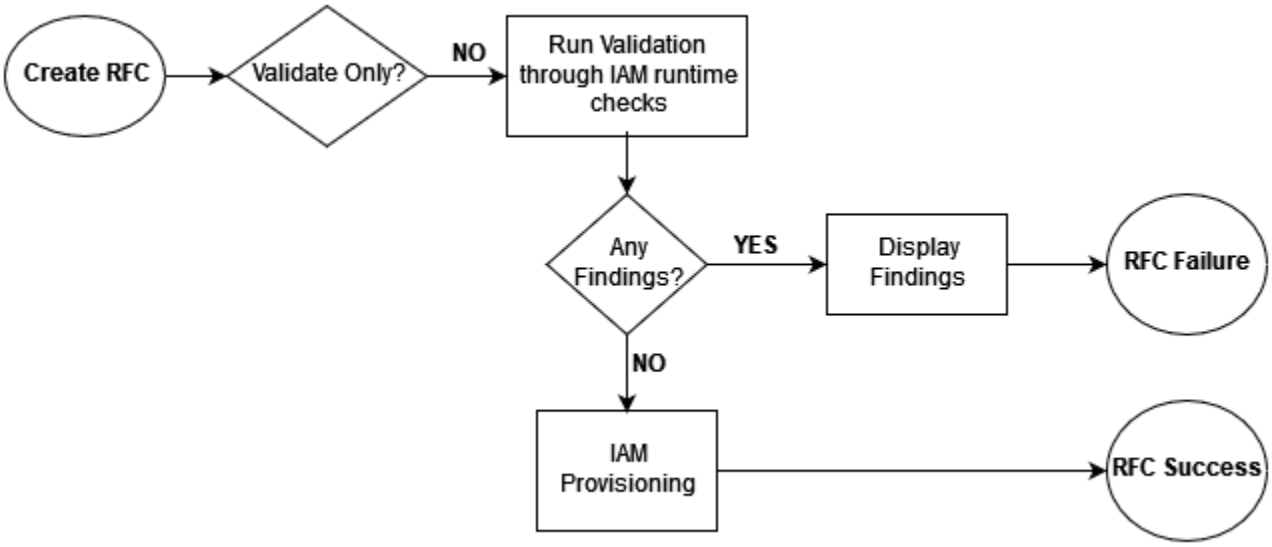

> **Note**
>
> This feature is suitable for teams that are experienced with AWS and IAM resources, and we do not recommend it for teams that are new to AWS. The automated validation process is designed to catch most errors and is helpful for teams to get quick reviews for changes to

IAM, when they understand the permissions that they need. To use the new change types safely and effectively, we recommend you to have a good understanding of AWS IAM, and the run-time checks offered by the change types to determine whether they are suitable for your team.

Onboarding to AMS Automated IAM Provisioning in AMS

To use the new change types, first enable AMS Automated IAM Provisioning by submitting an RFC using the following change type: Management | Managed account | AMS Automated IAM Provisioning with read-write permissions | Enable (review required) (ct-1706xvvk6j9hf). AWS requires that your organization go through a customer security risk management (CSRM) process to ensure that the use of these change types are aligned with your organizational policies. The AWS operations team works with you to acquire explicit approval from your security team contact in the form of risk acceptance as part of the required review. To learn more, see the RFC customer risk management (CSRM) process.

After the RFC for turning on AMS Automated IAM Provisioning with read-write permissions feature is successful, AMS enables the AMS Automated IAM Provisioning change types in the account used to submit the enable RFC. To confirm that an account has AMS Automated IAM Provisioning turned on, check the IAM console for the `AWSManagedServicesIAMProvisionAdminRole` role.

As part of onboarding, AMS provisions IAM Access Analyzer in the same AWS Region of the account to leverage its access preview capability. IAM Access Analyzer helps identify resources in your organization and accounts that are shared with an external entity, validates IAM policies against policy grammar and best practices, and generates IAM policies based on access activity in your AWS CloudTrail logs. To learn more, see Using AWS Identity and Access Management Access Analyzer.

Once onboarded, the `AWSManagedServicesIAMProvisionAdminRole` is deployed to the enabled accounts. If you choose to use this role through SAML federation, then you must onboard the role to your federation solution.

As part of onboarding, you can request to update AWSManagedServicesIAMProvisionAdminRole's trust policy to grant another IAM role ARN to assume this role using AWS Security Token Service.

Using AMS Automated IAM Provisioning in AMS

You can create RFCs with the following AMS Automated IAM Provisioning change types.

Note

- Only provisioning on roles and policies are supported.

  While updating roles, the Update CT replaces the existing list of managed policy Amazon resource names (ARNs) and the "assume role" policy document, with the provided list of managed policy ARNs and "assume role" policy document. In a partial update; for example, adding or removing an ARN in the existing list of managed policy ARNs, adding or removing individual policy statements to the "assume role" policy document is not allowed. Similarly, while updating policies, the Update CT replaces the existing policy document and does not allow adding or removing individual policy statement in the existing policy document.
- When the “validate only” option is selected, run-time checks are performed without provisioning any IAM entity or policy. Regardless of any findings, the RFC status is “success”. The "success" status indicates a successful validation against the provided IAM entity or policy.

- Deployment | Advanced Stack Components | Identity and Access Management (IAM) | Create entity or policy (read-write permissions)(ct-1n9gfnog5x7fl): A new IAM entity or policy is validated and provisioned automatically.
- Management | Advanced Stack Components | Identity and Access Management (IAM) | Update entity or policy (read-write permissions)(ct-1e0xmuy1diafq): An existing IAM entity or policy is updated and validated automatically.
- Management | Advanced Stack Components | Identity and Access Management (IAM) | Delete entity or policy (read-write permissions)(ct-17cj84y7632o6): An existing IAM entity or policy that's provisioned using the automated create entity or policy change type is deleted.

You can only call the preceding three CTs using a dedicated IAM role: `AWSManagedServicesIAMProvisionAdminRole`. This role is available only in the accounts that have been onboarded to the feature using the Management | Managed account | AMS Automated IAM Provisioning read-write permissions | Enable (review required) (ct-1706xvvk6j9hf).

> **Important**
>
> The Create, Update, and Delete change types are always visible in your account, but they aren't turned on by default. If you try submit an RFC using one of these change types without first enabling the AMS Automated IAM Provisioning feature, then an "unauthorized" error displays.

Limitations:

- The Create CT might allow you to create an IAM role or policy with permission to create AWS resources. However, AWS resources created by these roles and policies aren't managed by AMS. It's a best practice to adhere to your organizational control to limit creation of such roles or policies.
- The Update CT can not modify IAM roles and policies created with CFN Ingest, Direct Change Mode, Developer Mode, or, in some cases, through existing AMS Advanced manual or automated CTs.
- The Delete CT can not delete existing roles or policies that are not created with the AMS Automated IAM Provisioning Create CT.
- The AMS Automated IAM Provisioning with read-write permissions feature isn't supported in Direct Change Mode roles. This means that you can't provision or update IAM roles and policies with read-write permissions using these roles.
- AMS Automated IAM Provisioning with read-write permissions Create, Update, and Delete change types are not compatible with the ServiceNow Connector.

Runtime checks for AMS Automated IAM Provisioning in AMS

Automated IAM Provisioning leverages checks from AWS Identity and Access Management Access Analyzer, and performs additional checks and validations against the AMS boundary policy. AMS defined the additional checks and validations based on IAM best practices, experience operating customer workload in the cloud, and the collective AMS IAM manual evaluation experience.

You can view policy run-time check findings in the request for change (RFC) output. The findings include the resource identifier, location within the role and/or policy that generated the findings, and a message outlining the check that the IAM entity or resource failed to pass. These findings help you author policies that are functional and conform to security best practices.

> **Note**
>
> Automated IAM Provisioning attempts to be specific about the location within the entity or policy definition that fails to pass the check. Depending on the type, the location might include the resource name or ARN, or index within an array. For example, a statement to help you adjust the entity or policy for a successful outcome.

For a smooth AMS Automated IAM Provisioning experience, it's a best practice to use the “validate only” option to run the validation checks until there are no findings from the validation checks reported in the RFC outputs. When the validation checks report no findings, choose **Create copy** from the AMS Console to quickly create a copy of the existing RFC. When you are ready to provision, in the **Parameters** section, switch the **Validate only** value from **Yes** to **No**, and then proceed.

These are the run-time checks that AMS Automated IAM Provisioning performs to ensure that your IAM resources are secure:

> **Note**
>
> To provision IAM policies that contain actions denied by these automated change types, you must follow the RFC customer security risk management (CSRM) process. Use the following change type: Deployment | Advanced stack components | Identity and Access Management (IAM) | Create entity or policy (review required) (ct-3dpd8mdd9jn1r).

- **IAM Access Analyzer policy check and validation:** See also Access Analyzer policy check reference and IAM Access Analyzer policy validation.
- **AMS permissions boundary policy checks:** Actions on a set of services that are denied by default. For more information, see Automated IAM Provisioning permission boundary check.
- **Customer-defined permissions boundary policy checks:** Additional restricted actions on a set of services that are denied. For more information, see Automated IAM Provisioning permission boundary check.
- **AMS-defined custom checks**: Checks that identify various insecure and overly permissive policies or access patterns within a requested IAM entity or policy, and denies the request if found one. For for information, see AWS JSON policy elements: Principal.

Finding	Description
The role can be accessed from an external account that is outside of your zone of trust.	This finding refers to a principal listed in the role trust policy that is outside of your zone of trust. A zone of trust is defined as the account where the role is being created or the AWS organization that the account belongs to. An entity that does not belong to the account or to the same AWS Organization is an external entity. To resolve the finding, review the account ID in the principal ARNs and make sure that they belong to you and is an AMS onboarded account.
The role can be accessed by an external entity owned by account *External_Account_ID* that is not owned by the AMS customer-owning account *Account_ID* .	This finding is generated if the role trust policy includes a principal ARN that has an account ID not owned by you and an AMS onboarded account. To resolve this finding remove any such principal from the role trust policy.
The canonical user ID is not a supported principal in IAM trust policy.	Canonical principal IDs are not supported in IAM trust policy. To resolve the finding remove any such principal from the role trust policy.
The role can be accessed by an external web identity that is outside of your zone of trust.	This finding is generated if the role trust policy allows an external Web identity provider (IdP) other than SAML IdP. To resolve this finding, review the role trust policy and remove statements that allow the `sts:AssumeRoleWithWebIdentity` operation.
The role can be accessed through SAML federation; however, the provided SAML identity provider (IdP) does not exist.	This finding is generated if the role trust policy contains SAML IdP that does not exist in your account. To resolve ensure you all the listed SAML IdP exists in your account.
Policy contains privileged actions equivalent to administrator or power user access.	It's a best security practice in AWS Identity and Access Management to grant only the

Finding	Description
Consider reducing the permission scope to a specific service, action, or resource. If advanced policy elements such as **NotAction** or **NotResource** are used, make sure that they are not granting more access than you intend, particularly in **Allow** statements.	permissions required to perform a task when you set permissions with IAM policies. Do this by defining the actions that can be taken on specific resources under specific conditions, also known as least-privilege permissions. This finding is generated when automation detects the policy grants broad permissions and does not adhere to the principle of least privilege. To resolve the finding, review and reduce the permissions.
Statement contains privileged actions for *Service_Name* . Consider excluding these actions with a deny statement. Refer to the boundary policy reference in the AMS documentation for a list of privileged actions.	AMS identified certain actions for a given service as risky and require further risk review and acceptance by the customer security team. This finding is generated when automation detects the given policy granting such permissions. To resolve this finding, deny these actions in your policy. For a list of actions refer to the AMS boundary policy. For details on AMS boundary policy, see AMS Automated IAM Provisioning permission boundary check.
Statement grants access to privileged RFC Change Types: ct-1n9gfnog5x7fl, ct-1e0xmuy1diafq, and ct-17cj84y7632o6 for service *Service_Name* . Consider scoping the permissions to specific change types or exclude these change types with a deny statement.	This finding is generated if the policy grants permissions to perform RFC-related actions using Automated IAM Provisioning change types (CTs). The CTs are subject to risk acceptance and must only be used through onboarded roles. So, you can't granting permission to these CTs. To resolve this finding, deny RFC actions using these CTs.

Finding	Description
Statement contains privileged actions that are not scoped to your resources for service *Service_Name* . Consider scoping the actions to specific resources or exclude resources with AMS namespace prefixes. If wildcards are used ensure they restrict the scope to your resources.	This finding is generated if the policy grants privileged actions that are not scoped to your resources of the given service. Wild cards often create overly permissive policies that bring a broad set of resources or actions into the permission's scope. To resolve the finding, either reduce the scope of permissions to resources you own or exclude resources that are in the AMS namespace. For a list of AMS namespace prefixes, see the boundary policy in AMS documentation. Note that not all prefixes apply to all services. For details on the AMS boundary policy, see AMS Automated IAM Provisioning permission boundary check.
Invalid account Id or Amazon Resource Name (ARN).	This finding is generated if any ARN or account ID specified in the policy or role trust policy is invalid. To review valid resource ARN's resources for services, see the Service Authorization Reference. Make sure that the account ID is a 12-digit number and that the account is active in AWS.
Use of wildcard (*) for account id in ARN is restricted..	This finding is generated if a wild card (*) is specified in the account ID field of an ARN. A wild card in an account ID field matches any account and potentially grants unintende d permission to resources. To resolve this, replace the wild card with a specific account ID.

Finding	Description
Specified resource account not owned by same AMS customer owning account *Account_ID* .	This finding is generated if an account ID specified in a resource ARN does not belong to you and is not managed by AMS. To resolve this, make sure that all resources (as specified by their ARN in the policy) belong to your accounts that are managed by AMS.
The role name is in AMS restricted namespace.	This finding is generated if you try to create a role with a name that starts with an AMS reserved prefix. To resolve this, use a name for the role that is specific to your use case. For a list of AMS reserved prefixes, see AMS reserved prefixes
The policy name is in AMS restricted namespace.	This finding is generated if you try to create a policy with a name that starts with an AMS reserved prefix. To resolve this, use a name for the policy that is specific to your use case. For a list of AMS reserved prefixes, see AMS reserved prefixes.
The resource ID in the ARN is in AMS restricted namespace.	This finding is generated if you try to create a policy that grants permission to named resources that are in the AMS namespace . To resolve this, make sure that you scope the permissions to your resources or deny permissions to resources that are in the AMS namespace. For more information on AMS namespaces, see AMS restricted namespaces.

Finding	Description
Invalid policy variable case. Update the variable to *Variable_Names* .	This finding is generated if try to create a policy that contains an IAM global policy variable in the incorrect case. To resolve this, use the correct case for global variables in your policy. For a list of global variables, see AWS global condition context keys. For more information on the policy variables, see IAM policy elements: Variables and tags
Statement contains privileged actions that are not scoped to your KMS keys. Consider scoping these permissions to specific keys or exclude AMS owned keys.	This finding is generated if the policy contains permissions that are not scoped to specific KMS keys that you own. To resolve this, scope the permission to specific keys or exclude the keys that are AMS owned. AMS owned keys have specific alias sets. For a list of AMS owned key aliases, see AMS Automated IAM Provisioning permission boundary check.
Statement contains privileged actions that are not scoped to your KMS keys aliases. Consider scoping these permissions to your keys or aliases, or exclude AMS-owned key aliases.	This finding is generated if the policy contains permissions that are not scoped to specific KMS keys alias that you own. To resolve this, scope the permission to specific keys or exclude the keys that are AMS owned. AMS owned keys have specific alias sets. For a list of AMS owned key aliases, see AMS Automated IAM Provisioning permission boundary check.

Finding	Description
Statement contains privileged actions that are not adequately scoped to your KMS keys using the `kms:ResourceAliases condition` . Consider using specific alias names along with the appropriate set operator for the condition key. If wildcards are used in the alias names ensure they restrict the scope to a limited set of your KMS keys.	This finding is generated if you are scoping permissions to your KMS keys using condition s and not using `kms:ResourceAliases` to scope down to aliases for your KMS keys. Or, if the `kms:ResourceAliases` condition key has a value that also includes AMS owned KMS keys aliases. To resolve this, update the condition to scope down permission only to aliases of your KMS keys or exclude aliases for AMS owned KMS keys. For a list of AMS owned key aliases, see AMS Automated IAM Provision ing permission boundary check.
The role must have customer_deny_policy attached. Include the policy ARN in the list of managed policy ARNs.	This finding is generated if the role that you are creating does not have the `customer_ deny_policy` attached to it. To resolve this, include the `customer_deny_policy` in the managed policy ARNs list.
The AWS managed policy is overly permissiv e or grants permissions restricted by AMS boundary policy.	This finding is generated if the **ManagedPo licyArns** value for the role contains any AMS managed policy that provides full or administr ator level access to the relevant service. To resolve this, review use of the AWS managed policy and use a policy that provides scope down permission or define your own policy that follows the principle of least privilege.
The customer managed policy is in restricted AMS namespace.	This finding is generated if any customer managed policy with name prefixed in the AWS namespace is attached to the role. To resolve this, remove the policy from the **ManagedPolicyArn** list for the role.

Finding	Description
The customer_deny_policy can not be detached from the role. Include the policy ARN in the list of managed policy ARNs.	This finding is generated if the `customer_deny_policy` is detached from the role during an update. To resolve this, add the `customer_deny_policy` to the **ManagedPolicyArns** field of the role and try again.
The customer managed policies were provisioned outside AMS Change Management service or without prior validation.	This finding is generated if one or more existing customer managed policy ARNs are attached to a role and the policies are not provisioned through the AMS Change Management service (through an RFC). For example, Developer Mode or Direct Change Mode allow customers to provision IAM policies without an RFC. To resolve this, remove the customer managed policy ARNs from the **ManagedPolicyArns** list for the role.
The count of provided managed policy ARNs exceed attached policy per role quota.	This finding is generated if the total number of managed policies attached to the role exceeds the policy per role quota. For more information on IAM quotas, see IAM and AWS STS quotas, name requirements, and character limits. Use this information to reduce the number of policies that you attach to the role.
The trust policy size ({trust_policy}) exceeds assume role policy size quota of {size}.	This finding is generated if the size of the assume role policy document exceeds the policy size quota. For more information on IAM quotas, see IAM and AWS STS quotas, name requirements, and character limits.

Finding	Description
Statement contains all mutative actions for Amazon S3. Consider scoping these permissio ns to required actions only. If wild cards are used ensure they scope limited set of mutative actions.	This finding is generated if the given policy grants all Amazon Simple Storage Service mutative permissions irrespective of one or more resources. To resolve this, include only required Amazon S3 mutative actions against your buckets.
Statement contains privileged actions that are not allowed against any bucket in Amazon S3. Consider adding a statement denying these actions.	This finding is generated if the policy grants privileged actions on any bucket. For a list of privileged actions, see AMS Automated IAM Provisioning permission boundary check To resolve this finding, remove, or deny these actions in your policy.
Statement contains privileged actions that are not scoped to your buckets in Amazon S3. Consider including your buckets or exclude buckets with AMS namespace prefixes. If wild cards are used, make sure that they match buckets within your namespaces.	This finding is generated if the policy grants Amazon S3 actions that are not scoped to your buckets only. This is often occurs if wild cards are used when specifying bucket resources. To resolve this, specify bucket names or ARNs that you own or exclude the buckets that have AMS namespace prefixes.
Statement contains privileged actions that are not scoped to your buckets in Amazon S3. Consider avoiding use of wild cards (*) that scopes all buckets in the account.	This finding is generated if the policy grants Amazon S3 actions that are not scoped to your bucket. This is often occurs if wild cards are used when specifying bucket resources. To resolve this, specify bucket names or ARNs that you own or exclude the buckets that have AMS namespace prefixes.

Finding	Description
Statement contains a resource wildcard which is scoped to all Amazon S3 buckets, including non-existent buckets and buckets you do not own . Consider scoping the permissions using a condition and `s3:ResourceAccount` condition key.	This finding is generated if the policy grants permission to buckets specified using wild cards. Use of wild cards often brings non-existing or non-owner buckets in scope. To resolve this, use condition and the `aws:ResourceAccount` condition key to scope the permission to buckets within the current account only. For more details, see Limit access to Amazon S3 buckets owned by specific AWS accounts.
Statement contains a `NotResource` policy element, which may be scoped to a large number of buckets, including non-existent buckets and buckets you do not own. Consider scoping the permissions using a condition and `s3:ResourceAccount` condition key.	This finding is generated if the policy utilizes the `NotResources` policy element to specify bucket resources. The use of the `NotResource` element might scope a large number of buckets, including non-existent or non-owner buckets. To resolve this, use conditions and the `aws:ResourceAccount` condition key to scope the permission to buckets only within the current account.
Statement contains Amazon S3 action to buckets *`Bucket_Name`* that either do not exist, are not owned by account *`Account_ID`* , or name contains a wild card that might be scoped to a large number of buckets, including non-existent buckets and buckets you do not own. Consider scoping the permissions using a condition and the `s3:ResourceAccount` condition key	This finding is generated if the policy grants permission to buckets that either do no exist, are not owned by you, or have wild cards in the bucket names covering a large number of buckets and access is not scoped to the current account only. To resolve this, use condition and the `aws:ResourceAccount` condition key to scope the permission to buckets within the current account only.

Finding	Description
Statement contains Amazon S3 action to buckets *Bucket_Name* that either do not exist, not owned by account *Account_ID* , or the name contains a wild card that might be scoped to a large number of buckets, including non-existent buckets and buckets you do not own. Access is not restricted using `s3:ResourceAccount` or specified resource account in the condition does not belong to you.	This finding is generated if the policy grants permission to buckets that either do no exist, are not owned by you, or have wild cards in the bucket names covering a large number of buckets and access is scoped to a specific account only. However, the account specified in the `aws:ResourceAccount` condition key does not belong to you and is managed by AMS. To resolve this, update the `aws:ResourceAccount` condition key and set the appropriate account ID that you own and is managed by AMS.
Statement contains privileged actions that are not scoped to your instances for Amazon EC2. Consider scoping the actions to specific instance ARNs or exclude instances that have Name tag key with value in AMS namespace prefixes. If wild cards are used, ensure they match namespaces that you own.	This finding is generated if the policy grants privileged actions against Amazon EC2 instances that AMS owns. AMS instances are tagged with the **Name** tag key with values in AMS namespace. To resolve this, specify your resources or exclude AMS instances with a condition that has the `aws:ResourceTag/Name` key that excludes values in the AMS namespace using the `StringNotLike` operator
Statement contains privileged actions that are not scoped to your resources in AWS Systems Manager parameter store. Consider specifying ARNs of your parameters or exclude parameters with AMS namespace prefixes. If wild cards are used, ensure they scope only your parameters.	This finding is generated if the policy grants permissions to parameters that you do not own. This is usually when wild cards are used or parameters with AMS namespace prefixes are listed under resources in a policy statement. To resolve this, specify parameters that are within your namespace or exclude AMS parameters with a deny statement.

Finding	Description
Statement contains privileged actions against resources in AWS Systems Manager. Consider scoping the permissions to read only actions or actions against your resources.	This finding is generated if the policy grants permissions other than parameter store or readonly actions against Systems Manager resources. To resolve this finding reduce the permissions to readonly actions or parameter store only.
Statement contains privileged actions that are not scoped to {message} in *Service_Name* that you own. Consider scoping these permissions to specific resource types as appropriate or exclude AMS owned resources . If wild cards are used, ensure they match *Resources* .	This finding is generated if the policy allows privileged actions that are not granted against your resources, especially for named resources . To resolve this finding review your resource list and see if they only scope resource that is in your namespace. Alternatively exclude resources that are in AMS namespace.
Statement contains tagging actions of {*Service_Name* } that are not scoped to specific values for Name tag key. Consider scoping these actions by setting `aws:RequestTag/Name` condition key with values in your namespace or restrict these actions by setting `aws:RequestTag/Name` condition key with the `StringNotLike` operator with values in the AMS namespace prefixes.	This finding is generated if the policy grants tagging permission for given service and the permission is not scoped to specifc tag keys/values. To scope down what key or value can be used in tag actions, for example, when making request to perform the actions, use the `aws:RequestTag/tag key` condition . So, to resolve this, use this condition key to restrict key or values in your name space. Or, deny the `Name` tag key (`aws:RequestTag/Name`) with values in AMS namespace.
Internal error validating IAM role trust policy.	This finding is generated when CT automation encounters an error performing validation on the IAM role trust policy through the IAM Access Analyzer service. To resolve this, resubmit the RFC. If the error persists, then contact AMS Operations to troubleshoot the error.

Finding	Description
Internal error validating customer managed policy.	This finding is generated when CT automatio n encounters an error performing ovalidation on the customer managed policy through the IAM Access Analyzer service. To resolve this, resubmit the RFC. If the error persists, then contact AMS Operations to troubleshoot the error.
Access analyzer not found in *AWS Region*. Unable to perform access preview check for role trust policy.	This finding is generated when the IAM Access Analyzer resource is not found in the AWS Region. Contact AMS Operations to troublesh oot and create IAM Access Analyzer resource in the AWS Region.
Invalid trust policy for role *Role_Name*	This finding is generated when provided IAM role contains an invalid trust policy. To resolve review the trust policy to verify that it is valid.
IAM Access Analyzer encountered an internal error. Failed to create access preview for role *Role_Name*	This finding is generated when automation encounters an error while creating an access preview for a role through the IAM Access Analyzer. To resolve this, resubmit the RFC. If the error persists, then contact AMS Operation s to troubleshoot the error.
Failed to create access preview for trust policy of role *Role_Name*	This finding is generated when automation encounters an error while creating an access preview for a role through the IAM Access Analyzer. To resolve this, resubmit the RFC. If the error persists, then contact AMS Operation s to troubleshoot the error.

Finding	Description
Internal error validating listed SAML IdP.	This finding is generated when automatio n encounters an error while validating the provided SAML IdPs listed in the role trust policy. To resolve this, resubmit the RFC. If the error persists, then contact AMS Operations to troubleshoot the error.
Internal error validating permissions against AWS Key Management Service.	This finding is generated when automation encounters an error while validating the AWS KMS key permissions in the provided policy. To resolve this, resubmit the RFC. If the error persists, then contact AMS Operations to troubleshoot the error.
Internal error validating listed managed policy ARNs.	This finding is generated when automatio n encounters an error while validating listed managed policy ARNs. To resolve this, resubmit the RFC. If the error persists, then contact AMS Operations to troubleshoot the error.
Internal error validating default `customer_ deny_policy` attachment.	This finding is generated when automation encounters an error while validating that the `customer_deny_policy` is attached to the role. To resolve this, resubmit the RFC. If the error persists, then contact AMS Operation s to troubleshoot the error.
Internal error validating managed policy arns for the role *`Role_Name`*	This finding is generated when automation encounters an error while validating managed policy ARNs for the role. To resolve this, resubmit the RFC. If the error persists, then contact AMS Operations to troubleshoot the error.

Finding	Description
Internal error validating *Policy_name* against customer-defined boundary policy AWSManagedServicesIAMProvisionCustomerBoundaryPolicy	This finding is generated when automation encounters an error while validating the policy that cotains your custom deny list. To resolve this, resubmit the RFC. If the error persists, then contact AMS Operations to troubleshoot the error.
Customer-defined boundary policy AWSManagedServicesIAMProvisionCustomerBoundaryPolicy exists in the account. However, the policy contains allow statements that grant permissions. The policy must only contain deny statements.	This finding is generated when the policy that contains your custom deny list includes a statement that grants permission. Although the custom deny list exists within your account as an IAM managed policy, it can't be used for permission management. The policy must only contain deny statements that indicate that you want AMS Automated IAM Provisioning to validate and deny those actions in your IAM policies that AMS Automated IAM Provisioning creates.
Statement contains privileged actions defined by your organization for *Service_Name* . Consider excluding these actions with a deny statement. Refer to the policy named in your account for reference to the restricted list of actions.	This finding is generated when automatio n detects any action in your policy that you defined in the custom deny list. To resolve the finding, review your policy statement and remove any actions that are defined in your custom deny list or add a deny statement that denies those actions.

AMS Automated IAM Provisioning permission boundary check

AMS permission boundary checks help you adhere to the default permission boundary policy provided by AMS. This policy is a list of actions denied by AMS Automated IAM Provisioning. Provisioning policies that contain these restricted actions require additional explicit risk acceptance. Download the policy here: boundary-policy.zip.

Use customer-defined permission boundary policy checks to customize deny actions beyond the AMS permission boundary policy defaults. When you onboard to AMS Automated IAM Provisioning using the following change type: Management | Managed account | AMS Automated IAM Provisioning with read-write permissions | Enable (review required) (ct-1706xvvk6j9hf), you can include a list of custom deny actions that specify additional restricted actions.

You can update the list of deny actions using the change type: Management | Managed account | Automated IAM provisioning with read-write permissions | Update custom deny list (ct-2r9xvd3sdsic0). You must use the dedicated IAM role `AWSManagedServicesIAMProvisionAdminRole` to run this change type.

Note

- You must provide a comprehensive list of deny actions for each update. The previous list is replaced by the new list.
- The list of deny actions must contain only actions to be denied. Allow actions aren't supported.
- The list of deny actions resides within the account as an IAM managed policy named `AWSManagedServicesIAMProvisionCustomerBoundaryPolicy`. The policy must not be attached to any role.
- The term *permission boundary* used to denote denied actions in AMS Automated IAM Provisioning has a different contextual meaning compared to the IAM permission boundary. The IAM permission boundary sets the maximum permission that a policy can grant at runtime to an IAM entity. For more information on IAM permission boundary see Policy types in the *AWS Identity and Access Management User Guide*. The permission boundary in AMS Automated IAM Provisioning prevents you from provisioning an IAM policy that contains a certain set of permissions, for example, a denied list of actions.

Troubleshooting AMS Automated IAM Provisioning fndings and errors

There are three ways you might run into problems when using AMS Automated IAM Provisioning:

- RFC errors: These can happen for a variety of reasons; for example, incorrect input. For more information, see Troubleshooting RFC errors in AMS.
- SSM errors: These can happen for a variety of reasons; for example, poor formatting. For more information, see Troubleshooting Systems Manager Automation.

- Validation check findings: These occur when one of the many validation checks that Automated IAM Provisioning runs finds a problem. For a list of validation checks, and recommended actions to fix, see Runtime checks for AMS Automated IAM Provisioning in AMS.

Setting permissions in AMS with IAM roles and profiles

AMS uses AWS Identity and Access Management (IAM) to manage users, security credentials such as access keys, and permissions that control which AWS resources users and applications can access. AMS provides a default IAM user role and a default Amazon EC2 instance profile (which includes a statement allowing the resource access to the default IAM user role).

Requesting a new IAM user role or instance profile

AMS uses an IAM role to set user permissions through your federation service and an IAM instance profile as a container for that IAM role.

You can request, with an AMS service request, or a Management | Other | Other | Create CT, a custom IAM role or instance profile. See the descriptions of each in this section.

> **Note**
>
> AMS has an IAM policy, `customer_deny_policy` that blocks out dangerous namespaces and actions. This policy is attached to all AMS customer roles by default and is rarely a problem for users. Your IAM user and role requests don't include this policy, but automatic inclusion of the `customer_deny_policy` in requests for IAM roles helps AMS deploy new IAM instance profiles more quickly. You can request the exclusion of the `customer_deny_policy` policy. However, this request will go through a weighty security review and is likely to be declined due to security reasons.

Restrict permissions with IAM role policy statements

AMS uses an IAM role to set user permissions through your federation service.

Single-Account Landing Zone AMS: See SALZ: Default IAM User Roles.

Multi-Account Landing Zone AMS: See MALZ: Default IAM User Roles.

An IAM role is an IAM entity that defines a set of permissions for making AWS service requests. IAM roles are not associated with a specific user or group. Instead, trusted entities assume roles, such as IAM users, applications, or AWS services such as Amazon EC2. For more information, see IAM Roles.

You can scope down the desired policy for a user assuming the AMS IAM user role by using the AWS Security Token Service (STS) API operation AssumeRole by passing a more restrictive IAM policy under the `Policy` request field.

Example policy statements that you can use to restrict CT access are provided next.

Using your configured Active Directory (AD) groups, and the AWS Security Token Service (STS) API operation AssumeRole, you can set permissions for certain users or groups, including restricting access to certain change types (CTs). You can use the policy statements shown below to restrict CT access in various ways.

AMS change type statement in the default IAM instance profile that allows access to all AMS API calls (amscm and amsskms) and all change types:

```
{
    "Sid": "AWSManagedServicesFullAccess",
    "Effect": "Allow",
    "Action": [
        "amscm:*",
        "amsskms:*"
    ],
    "Resource": [
        "*"
    ]
}
```

1. Statement to allow access and all actions for only two specified CTs, where "Action" is the AMS API operations (either `amscm` or `amsskms`), and "Resource" represents existing change type IDs and version number:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action":amscm:*,
      "Resource": [ "arn:aws:amscm:*:*:changetype/ct-ID1:1.0",
```

```
                    "arn:aws:amscm:*:*:changetype/ct-ID2:1.0"]
    }
  ]
}
```

2. Statement to allow access for CreateRfc, UpdateRfc, and SubmitRfc on only two specified CTs:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action":["amscm:CreateRfc",
                "amscm:UpdateRfc",
                "amscm:SubmitRfc"],
      "Resource": ["arn:aws:amscm:*:*:changetype/ct-ID1:1.0",
                   "arn:aws:amscm:*:*:changetype/ct-ID2:1.0"]
    }
  ]
}
```

3. Statement to allow access for CreateRfc, UpdateRfc, and SubmitRfc on all available CTs:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action":["amscm:CreateRfc",
                "amscm:UpdateRfc",
                "amscm:SubmitRfc"],
      "Resource":"*"
    }
  ]
}
```

4. Statement to deny access for all actions on restricted CT and allow on other CTs:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
```

```
      "Action":amscm:*"
      "Resource":"arn:aws:amscm:*:*:changetype/ct-RetrictedID:1.0"
    },
    {
      "Effect": "Allow",
      "Action":amscm:*"
      "Resource":"*"
    }
  ]
}
```

Restrict permissions with Amazon EC2 IAM instance profiles

An IAM instance profile is a container for an IAM role that you can use to pass role information to an Amazon EC2 instance when the instance starts.

Currently there is one AWS Managed Services (AMS) default instance profile, `customer-mc-ec2-instance-profile`, that grants permissions to the applications running on the instance, not to users logging into the instance. You might want to modify the default instance profile, or create a new one, if you want to give an instance access to something, without granting other instances access as well. You can request a new IAM instance profile with the Management | Applications | IAM instance profile | Create change type (ct-0ixp4ch2tiu04). When submitting the RFC, you could fashion your own instance profile and include that as the InstanceProfileDescription, or you could just inform AMS (using the same field) of what changes you want. Because this is a Manual CT, AMS must approve the change and will be in contact with you about it.

If you're unfamiliar with Amazon IAM policies, see Overview of IAM Policies for important information. There is also a good blog post, Demystifying Amazon EC2 Resource-Level Permissions. Note that AMS does not currently support Resource-based access control, but does support Resource-level controls using IAM role policies (for an explanation of the difference, see AWS Services That Work with IAM.

Single-Account Landing Zone AMS:

To see a table of permissions that the default AMS IAM instance profile grants, go to EC2 IAM Instance Profile.

AD FS claim rule and SAML settings

ActiveDirectory Federation Services (AD FS) claim rule and SAML settings for AWS Managed Services (AMS)

For detailed step-by-step instructions on how to install and configure AD FS see Enabling Federation to AWS Using Windows Active Directory, ADFS, and SAML 2.0.

ADFS claim rule configurations

If you already have an ADFS implementation, configure following:

- Relying party trust
- Claims rules

The relying party trust and claims rules steps are taken from Enabling Federation to AWS Using Windows Active Directory, AD FS, and SAML 2.0blog

- Claims rules:
  - **Nameid**: Configuration per blog post
  - **RoleSessionName**: Configure as follows
    - **Claim rule name**: **`RoleSessionName`**
    - **Attribute store**: **`Active Directory`**
    - **LDAP Attribute**: **`SAM-Account-Name`**
    - **Outgoing Claim Type**: **`https://aws.amazon.com/SAML/Attributes/RoleSessionName`**
    - **Get AD Groups**: Configuration per blog post
    - **Role claim**: Configure as follows

```
c:[Type == "http://temp/variable", Value =~ "(?i)^AWS-([^d]{12})-"]
```

```
=> issue(Type = "https://aws.amazon.com/SAML/Attributes/Role", Value =
 RegExReplace(c.Value, "AWS-([^d]{12})-", "arn:aws:iam::$1:saml-provider/
customer-readonly-saml,arn:aws:iam::$1:role/"));
```

Web console

You can access the AWS Web console by using the link below replacing *[ADFS-FQDN]* with the FQDN of your ADFS implementation.

https://*[ADFS-FQDN]*/adfs/ls/IdpInitiatedSignOn.aspx

Your IT department can deploy the above link to the user population via a Group Policy.

API and CLI access with SAML

How to configure API and CLI access with SAML.

The python packages are sourced from the blog posts below:

- NTLM: How to Implement Federated API and CLI Access Using SAML 2.0 and AD FS
- Forms: How to Implement a General Solution for Federated API/CLI Access Using SAML 2.0
- PowerShell: How to Set Up Federated API Access to AWS by Using Windows PowerShell

Script configuration

1. Using Notepad++, change the default region to the correct region
2. Using Notepad++, disable SSL verification for test and dev environments
3. Using Notepad++, configure idpentryurl

   ```
   https://[ADFS-FDQN]/adfs/ls/IdpInitiatedSignOn.aspx?
   loginToRp=urn:amazon:webservices
   ```

Windows configuration

The instructions below are for the python packages. The credentials generated will be valid for 1 hour.

1. Download and install python (2.7.11)
2. Download and install AWS CLI tools
3. Install the AMS CLI:

a. Download the AMS distributables zip file provided by your cloud service delivery manager (CSDM) and unzip.

   Several directories and files are made available.

b. Open either the **Managed Cloud Distributables -> CLI -> Windows** or the **Managed Cloud Distributables -> CLI -> Linux / MacOS** directory, depending on your operating system, and:

   For **Windows**, execute the appropriate installer (this method only works on Windows 32 or 64 bits systems):

   - 32 Bits: ManagedCloudAPI_x86.msi
   - 64 Bits: ManagedCloudAPI_x64.msi

   For **Mac/Linux**, execute the file named: **MC_CLI.sh**. You can do this by running this command: `sh MC_CLI.sh`. Note that the **amscm** and **amsskms** directories and their contents must be in the same directory as the **MC_CLI.sh** file.

c. If your corporate credentials are used via federation with AWS (the AMS default configuration) you must install a credential management tool that can access your federation service. For example, you can use this AWS Security Blog How to Implement Federated API and CLI Access Using SAML 2.0 and AD FS for help configuring your credential management tooling.

d. After the installation, run `aws amscm help` and `aws amsskms help` to see commands and options.

4. Download the required SAML script

   Download to c:\aws\scripts

5. Download PIP

   Download to c:\aws\downloads

6. Using PowerShell, install PIP

   <pythondir>.\python.exe c:\aws\downloads\get-pip.py

7. Using PowerShell, install boto module

   <pythondir\scripts>pip install boto

8. Using PowerShell, install requests module

   <pythondir\scripts>pip install requests

9. Using PowerShell, install requests security module

   <pythondir\scripts>pip install requests[security]

10. Using PowerShell, install beautifulsoup module

    <pythondir\scripts>pip install beautifulsoup4

11. Using PowerShell, create a folder called .aws in the users profile (%userprofile%\.aws)

    mkdir .aws

12. Using PowerShell, create a credential file in the .aws folder

    New-Item credentials -type file –force

    The credentials file mustn't have a file extension

    The filename must be all lowercase and have the name credentials

13. Open the credentials file with notepad and paste in the following data, specifying the correct region

    ```
    [default]
    output = json
    region = us-east-1
    aws_access_key_id =
    aws_secret_access_key =
    ```

14. Using PowerShell, the SAML script and logon

    <pythondir>.\python.exe c:\aws\scripts\samlapi.py

    Username: [USERNAME]@upn

    Choose the role you would like to assume

Linux configuration

The credentials generated will be valid for 1 hour.

1. Using WinSCP, transfer the SAML script
2. Using WinSCP, transfer the Root CA certificate (ignore for test and dev)
3. Add the ROOT CA to the trusted root certificates (ignore for test and dev)

   $ openssl x509 -inform der -in [certname].cer -out certificate.pem (ignore for test and dev)

   Add contents of certificate.pem to end of /etc/ssl/certs/ca-bundle.crt file ((ignore for test dev)
4. Create .aws folder in home/ec2-user 5

   ```
   [default]
   output = json
   region = us-east-1
   aws_access_key_id =
   aws_secret_access_key =
   ```
5. Using WinSCP, transfer the credentials file to .aws folder
6. Install boto module

   $ sudo pip install boto
7. Install requests module

   $ sudo pip install requests
8. Install beautifulsoup module

   $ sudo pip install beautifulsoup4
9. Copy the script to home/ec2-user

   Set the required permissions

   Execute the script: samlapi.py

Restrict with network ACL

A network access control list (NACL) is an optional layer of security for your VPC that acts as a firewall for controlling traffic in and out of one or more subnets. You might set up network ACLs with rules similar to your security groups in order to add an additional layer of security to your

VPC. For more information about the differences between security groups and network ACLs, see Comparison of security groups and network ACLs.

However, in AMS Managed Multi-Account Landing Zone, in order for AMS to effectively manage and monitor infrastructure, the use of NACLs is limited to following scope:

- NACLs are not supported in the Multi-Account Landing Zone Core accounts, i.e. Management account, Networking, Shared-Services, Logging and Security.
- NACLs are supported in Multi-Account Landing Zone Application accounts as long as they are only used as a "Deny" list and have "Allow All" to allow AMS monitoring and management operations.

In large scale multi-account environments, you can also leverage features like centralized egress firewalls to control outbound traffic and/or AWS Transit Gateway routing tables in AMS Multi-Account Landing Zone to segregate network traffic among VPCs.

AMS on Outposts

AWS Outposts is a managed hardware solution that extends AMS managed landing zones to customer data centers. With AMS support on AWS Outposts, customers seeking the cloud expertise, cost savings and standardized platform offered by AMS, are no longer limited to hosting resources inside AWS Regions. With AMS on AWS Outposts, customers with on-premise requirements can now modernize on AWS, while enjoying the patching, backup, provisioning, incident management, business continuity, and cost optimization services offered by AMS.

Once an AWS Outposts is activated in your AMS Multi-Account Landing Zone or Single-Account Landing Zone account, you can follow existing AMS change management processes to provision and manage AWS resources. AMS-hosted infrastructure can be managed by specifying your AWS Outposts-specific subnet. AWS Outposts lifecycles can be managed directly in the AWS Outposts console using the AWS Outposts self-provision services role.

For information on the role, see AWS Outposts.

AWS Outposts installation and operational management

The onboarding to AMS on AWS Outposts process is comprised of:

1. Outposts Planning
2. Order Validation

3. Outposts Onboarding to AMS
4. Lifecycle Management

AWS Outposts planning

During AWS Outposts planning, you identify AMS on AWS Outposts use cases and engage key stakeholders, including your AMS account team and AWS Outposts representatives, to align on capacity strategy.

1. Once use cases requiring AMS on AWS Outposts have been identified, engage your AMS account team to discuss capacity planning.
2. Once your AWS Outposts capacity requirements have been determined, your AMS account team engages the AWS Outposts service team to discuss AWS Outposts onboarding plan, roles and responsibilities. During this time an AWS Outposts single point of contact (SPOC) is assigned to you. The AWS Outposts SPOC assists in finalizing AWS Outposts sizing requirements.

AWS Outposts order validation

During order validation, you create an AWS Outposts site, and order your required capacity directly in the AWS Outposts console or through your AWS Outposts account representative.

Once you, the AMS account team, the AWS Outposts team are aligned, you can request the AWS Outposts self-provisioned service role using change type ID ct-3qe6io8t6jtny, to create your site and AWS Outposts order directly in the AWS Outposts console.

Alternatively, you can work through the AWS Outposts SPOC to create Outpost Sites and orders. Your AWS Outposts SPOC remains to provide status updates to you and the AMS account team during site and order validation, and AWS Outposts installation.

AWS Outposts onboarding to AMS

Once your AWS Outposts unit is activated in your AMS managed VPC, you can request that monitors be created to track availability, capacity, exceptions and network connectivity for your Outposts hardware. By following the monitoring deployment steps described next, your AWS Outposts hardware is actively monitored by AMS.

1. Once your AWS Outposts has been installed and activated, you can request AWS Outposts-specific monitoring by submitting the following template with an RFC using the Management |

Other | Other | Create (ct-1e1xtak34nx76) change type. AMS operations ensures that the AWS Outposts subnet is tracked in AMS internal tooling.

- AWS Outposts ID
- Subnet CIDR
- Recommended AWS Outposts alarms:
  - InstanceFamilyCapacityAvailabilityAlert
  - InstanceTypeCapacityAvailabilityAlert
  - EBSVolumeTypeCapacityAvailabilityAlert
  - CapacityExceptionsAlert
  - Direct Connect ConnectionAlert
- For each of the above alerts, specify the following parameters:
  - Statistic ("Average" is recommended. Other options include sum, maximum, minimum, sample count and p90)
  - Period ("5 minutes" is recommended. Other options include 10 and 30 seconds, 1, 5, and 15 minutes, 1 and 6 hours, and 1 day)
  - Threshold type ("Static" is recommended. "Anomaly" are also options.)
  - Condition ("Whenever call count is greater than", "equal to", "less than" are also options.)
  - Condition Value ("25%" is configured by default. Another other positive integer is allowable.)
  - Notification topic (AMS operations topics are automatically assigned. However any other, or custom, topic can also be added.)

2. Monitoring and operations Support
   - AMS operations monitors AWS Outposts metrics for network disconnection or component failures. AMS operations provides first response services for AWS Outposts issues, and escalates, if needed, to Premium support or EC2 support.
   - AMS operations is available to address issues related to your AWS Outposts unit.
3. When EC2 instance status or system status checks fail, AMS operations follows existing processes to bring the instance back online. If the restart fails or AWS Outposts capacity is insufficient, then an AMS operations team member notifies you directly to determine next steps.

AWS Outposts life cycle management

Once AWS Outposts has been onboarded to your AMS account, you are notified if any availability, capacity, or network exceptions, occur. You can decommission AWS Outposts directly through the AWS Outposts console or the AWS Outposts SPOC.

You can manage AWS Outposts directly in the AWS Outposts console using the AWS Outposts self-service provisioning service role or developer mode. You can also request AWS Outposts through your CSDM, or AWS Outposts single point of contact, (SPOC).

High-availability on AWS Outposts can be achieved by deploying two or more AWS Outposts. Configuring two or more AWS Outposts enables the multi-availability zone option for your Amazon Relational Database Service instances.

Provisioning AMS managed resources on AWS Outposts

Provisioning AWS resources hosted on AWS Outposts (for example, Amazon EC2, Amazon EMR, Amazon EKS, Amazon ECS, Amazon EBS, and Amazon S3) in AMS accounts (Single-Account Landing Zone, Multi-Account Landing Zone, and Accelerate accounts) are subject to the same AMS support levels as resources in AWS Regions. You can use AMS change management, self-service provisioning services, or developer mode to create and modify the resources created on AWS Outposts.

Currently, all instance types (M5/M5d, C5/C5d, R5/R5d, I3en, G4dn), Amazon Elastic Block Store, Amazon Elastic Container Service, Amazon Elastic Kubernetes Service, Amazon EMR, Amazon Relational Database Service DBs, Application Load Balancers, and App Mesh Envoy proxy are available directly on AWS Outposts. These resources are eligible for the same AMS operations support as resources in existing regions.

Limitations of AMS on AWS Outposts

- Operational support for AWS Outposts-hosted resources is dependent on consistent network connectivity. AWS Outposts network disconnection prevents AMS operations from being able to troubleshoot any incidents or problems that occur on the disconnected AWS Outposts unit. For AMS on AWS Outposts service level contingencies, see the updated AWS Service Level Agreements (SLAs).
- Amazon Relational Database Service:
  - The create RDS change type (ct-2z60dyvto9g6c), by default, enables multi-AZ and requires a DB subnet group. DB subnet groups require two subnets in two separate Availability

Zones (AZ). If you have only one AWS Outposts, creating a DB subnet group is an issue since AWS Outposts are only assigned to a single AZ. To circumvent this limitation, follow these instructions:

1. Request a DB subnet group through an RFC with a Management | Other | Other CT, and specify the subnet on the AWS Outposts.
2. Create a custom CFN template to deploy RDS on AWS Outposts, and specify the subnet group created in the previous step. To learn more about doing this, see Custom resources.
3. Request that AMS deploy the CFN template containing the target RDS instance through the AMS CFN ingest CT (ct-36cn2avfrrj9v).
4. Note that currently, RDS for AWS Outposts does not provide metrics and logs due to a limitation of RDS Service.

- Workload ingest (WIGs): Linux WIGs only works if the pre-WIGs EC2 instance is on a non-AWS Outposts subnet. The reason is because Linux WIGs creates a WIGs node in the subnet of the first EC2 instance using m4.large, by default. As AWS Outposts doesn't have that instance type, WIGs is not able to launch its worker node. The workaround for this is to create the initial EC2 instance in a non-AWS Outposts subnet, then the target instance can be created on AWS Outposts. Moreover, currently, only Nitro-based EC2 instance types including C5, C5d, M5, M5d, R5, R5d, G4, and I3en are supported on AWS Outposts.
- Amazon Elastic Block Store (EBS): Create EBS Volume CT (ct-16xg8qguovg2w) does not work, as volumes get created in AWS instead of AWS Outposts as we do not provide the AWS Outposts Amazon resource number (ARN) as an input parameter to the CT.
- Network connectivity: Network connectivity is your responsibility per the AWS Outposts team.
- Brownfield and account takeover: AWS Outposts activated in non-AMS accounts cannot be transitioned into AMS, due to the nature of AWS Outposts billing and enterprise support requirements.

AMS on AWS Outposts compliance

AMS on AWS Outposts compliance attestation

AWS Outposts control plane has been attested to HIPAA eligible, PCI and ISO compliance. However, AMS on AWS Outposts control plane has not been attested for AWS Outposts. For this reason, customers are encouraged to pursue compliance attestation AMS on AWS Outposts environment.

For controlling resource creation on the Outpost unit, customers are encouraged to segregate developer access to the Outpost, to prevent excess developer access in standard AMS managed accounts.

AMS Managed Workloads requiring FedRAMP compliance

Foremost, AMS management accounts must first be assessed for regulatory compliance, since control plane data would flow out of the AWS Outposts to AMS management accounts.

If FedRAMP certification is required and the AMS account structure is compliant, then it is recommended that you either utilize a datacenter vendor that already has the required certification and owns all of the service link appliance (or already encrypts egress data).

Finally, additional data protection can be put in place by working with your account team to deploy an SCP that restricts data to the AWS Outposts and prevents the creation of any in-region resources in the account hosting the Outpost.

Impact on existing compliance for AMS accounts

An account utilizing AWS Outposts does not need to be retested for compliance as long as no regulated data is being consumed and the account is logically separated. AMS management accounts can manage non-regulated and regulated accounts as long as cross account authentication/authorization and ingress/ egress data flows are segregated between VPCs. Therefore, even though both the non-compliant Outpost account and existing compliant application accounts are in the same organization (including shared services, networking, logging, master, security AMS services), the compliance application account still retains compliance since data is logically separated.

AMS on AWS Outposts FAQs

Which use cases qualify for AMS support on AWS Outposts?

AMS on AWS Outposts can be leveraged by enterprises needing a proven cloud operating model have workloads requiring low latency (e.g., factory robot management and mainframe migration), edge computing (e.g., remote workstations and edge data streamlining), and large data transfer loads.

Why should I use this feature?

AMS provides monitoring of AWS Outposts hardware and first response to any AWS Outposts hardware issue. Moreover, the following support features for all managed resources hosted on AWS Outposts:

- Logging, Monitoring, Guardrails, and Event Management
- Continuity Management
- Security and Access Management
- Patch Management
- Change Management
- Automated and Self-Service Provisioning Management
- Incident and Problem Management
- Reporting (Reporting for AWS Outposts hardware will not be initially supported with AMS on AWS Outposts)
- Service Request Management
- Developer Mode
- Enterprise Support

How do I use this feature?

AWS Outposts planning: During AWS Outposts planning, you have identified AMS on AWS Outposts use cases and will engage key stakeholders, including the AMS account team and AWS Outposts representatives, to align on capacity strategy.

Order validation: During order validation, you create an AWS Outposts site, and order your required capacity directly in the AWS Outposts console, or through your AWS Outposts account representative.

AWS Outposts onboarding to AMS: Once your AWS Outposts unit is activated in your AMS managed VPC, you can request that your AWS Outposts be onboarded to your AMS account by submitting a request for change (RFC) using the template in the AMS User Guide (AWS Outposts). AMS operations then creates a subnet and monitors for your Outpost using the inputs provided on the RFC.

Lifecycle management: Once AWS Outposts has been onboarded to your AMS account, you are notified of any availability, capacity, or network exceptions. You can decommission AWS

Outposts directly through the AWS Outposts console or your AWS Outposts single point of contact (SPOC).

What are the limitations of AMS on AWS Outposts?

Data residency (e.g., country-specific data localization laws, etc.) use cases have not yet been validated for AMS on AWS Outposts.

AWS Outposts activated in non-AMS accounts cannot be transitioned into AMS, due to the nature of AWS Outposts billing and Enterprise Support requirements.

AWS Outposts control plane has been attested to HIPAA eligible, PCI and ISO compliance. However, AMS on AWS Outposts control plane has not been attested for AWS Outposts. For this reason, customers are encouraged to pursue compliance attestation AMS on AWS Outposts environment.

Can I opt out of this feature?

Provisioning AWS Outposts into your AMS environment is optional. Once deployed into your AMS account, AWS Outposts can be deprovisioned via the AWS Outposts console at any time, if no longer needed.

How will AMS on AWS Outposts be billed?

AMS uplift on AWS Outposts charges will be applied at the Group B tier.

How will the AMS Service Level Agreement change to accommodate AWS Outposts?

Incident management will be contingent on AWS Outposts availability. AWS Outposts availability is subject to customer network availability, which is the responsibility of the customer. AWS Outposts availability is also subject to AWS Outposts hardware uptime, which is dependent on AWS Outposts Service Level Agreements.

See also AWS Outposts FAQs.

Using tags in AMS

Topics

- AMS infrastructure automatic tagging
- AMS recommended tags

- Tag bulk update notes

Providing tags can be of great value. For in-depth information, read Tagging Your Amazon EC2 Resources.

RFCs that create instances, such as an Amazon S3 bucket or an AWS Elastic Load Balancer (ELB), generally provide a schema that enables you to add up to seven tags (key/value pairs); you can add more tags to your S3 bucket by submitting a Management | Advanced stack components | Amazon S3 storage | Update CT. Amazon EC2, Amazon EFS, Amazon RDS, and the multi-tiered (HA Two-Tiered and HA One-Tiered) schemas allow up to fifty tags. Tags are specified in the ExecutionParameters part of the schema.

When using the AMS console, you must enable the **Additional configuration** view in order to add tags.

> **Tip**
>
> Many CT schemas have a `Description` and `Name` field near the top of the schema. Those fields are used to name the stack or stack component, they do not name the resource you are creating. Some schemas offer a parameter to name the resource you are creating, and some do not. For example, the CT schema for Create Amazon EC2 stack does not offer a parameter to name the Amazon EC2 instance. In order to do so, you must create a tag with the key "Name" and the value of what you want the name to be. If you do not create such a tag, your Amazon EC2 instance displays in the Amazon EC2 console without a name attribute.

Are there tag restrictions? Yes:

- Tags are case sensitive.
- The maximum key length is 128 Unicode characters.
- The maximum value length is 256 Unicode characters.
- The maximum number of tags per resource is 50.
- The reserved prefix is `aws:`.
- AWS-generated tag names and values are automatically assigned the `aws:` prefix, which you can't assign. User-defined tag names have the prefix `user:` in the cost allocation report.

- Use each key only once for each resource. If you attempt to use the same key twice on the same resource, your request is rejected.
- Allowed characters are Unicode letters, white space, numbers, and the following special characters: + - = . _ : /

Which AWS resource types support tags? See Now Organize Your AWS Resources by Using up to 50 Tags per Resource.

AMS infrastructure automatic tagging

AMS can tag all resources created by AMS for management purposes, in your multi-account landing zone (MALZ) and single-account landing zone (SALZ) accounts through a request for change (RFC) with the Management | Other | Other | Create change type (ct-1e1xtak34nx76). This can help you in identifying resources created by AMS for management purposes.

AMS can automatically identify AMS-created resources based on the naming standards and check if the resource has the following tag keys and values - "AppName", "AppId", "AMSResource", and "EnvironmentType". If the tag key does not exist, or the value is empty, those tag-keys can be created automatically by AMS with tag-value "AMSInfrastructure".

You can customize the tags you want on AMS-created resources based on your organization's tagging standards. You can include your own tag-keys and tag-values when you submit the request to AMS. Follow these AWS tag naming standards: Tagging Best Practices

> **Note**
>
> For MALZ accounts, custom tagging of AMS infrastructure is supported on Application accounts only. Custom tagging on core accounts is currently not supported.
> If the tag-key name you provide in your RFC, already exists on the resource, then the tag-value gets replaced with the new tag-value that you provided in the RFC.
> Total length of tag key:value pairs must not exceed 256 characters.

Include the following information in your RFC with the Management | Other | Other | Create change type (ct-1e1xtak34nx76) for tagging AMS-created resources.

1. List of multi-account landing zone or single-account landing zone accounts where you would like to tag AMS-created resources for management purposes.

2. Required tag-key name and tag-value (if needed). By default, AMS can tag with tag-key name as "EnvironmentType" and tag-value as "AMSInfrastructure". If you need a custom tag-key name and tag-value, follow AWS tag naming standards: Tagging Best Practices

These resources are currently supported by AMS infrastructure tagging:

```
API Gateway
Amazon CloudFront
Amazon DynamoDB
Amazon EBS
Amazon EC2
Amazon OpenSearch Service
Amazon Quantum Ledger Database (QLDB)
Amazon Redshift
Amazon RDS
Amazon S3 (specific buckets only*)
Amazon Simple Queue Service (SQS)
Amazon Simple Notification Service (SNS)
Amazon VPC
AWS Certificate Manager
AWS CloudFormation
AWS CloudTrail
AWS CodeBuild
AWS CodePipeline
AWS Elastic Beanstalk
AWS Lambda
AWS Secrets Manager
AWS Service Catalog
AWS Systems Manager
AWS WAF Regional
Elastic Load Balancing
```

* "arn:aws:s3:::awsms-a*-patch-data-*", "arn:aws:s3:::ams-a*-log-management-*", "arn:aws:s3:::cf-templates-*", "arn:aws:s3:::mc-a*", "arn:aws:s3:::ams-a*-backup-reports-*", "arn:aws:s3:::ams-a*-patch-data-customer-reports-*", "arn:aws:s3:::ams-a*-patch-data-raw-*", "arn:aws:s3:::ams-a*-patch-data-reporting-*", "arn:aws:s3:::ams-a*-release-assets-*", "arn:aws:s3:::ams-cfn-drift-remediation-*", "arn:aws:s3:::ams-reporting-data-a*"

AMS recommended tags

AMS recommends the following tags on supported resources. Starred (*) tags are highly recommended.

> **Note**
>
> You can use tags to schedule patching. For information, see AMS Advanced Patch Orchestrator: a tag-based patching model.

AMS reserved prefixes

AMS resource attributes must comply with certain patterns; for example, IAM instance profile names, BackupVault names, tag names, and so forth, must not start with AMS reserved prefixes. Those reserved prefixes are:

```
*/aws_reserved/*
ams-*
/ams/*
ams*
AMS*
Ams*
aws*
AWS*
AWS_*
AWSManagedServices*
codedeploy_service_role
CloudTrail*
Cloudtrail*
customer-mc-*
eps
EPSDB*
IAMPolicy*
INGEST*
LandingZone*
Managed_Services*
managementhost
mc*
MC*
Mc*
```

```
MMS*
ms-
NewAMS*
Root*
sentinel*
Sentinel*
sentinel.int.
StackSet-ams*
StackSet-AWS-Landing-Zone
StateMachine*
TemplateId*
VPC_*
UnhealthyInServiceBastion
```

AMS recommended tags

Tag key	Supported values	Notes
AppName*	Unconstrained.	Identify the applications that will reside on, or require access to, the resource. This facilitates tracking and communications between AMS and you.
AppId*		
EnvironmentType *		Distinguish between development, test, and production infrastructure as the environment for the resource.
OwnerTeamEmail *	Distribution list email address	Identify the distribution list email address for the team responsible for the resource. The email should not be a personal email; it must be an anonymous email like a distribution list.
ComplianceFramework	Unconstrained	Identify which controls and policies should be applied to the resource.

Tag key	Supported values	Notes
CostCenter		Identify the cost center or business unit associated with a resource (typically for cost allocation and tracking).
Customer		Used by AMS customers that have resources from multiple customers (AMS sub-customers). To group resources in the managed environme nt into the specific customer they are serving. Identify a specific client on a particular group of resources or services.
DataClassification		Identify the specific data-confidentiali ty level a resource supports. Identify which controls and policies should be applied to the resource.
HoursOfOperation		Identify the date or time a resource should be started, stopped, deleted, or rotated.
OwnerTeam		Identify the team responsible for the resource. Facilitates communication with the team responsible for the resource.

Tag key	Supported values	Notes
`Patch Group`	You have two options: • You can submit a service request with the informati on outlined in AMS Advanced Patch Orchestrator: a tag-based patching model, and AMS creates on your specified resources, and using the information you provide, a `Patch Group` tag for you. • If you have already created a `Patch Group` tag on your resources, the supported values are unconstrained.	What resources to include in an automated patching maintenance window.
`ProjectId`	Unconstrained	Identify the projects the resource supports.

Tag key	Supported values	Notes
SupportPriority	One of six possible acronyms for Confident iality, Integrity, or Availability; the order of the acronym is the priority. For example, I.C.A. would mean the order is Integrity, Confidentiality, Availabil ity. Other acceptable values are C.I.A., C.A.I., I.A.C., A.C.I., and A.I.C.	Identify which type of support should be priority: Confidentiality, Integrity, or Availability.

Tag bulk update notes

These notes are for use with the Tag | Bulk Update and Tag | Bulk Update (Review Required) change type example walkthroughs.

Notes on bulk update tags:

- Supported services:
  - For Tag | Bulk Update (Review Required): All.
  - For Tag | Bulk Update:
    - Auto Scaling
    - Amazon EC2
    - Elastic Load Balancing
    - Amazon RDS
    - Amazon S3 buckets
- To generate the comma-separated values (CSV) file, log in to **Resource Groups** > **Tag Editor** in the AWS console, find the resources you want to tag, and then export them to a CSV file. For more information, see Export Results to CSV.

  NOTE: We recommend using an S3 pre-signed URL for providing the CSV object location.

- There must be columns with the headers **Service**, **Type**, and **Identifier**. These columns are mandatory.
- Rename the column **Identifier** to **ID**.
- The values for columns **Service**, **Type**, and **Identifier** (changed to **ID**) are as provided by the **Tag Editor** under **Resource Groups**.
- If a resource to be tagged is not available under the **Tag Editor**, use the Amazon Resource Name (ARN) of the resource for the column **Identifier** (changed to **ID**) and keep **Service** and **Type** blank.
- The column headers of the tags to be added or modified must have the format **Tag: *TagKey***.
- The column headers of the tags to be deleted must have the format **Untag: *TagKey***.
- Each row is for a resource to be tagged or untagged.
- If a tag is to be added to a resource, specify the value of the tag under the appropriate **Tag: *TagKey*** column.
- If a tag is to be deleted from a resource, specify **True** under the appropriate **Untag: *TagKey*** column.
- If a **Tag: *TagKey*** or **Untag: *TagKey*** is not applicable for a resource, add a dash (-) to the respective column or keep it blank.

  NOTE: The Untag overrides the Tag - this is because, unlike Tag, Untag needs to be manually added to the CSV file, so we assume the intention was to remove, not modify, the tag.
- Any additional columns that don't fit the criteria are ignored.

> **Important**
>
> The Tag Editor export populates a matrix of all tags against all resources, missing tags are populated with a value of 'not tagged'. Re-using this export CSV as input to the RFC results in all the previously missing tags being created, with literal values of 'not tagged'.

	A	B	C	D	E	F	G	H	I	J
1	Name	Service	Type	Region	ID	Tag: Name	Tag: App	Tag: Backup	Untag: Version	Untag: App ID
2	S3 Bucket app-a-bucket	S3	Bucket	us-east-1	app-a-bucket	app-a-bucket	app-a		TRUE	TRUE
3	EC2 Image ami-1234567890abcdef1	EC2	Image	us-east-1	ami-1234567890abcdef1	app-a-2019-09-24	app-a		TRUE	TRUE
4	EC2 Instance i-1234567890abcdef1	EC2	Instance	us-east-1	i-1234567890abcdef1	app-a-1	app-a	TRUE		
5	EC2 SecurityGroup sg-1234567890abcdef1	EC2	SecurityGroup	us-east-1	sg-1234567890abcdef1	app-b-sg	app-b			
6	EC2 Snapshot snap-1234567890abcdef1	EC2	Snapshot	us-east-1	snap-1234567890abcdef1	app-a-sda1-2019-09-24	app-a			
7	EC2 Volume vol-1234567890abcdef1	EC2	Volume	us-east-1	vol-1234567890abcdef1	app-a-sda1	app-a		TRUE	TRUE
8	RDS DBInstance app-a-db	RDS	DBInstance	us-east-1	app-a-db	app-a-db	app-a			
9	RDS DBSnapshot app-a-db-snapshot-2019-09-24	RDS	DBSnapshot	us-east-1	app-a-db-snapshot-2019-09-24	app-a-db-snapshot-2019-09-24	app-a			
10					arn:aws:s3:::app-b-logs	app-b-logs	app-b			
11										

> **Note**
>
> For information about exporting tags to a CSV file, see Find Resources to Tag -> Export Results to CSV.

AWS Managed Services Resource Scheduler

Use AWS Managed Services (AMS) Resource Scheduler to schedule the automatic start and stop of AutoScaling groups, Amazon EC2 instances, and RDS instances in your account. This helps reduce infrastructure costs where the resources are not meant to be running 24/7. The solution is built on top of Instance Scheduler on AWS, but contains additional features and customizations specific to AMS needs.

> **Note**
>
> AMS Resource Scheduler doesn't interact with resources that aren't part of a AWS CloudFormation stack. The resource must be part of a stack that starts with "stack-" , "sc-" or "SC-".

AMS Resource Scheduler uses periods and schedules:

- *Periods* define the times when Resource Scheduler runs, such as start time, end time, and days of the month.
- *Schedules* contain your defined periods, along with additional configurations, such as SSM maintenance window, timezone, hibernate setting, and so forth; and specify when resources should run, given the configured period rules.

You can configure these periods and schedules using AMS Resource Scheduler's automated change types (CTs).

For full details on the settings available for AMS Resource Scheduler, see the corresponding AWS Instance Scheduler documentation at Solution components. For an architectural view of the solution, see the corresponding AWS Instance Scheduler documentation at Architecture overview.html.

Deploying AMS Resource Scheduler

To deploy AMS Resource Scheduler, use the automated change type (CT) : Deployment | AMS Resource Scheduler | Solution | Deploy (ct-0ywnhc8e5k9z5) to raise an RFC that then deploys the solution in your account. Once the RFC is executed, a CloudFormation stack containing AMS Resource Scheduler resources with default configuration, is automatically provisioned into your account. For more on Resource Scheduler change types, see AMS Resource Scheduler.

> **Note**
>
> To find out if AMS Resource Scheduler is already deployed in your account, check the AWS Lambda console for that account and look for the **AMSResourceScheduler** function.

After the AMS Resource Scheduler is provisioned in your account, we recommend you review the default configuration and, if required, customize configurations such as tag key, timezone, scheduled services, and so forth, based on your preferences. For details on the recommended customizations, see Customizing AMS Resource Scheduler, next.

To make the custom configurations, or just confirm the Resource Scheduler configuration,

Customizing AMS Resource Scheduler

We recommend you customize the following properties of AMS Resource Scheduler using the update AMS Resource Scheduler change types, see AMS Resource Scheduler.

- **Tag name**: The name of the tag that Resource Scheduler will use to associate instance schedules with resources. The default value is Schedule.
- **Scheduled Services**: A comma-separated list of services that Resource Scheduler can manage. The default value is "ec2,rds,autoscaling". Valid values are "ec2", "rds" and "autoscaling"
- **Default timezone**: Specify the default time zone for the Resource Scheduler to use. The default value is UTC.
- **Use CMK**: A comma-separated list of Amazon KMS Customer Managed Key (CMK) ARNs that Resource Scheduler can be granted permissions to.
- **Use LicenseManager**: A comma-separated list of AWS Licence Manager ARNs to that Resource Scheduler can be granted permissions to.

> **Note**
>
> AMS may, time to time, release features and fixes to keep AMS Resource Scheduler up to date in your account. When this happens, any customization that you make to the AMS Resource Scheduler are preserved.

Using AMS Resource Scheduler

To configure AMS Resource Scheduler after the solution is deployed, use the automated Resource Scheduler CTs to create, delete, update, and describe (get details on) AMS Resource Scheduler periods (the times when Resource Scheduler runs) and schedules (the configured periods and other options). For an example of using the AMS Resource Scheduler change types, see AMS Resource Scheduler.

To select resources to be managed by AMS Resource Scheduler, following deployment and schedule creation, you use the AMS Tag Create CTs to tag Auto Scaling groups, Amazon RDS stacks, and Amazon EC2 resources with that tag key you provided during deployment, and the defined schedule as the tag value. After the resources are tagged, the resources are scheduled for start or stop per your defined Resource Scheduler schedule.

There is no additional cost to using AMS Resource Scheduler. However the solution makes use of several AWS services and you're charged for these resources as they are used. For more details, see Architecture overview.

To opt out temporarily or completely of AMS Resource Scheduler, submit a Management | Other | Other | Update RFC requesting to disable or remove the solution from our release automation system.

AMS Resource Scheduler cost estimator

In order to track cost savings, AMS Resource Scheduler features a component that hourly calculates the estimated cost savings for Amazon EC2 and RDS resources that are managed by scheduler. This cost savings data is then published as a CloudWatch metric (`AMS/ResourceScheduler`) to help you track it. The cost savings estimator only estimates savings on instance running hours. It does not account any other cost, such as data transfer costs associated with a resource.

The cost savings estimator is enabled with Resource Scheduler. It runs hourly and retrieves cost and usage data from AWS Cost Explorer. From that data it calculates the average cost per hour for each

instance type and then projects the cost for a full day if it was running without being scheduled. The cost savings is the difference between the projected cost and the actual reported cost from Cost Explorer for a given day.

For example, if instance A is configured with Resource Scheduler to run from 9 a.m. to 5 p.m., that is eight hours on a given day. Cost Explorer reports the cost as $1 and usage as 8. The average cost per hour is therefore $0.125. If the instance was not scheduled with Resource Scheduler, then the instance would run 24 hours on that day. In that case, the cost would have been 24x0.125 = $3. Resource Scheduler helped you achieve a cost savings of $2.

In order for the cost savings estimator to retrieve cost and usage only for resources managed by Resource Scheduler from Cost Explorer, the tag key that Resource Scheduler uses to target resources needs to be activated as the **Cost allocation** tag in the Billing Dashboard. If the account belongs to an organization, the tag key needs to be activated in the Management account of the organization. For information on doing this, see Activating User-Defined Cost Allocation Tags and User-Defined Cost Allocation Tags

After the tag key is activated as Cost Allocation Tag, AWS billing starts tracking cost and usage for resources managed by Resource Scheduler, and after that data is available, the cost savings estimator starts to calculate the cost savings and publish the data under the `AMS/ResourceScheduler` metric namespace in CloudWatch.

Cost estimator tips

Cost Savings Estimator does not accept discounts such as reserved instances, savings plans, and so forth, into consideration in its calculation. The Estimator takes usage costs from Cost Explorer and calculates the average cost per hour for the resources. For more details, see Understanding your AWS Cost Datasets: A Cheat Sheet

In order for the cost savings estimator to retrieve cost and usage only for resources managed by Resource Scheduler from Cost Explorer, the tag key that Resource Scheduler uses to target resources needs to be activated as the **Cost Allocation** tag in the Billing Dashboard. If the account belongs to an organization, the tag key needs to be activated in the management account of the organization. For information on doing this, see User-Defined Cost Allocation Tags. If the cost allocation tag is not activated, the estimator is not able to calculate the savings and publish the metric, even if it is enabled.

AMS Resource Scheduler best practices

Scheduling Amazon EC2 Instances

- Instance shut down behavior must be set to `stop` and not to `terminate`. This is pre-set to `stop` for instances that are created with the AMS Amazon EC2 Create automated change type (ct-14027q0sjyt1h) and can be set for Amazon EC2 instances created with AWS CloudFormation ingestion, by setting the `InstanceInitiatedShutdownBehavior` property to `stop`. If instances have shut down behavior set to `terminate`, then the instances will end when the Resource Scheduler stops them and the scheduler won't be able to start them back up.
- Amazon EC2 instances that are part of an Auto Scaling group aren't processed individually by AMS Resource Scheduler, even if they are tagged.
- If the target instance root volume is encrypted with a KMS customer master key (CMK), an additional `kms:CreateGrant` permission needs to be added to your Resource Scheduler IAM role, for the scheduler to be able to start such instances. This permission is not added to the role by default for improved security. If you require this permission, submit an RFC with the Management | AMS Resource Scheduler | Solution | Update change type, and specify a comma separated list of ARNs of the KMS CMKs.

Scheduling Auto Scaling groups

- AMS Resource Scheduler starts or stops the auto scaling of Auto Scaling groups, not individual instances in the group. That is, the scheduler restores the size of the Auto Scaling group (start) or sets the size to 0 (stop).
- Tag AutoScaling group with the specified tag and not the instances within the group.
- During stop, AMS Resource Scheduler stores the Auto Scaling group's Minimum, Desired, and Maximum capacity values and sets the Minimum and Desired Capacity to 0. During start, the scheduler restores the Auto Scaling group size as it was during the stop. Therefore, Auto Scaling group instances must use an appropriate capacity configuration so that the instances' termination and relaunch don't affect any application running in the Auto Scaling group.
- If the Auto Scaling group is modified (the minimum or maximum capacity) during a running period, the scheduler stores the new Auto Scaling group size and uses it when restoring the group at the end of a stop schedule.

Scheduling Amazon RDS instances

- The scheduler can take a snapshot before stopping the RDS instances (does not apply to Aurora DB cluster). This feature is turned on by default with the **Create RDS Instance Snapshot** AWS

CloudFormation template parameter set to **true**. The snapshot is kept until the next time the Amazon RDS instance is stopped and a new snapshot is created.

Scheduler can start/stop Amazon RDS instance that are part of a cluster or Amazon RDS Aurora database or in a multi availability zone (Multi-AZ) configuration. However, check Amazon RDS limitation when the scheduler won't be able to stop the Amazon RDS instance, especially Multi-AZ instances. To schedule Aurora Cluster for start or stop use the **Schedule Aurora Clusters** template parameter (default is **true**). The Aurora cluster (not the individual instances within the cluster) must be tagged with the tag key defined during initial configuration and the schedule name as the tag value to schedule that cluster.

Every Amazon RDS instance has a weekly maintenance window during which any system changes are applied. During the maintenance window, Amazon RDS will automatically start instances that have been stopped for more than seven days to apply maintenance. Note that Amazon RDS will not stop the instance once the maintenance event is complete.

The scheduler allows specifying whether to add the preferred maintenance window of an Amazon RDS instance as a running period to its schedule. The solution will start the instance at the beginning of the maintenance window and stop the instance at the end of the maintenance window if no other running period specifies that the instance should run, and if the maintenance event is completed.

If the maintenance event is not completed by the end of the maintenance window, the instance will run until the scheduling interval after the maintenance event is completed.

> **Note**
>
> The Scheduler doesn't validate that a resource is started or stopped. It makes the API call and moves on. If the API call fails, it logs the error for investigation.

AWS Systems Manager in AMS Advanced

An AWS Systems Manager document (SSM document) defines the actions that Systems Manager performs on your AWS resources. Systems Manager includes more than a dozen pre-configured documents that you can use by specifying parameters at runtime. Documents use JavaScript Object Notation (JSON) or YAML, and they include steps and parameters that you specify.

AWS Managed Services (AMS) is a trusted publisher for SSM documents. SSM documents owned by AMS are shared only with onboarded AMS accounts, always begin with a reserved prefix (AWSManagedServices-*), and show up in the Systems Manager console, as owned by AWS. The AMS process for SSM document development and publishing follows AWS best practices and requires multiple peer reviews throughout the document life cycle. For more information on AWS best practices for sharing SSM documents, see Best practices for shared SSM documents.

Available AMS Advanced SSM documents

AMS Advanced SSM documents are available exclusively to AMS Advanced customers, and are used to automate operational workflow to operate your account.

To see the available AMS Advanced SSM documents from the AWS Management Console:

1. Open the Systems Managerconsole at AWS Systems Manager console.
2. Choose **Shared with me**.
3. In the search bar, filter by **Document name prefix**, then **Equals**, and set the value to **AWSManagedServices-**.

   For AWS CLI instructions, see Using shared SSM documents.

AMS Advanced SSM document versions

SSM documents support versioning. AMS Advanced SSM documents can't be modified from the customer's account and can't be re-shared. They're centrally managed and maintained by AMS Advanced in order to operate the account.

Version numbers are incremented with each document update in a specific AWS Region. As new Regions become available, the same document content in two Regions can have different version numbers; this is typical and doesn't mean their behavior will be different. If you want to compare two AMS Advanced SSM documents, we recommend comparing their hashes with the AWS CLI:

```
aws ssm describe-document \
--name AWSManagedServices-DOCUMENTNAME \
--output text --query "Document.Hash"
```

Two SSM documents are identical if their hashes match.

Systems Manager pricing in AMS

There is no cost associated with AMS Advanced SSM document access. Runtime cost varies based on the type of SSM document, its steps, and runtime duration. For more information, refer to AWS Systems Manager pricing.

Offboard AMS accounts

AMS offers off-boarding assistance for single-account landing zone and multi-account landing zone accounts.

AMS does not unshare any AMIs from you during offboarding to avoid impact for any of your depedencies. If you want to remove AMS AMIs from your account, you can use the `cancel-image-launch-permission` API to hide specific AMIs. For example, you can use the script below to hide all of the AMS AMIs that were shared with your account earlier:

```
for ami in $(aws ec2 describe-images --executable-users self --owners 027415890775 --
query 'Images[].ImageId' --output text) ;
    do
    aws ec2 cancel-image-launch-permission --image-id $ami ;
    done
```

You must have the AWS CLI v2 installed for the script to execute without any errors. For AWS CLI installation steps, see Installing or updating the latest version of the AWS CLI. For details on the `cancel-image-launch-permission` command, see `cancel-image-launch-permission`.

Topics

- Offboard from AMS single-account landing zone accounts
- Offboard from AMS multi-account landing zone accounts

Offboard from AMS single-account landing zone accounts

AMS offers off-boarding assistance within 30 days prior to termination of AMS.

You must request off-boarding assistance at least 7 days before such assistance can be provided. Off-boarding assistance can be offered in two forms:

- Control hand-over: AMS transfers account control back to you along with access credentials for all AMS-managed applications, or

- Resource termination for account closure: AMS deletes all of the data in your AMS-managed environment and de-provisions any active resources in the account. When submitting the offboarding request, customers can request that AMS:
  - Delete or retain the data objects (including logs) that are stored on Amazon S3 buckets
  - Remove or retain Amazon S3 buckets
  - Remove or retain AWS Backup restore points

> **Important**
>
> Any other specific requests (subject to plausibility) must be communicated to AMS before initialization of offboarding.

Optional Prerequisites (if required):

> **Note**
>
> Prior to the offboarding request, customers can request AMS assistance to transfer your data in the existing format using AWS Snowball or any other media that AWS interfaces with.

In addition to data backups, the following customer data can be provided as part of off-boarding assistance:

- Data stored in storage services including logs
- Customer-specific change type schemas
- CloudFormation templates for change type schemas

If off-boarding activities are not completed upon the termination of AMS, we hand over the controls of the account(s) to enable you to complete any pending activity.

Function	What was removed	Impact	Actions needed
Monitoring, Logging, Alerting	AMS Monitoring removed	AMS no longer has access or visibility	Contingencies for removed and

Function	What was removed	Impact	Actions needed
	MMS (Managed Monitoring System) unsubscribed Resource Tagger and Alarm Manager removed Baseline CloudWatch alerts remain on existing resources GuardDuty and Macie: Ownership reverts to you	into your resources and environment.	unsubscribed services are owned by you.
Backup management	AMS Backup automation is removed although the AWS Backup service remains available for use. Backup vaults and data are retained unless deletion is requested.	AMS no longer monitors the backup jobs or performs restoration actions during incidents. Alarms and alerts are disabled. Deletion of the IAM backup role and KMS keys render your AMS backups inoperable.	AMS Backup Plans must be reconfigured. All monitoring and remediation ownership returns to you.
AMS automations for service management	AMS-curated AWS SSM automation runbooks, Amazon Simple Notification Service (SNS), and AWS Lambda functions are no longer available.	No AMS access to your accounts. All automation disabled.	All automation including SSM, SNS, and Lambda functions need to be recreated, if required.

Function	What was removed	Impact	Actions needed
Compliance	AMS visibility into and monitoring for all GuardDuty and AWS Config rules removed, although these rules remain on the accounts.	All monitoring, reporting, and remediation from Amazon GuardDuty and AWS Config Rules is no longer managed by AMS.	Monitoring and remediation for all security and compliance tools to be assumed by you.
On-instance agents	Access to Resource Scheduler, Resource Tagger or automated instance configura tion to install required agents in your EC2 instances is removed.	CloudWatch and SSM Agents on instances are left in place with existing configura tions however, AMS no longer assists with these configurations.	You manage tagging and on-instance CloudWatch and SSM agent configurations.
Patch and reporting infrastructure	AMS no longer manages pre- and post- patching activities, and access and visibility to these services are removed.	AMS no longer creates a snapshot of the instance prior to patching, no longer installs and monitors the patch installation, and no longer notifies you of the outcome. Reports and "audit" S3 buckets are left in your accounts at your request. AMS no longer generates service metric reports.	You retain the Patch baselines and snapshots created in the past. Additionally, the configuration of the patch maintenan ce windows remains but the patches are no longer installed or remediated by AMS. All reporting on infrastructure operational metrics are now your responsibility.

Function	What was removed	Impact	Actions needed
Process management	All accounts are offboarded from the service managemen t provided for incidents, including service requests, problem, and change, management.	All service disruptio n formerly remediate d by AMS through incidents and service requests, and changes to the environment, as well as Root Cause investigations, are longer managed by AMS.	You regain full ownership of all process management.

Offboard from AMS multi-account landing zone accounts

There are two types of AWS accounts that you can offboard from AMS Advanced multi-account landing zone:

- Application accounts
- Core accounts

To offboard all accounts from your AMS Multi-account landing zone, you must offboard all Application accounts before you offboard Core accounts.

To take over and continue operating workloads in offboarded Application or Core accounts, make sure that you review this documentation with your AMS account team. This documentation outlines the changes that AMS performs during the offboard process.

Tasks to complete for continued operation of offboarded accounts

The following tasks are required for continued operation of accounts that you offboarded from AMS multi-account landing zone:

- **Turn on Developer mode:** To gain more permission to your accounts, turn on Developer mode before you offboard Application accounts from AMS. When you turn on Developer mode, you can more easily make the necessary changes to prepare for offboarding. Don't try to remove

or modify AMS infrastructure resources. If you delete AMS infrastructure resources, then AMS might not be able to successfully offboard your account. For information about how to enable Developer mode, see Getting started with AMS Advanced Developer mode.

If you can't complete the necessary changes to prepare for offboarding after you turn on Developer mode, then contact your AMS account team to discuss your requirements.

- **Choose an alternate method for EC2 stack access:** After you offboard Application accounts from AMS, you can't use RFCs to access your stack resources. Review Offboarding changes, and then choose an alternate access method so that you retain access to your stacks. For more information, see Access Alternatives.

Offboard AMS Application accounts

To offboard Application accounts from your multi-account landing zone environment, complete the following steps for each account:

1. Verify that there are no open RFCs in the account. For more information, see Creating, cloning, updating, finding, and canceling RFCs.
2. Verify that you can access the primary or root user email address for the account.
3. From the Application account, submit an RFC with the Application Account | Confirm offboarding (ct-2wlfo2jxj2rkj) change type. In the RFC, specify the Application account to offboard.
4. From the Management account, submit an RFC with the Management account | Offboard Application Account (ct-0vdiy51oyrhhm) change type. In the RFC, specify the Application account to offboard. Also, indicate if you want to delete or retain the transit gateway attachment to the landing zone.
5. To make sure that AMS billing is stopped, notify your CSDM that you offboarded the account.

The following occurs after the Application account is offboarded:

- All components are disassociated from AMS services, but your created resources remain in the account. You can choose to keep or close the AMS offboarded account.
- Core accounts and other remaining Application accounts function normally after an Application account is offboarded.
- AMS billing is stopped, but AWS billing isn't stopped until you close the account. For more information, see What you need to know before closing your account.

- If an account is closed, then the account is visible in your organization in the `suspended` state for 90 days. After 90 days, the closed account is permanently removed and no longer visible in your organization.
- After the account is closed, you can still sign in and file a support case or contact AWS Support for 90 days.
- After 90 days, any content that remains in the account is permanently deleted and the remaining AWS services are terminated.

Application account offboarding FAQ

Q: Can I use my federated IAM roles to continue to access an Application account that I offboarded from my AMS multi-account landing zone?

Yes. AMS created default AWS Identity and Access Management (IAM) roles remain available in the account after AMS offboarding. However, these roles and policies are designed for use with AMS access management. To provide the necessary access for your users, you might need to deploy your own IAM resources.

Q: How do I gain full access to an Application account that I offboarded from my AMS multi-account landing zone?

Offboarded Application accounts are moved to the **Deprecated** Organizational Unit (OU) in the AWS Organizations account structure. This move lifts SCP access restrictions that previously blocked root user access. For information about how to reset root user credentials, see Resetting a lost or forgotten root user password.

Q: What changes are made during Application account offboarding?

For information about actions that AMS takes when the service offboards accounts, see Offboarding changes.

Q: Can I offboard an Application account without detaching it from the transit gateway?

Yes. Use the Management account | Offboard Application Account (ct-0vdiy51oyrhhm) change type to submit the RFC, and specify the `DeleteTransitGatewayAttachment` parameter as `False`.

Q: How long does it take to offboard an Application account?

When you use the Management account | Offboard Application Account (ct-0vdiy51oyrhhm) change type, RFCs complete within 1 hour.

Q: Is it mandatory that I close the offboarded account?

No. Account closure after AMS offboarding isn't mandatory. During the offboarding process, AMS removes its access and management of your AWS account, but your account and your resources within the account remain. It's important to note that after AMS offboarding, you're solely responsible for managing and maintaining your AWS account and resources. AMS isn't responsible for any issues, incidents, or service disruptions that might occur in your account after the offboarding process is complete. For more information, see How do I close my AWS account?.

Q: If I submit an account closure request, are all existing resources deleted immediately?

No. Account closure doesn't terminate your resources. Resources in the account automatically terminate 90 days after the closure request. AMS billing stops, but AWS resource billing doesn't stop until you close the account. For more information, see What you need to know before closing your account.

Q: Can I schedule the offboarding of an Application account?

Yes. You can schedule the RFCs to run at a specific time. However, the Application Account | Confirm offboarding RFC must complete before you can schedule the Management account | Offboard Application Account RFC. For more information, see RFC scheduling.

Application account offboarding RACI

- **R**: Responsible party. The party responsible for completing the listed task.
- **A**: Accountable party. The party that approves the completed task.
- **C**: Consulted party. A party whose opinions are sought, typically as subject matter experts, and with whom there is bilateral communication.
- **I**: Informed party. A party that's informed on progress, often only on completion of the task or deliverable.

Activity	Custor	AWS Managed Services (AMS)
Prerequisites		

Activity	Custor	AWS Managed Services (AMS)
Verify access to the root email address for each AWS account ID that will be offboarded	R	C
Review AMS documentation on recommended customer actions and prepare accounts for AMS offboarding	R	C
If needed, submit an RFC to enable Developer mode to prepare accounts for AMS offboarding	R	I
If needed, choose an alternate method for EC2 stack access.	R	I
Offboarding		
Submit RFCs to confirm and request offboarding of Application accounts	R	I
Offboard AMS components from Application accounts	I	R
Notify AMS CSDM of offboarded accounts to stop AMS billing	R	I
Post-offboarding		
Reset root user account password and verify root access in offboarded accounts	R	C
Close the account or follow AMS guidance on recommended customer actions in the AMS offboarding documentation to continue operating the accounts	R	C

Offboard Core accounts

To offboard multi-account landing zone Core accounts, complete the following steps:

1. Verify that all Application accounts in the landing zone were offboarded from AMS.

2. Verify that you have no open RFCs in the accounts. For more information, see Creating, cloning, updating, finding, and canceling RFCs.
3. Verify that you can access the primary or root user email address for all Core accounts. For more information, see Multi-Account Landing Zone accounts.
4. Verify that you can access the primary or root user phone number for the management account. Use the `AWSManagedServicesBillingRole` IAM role to update the phone number. For more information, see How do I update my telephone number that's associated with my AWS account?.
5. Log in to your AMS landing zone management account and submit an AMS service request. In the service request, specify to offboard your entire landing zone.

The following occurs after the Core accounts are offboarded:

- All components are disassociated from AMS services, but some AWS resources remain in the account. You can choose to keep or close the AMS offboarded Core accounts.
- AMS billing is stopped, but AWS billing isn't stopped until you close the account. For more information, see What you need to know before closing your account.
- If an account is closed, then the account is visible in your organization in the `suspended` state for 90 days. After 90 days, the closed member account is permanently removed and is no longer visible in your organization.
- After the account is closed, you can still sign in and file a support case or contact AWS Support for 90 days.
- After the account is closed for 90 days, any content that remains in the account is permanently deleted, and the remaining AWS services are terminated.

Core account offboarding FAQ

Q: Can I use my federated IAM roles to continue to access the offboarded Core accounts?

Yes. AMS created default AWS Identity and Access Management (IAM) roles remain available in the offboarded account. However, these roles and policies are designed for use with AMS access management. To provide the necessary access for your users, you might need to deploy your own IAM resources.

Q: How do I gain full access to the Management, Shared Services, Networking, or other non-Application MALZ account after offboarding from AMS multi-account landing zone?

After offboarding, follow the instructions in Resetting a lost or forgotten root user password to use primary (root) user credentials to access Core accounts other than the management account. Unlike other account types, the management account retains an inaccessible multi-factor authentication (MFA) device that's associated to the root user to prevent usage. To regain root access you must follow the lost MFA device recovery process.

Q: What changes are made during Core account offboarding?

For information about actions that AMS takes when the service offboards accounts, see Offboarding changes.

Q: How long does Core account offboarding take to complete?

The Core account offboarding process typically takes up to 30 days to complete. However, to make sure that all required steps are correctly completed, you must initiate the offboarding request at least 7 days before offboarding starts. To facilitate an easy transition, plan ahead and submit your offboarding request in advance.

Q: How do I manage shared components after AMS offboarding?

AMS Managed Active Directory and other shared services infrastructure components are designed for AMS operator access. You might need to update Amazon Elastic Compute Cloud (Amazon EC2) security groups, AWS Organizations service control policy (SCP), or make other changes to retain full access to these components.

Q: Can I close offboarded Core accounts?

By default, Application accounts have multiple dependencies on MALZ Core accounts, such as AWS Organizations membership, transit gateway network connectivity, and DNS resolution through AMS Managed Active Directory. After you resolve these dependences, you can decommission and close the offboarded Core account. For more information, see Multi-Account Landing Zone accounts.

Core account offboarding RACI

- **R**: Responsible party. The party responsible for completing the listed task.
- **A**: Accountable party. The party that approves the completed task.
- **C**: Consulted party. A party whose opinions are sought, typically as subject matter experts, and with whom there is bilateral communication.

- **I**: Informed party. A party that's informed on progress, often only on completion of the task or deliverable.

Activity	Custor	AWS Managed Services (AMS)
Prerequisites		
Verify access to the root email address for each AWS account ID that will be offboarded	R	C
Verify access to and update the root user phone number for the management account	R	C
Review AMS documentation on recommended customer actions and prepare accounts for AMS offboarding	R	C
Offboarding		
Submit service request to request offboarding of the landing zone	R	I
Offboard AMS components from Core accounts	I	R
Post-offboarding		
Reset root user account password and verify root access in offboarded accounts	R	C
Close the accounts or follow AMS guidance on recommended customer actions in AMS offboarding documentation to continue to operate the accounts	R	C

Offboarding changes

The following table describes the actions that AMS takes for multi-account landing zone offboarding, potential impacts, and recommended actions.

Componer	Account type	Actions taken to offboard	Potential impacts	Recommended customer action
Access managem t	Applicati on accounts	After offboarding, stack access RFCs for just-in-time, time-bound access can no longer be submitted to access EC2 stacks through AMS bastion hosts AMS no longer manages access-related component s on any existing EC2 resource stacks (PBIS Open agent, domain join scripts)	Can't use AMS bastions through RFS to access EC2 instances EC2 instances launched from non AMS provided AMIs aren't joined to Managed Active Directory domain If not removed, AMS launch scripts in existing resource stacks might produce errors because of missing AMS dependenc ies, and prevent rejoining to a different domain	Use alternate methods to access EC2 instance (see Access Alternatives) Remove AMS launch scripts from existing EC2 resource stacks (see Disable AMS EC2 launch scripts)
Access managem t (cont'd)	Core accounts	If you migrated from PBIS Open to PBIS Enterprise (AD Bridge), then AMS no longer renews licensing after core account offboarding	If PBIS Enterprise license is allowed to expire, Active Directory credentia ls aren't valid for existing Linux-based EC2 instance stacks	If you migrated to PBIS Enterprise (AD Bridge), then decide whether to maintain licensing or decommission (see PBIS Open/ Enterprise (AD Bridge))

Componer	Account type	Actions taken to offboard	Potential impacts	Recommended customer action
Logging, Monitoring, Incident/ Event Managemt	Application and Core accounts	AMS components to deploy	Newly created resources don't have AMS baseline monitoring and alarms applied Infrastructure metric alarms and security events no longer generate AMS incidents AWS Config is no longer aggregated in a central account	Define, capture, and analyze operations metrics to view workload events and take appropriate action. Implement any required alert workflow to continue to apply required operational monitoring and alarms on new resources and to for receive security alerts from AWS Config and Amazon GuardDuty.

Componer	Account type	Actions taken to offboard	Potential impacts	Recommended customer action
		Learn about AMS monitoring defaults; for more information, see Monitoring and event management in AMS. The following table shows what is monitored, and the default alerting threshold s. You can change the alerting thresholds with a Management | Other | Other | Update (ct-0xdaw ir96cy7k) RFC after determining what changes you want and subscribing to the relevant CloudWatc h Amazon SNS topic. For information about creating and subscribing to topics, see Subscribe to a Topic. For general information, see Amazon SNS FAQs. To be notified directly when alarms cross their threshold, in addition to AMS's standard alerting process, follow these instructions about how to overwrite alarm configurations, Receiving alerts generated by AMS. Amazon CloudWatc h provides extended retention of metrics. For more information, see CloudWatch Limits.		

Componer	Account type	Actions taken to offboard	Potential impacts	Recommended customer action
		are removed Existing deployed Amazon CloudWatch alarms remain but no longer create AMS incidents AWS Config aggregation authorizations from AMS and the MALZ Core security account are removed AWS Config rules remain deployed and Amazon GuardDuty remains enabled, but no longer creates AMS incidents		
Continuity management (Backup and Restore)	Application accounts	AMS no longer monitors backup jobs or performs backup and restore requests AMS default backup vaults, backup encryption key, and backup role remain	Backup operation failures might not be noticed	Monitor and review backup plan configurations

Componer	Account type	Actions taken to offboard	Potential impacts	Recommended customer action
Patch managem t	Applicati on and Core accounts	AMS no longer monitors patching operations for successful execution or performs manual patching AMS no longer updates AMS infrastructure components AMS provided patch baselines are retained AMS provided AWS Systems Manager patch automation runbooks are unshared and no longer available for use	Patching operation failures might not be noticed Existing patch configurations that depend on AMS provided Systems Manager Automatio n runbooks must be reconfigured to continue uninterru pted	Review and reconfigure patching configura tions as needed
Network managem t	Applicati on accounts	If specified, the transit gateway attachment in the offboarded Application account is removed	Offboarded Application account can no longer use a transit gateway to access shared services, such as Managed Active Directory or other Application accounts	Specify `DeleteTra nsitGatew ayAttachment` as `False` to retain the transit gateway connectivity

Componer	Account type	Actions taken to offboard	Potential impacts	Recommended customer action
Security managem t	Applicati on accounts	Account is detached from central Trend Micro DSM console. Also, endpoint agents no longer forward alerts through the AMS incident process Trend Micro agents remain installed but are no longer managed or updated by AMS AMS provided AMI customizations that deploy Trend Micro agent are no longer maintained or updated by AMS	EC2 instance endpoint malware detections might not be noticed The Trend micro agent isn't deployed on EC2 instances that are launched from non AMS provided AMIs	Consider options for continuing or discontinuing Trend Micro (see Trend Micro Deep Security)
Security managem t (cont'd)	Core accounts	Trend Micro DSM infrastru cture is left in place within Shared Services accounts but no longer maintained or updated by AMS Trend Micro DSM no longer forwards alerts through the AMS incident process	EC2 instance endpoint malware detections might go unnoticed EC2 instance endpoint protection might be impacted if infrastructure isn't maintained (definition updates, licensing, and so on)	Decide whether to continue or discontinue Trend Micro (see (see Trend Micro Deep Security)

Componer	Account type	Actions taken to offboard	Potential impacts	Recommended customer action
Change managem t	Applicati on and Core accounts	AMS RFC console and API is removed AMS custom service control policies (SCPs) that contain account-l evel access restrictions are detached during Applicati on account offboarding, and deleted during Core account offboarding	You must use a native AWS API to create new resources, change existing resources, or update existing AWS CloudForm ation stacks Access restricti ons are no longer imposed at the account-level through AMS provided SCPs	Make sure that user roles provide sufficient access to use AWS services Create SCPs to provide account-level permission restrictions
AMS OS images and automatio ns for service managem t	Applicati on and Core accounts	AMS no longer provides support on customiza tions and launch scripts included in AMS provided EC2 AMIs AMS provided EC2 AMIs remain available in your offboarded accounts AMS provided Systems Manager Automation runbooks are unshared and no longer available for use	After offboardi ng, AMS provided AMIs launched with AWS CloudForm ation send cfn-signal `FAILURE` because of missing dependencies on AMS components Operational processes that depend on AMS provided Systems Manager Automatio n runbooks might fail	Review and update any build or operational processes that are dependent on AMS provided AMIs or Systems Manager Automatio n runbooks

Componer	Account type	Actions taken to offboard	Potential impacts	Recommended customer action
Shared services infrastructure	Core accounts	AMS access is removed and AMS no longer manages shared components, including AMS Managed Active Directory, AWS Transit Gateway, and AWS Organizations	Loss of management over shared infrastructure	Reset admin access to AMS Managed Active Directory and assume management of shared services components
Reporting	Application and Core accounts	AMS no longer collects account or resource-level details for aggregate reporting	Loss of insight into operational and business metrics (backup and patching coverage, change management, and incident activity)	Replace any needed aggregate data reporting across accounts with their own solution
AMS account team and service desk	Application and Core accounts	AMS account team (CSDM, CA) and AMS operations service desk no longer support the offboarded accounts	Loss of operational support with expertise in the AMS designed multi-account landing zone architecture and related components	Make sure that there's sufficient personnel and familiarity with account structure and resources to support operations in the environment

Access Alternatives

The following are alternatives methods for retaining access to your EC2 stack after you offboard AMS accounts:

- **Use Session Manager** to access EC2 instances with elevated permissions without requiring bastions or inbound network access. For more information, see AWS Systems Manager Session Manager.
- **Rejoin EC2 instances to a different Active Directory domain** with new domain credentials. If you use AWS Directory Service, then see Join an EC2 instance to your AWS Managed Microsoft AD directory.
- **Use local user accounts** that you created through one of the other access methods or through AWS Systems Manager Run Command.

Disable AMS EC2 launch scripts

Linux operating systems

Use your distribution's package manager to uninstall the `ams-modules` package. For example, for Amazon Linux 2 use `yum remove ams-modules`.

Windows operating systems

To disable EC2 launch scripts in Windows, complete the following steps:

1. Windows Server 2008/2012/2012r2/2016/2019:

   Disable or remove the Managed Services Startup scheduled task from Task Scheduler. To list scheduled tasks, run the `Get-ScheduledTask -TaskName '*Ec2*'` command.

   Windows Server 2022:

   Remove the EC2Launch v2 task. This task runs `Initialize-AMSBoot` in `postReady` stage in C:\ProgramData\Amazon\EC2Launch\config\agent-config.yml on the instance. The following is a snippet from an example `agent-config.yml`:

```
{
 "task": "executeScript",
 "inputs": [
  {
   "frequency": "always",
   "type": "powershell",
   "runAs": "localSystem"
  }
 ]
```

```
}
```

2. (Optional) Remove the following file contents:

```
C:\Program Files\WindowsPowerShell\Modules\AWSManagedServices.*
C:\Windows\System32\WindowsPowerShell\v1.0\Modules
\AWSManagedServices.Build.Utilities\*
```

PBIS Open/Enterprise (AD Bridge)

To determine if you use PBIS Open or PBIS Enterprise (AD Bridge) edition, run the following command in a Linux EC2 managed instance:

```
yum info | grep pbis
```

The following is example output that shows PBIS Enterprise (AD Bridge):

```
Name        : pbis-enterprise
From repo   : pbise
Name        : pbis-enterprise-devel
Repo        : pbise
Description : The pbis-enterprise-devel package includes the development
```

PBIS Open

PBIS Open is a deprecated product that BeyondTrust no longer supports. For more information, see the BeyondTrust pbis-open documentation.

AD Bridge (PBIS Enterprise)

You can do one of the following:

- **Renew licensing and continue operating AD Bridge.** Contact BeyondTrust to discuss licensing and support.
- **Discontinue use of AD Bridge.** Run the following Shell command to remove PBIS-Enterprise package from Linux managed instances. For more information, see the BeyondTrust documentation Leave a Domain and Uninstall the AD Bridge Agent.

```
$ sudo /opt/pbis/bin/uninstall.sh purge
```

Leave the AMS managed Active Directory domain without removing PBIS agent

You have the option to leave the AMS managed Active Directory without removing the PBIS agent. Use one of the following solutions, depending on your operating system:

Linux operating systems

Use PBIS from the AMS managed AD to run the following shell command to unjoin a Linux EC2 instance. For more information, see the BeyondTrust pbis-open or BeyondTrust AD Bridge documentation, depending on which version you use.

```
$ sudo /opt/pbis/bin/domainjoin-cli leave
```

You might see an error message similar to the following:

```
Error: LW_ERROR_KRB5_REALM_CANT_RESOLVE [code 0x0000a3e1]
Cannot resolve network address for KDC in requested realm
```

If this error occurs, then run the following commands to delete AD provider registry and restart `lwsm` services:

```
$ /opt/pbis/bin/regshell dir '[HKEY_THIS_MACHINE\Services\lsass\Parameters\Providers
\ActiveDirectory\DomainJoin]'
```

Use the directory ID output that you received from the previous command (for example, **`A123EXAMPLE.AMAZONAWS.COM`**) to run the following commands:

```
$ /opt/pbis/bin/regshell delete_tree \
'[HKEY_THIS_MACHINE\Services\lsass\Parameters\Providers\ActiveDirectory\DomainJoin
\DIRECTORYID]'

$ /etc/pbis/redhat/lwsmd restart

$ /opt/pbis/bin/lwsm restart lwreg
```

Windows operating systems

```
# Collect hostname and domain name using:
```

```
Test-ComputerSecureChannel -verbose

# Disjoin computer from the domain:

netdom remove hostname /domain:domain name /force
```

> **Note**
>
> Make sure to disable or remove Managed Services Startup scheduled task as mentioned in Disable AMS EC2 launch scripts.

Trend Micro Deep Security

Use one of the following options to continue or discontinue the use of Trend Micro Deep Security:

Continue usage

- **(If offboarding the entire MALZ) After Core account offboarding, reconnect offboarded Application accounts to the existing Trend Micro Deep Security Manager (DSM) and maintain licensing in shared Services account.** For more information, see Add AWS cloud accounts and Check your license information.

  1. Log in to the shared services account and navigate to the Secrets Manager console.
  2. Retrieve DSM console admin credentials that are stored in the `/ams/eps/` path.
  3. Log in to the DSM console at https://dsm.sentinel.int.
  4. Choose **Use Cross Account Role**, and then enter **`arn:aws:iam::ACCOUNTID:role/mc_eps_cross_account_role`**. Replace **ACCOUNTID** with the offboarded Application account ID.
  5. Choose **Next**.
  6. Wait several minutes for DSM to process the account discovery and show that sync was successful.

- **Reconnect offboarded Application accounts to a new Trend Micro DSM installation.** For more information, see Activate and protect agents and Activate the agent.
- **Reconnect offboarded Application accounts to Trend Micro Cloud One.** For more information, see Migrate from Deep Security to Workload Security and Migrate from an on-premises DSM.

Discontinue usage

- **Uninstall Trend Micro Deep Security Agents from offboarded Application accounts.** For more information, see Uninstall Deep Security.

Change management modes

AWS Managed Services (AMS) change Management is the process that the service uses to guardrail changes in AMS Advanced. The change management system helps you control risk, and prevent adverse impact, while maintaining high operational standards for the environment. AMS Advanced has different modes that provide different levels of control and risk (all except Customer-Managed mode are managed by AMS):

- RFC mode (formerly Standard CM mode): Provides a "request for change" (RFC) system and AMS-custom change types (CTs)
- Direct Change mode: Same as RFC mode plus use of AWS APIs and consoles to create AMS-managed resources
- AWS Service Catalog on AMS: Similar to Direct Change mode, but instead of using the AMS change management system (RFCs), you use AWS Service Catalog to create resources that AMS then manages.
- Developer mode: Same as Direct Change mode only the resources you create with AWS APIs and consoles are not AMS-managed—you are responsible for their management
- Self Service Provisioning (SSP) mode: Same as Developer mode except there is no access to the AMS change management system (no RFCs)
- Customer Managed mode: AMS provides you with a multi-account landing zone landing zone but all resource management is your responsibility

The AWS Managed Services (AMS) change management system, using the change management (CM) API, provides operations to create and manage requests for change (RFCs) for both multi-account landing zone (MALZ) and single-account landing zone (SALZ) accounts.

A request for change (RFC) is a request created by either you or AMS through the AMS interface to make a change to your managed environment and includes a change type (CT) ID for a particular operation.

The AMS change management (CM) API provides operations to create and manage requests for change (RFCs). You can create, update, submit, approve, reject, and cancel RFCs. The AMS operators can create, update, submit, approve, reject, cancel, and mark RFCs as closed.

For a list of AMS reserved prefixes not to be used in tag or other names, see Reserved prefixes.

For information on each change type, including schemas and examples, see the AMS Change Type Reference.

> **Note**
>
> All change management API calls are recorded in AWS CloudTrail. For more information, see Accessing your logs.

Modes overview

Use this information to help you select the appropriate AWS Managed Services (AMS) mode for hosting your applications, based on your desired combination of flexibility and prescriptive governance to achieve your business outcomes.

The intended audience for this information is:

- Customer teams responsible for the strategy and governance of their landing zone. This information will help the team lay out the foundation of an AMS-managed landing zone, with the AMS modes they'd like to offer to their internal and external customers.
- Business and application owners tasked with migrating their application to AMS. This information will help with planning application migration, with the appropriate AMS mode to migrate/host their application. Note, the same application can be hosted in more than one AMS mode during different phases of its Software Development Life Cycle (SDLC) lifecycle.
- AMS partners tasked with guiding customers on the different options to build and migrate to AMS.

This information is most useful during the foundation phase of setting up your AMS-managed platform, and when you are transitioning from the foundation to the migration phase of your cloud adoption journey, just after onboarding to AMS is complete and you're focusing on application governance and operations.

Types of modes and accounts in AMS

AWS Managed Services (AMS) modes can be defined as the ways of interacting with the AMS service under the specific governance framework for each mode. The landing zone differences, multi-account landing zone or MALZ and single-account landing zone or SALZ are noted.

Note

For details about application deployment and choosing the right AMS mode, see AMS modes and applications or workloads.
For real-world use cases of the different modes, see Real world use cases for AMS modes

The following table provides descriptions of the modes per AMS service.

AMS feature	RFC mode (formerly Standard CM mode) / OOD*	Direct Change mode	AWS Service Catalog	Self-service provisioning / Developer mode	Customer Managed
Landing Zone Configuration	MALZ and SALZ	MALZ and SALZ	MALZ and SALZ		
Change Management	Change scheduling, review of manual changes, and change record	Same as RFC mode for high-risk changes like IAM or security groups	None		
Logging, Monitoring, Guardrails, and Event Management	Yes (supported resources)			No	

AMS feature	RFC mode (formerly Standard CM mode) / OOD*	Direct Change mode	AWS Service Catalog	Self-service provisioning / Developer mode	Customer Managed
Continuity management	Yes (supported resources)			Not applicable / No	No
Security management	Instance level security controls and account level controls			Account level controls	AWS Org level controls
Patch management	Yes			Not applicable / No	No
Incident and problem management	Response and resolution SLA for AMS supported resources			Response SLA for resulting resources	No
Reporting	Yes			No	
Service request management	Yes			Support requests only	No

*Operations On Demand (OOD) has an offering for customers using the RFC mode to manage their changes through dedicated resourcing. For more details, see the Operations on Demand catalog of offerings and talk to your cloud service delivery manager (CSDM).

Note

Self-Service Provisioning mode and AMS Advanced Developer mode may both appear to be a suitable fit for an application that has complex architecture rooted in native

AWS Services. When architecting workloads, you make trade-offs between operational excellence and agility, based on your business context. This is a good way to think about selecting SSP mode or Developer mode for your application. The selection may also change based on the SDLC phase of the application. For example: When the application is production-ready, then SSP mode maybe a more appropriate option due to stricter AMS guardrails in this mode. The guardrails are enforced in the form of preventative controls like RFC-based change control for IAM updates and SCPs at the application OU level. These business decisions can drive your engineering priorities. You might optimize to increase flexibility for application owners in "pre-prod" phase at the expense of governance and operational support.

MALZ architecture and associated AMS modes

AMS multi-account landing zone (MALZ) gives you the option to automatically provision application accounts (or resource accounts) under the default Organizational Units (OU): Customer Managed OU, Managed OU, or Development OU. The infrastructure provisioned in the application accounts created under each of these OUs is subject to the specific AMS mode offered by those foundational OUs. It is common to find a mix of two or more modes in the same application account. For example: RFC mode and SSP mode can coexist in an AMS managed account that hosts pipeline architecture consisting of API Gateway and Lambda for trigger functions, and EC2, S3, and SQS for ingestion and orchestration. In this case, SSP mode would apply to Lambda and API Gateway.

Figure 1 presents how different modes are offered through the foundational OUs in AMS. When requesting a new application account in AMS, you must select the OU for the account.

MALZ architecture and associated AMS modes

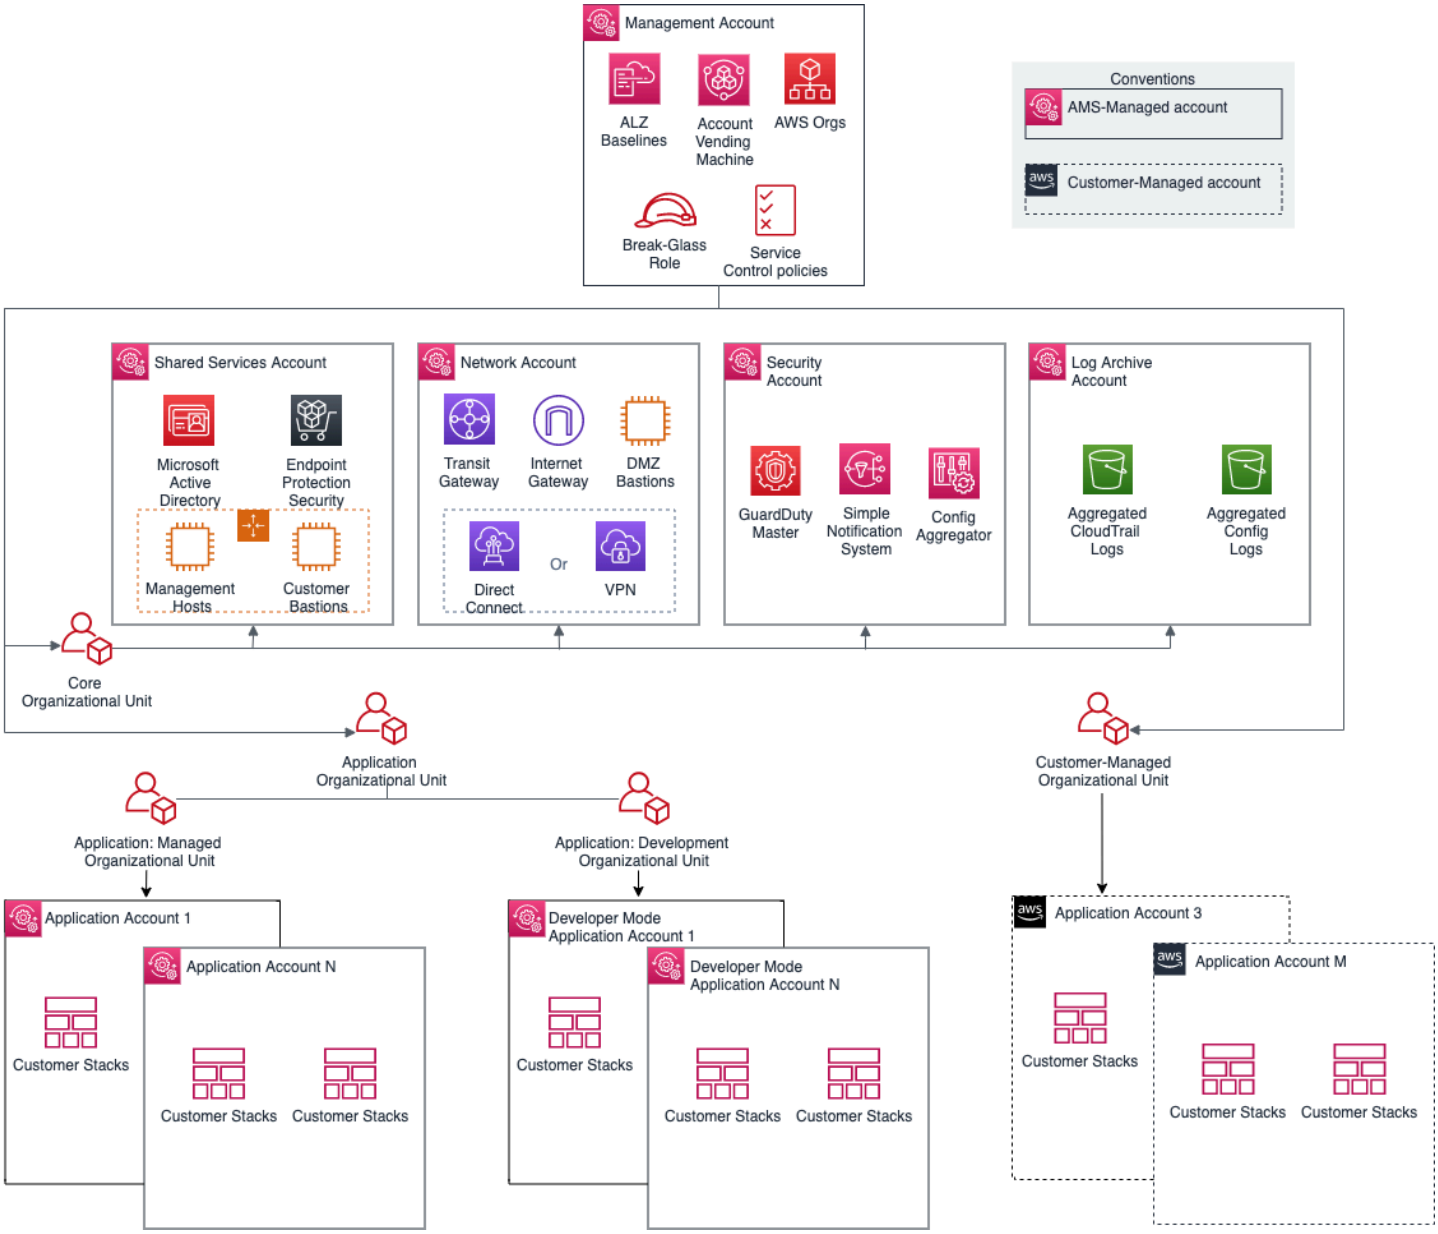

AMS leverages the foundational OUs based on AWS best practices as a way to logically manage accounts using Service Control Policies (SCPs). This serves as a way to enforce the governance framework with each AMS mode. Any governance and security guardrails (in the form of SCPs) applied to the foundational OUs also get applied to the custom/child OUs automatically. Additional SCPs can be requested for the child OUs. It is important to understand that application accounts are not the same as modes. Modes are applied to the infrastructure provisioned within the accounts and define the operational responsibilities between AMS and customers.

Figure 1: MALZ architecture and associated AMS modes

AMS Modes	Default Governance controls (Guardrails)		Support for Customer Added Governance Controls
	Preventative Controls	Detective Controls	
AMS Managed – Standard CM Mode and OOD			Yes (Restrictive)*
AMS Managed - Direct Change Mode (DCM)			Yes (Restrictive)*
AMS Managed – AWS Service Catalog			Yes (Restrictive)*
AMS Managed – Self Service Provisioning (SSP)			Yes (Restrictive)*
AMS Managed – Developer Mode			Yes
Customer Managed			Yes

Note

"Restrictive" implies that you can request custom policies for these OUs, they are approved by AMS on a case-by-case basis to ensure they don't interfere in AMS's capabilities to provide operational excellence. For a detailed list of AMS guardrails see AMS Guardrails in the user guide.

AMS modes and applications or workloads

Consider operational and governance requirements for your applications when selecting the right mode, either by requesting a new application account or hosting the application in an existing application account. The selection of the appropriate AMS mode for each application or workload depends on the following factors:

- The type of SDLC lifecycle function that the environment will provide (e.g., sandbox with unmoderated changes, UAT with some frequent changes, production with minimal changes and highly regulated)

- The governance policies needed (enforced through SCPs at the OU level)
- Operational Model (if you want to own the operational responsibility or want to outsource that to AMS)
- The desired business outcomes, like time to operate in the cloud, and cost of operations.

> **Note**
>
> For a descriptions of the mode types per AMS service, see Types of modes and accounts in AMS.
> For real-world use cases of the different modes, see Real world use cases for AMS modes

The following table outlines key considerations for application owners to help decide on the most suitable AMS mode. Application owners should include an assessment phase ahead of application migration to fully understand which mode applies to their specific application. Example: For applications based on cloud-native services or serverless architecture, the best option could be to start building and iterating in Developer mode and deploy the final Infrastructure as Code using AMS Managed – SSP mode. In this case light re-factoring may be required to ensure that any CloudFormation templates created for automated deployment meet the ingest guidelines laid out by AMS. Additionally, any IAM permissions need to be approved by AMS Security to ensure they follow the least privilege model.

The AMS mode selected to host the application, can help enable you to build towards you desired cloud operating model.

> **Note**
>
> More than one cloud operating model can existing in a single AMS Managed Landing Zone based on the different AMS modes selected to host the applications.

Decision issues	Standard CM mode / OOD*	AWS Service Catalog	Direct Change mode	Self-service provisioning	Developer mode	Customer Managed
Operational readiness						
Logging, Monitoring and Event Management	AMS responsible for all managed infrastructure			Customer responsible for Self-Service Provisioned Services (SSP)	Customer responsible for resources provisioned using developer IAM role outside AMS CM system	Customer responsible
Continuity Management	AMS responsibility to execute backup plan selected by customer			Customer responsible for Self-Service Provisioned Services (SSP)	Customer responsible for resources provisioned using developer IAM role outside AMS CM system	
Instance Level Access Management	AMS-managed through one-way AD trust with on-prem domain. Requires managed infrastructure to join AMS domain			Not applicable	Customer responsible for resources provisioned using	

Decision issues	Standard CM mode / OOD*	AWS Service Catalog	Direct Change mode	Self-service provisioning	Developer mode	Customer Managed
					developer IAM role outside AMS CM system	
Security Management and Account Level Access Management	AMS responsibility for all managed accounts			AMS responsible for all managed accounts	Customer responsible for resources provisioned using developer IAM role outside AMS CM system	
Patch Management	AMS responsibility for all managed accounts			Customer responsible for Self-Service Provisioned Services (SSP)	Customer responsible for resources provisioned using developer IAM role outside AMS CM system	

Decision issues	Standard CM mode / OOD*	AWS Service Catalog	Direct Change mode	Self-service provisioning	Developer mode	Customer Managed
Change Management	AMS responsibility for all managed accounts			Customer responsible for Self-Service Provisioned Services (SSP)	Customer responsible for resources provisioned using developer IAM role outside AMS CM system	
Provisioning Management	Prescriptive and standardized for the provisioning options offered in AMS	Flexibility to directly use AWS service API for AWS Service Catalog following AMS prescriptive standards	Flexibility to directly use AWS service API following AMS prescriptive standards	Flexibility to directly use AWS service APIs for SSP services	Flexibility to directly use AWS service API for provisioning	

Decision issues	Standard CM mode / OOD*	AWS Service Catalog	Direct Change mode	Self-service provisioning	Developer mode	Customer Managed
Incident Management and Audit	AMS responsibile for all managed accounts				Customer responsible for resources provisioned using developer IAM role outside AMS Change Management System	
GuardRails and Shared infrastructure (Network) and Security Framework	Prescriptive and standardized leveraging AMS Core Accounts					Flexible and bespoke leveraging AMS Core Accounts
Application readiness						

Decision issues	Standard CM mode / OOD*	AWS Service Catalog	Direct Change mode	Self-service provisioning	Developer mode	Customer Managed
Application refactoring	Light refactoring is needed				Light refactoring is needed (if provisioned using AMS Standard CM)	No need for refactoring
Support for AWS services	Limited to what is supported by AMS					Not limited
Business considerations						
Time to operational readiness	Three to six months			6 months + dependent on customer application operations competencies		6-18 months dependent on customer infrastructure and application operations competencies
Costs	$$$$			$$$	$$	$

Decision issues	Standard CM mode / OOD*	AWS Service Catalog	Direct Change mode	Self-service provisioning	Developer mode	Customer Managed
Application examples	Webserver with 3 tier stack, apps with compliance and regulatory requirements			Webserver using API Gateway, containerized application leveraging ECS/EKS	Iterating /optimizing on Data Lake application that uses Lambda, Glue, Athena, etc	De-centralized accounts/ applications like sandbox, third party managed applications

*Operations On Demand (OOD) has an offering for customers using the Standard CM mode to manage their changes through dedicated resourcing. For more details, see the Operations on Demand catalog of offerings and talk to your cloud service delivery manager (CSDM).

> **Note**
>
> The price comparison between SSP mode and Developer mode assumes that the same AWS services are provisioned.

Comparing AMS Modes against business and IT objectives

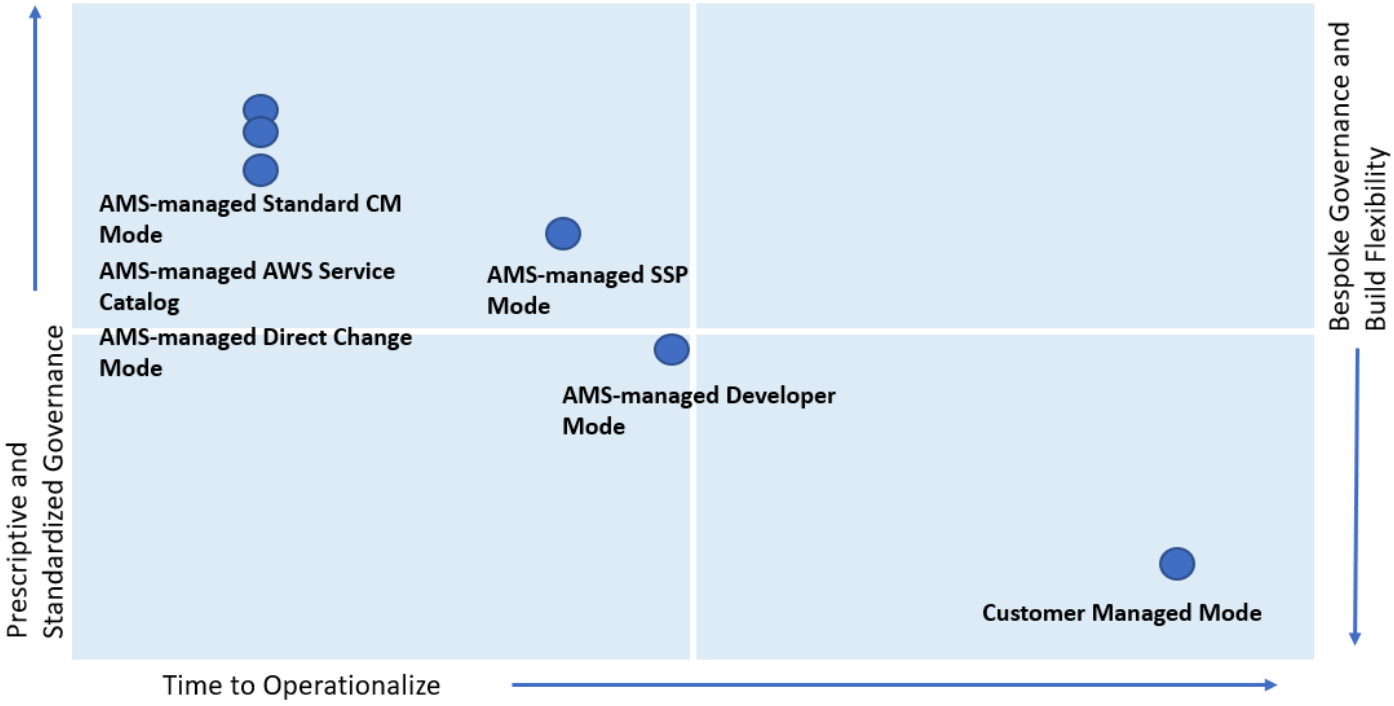

As shown, if you are looking for a highly controlled and standardized governance model for you applications, then AMS-managed Standard Change, AWS Service Catalog, or Direct Change modes are the best fit. If you require a bespoke governance model with a focus on application innovation without the need for operational readiness, select Customer Managed mode. With Customer Managed mode, it could take you a longer time to operationalize you applications as you bear the responsibility to establish people, processes, and tools to support operational capabilities such as Incident Management, Configuration Management, Provisioning Management, Security Management, Patch Management, etc.

Real world use cases for AMS modes

Examine these to help determine how to use AMS modes.

- **Use Case 1, business imperative to lower costs with a time-sensitive data center exit**: An enterprise with a compelling business event, like a data center exit, is interested in re-hosting their on-prem applications on the cloud. Most of the on-prem inventory consists of Windows and Linux servers with a mix of operating system versions. In doing so, the customer also wants to take advantage of cost savings that moving to the cloud offers and improving the technical and security posture of their applications. The customer wants to move fast but does not have the in-house cloud operations expertise built out yet. The customer has to find a balance of refactoring, too much refactoring can be risky against a tight timeline. However, with some refactoring,

like updating OS versions and optimizing databases, applications can achieve the next level of performance. In this example, the customer can select AMS-managed RFC mode to re-host most of their applications. AMS provides infrastructure operations, while also guiding the customer operations teams on best practices on securely operating in the cloud.

AMS-managed AWS Service Catalog and AMS-managed Direct Change mode gives the customer an extra flexibility while achieving the same business outcomes and objectives. In addition, the customer can use the AMS Operations On Demand (OOD) offering to have dedicated AMS operations engineers to prioritize the execution of requests for change (RFCs).

While offloading the undifferentiated infrastructure operational tasks (patching, backups, account management, etc) to AMS, the customer can continue to focus on optimizing their application and ramp-up their internal teams on cloud operations. AMS provides monthly reports to the customer on cost savings, and makes recommendations on resource optimizations. In this use case, if there were end-of-life applications hosted on legacy OS versions like Windows 2003 and 2008, that the customer decided not to re-factor, those can also be migrated to AMS and hosted in an account that leverages Customer Managed mode.

- **Use Case 2, building a data lake with Lambda, Glue, Athena within the secure AMS boundary**: An enterprise is looking to set up a Data Lake to meet the reporting needs for multiple applications in AMS. The customer wants to use S3 buckets for the storage of datasets and AWS Athena to query against the dataset for each report. S3 and AWS Athena will be deployed in separate AMS Managed accounts. The account with S3 also has other services like Glue, Lambda, and Step Functions to build a data ingestion pipeline. Glue, Lambda, Athena, and Step Functions are considered Self-Service Provisioning (SSP) services in this case. The customer also deployed an EC2 instance in the account that acts as an ad hoc tooling/scripting server. The customer starts by requesting AMS to enable the SSP services in their AMS Managed account. AMS provisions an IAM role for each service that the customer can assume, once the role is onboarded to the customer's federation solution. For ease of management, the customer can also combine the policies for the separate IAM roles into one custom role, alleviating the need to switch roles when working between the AWS services. Once the role is enabled in the account, the customer is able to configure the services as per their requirements. However, the customer must work with the AMS change management system to request additional permissions, depending on their use case.

  For example, for access to Glue Crawlers, additional permissions are needed by Glue. Additional permissions will also be needed to create event sources for Lambda. The customer will work with AMS to update IAM roles to allow cross-account access for Athena to query S3 buckets. Updates

to service roles or service-linked roles will also be needed through AMS change management for Lambda to call the Step Functions service, and Glue to read and write to all S3 buckets. AMS works with customers to ensure that the least-privilege access model is followed and the IAM changes requested are not overly permissive and opening up the environment to unnecessary risk. The customer's data lake team spends time planning for all IAM permissions needed for the services specific to the customer's architecture and requests AMS to enable them. This is because all IAM changes are processed manually and undergo review from the AMS Security team. Time to process these requests should be accounted for in the application deployment schedule.

As the SSP services are operational in the account, the customer can request support and report issues through AMS incident management and service requests. However, AMS will not actively monitor performance and concurrency metrics for Lambda, or job metrics for Glue. It is the customer's responsibility to ensure appropriate logging and monitoring is enabled for SSP services. The EC2 instance and S3 bucket in the account are fully managed by AMS.

- **Use Case 3, quick and flexible set up of a CICD deployment pipeline in AMS**: A customer is looking to set up a Jenkins-based CICD pipeline to deploy code pipeline to all application accounts in AMS. The customer may find it most suitable to host this CICD pipeline in the AMS-managed Direct Change mode (DCM) or AMS-managed Developer mode because it gives them flexibility to set up the Jenkins server with required custom configuration on EC2, with the desired IAM permissions to access CloudFormation and S3 buckets that host the artifact repository. While this can also be done in the AMS-managed RFC mode, the customer team would need to create multiple manual RFCs for IAM roles to iterate on the least permissive set of approved permissions, which are manually reviewed by AMS. DCM allows the customers to achieve their operational goals on AWS while avoiding the need to create multiple manual RFCs for IAM roles, when using AMS-managed RFC mode, to iterate on the least permissive set of approved permissions, which are manually reviewed by AMS. This would take time as well as education on the customer's part to ramp up AMS processes and tools. Working with Developer mode, the customer can start with a "developer role" to provision infrastructure using native AWS APIs. The quickest and most flexible way to set up this pipeline would be to use AMS Managed-Developer mode. Developer mode gives the quickest and easiest way, while compromising on operational integration, while DCM is less flexible but does provide the same level of operational support as RFC mode.
- **Use Case 4, bespoke operating model within the AMS foundation**: A customer is looking at a deadline-driven data center exit and one of their enterprise applications is fully managed by a third party MSP, including application operations and infrastructure operations. Assuming that the customer does not have time in the schedule to re-factor this application so that it

can be operated by AMS, Customer Managed mode is a suitable option. The customer can take advantage of the automated and quick set up of AMS managed Landing Zone. They can leverage the centralized account management that controls account vending and connectivity through the centralized networking account. It also simplifies their billing by consolidating charges for all customer managed accounts through the AMS Payer account. The customer has flexibility to set up their bespoke access management model with the MSP separate from standard access management used for AMS Managed accounts. This way, using Customer Managed mode, they can set up an AMS managed environment while meeting their business requirement of vacating their on-prem environment. In this case, if the customer also has Windows-based applications that they are migrating to the cloud, and choose to move them to a Customer Managed account, the customer is responsible for creating a cloud operating model. This can be complex, expensive, and time consuming depending on the customer's ability to transform traditional IT processes and train people. The customer can save time and cost by "lift and shift" of such workloads to an AMS Managed account and offload infrastructure operations to AMS.

Note

Customers may sometimes feel the need to move application accounts between the governance framework of RFC or SSP mode and Developer mode. For example, customers may host an application in AMS-managed mode as part of initial lift and shift migration, but overtime want to re-write the application to optimize it for cloud-native AWS services. They could change the mode of the pre-prod account from AMS-managed RFC to AMS-managed Developer mode, giving them the flexibility and agility for provisioning infrastructure. However, once infrastructure provisioning changes have been made using the "developer role", the same infrastructure cannot be moved back to AMS-managed RFC mode. This is because AMS cannot guarantee operations of infrastructure that was provisioned outside of the AMS change management system. Customers may need to create a new application account that offers AMS-managed RFC mode and then re-deploy the "optimized" infrastructure configuration through CloudFormation templates or custom AMIs ingested into an AMS-managed account. This is a clean way to deploy a production ready configuration. Once deployed, the application will be under prescriptive AMS governance and operations. The same applies to switching modes between Customer Managed mode and AMS-managed.

RFC mode

RFC mode is the AMS mode that provides the change management system, including requests for change (RFCs) and change types (CTs).

For example walkthroughs for each change type, see the "Additional Information" section for the relevant change type in the *AMS Advanced Change Type Reference* Change Types by Classification section.

> **Note**
>
> RFC mode is the new name for what was previously called "Change Management mode" or "Standard CM mode."

Topics

- Getting started with RFC mode
- What are RFCs?
- How RFCs work
- What are change types?
- Troubleshooting RFC errors in AMS

Getting started with RFC mode

The AMS RFC mode is the default mode for AMS Advanced operations plan customers. It includes a change management system with requests for change or RFCs and a catalog of change types to use to request the addition or change that you need to your accounts. This change management system provides a level of security in limiting who can make changes to your accounts.

For details on AMS Advanced change types, see What Are AMS Change Types?.

For details about onboarding to AMS Advanced, see AWS Managed Services Onboarding Introduction.

What are RFCs?

A request for change, or RFC, is how you make a change in your AMS-managed environment, or ask AMS to make a change on your behalf. To create an RFC, you choose from AMS change types,

choose RFC parameters (such as schedule), and then submit the request using either the AMS console or the API commands CreateRfc and SubmitRfc.

An RFC consists of two specifications, one for the RFC itself, and one for the change type (CT) parameters. At the command line, you can use an Inline RFC command, or a standard CreateRfc template in JSON format, which you fill out and submit along with the CT JSON schema file that you create (based on the CT parameters). The CT name is an informal description of the CT. A CSIO (category, subcategory, item, operation) is a more formal description of a CT. Only the CT ID must be specified when creating an RFC.

You're notified when the change has completed, successfully (Success) or unsuccessfully (Failure).

> **Note**
>
> For information about troubleshooting RFC failures, see Troubleshooting RFC errors in AMS.

The following graphic depicts the workflow of an RFC submitted by you.

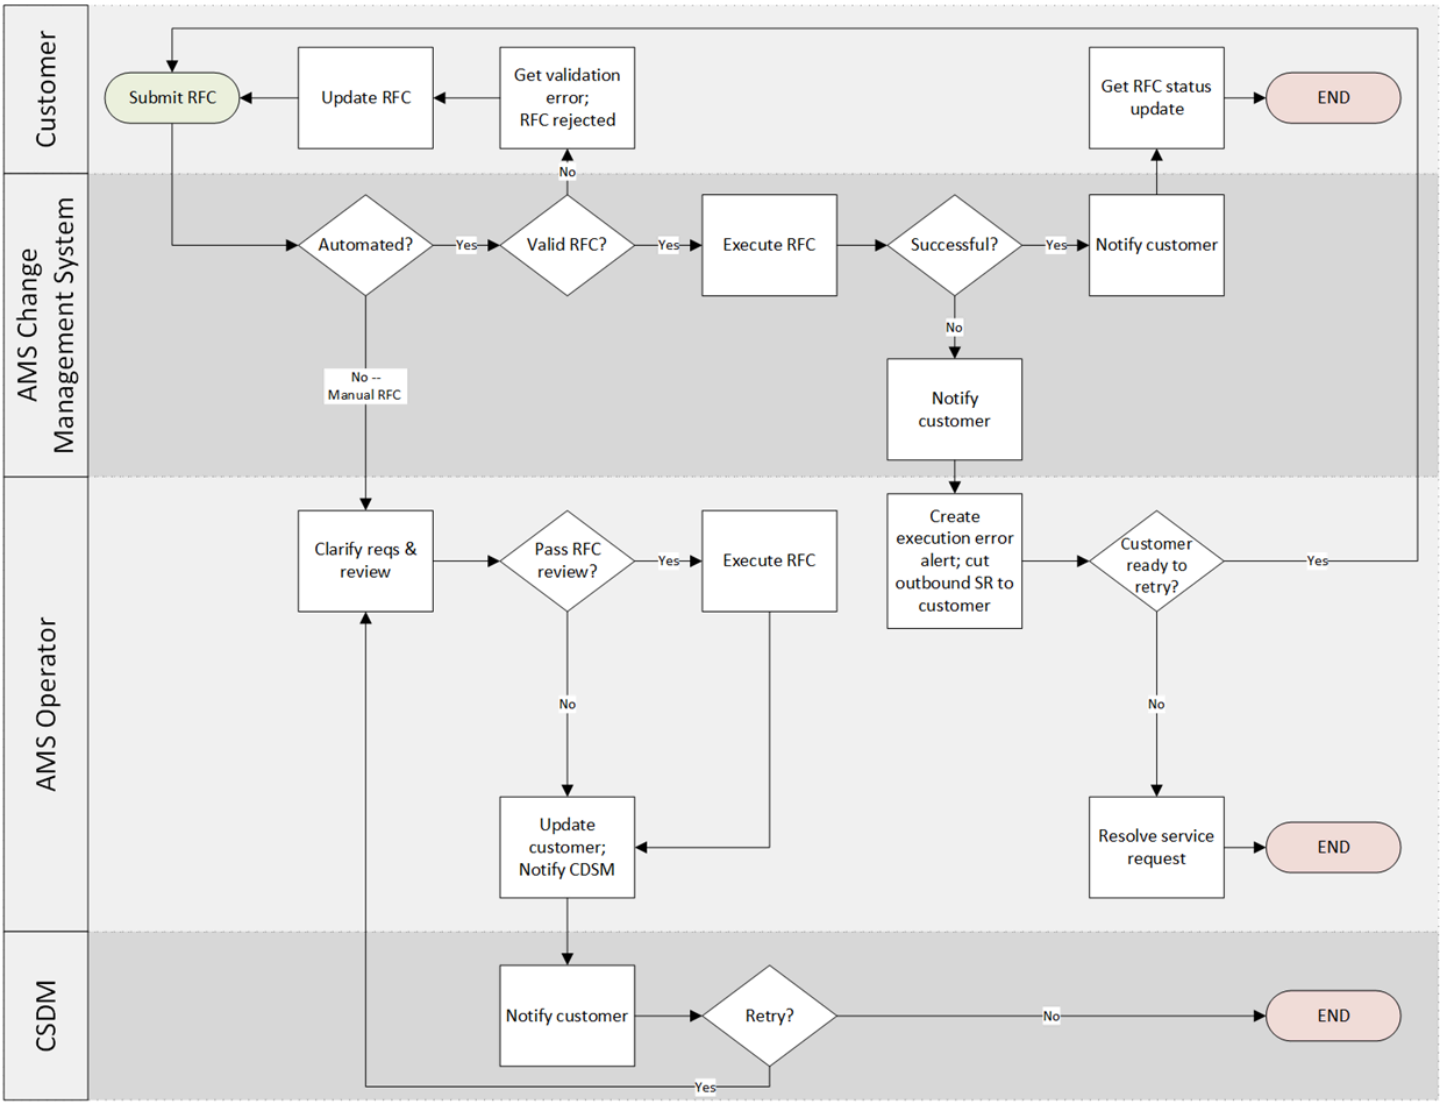

How RFCs work

In AWS Managed Services (AMS) requests for change, or RFCs, work in a two-fold manner. First, there are parameters required for the RFC itself (the options in the `CreateRfc` API); second, there are parameters required for the action of the RFC (the execution parameters). To learn about the `CreateRfc` options, see the CreateRfc section of the AMS API Reference. These options typically appear in the **Additional configurations** area of the Create RFC pages.

You can create and submit an RFC with the `CreateRfc` API, `aws amscm create-rfc` CLI, or using the AMS console Create RFC pages. For a tutorial on creating an RFC, see Creating an RFC.

Topics

- RFC authentication

- RFC security reviews
- RFC change type classification
- RFC action/activity states
- RFC status codes
- RFC update CTs and CloudFormation template drift detection
- RFC scheduling
- Approving or rejecting RFCs
- RFC restricted run periods
- Creating, cloning, updating, finding, and canceling RFCs
- RFCs and the AMS console
- RFC common parameters
- RFC daily email

RFC authentication

When using the AMS API/CLI you need to authenticate with temporary credentials. You can request temporary security credentials for federated users by calling GetFederationToken, AssumeRole, AssumeRoleWithSAML, or AssumeRoleWithWebIdentity AWS security token service (STS) APIs.

A common choice is SAML, which, once set up, means that you add an argument to each operation that you call. For example: `aws --profile saml amscm list-change-type-categories`.

A shortcut for SAML 2.0 profiles is to set the profile variable at the start of each API/CLI with `set AWS_DEFAULT_PROFILE=saml` (for Windows; for Linux it would be `export AWS_DEFAULT_PROFILE=saml`). For information about setting CLI environment variables, see Configuring the AWS Command Line Interface, Environment Variables.

RFC security reviews

The AWS Managed Services (AMS) change management approval process ensures that we perform a security review of changes we make in your accounts.

AMS evaluates all the requests for change (RFCs) against our AMS technical standards. Any change that could potentially lower your account's security posture by deviating from the technical standards, goes through a Security Review process where relevant risk is highlighted by AMS and, in case of a high or very high security risk, the RFC execution is accepted or rejected by your

authorized security personnel. All such changes are also evaluated to assess if there would be any adverse impact on AMS's ability to operate the account and, if so, additional reviews and approvals are required within AMS.

AMS technical standards

AMS Technical Standards define the minimum security criteria, configurations, and processes to establish the baseline security of your accounts. These standards must be followed by both AMS and you.

Any change that could potentially lower the security posture of your account by deviating from the technical standards, goes through a Risk Acceptance process, where relevant risk is highlighted by AMS and accepted or rejected by the authorized security personnel from your end. All such changes are also evaluated to assess if there would be any adverse impact on AMS's ability to operate the account and, if so, additional reviews and approvals are required within AMS.

RFC customer security risk management (CSRM) process

When someone from your organization requests a change to your managed environment, AMS reviews the change to determine whether the request might deteriorate the security posture of your account by falling outside the technical standards. If the request does lower the security posture of the account, AMS notifies your security team contact with the relevant risk, and executes the change; or, if the change introduces high or very high security risk in the environment, AMS seeks explicit approval from your security team contact in the form of risk acceptance (explained next). The AMS Customer Risk Acceptance process is designed to:

- Ensure risks are clearly identified and communicated to the right owners
- Minimize identified risks to your environment
- Obtain and document approval from the designated security contacts who understand your organization's risk profile
- Reduce ongoing operational overhead for identified risks

How to access technical standards and high or very high risks

We have made AMS Technical Standards documentation available for your reference in the https://console.aws.amazon.com/artifact/ as a report. Use the AMS Technical Standards documentation to understand whether a change would require risk acceptance from your authorized security contact prior to submitting a request for change (RFC).

Find the Technical Standards report by searching on "AWS Managed Services (AMS) Technical Standards" in the AWS Artifact **Reports** tab search bar after logging in with the default **AWSManagedServicesChangeManagementRole**.

> **Note**
>
> The AMS technical standard document is accessible for the Customer_ReadOnly_Role in single-account landing zone. In multi-account landing zone, the AWSManagedServicesAdminRole used by security admins and AWSManagedServicesChangeManagementRole used by application teams, can be used to access the document. If your team uses a custom role, create an Other | Other RFC to request access and we will update the specified custom role.

RFC change type classification

The change types that you use when submitting an RFC are divided into two broad categories: **Deployment** for creating resources, **Management** for updating or deleting resources. The **Management** category is also where you'll find change types for accessing instances, encrypting or sharing AMIs, and starting/stopping/rebooting/deleting stacks.

RFC action/activity states

`RfcActionState` (API) / **Activity State** (console) help you understand the status of human intervention, or action, on an RFC. Used primarily for manual RFCs, the `RfcActionState` helps you understand when there is action needed by either you or AMS operations, and helps you see when AMS Operations is actively working on your RFC. This provides increased transparency into the actions being taken on an RFC during its lifecycle.

`RfcActionState` (API) / **Activity State** (console) definitions:

- **AwsOperatorAssigned**: An AWS operator is actively working on your RFC.
- **AwsActionPending**: A response or action from AWS is expected.
- **CustomerActionPending**: A response or action from the customer is expected.
- **NoActionPending**: No action is required from either AWS or the customer.
- **NotApplicable**: This state can't be set by AWS operators or customers, and is used only for RFCs that were created prior to this functionality being released.

RFC action states differ depending on whether the change type submitted requires manual review and has scheduling set to **ASAP** or not.

- RFC **ActionState** changes during the review, approval, and start of a manual change type with deferred scheduling:
  - Once you submit a manual, scheduled, RFC, the **ActionState** automatically changes to **AwsActionPending** to indicate that an operator needs to review and approve the RFC.
  - When an operator begins actively reviewing your RFC, the **ActionState** changes to **AwsOperatorAssigned**.
  - When the operator approves your RFC, the RFC Status changes to Scheduled, and the **ActionState** automatically changes to **NoActionPending**.
  - When the scheduled start time of the RFC is reached, the RFC Status changes to **InProgress**, and the **ActionState** automatically changes to **AwsActionPending** to indicate that an operator needs to be assigned for review of the RFC.
  - When an operator begins actively running the RFC, they change the **ActionState** to **AwsOperatorAssigned**.
  - Once completed, the Operator closes the RFC. This automatically changes the **ActionState** to **NoActionPending**.

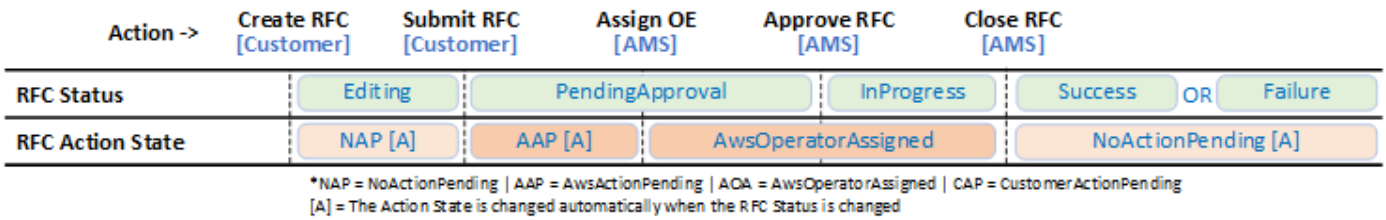

Important

- Action states can't be set by you. They are either set automatically based on changes in the RFC, or set manually by AMS operators.
- If you add correspondence to an RFC, the **ActionState** is automatically set to **AwsActionPending**.
- When an RFC is created, the **ActionState** is automatically set to **NoActionPending**.
- When an RFC is submitted, the **ActionState** is automatically set to **AwsActionPending**.
- When an RFC is Rejected, Canceled, or completed with a status of Success or Failure, the **ActionState** is automatically reset to **NoActionPending**.

- Action states are enabled for both automated and manual RFCs, but mostly matter for manual RFCs because those type of RFCs often require communications.

RFC action states usage examples

Use Case: Visibility on Manual RFC Process

- Once you submit a manual RFC, the RFC action state automatically changes to `AwsActionPending` to indicate that an operator needs to review and approve the RFC. When an operator begins actively reviewing your RFC, the RFC action state changes to `AwsOperatorAssigned`.
- Consider a manual RFC that has been approved and scheduled and is ready to begin running. Once the RFC status changes to `InProgress`, the RFC action state automatically changes to `AwsActionPending`. It changes again to `AwsOperatorAssigned` once an operator starts actively running the RFC.
- When a manual RFC is completed (closed as "Success" or "Failure"), the RFC Action state changes to `NoActionPending` to indicate that no further actions are necessary from either the customer or operator.

Use case: RFC correspondence

- When a manual RFC is `Pending Approval`, an AMS Operator might need further information from you. Operators will post a correspondence to the RFC and change the RFC action state to `CustomerActionPending`. When you respond by adding a new RFC correspondence, the RFC action state automatically changes to `AwsActionPending`.
- When an automated or manual RFC has failed, you can add a correspondence to the RFC details, asking the AMS Operator why the RFC failed. When your correspondence is added, the RFC action state is automatically set to `AwsActionPending`. When the AMS operator picks up the RFC to view your correspondence, the RFC action state changes to `AwsOperatorAssigned`. When the operator responds by adding a new RFC correspondence, the RFC action state may be set to `CustomerActionPending`, indicating that there is another response from the customer expected, or to `NoActionPending`, indicating that no response from the customer is needed or expected.

RFC status codes

RFC status codes help you track your requests. You can observe these status codes during an RFC run in the CLI output, or by refreshing the RFC list page in the console.

You can also see the codes for an RFC on the details page for that RFC, which might look like this:

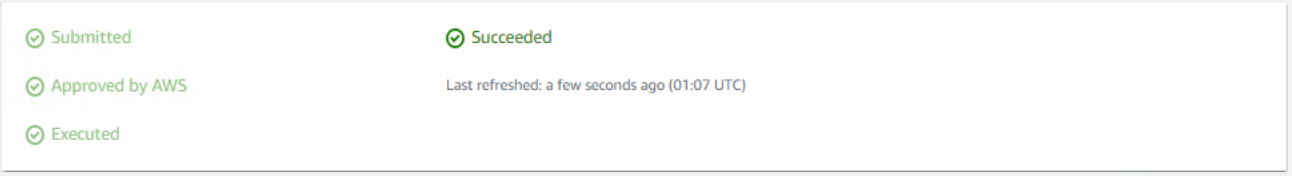

You may sometimes see an RFC in your list that you did not submit. When AMS operators use an internal-only CT, they submit it in an RFC and it displays in your RFC list. For more information, see Internal-only change types.

Important

You can request notifications of RFC state changes. For details, see RFC State Change Notifications.

RFC status codes

Success	Failure
Editing: the RFC has been created but not submitted	Rejected: RFCs are rejected typically because they fail validation; for example, an unusable resource, i.e. a subnet, is specified
PendingApproval / Submitted: The RFC has been submitted and the system is determining if it requires approval, and obtaining that approval, if required	Canceled: RFCs are canceled typically because they do not pass validation before the configured start time has passed
Approved by AWS / Approved by customer: the RFC has been approved. Automated RFCs are approved by AWS, manual RFCs are approved by Operators and, sometimes, customers	Failure: The RFC has failed; see the StatusReason in the output for failure reasons, and AMS operations automatically creates a trouble ticket and communicates with you as needed

Success	Failure
Scheduled: the RFC has passed syntax and requirement checks and is scheduled for running InProgress: the RFC is being run, note that RFCs that provision multiple resources or have long-running UserData, take longer to run Executed: The RFC has been run Success / Succeeded: The RFC has been successfully completed	

> **Note**
>
> Canceled or rejected RFCs can be re-submitted using UpdateRfc; see also Updating RFCs.

If the RFC passes all the necessary conditions (for example, all required parameters are specified), the status changes to `PendingApproval` (even automated CTs require approval, which happens automatically if syntax and parameter checks pass). If it does not pass, the status changes to `Rejected`. The `StatusReason` provides information about rejections; the `ExecutionOutput` fields provide information about approval and completion. Error codes include:

- InvalidRfcStateException: The RFC is in a status that doesn't allow the operation that was called. For example, if the RFC has moved to the Submitted state, it can no longer be modified.
- InvalidRfcScheduleException: The StartTime, EndTime, or TimeoutInMinutes parameters were breached.
- InternalServerError: A difficulty with the system was encountered.
- InvalidArgumentException: A parameter is incorrectly specified; for example, an unacceptable value is used.
- ResourceNotFoundException: A value, such as the stack ID, cannot be found.

If the scheduled requested start and end times (also known as the change run window) occur before the change is approved, the RFC status changes to `Canceled`. If the change is approved, the RFC status changes to `Scheduled`. The change run window for ASAP RFCs is the submitted time plus the `ExpectedExecutionDuration` value for the CT.

At any time before the arrival of the change run window, a scheduled change (submitted with a `RequestedStartTime` in the CLI) can be modified or canceled. If the scheduled change is modified, it must then be re-submitted.

When the change start time arrives (scheduled or ASAP) and after approvals are complete, the status changes to `InProgress` and no modifications can be made. If the change is completed within the specified change run window, the status changes to `Success`. If any part of the change fails, or if the change is still in progress when the change run window ends, the status changes to `Failure`.

> **Note**
>
> During the `InProgress`, `Success`, or `Failure` change states, the RFC cannot be modified or canceled.

The following diagram illustrates the RFC statuses from the CreateRFC call through to resolution.

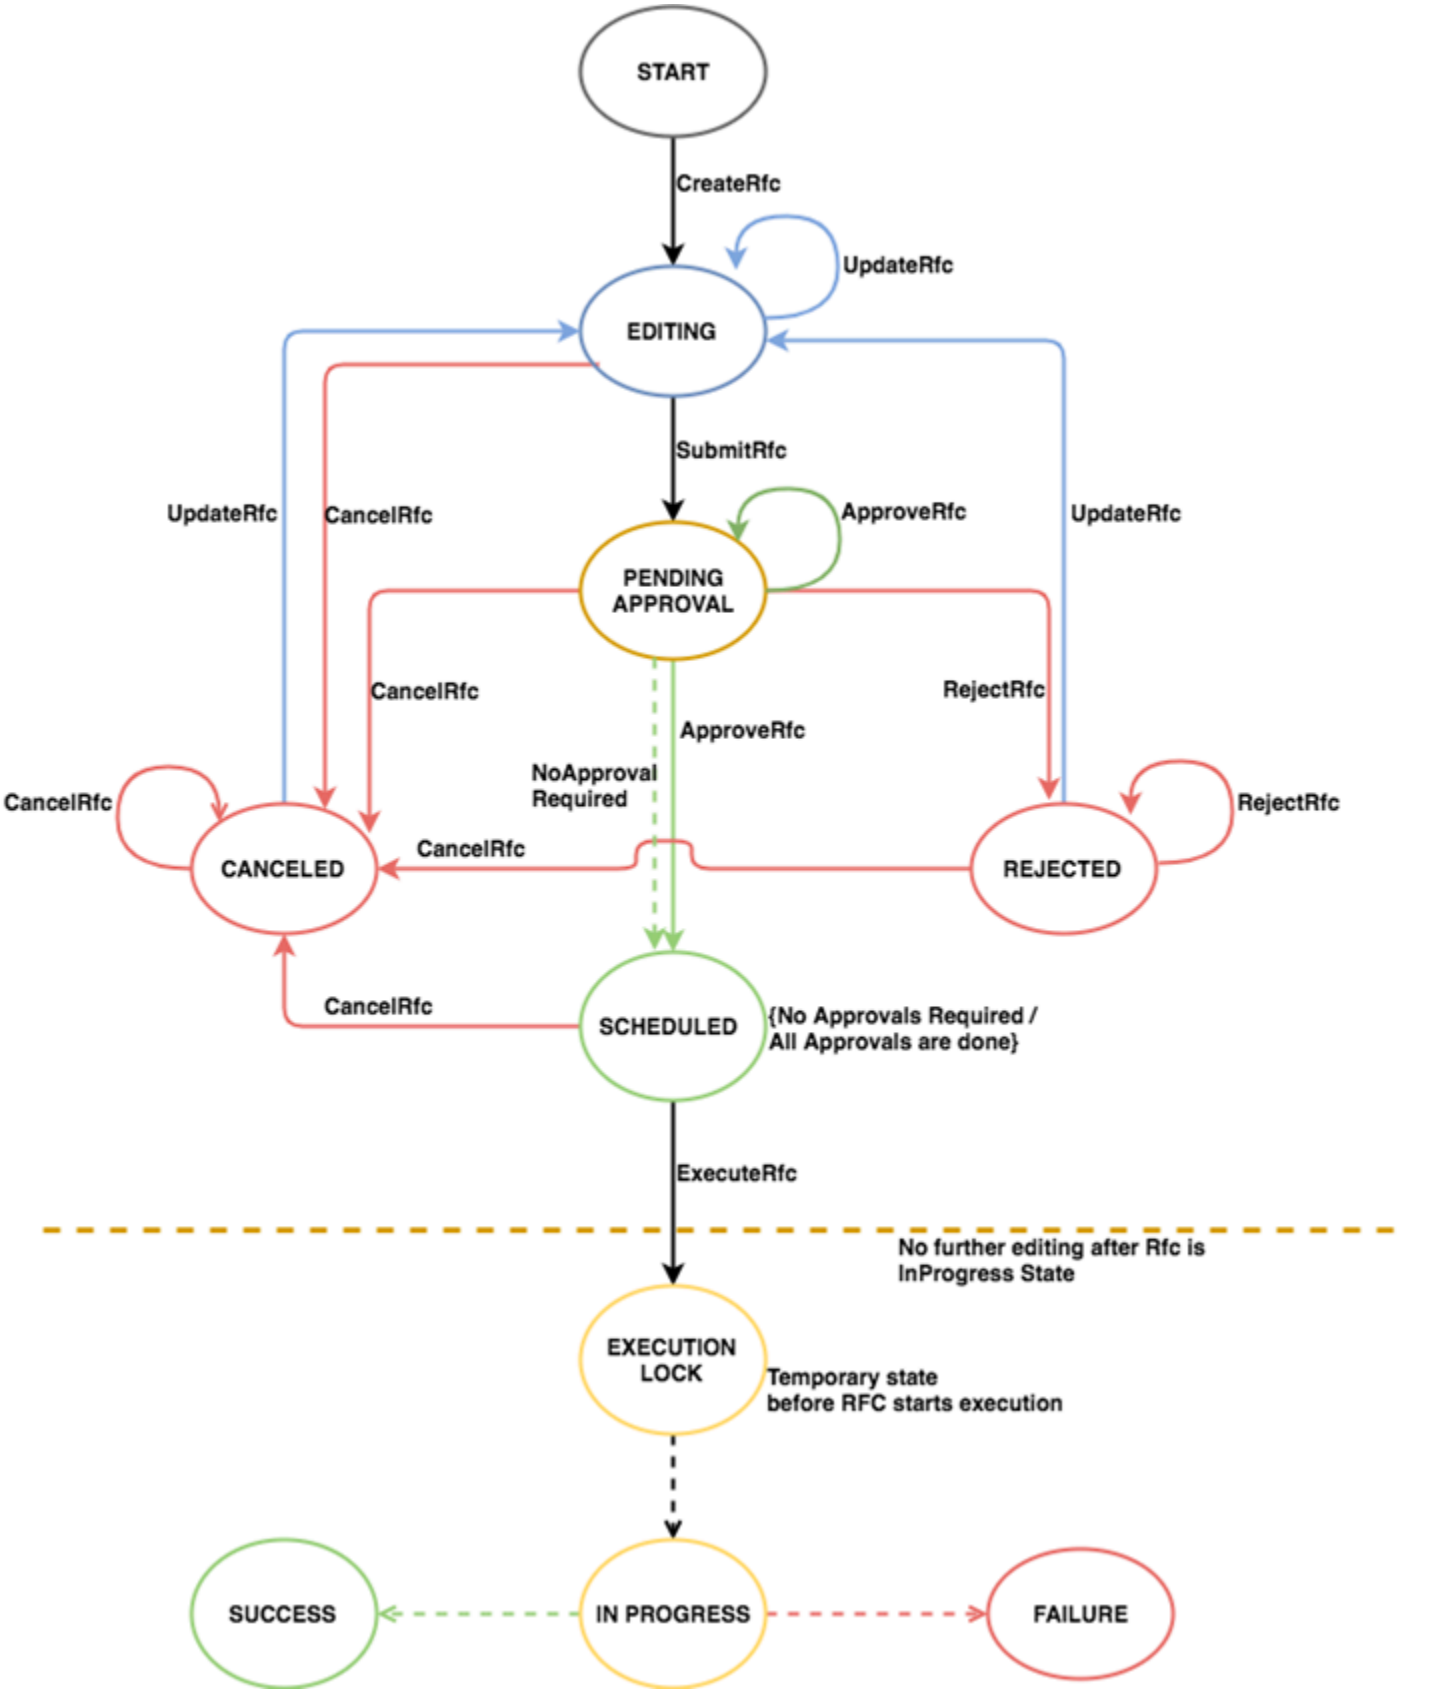
START
CreateRfc
UpdateRfc
EDITING
SubmitRfc
UpdateRfc
CancelRfc
ApproveRfc
UpdateRfc
PENDING APPROVAL
CancelRfc
RejectRfc
ApproveRfc
NoApproval Required
CancelRfc
CANCELED
CancelRfc
REJECTED
RejectRfc
CancelRfc
SCHEDULED
{No Approvals Required / All Approvals are done}
ExecuteRfc
No further editing after Rfc is InProgress State
EXECUTION LOCK
Temporary state before RFC starts execution
SUCCESS
IN PROGRESS
FAILURE

RFC update CTs and CloudFormation template drift detection

Resources provisioned in AMS use a modified CloudFormation template. If a resource has a parameter changed directly through an AWS Management Console, the CloudFormation creation record of that resource becomes out of- sync. If this happens and if you should attempt to use an AMS update change type to update the resource in AMS, the AMS system references the original resource configuration and potentially resets changed parameters. This reset could be damaging, so AMS disallows RFCs with update change types if any extra-AMS configuration changes are detected.

For a list of update change types, use the console filter.

Drift remediation FAQs

Questions and answers on AMS drift remediation. There are two change types that you can use to initiate drift remediation, one is execution mode=manual or "review required," the other is execution mode=automated.

Drift remediation supported resources (ct-3kinq0u4l33zf)

These are the resources that are supported by the drift remediation change type, (ct-3kinq0u4l33zf). All resources are supported for the "review required" (ct-34sxfo53yuzah) change type.

```
AWS::EC2::Instance
AWS::EC2::SecurityGroup
AWS::EC2::VPC
AWS::EC2::Subnet
AWS::EC2::NetworkInterface
AWS::EC2::EIP
AWS::EC2::InternetGateway
AWS::EC2::NatGateway
AWS::EC2::NetworkAcl
AWS::EC2::RouteTable
AWS::EC2::Volume
AWS::AutoScaling::AutoScalingGroup
AWS::AutoScaling::LaunchConfiguration
AWS::AutoScaling::LifecycleHook
AWS::AutoScaling::ScalingPolicy
AWS::AutoScaling::ScheduledAction
AWS::ElasticLoadBalancing::LoadBalancer
AWS::ElasticLoadBalancingV2::Listener
```

```
AWS::ElasticLoadBalancingV2::ListenerRule
AWS::ElasticLoadBalancingV2::LoadBalancer
AWS::CloudWatch::Alarm
```

Drift remediation change types

Questions and answers on using the AMS drift remediation change types.

For a list of supported resources for the drift remediation feature, see Drift remediation supported resources (ct-3kinq0u4l33zf).

> **Important**
>
> Drift remediation modifies the stack template and/or parameters and it is mandatory to update your local template repositories or any automation that is updating these stacks to use the latest stack template and parameters. Using old template and/or parameters without syncing can cause damaging changes to underlying resources.
> The no review required, automated, CT (ct-3kinq0u4l33zf) supports remediating only 10 resources per RFC. To remediate remaining resources in batches of 10 create new RFCs until all resources are remediated.

Which drift remediation change type should I use?

We recommend using the **no review required**, automated CT (ct-3kinq0u4l33zf) when:

- You attempt to perform an update to an existing stack resource using an automated CT and the RFC gets rejected as the stack is `DRIFTED`.
- You used an Update CT in the past and it failed as the stack was DRIFTED. You do not need to attempt an update again and can use the review required, manual, CT instead.

We recommend using the **review required**, manual CT (ct-34sxfo53yuzah) only when drifted resource types are not supported by the drift remediation no review required, automated, CT (ct-3kinq0u4l33zf), or when the drift remediation no review required, automated, CT fails.

What changes are performed to the stack during remediation?

Remediation requires updates to the stack template and/or parameters depending on the properties that are drifted. Remediation also updates the stack policy of the stack during remediation and restores the stack policy to its previous value once remediation is completed.

How can we see the changes performed to the stack template and/or parameters?

In the response to the RFC, a change summary is provided with the following information:

- `ChangeSummaryJson`: Contains change summary of Stack Template and/or Parameters as part of drift remediation. Remediation is performed in multiple phases. This change summary consists of changes for individual phases. If Remediation is successful check changes of the last phase. See ExecutionPlan in the JSON for phases executed in order. For example, RestoreReferences section when present is always executed at the end and contains JSON for post remediation changes. If remediation is run in DryRun mode none of these changes would have been applied to the stack.
- `PreRemediationStackTemplateAndConfigurationJson`: Contains configuration snapshot of CloudFormation Stack including Template, Parameters, Outputs, StackPolicyBody before remediation was triggered on the stack.

What do I need to do once remediation is performed?

> **Important**
>
> You need to update your local template repositories, or any automation, that would be updating the remediated stack, with the latest template and parameters provided in the RFC summary. It is very important to do this because using the old template and/or parameters can cause further destructive changes on the stack resources.

Will my application be effected during this remediation?

Remediation is an offline process that is performed only on the CloudFormation stack configuration. No updates are performed on the underlying resource.

Can I continue using Management | Other | Other RFCs to perform updates to resources after remediation?

We recommend that you always perform updates to stack resources using the available automated Update CTs. When the available Update CTs do not support your use case, use Management | Other | Other requests.

Does remediation create any new resources in the stack?

Remediation does not create any new resources in the stack. However, remediation creates new outputs and updates the stack template metadata section to store the remediation summary for your reference.

Will remediation always be successful?

Remediation requires careful analysis and validation of the template configuration to determine if it can be performed. In scenarios where these validations fail, the remediation process is stopped and no changes are performed to the stack template or parameters. Also, remediation can only be performed on supported resource types.

How can I perform updates to stack resources if remediation is not successful?

You can use the Management | Other | Other | Update CT (ct-0xdawir96cy7k) to request changes. AMS monitors such scenarios and works towards improving the remediation solution.

Can I remediate stacks that have both supported and unsupported resource types?

Yes. However, remediation is performed only if the supported resource types are found DRIFTED in the stack. If any unsupported resource types are DRIFTED, remediation does not continue.

Can I request remediation for stacks created through non-CFN Ingest CTs?

Yes. Remediation can be performed on stacks irrespective of the change type used for creating the stack.

Can I know the changes that would be performed to the stack before remediation?

Yes. Both change types provide a **DryRun** option that you can use to request changes that would be performed if the stack was remediated. However, the final remediation changes may differ depending on the drift present on the stack at the time of remediation.

RFC scheduling

AMS has a feature, **Scheduling**, that enables you to choose an RFC start time that is as soon as the RFC is approved, or schedule an RFC. The **Scheduling** option has two possible choices:

- **Execute this change ASAP**: This choice directs AMS to run the RFC as soon as it has been approved. Most CTs are automatically approved, so using this option is best if you're not submitting an RFC that you want to start at a specific time.
- **Schedule this change**: This choice sets a day, time, and timezone for the implementation of the RFC to begin. We recommend requesting a start time that is at least 10 minutes after you plan to submit the RFC for automated change types and at least 24 hours after you plan to submit the RFC for "review required" change types. If the RFC is not approved by the time the configured start time is reached, the RFC is rejected.

To set an RFC schedule:

1. For **Execute this change ASAP**:
   - Console: Do nothing, it is the default RFC schedule.
   - API/CL: Remove the `RequestedStartTime` and `RequestedEndTime` options in the Create RFC operation.

   ASAP "review required" RFCs are auto-rejected if they are not approved within thirty days of submission.
2. For **Schedule this change**:
   - Console: Select the radio button to open a **Start time** area for specifying the schedule. Then either type in, or use the Calendar widget, to pick a day. Also enter a time, in UTC, expressed in ISO 8601 format, and use the drop-down list to pick a location. By default, AMS uses the ISO 8601 format (YYYYMMDDThhmmssZ or YYYY-MM-DDThh:mm:ssZ, either format is accepted.
   - API/CLI: Submit values for the `RequestedStartTime` and `RequestedEndTime` parameters in the Create RFC operation. Passing a configured `RequestedEndTime` doesn't stop the run for an automated change type that has already started. For a "review required" change type, if the `RequestedEndTime` is reached while AMS Operations research is still ongoing, and you're in communication with AMS, you can request an extension, or you might be asked to re-submit the RFC.

> **Tip**
>
> A good UTC time readout can be found at Time.is. Example ISO 8601 format for a date/time value of 2016-12-05 at 2:20pm: **2016-12-05T14:20:00Z** or **20161205T142000Z**.

> **Note**
>
> To request a call back during the scheduled Manual (MOO) RFC's execution window, add a note on the execution parameters of the RFC, along with bridge information, if you have one. If you don't have a bridge, include a request for one in your note. Also, use RFC bi-directional communication or Live Contact features to communicate in the absence of an engineer in the bridge at the time of the scheduled window.

If you provide...

- only a `RequestedStartTime`, the RFC is considered scheduled and the `RequestedEndTime` is populated using the `ExecutionDurationInMinutes` value.
- only a `RequestedEndTime`, we throw an InvalidArgumentException.
- both `RequestedStartTime` and `RequestedEndTime`, we overwrite the `RequestedEndTime` with the specified start time plus the `ExecutionDurationInMinutes` value.
- neither `RequestedStartTime` nor `RequestedEndTime`, we keep those values as null and the RFC is treated as an ASAP RFC.

> **Note**
>
> For all scheduled RFCs, an unspecified end time is written to be the time of the specified `RequestedStartTime` plus the `ExpectedExecutionDurationInMinutes` attribute of the submitted change type. For example, if the `ExpectedExecutionDurationInMinutes` is "60" (minutes), and the specified `RequestedStartTime` is `2016-12-05T14:20:00Z` (December 5, 2016 at 4:20 AM), the actual end time would be set to December 5, 2016 at 5:20 AM. To find the `ExpectedExecutionDurationInMinutes` for a specific change type, run this command:
>
> ```
> aws amscm --profile saml get-change-type-version --
> change-type-id CHANGE_TYPE_ID --query "ChangeTypeVersion.
> {ExpectedDuration:ExpectedExecutionDurationInMinutes}"
> ```

RFC priority

Use the **Priority** option in `execution mode = manual` change types to alert AMS Operations to the urgency of the request.

Priority option in `execution mode = manual`:

You can specify the priority of a manual RFC as **High**, **Medium**, or **Low**. RFCs classified as **High** are reviewed and approved prior to RFCs classified as **Medium**, subject to RFC service level objectives (SLOs) and their submission times. RFCs with **Low** priority or no priority specified are processed in the order they are submitted.

Approving or rejecting RFCs

RFCs submitted with approval-required (manual) CTs are subject to approval by you or AMS; pre-approved CTs, conversely, are automatically processed. For more information, see CT approval requirements.

> **Note**
>
> When using "review required" CTs, AMS recommends that you use the ASAP **Scheduling** option (choose **ASAP** in the console, leave start and end time blank in the API/CLI) as these CTs require an AMS operator to examine the RFC, and possibly communicate with you before it can be approved and run. If you schedule these RFCs, be sure to allow at least 24 hours. If approval does not happen before the scheduled start time, the RFC is rejected automatically.

If an approval-required RFC has been successfully submitted by AMS, it must be explicitly approved by you (if you submit an approval-required RFC, it must be approved by AMS). If you're required to approve an RFC that AMS submitted, an email or other pre-determined communication is sent to you, requesting the approval and including the RFC ID. The following actions can be taken at that point:

- Console Approve or Reject: Use the RFC details page for the relevant RFC:

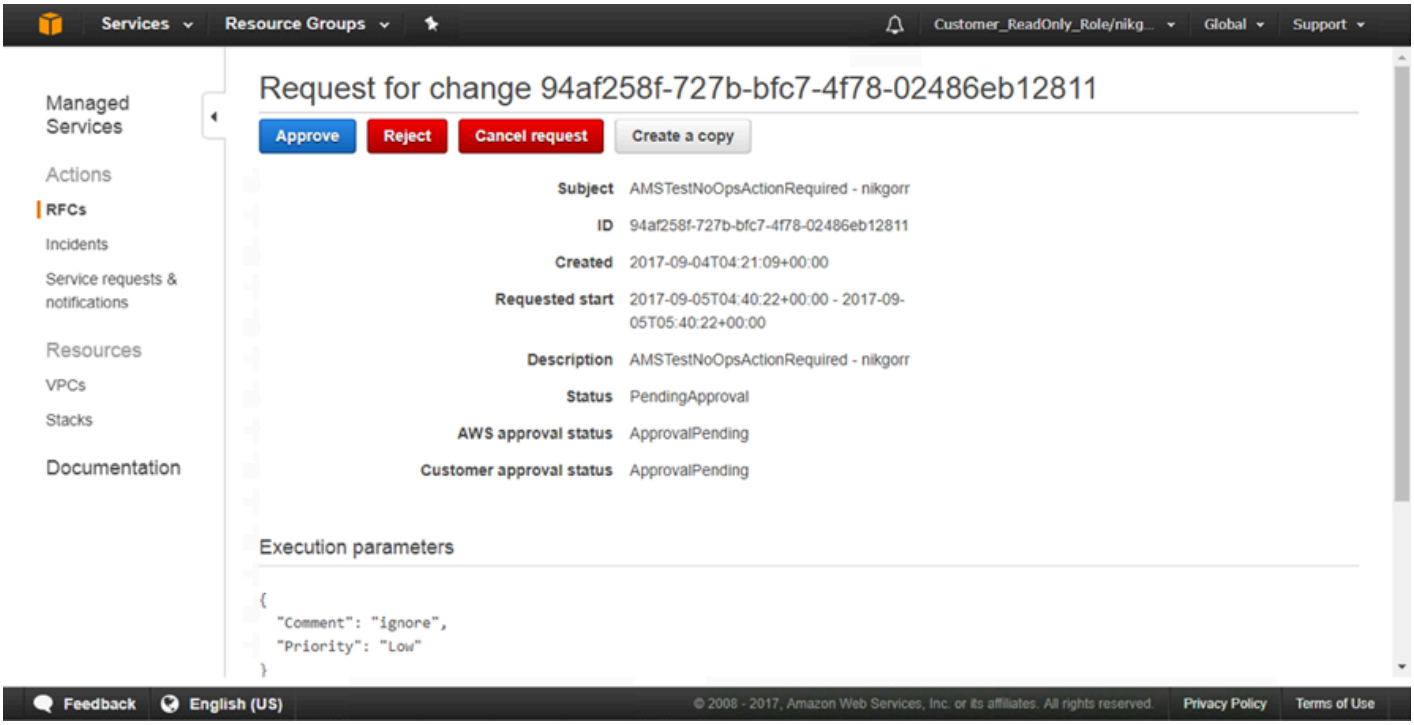

- API / CLI Approve: ApproveRfc marks a change as approved. The action must be taken by both the owner and operator, if both are required. The CLI command would look like this (with the appropriate RFC ID):

```
aws amscm approve-rfc --rfc-id RFC_ID
```

- API / CLI Reject: RejectRfc marks a change as rejected. The CLI command would look like this:

```
aws amscm reject-rfc --rfc-id RFC_ID --reason "no longer relevant"
```

RFC restricted run periods

Formerly known as blackout days, you can request that certain time periods be restricted so that no changes can be run during those times.

To set a restricted run period, use the UpdateRestrictedExecutionTimes API operation and set a specific time period, in UTC. The period that you specify overrides any previous periods that were specified. If you submit an RFC during the specified restricted run time, submission fails with the error Invalid RFC Schedule. You can specify up to 200 restricted time periods. By default, no restricted period is set. Example (with SAML authentication configured):

```
aws amscm  --profile saml update-restricted-execution-times --restricted-execution-
times="[{\"TimeRange\":{\"StartTime\":\"2018-01-01T12:00:00Z\",\"EndTime\":
\"2018-01-01T12:00:01Z\"}}]"
```

You can also view your current RestrictedExecutionTimes setting by running the ListRestrictedExecutionTimes API operation. Example:

```
aws amscm  --profile saml list-restricted-execution-times
```

If you want to submit an RFC during a specified restricted execution time, you can add the RestrictedExecutionTimesOverrideId with the value of OverrideRestrictedTimeRanges, and then submit RFC as you normally would. We recommend only using this for a critical or emergency RFC. For more information, see the API reference for SubmitRfc.

Creating, cloning, updating, finding, and canceling RFCs

These examples walk you through various RFC operations.

Topics

- Creating an RFC
- RFC cloning (re-creating) with the console
- Updating RFCs
- Finding RFCs
- Canceling RFCs

Creating an RFC

Creating an RFC with the console

The following is the first page of the RFC Create process in the AMS console, with **Quick cards** open and **Browse change types** active:

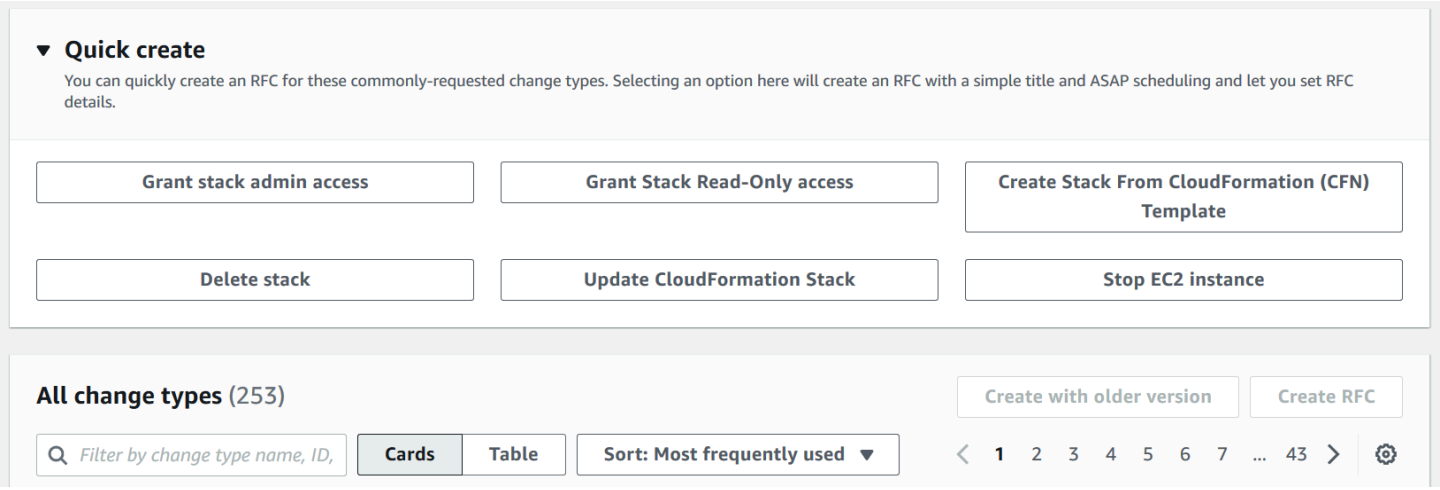

The following is the first page of the RFC Create process in the AMS console, with **Select by category** active:

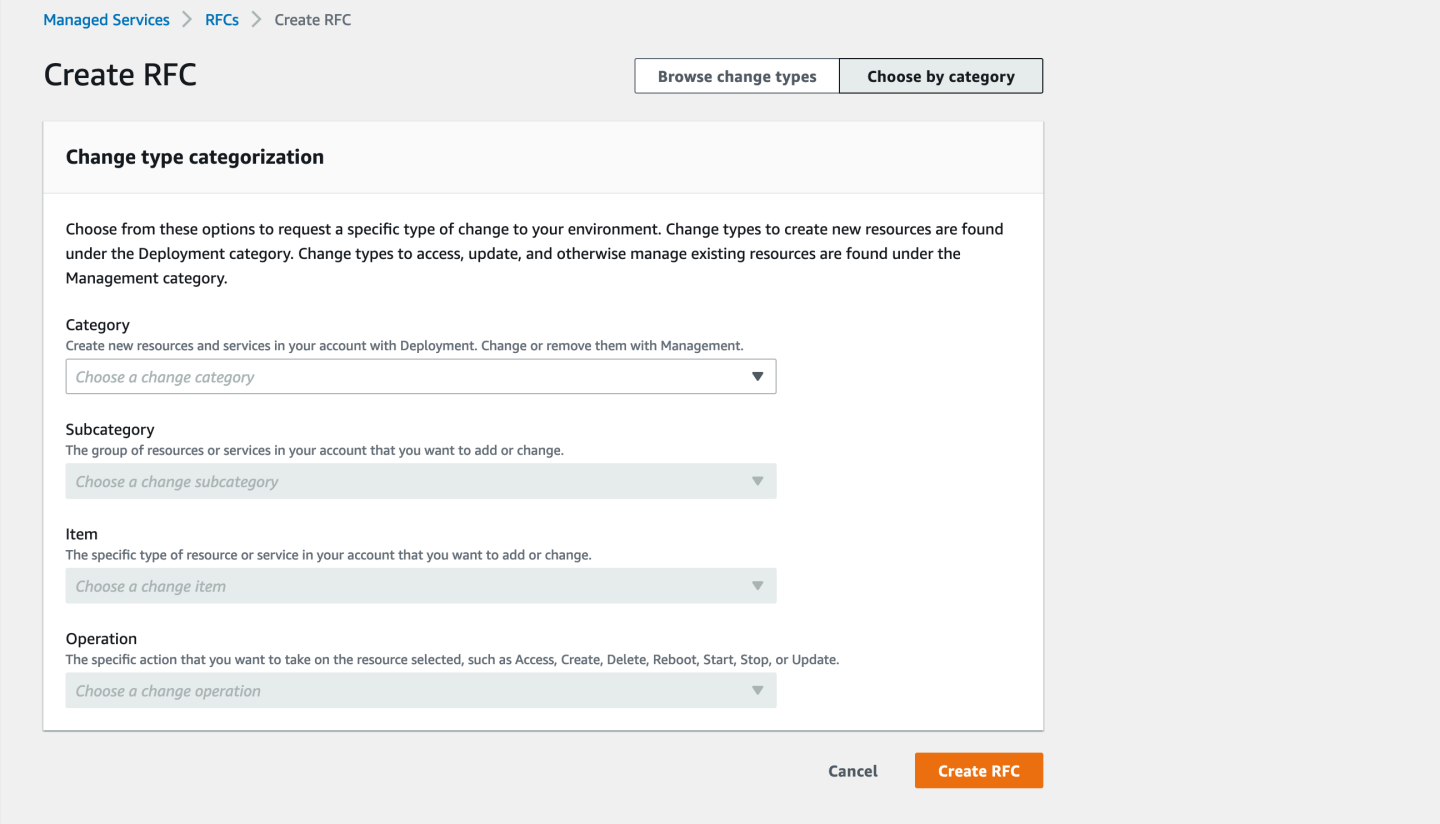

How it works:

1. Navigate to the **Create RFC** page: In the left navigation pane of the AMS console click **RFCs** to open the RFCs list page, and then click **Create RFC**.

2. Choose a popular change type (CT) in the default **Browse change types** view, or select a CT in the **Choose by category** view.
   - **Browse by change type**: You can click on a popular CT in the **Quick create** area to immediately open the **Run RFC** page. Note that you cannot choose an older CT version with quick create.

     To sort CTs, use the **All change types** area in either the **Card** or **Table** view. In either view, select a CT and then click **Create RFC** to open the **Run RFC** page. If applicable, a **Create with older version** option appears next to the **Create RFC** button.
   - **Choose by category**: Select a category, subcategory, item, and operation and the CT details box opens with an option to **Create with older version** if applicable. Click **Create RFC** to open the **Run RFC** page.
3. On the **Run RFC** page, open the CT name area to see the CT details box. A **Subject** is required (this is filled in for you if you choose your CT in the **Browse change types** view). Open the **Additional configuration** area to add information about the RFC.

   In the **Execution configuration** area, use available drop-down lists or enter values for the required parameters. To configure optional execution parameters, open the **Additional configuration** area.
4. When finished, click **Run**. If there are no errors, the **RFC successfully created** page displays with the submitted RFC details, and the initial **Run output**.
5. Open the **Run parameters** area to see the configurations you submitted. Refresh the page to update the RFC execution status. Optionally, cancel the RFC or create a copy of it with the options at the top of the page.

Creating an RFC with the CLI

How it works:

1. Use either the Inline Create (you issue a `create-rfc` command with all RFC and execution parameters included), or Template Create (you create two JSON files, one for the RFC parameters and one for the execution parameters) and issue the `create-rfc` command with the two files as input. Both methods are described here.
2. Submit the RFC: `aws amscm submit-rfc --rfc-id` *`ID`* command with the returned RFC ID.

   Monitor the RFC: `aws amscm get-rfc --rfc-id` *`ID`* command.

To check the change type version, use this command:

```
aws amscm list-change-type-version-summaries --filter
 Attribute=ChangeTypeId,Value=CT_ID
```

> **Note**
>
> You can use any `CreateRfc` parameters with any RFC whether or not they are part of the schema for the change type. For example, to get notifications when the RFC status changes, add this line, `--notification "{\"Email\": {\"EmailRecipients\" : [\"email@example.com\"]}}"` to the RFC parameters part of the request (not the execution parameters). For a list of all CreateRfc parameters, see the AMS Change Management API Reference.

INLINE CREATE:

Issue the create RFC command with execution parameters provided inline (escape quotes when providing execution parameters inline), and then submit the returned RFC ID. For example, you can replace the contents with something like this::

```
aws amscm create-rfc --change-type-id "CT_ID" --change-type-version "VERSION" --title
 "TITLE" --execution-parameters "{\"Description\": \"example\"}"
```

TEMPLATE CREATE:

> **Note**
>
> This example of creating an RFC uses the Load Balancer (ELB) stack change type.

1. Find the relevant CT. The following command searches CT classification summaries for those that contain "ELB" in the **Item** name and creates output of the Category, Item, Operation, and ChangeTypeID in table form (Subcategory for both is `Advanced stack components`).

   ```
   aws amscm list-change-type-classification-summaries --query
    "ChangeTypeClassificationSummaries[?contains(Item,'ELB')].
   [Category,Item,Operation,ChangeTypeId]" --output table
   ```

```
-------------------------------------------------------------------
|                            CtSummaries                          |
+-----------+--------------------------+--------------------------+
| Deployment| Load balancer (ELB) stack | Create | ct-123h45t6uz7jl |
| Management| Load balancer (ELB) stack | Update | ct-0ltm873rsebx9 |
+-----------+--------------------------+--------------------------+
```

2. Find the most current version of the CT:

   `ChangeTypeId` and `ChangeTypeVersion`: The change type ID for this walkthrough is `ct-123h45t6uz7jl` (create ELB), to find out the latest version, run this command:

```
aws amscm list-change-type-version-summaries --filter
 Attribute=ChangeTypeId,Value=ct-123h45t6uz7jl
```

3. Learn the options and requirements. The following command outputs the schema to a JSON file named CreateElbParams.json.

```
aws amscm get-change-type-version --change-type-id "ct-123h45t6uz7jl" --query
 "ChangeTypeVersion.ExecutionInputSchema" --output text > CreateElbParams.json
```

4. Modify and save the execution parameters JSON file. This example names the file CreateElbParams.json.

   For a provisioning CT, the StackTemplateId is included in the schema and must be submitted in the execution parameters.

   For TimeoutInMinutes, how many minutes are allowed for the creation of the stack before the RFC is failed, this setting will not delay the RFC execution, but you must give enough time (for example, don't specify "5"). Valid values are "60" up to "360," for CTs with long-running UserData: Create EC2 and Create ASG. We recommend the max allowed "60" for all other provisioning CTs.

   Provide the ID of the VPC where you want the stack to be created; you can get the VPC ID with the CLI command `aws amsskms list-vpc-summaries`.

```
{
"Description":      "ELB-Create-RFC",
"VpcId":            "VPC_ID",
"StackTemplateId":  "stm-sdhopv00000000000",
```

```
"Name":             "MyElbInstance",
"TimeoutInMinutes": 60,
"Parameters":   {
    "ELBSubnetIds":                     ["SUBNET_ID"],
    "ELBHealthCheckHealthyThreshold":   4,
    "ELBHealthCheckInterval":           5,
    "ELBHealthCheckTarget":             "HTTP:80/",
    "ELBHealthCheckTimeout":            60,
    "ELBHealthCheckUnhealthyThreshold": 5,
    "ELBScheme":                        false
    }
}
```

5. Output the RFC JSON template to a file in your current folder named CreateElbRfc.json:

```
aws amscm create-rfc --generate-cli-skeleton > CreateElbRfc.json
```

6. Modify and save the CreateElbRfc.json file. Because you created the execution parameters in a separate file, remove the `ExecutionParameters` line. For example, you can replace the contents with something like this:

```
{
"ChangeTypeVersion":    "2.0",
"ChangeTypeId":         "ct-123h45t6uz7jl",
"Title":                "Create ELB"
}
```

7. Create the RFC. The following command specifies the execution parameters file and the RFC template file:

```
aws amscm create-rfc --cli-input-json file://CreateElbRfc.json --execution-
parameters file://CreateElbParams.json
```

   You receive the ID of the new RFC in the response and can use it to submit and monitor the RFC. Until you submit it, the RFC remains in the editing state and does not start.

Tips

> **Note**
>
> You can use the AMS API/CLI to create an RFC without creating an RFC JSON file or a CT execution parameters JSON file. To do this, you use the `create-rfc` command and add the required RFC and execution parameters to the command, this is called "Inline Create". Note that all provisioning CTs have contained within the `execution-parameters` block a `Parameters` array with the parameters for the resource. The parameters must have quote marks escaped with a back slash (\).
> The other documented method of creating an RFC is called "Template Create." This is where you create a JSON file for the RFC parameters and another JSON file for the execution parameters, and submit the two files with the `create-rfc` command. These files can serve as templates and be re-used for future RFCs.
> When creating RFCs with templates, you can use a command to create the JSON file with the contents you want by issuing a command as shown. The commands create a file named "parameters.json" with the shown content; you could also use these commands to create the RFC JSON file.

RFC cloning (re-creating) with the console

You can use the AMS console to clone an existing RFC.

To clone, or recreate, an RFC by using the AMS console, follow these steps:

1. Find the relevant RFC. From the left navigation, click **RFCs**.

   The RFCs dashboard opens.

2. Scroll through the pages until you find the RFC you want to clone. Use the **Filter** option to narrow the list. Choose the RFC that you want to clone.

   The RFC details page opens.

3. Click **Create a Copy**.

   The **Create a request for change** page opens with all options set as in the original RFC.

4. Make the changes you want. To set additional options, change the **Basic** option to **Advanced**. After you have set all options, choose **Submit**.

The active RFC details page opens with a new RFC ID for the cloned RFC and the cloned RFC appears in the RFC dashboard.

Updating RFCs

You can resubmit an RFC that has been rejected or that has not yet been submitted, by updating the RFC and then submitting it, or re-submitting it. Note that most RFCs are rejected because the specified `RequestedStartTime` has passed before submission or the specified TimeoutInMinutes is inadequate to run the RFC (since TimeoutInMinutes does not prolong a successful RFC, we recommend always setting this to at least "60" and up to "360" for an Amazon EC2 or an Amazon EC2 Auto Scaling group with long-running UserData). This section describes how to use the CLI version of the `UpdateRfc` command to update an RFC with a new RFC parameter, or new parameters using either stringified JSON or an updated parameters file.

This example describes using the CLI version of the AMS UpdateRfc API (see Update RFC). While there are change types for updating some resources (DNS private and public, load balancer stacks, and stack patching configuration), there is no CT to update an RFC.

We recommend that you submit one UpdateRfc operation at a time. If you submit multiple updates, for example on a DNS stack, the updates might fail attempting to update the DNS at the same time.

REQUIRED DATA: `RfcId`: The RFC you're updating.

OPTIONAL DATA: `ExecutionParameters`: Unless you're updating a non-required field, like `Description`, you would submit modified execution parameters to address the issues that caused the RFC to be rejected or canceled. All submitted non-null values overwrite those values in the original RFC.

1. Find the relevant rejected or canceled RFC, you can use this command (you can substitute the value with `Canceled`):

```
aws amscm list-rfc-summaries --filter Attribute=RfcStatusId,Value=Rejected
```

2. You can modify any of the following RFC parameters :

```
{
    "Description": "string",
    "ExecutionParameters": "string",
```

```
    "ExpectedOutcome": "string",
    "ImplementationPlan": "string",
    "RequestedEndTime": "string",
    "RequestedStartTime": "string",
    "RfcId": "string",
    "RollbackPlan": "string",
    "Title": "string",
    "WorstCaseScenario": "string"}
```

Example command updating the Description field:

```
aws amscm update-rfc --description "AMSTestNoOpsActionRequired" --rfc-id "RFC_ID"
 --region us-east-1
```

Example command updating the ExecutionParameters VpcId field:

```
aws amscm update-rfc  --execution-parameters "{\"VpcId\":\"VPC_ID\"}" --rfc-id
 "RFC_ID" --region us-east-1
```

Example command updating the RFC with an execution parameters file that contains the updates; see example execution parameters file in step 2 of: EC2 stack | Create:

```
aws amscm update-rfc --execution-parameters file://CreateEc2ParamsUpdate.json --
rfc-id "RFC_ID" --region us-east-1
```

3. Resubmit the RFC using `submit-rfc` and the same RFC ID that you have from when the RFC was first created:

```
aws amscm submit-rfc --rfc-id RFC_ID
```

If the RFC succeeds, you receive no confirmation or error messages at the command line.

4. To monitor the status of the request and to view Execution Output, run the following command.

```
aws amscm get-rfc --rfc-id RFC_ID
```

Finding RFCs

Finding a request for change (RFC) with the console

To find an RFC by using the AMS console, follow these steps.

> **Note**
>
> This procedure applies only to scheduled RFCs, that is, RFCs that did not use the **ASAP** option.

1. From the left navigation, click **RFCs**.

   The RFCs dashboard opens.
2. Scroll through the list or use the **Filter** option to refine the list.

   The RFC list changes per filter criteria.
3. Choose the Subject link for the RFC you want.

   The RFC details page opens for that RFC with information including RFC ID.
4. If there are many RFCs in the dashboard, you can use the **Filter** option to search by RFC:

   - **Subject**: The subject line, or title (in the API/CLI) given to the RFC when it was created.
   - **RFC ID**: The identifier for the RFC.
   - **Activity state**: If you know the RFC state, you can choose between **AwsOperatorAssigned** meaning an operator is currently looking at the RFC, **AwsActionPending** meaning that an AMS operator must perform something before the RFC execution can proceed or **CustomerActionPending** meaning that you need to take some action before the RFC execution can proceed.
   - **Status**: If you know the RFC status, you can choose between:
     - **Scheduled**: RFCs that were scheduled.
     - **Canceled**: RFCs that were canceled.
     - **In progress**: RFCs in progress.
     - **Success**: RFCs that executed successfully.
     - **Rejected**: RFCs that were rejected.
     - **Editing**: RFCs that are being edited.

  - **Failure**: RFCs that failed.
  - **Pending approval**: RFCs that cannot proceed until either AMS or you approve. Typically, this indicates that you need to approve the RFC. You will have gotten a service notification of this in your Service Requests list.
- **Change type**: Pick the **Category**, **Subcategory**, **Item**, and **Operation**, and the change type ID is retrieved for you.
- **Requested start time** or **Requested end time**: This filter option lets you choose **Before** or **After**, and then enter a **Date** and, optionally, a **Time** (hh:mm and time zone). This filter operates successfully only on scheduled RFCs (not ASAP RFCs).
- **Status**: Either **Scheduled**, **Canceled**, **In progress**, **Success**, **Rejected**, **Editing**, or **Failure**.
- **Subject**: The subject (or title, if the RFC was created with the API/CLI) that you gave the RFC.
- **Change type ID**: Use the identifier for the change type submitted with the RFC.

The search allows you to add the filters, as shown in the following screenshot.

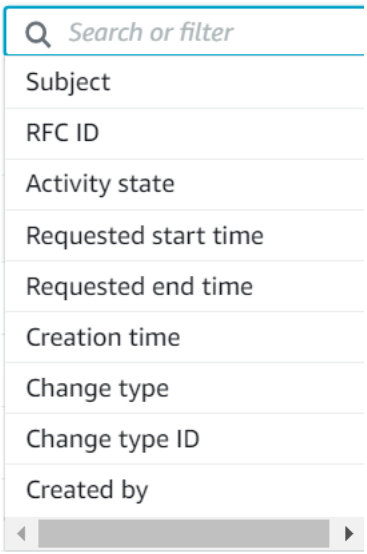

5. Click on the Subject link for the RFC you want.

   The RFC details page opens for that RFC with information including RFC ID.

Finding a request for change (RFC) with the CLI

You can use multiple filters to find an RFC.

To check the change type version, use this command:

```
aws amscm list-change-type-version-summaries --filter
 Attribute=ChangeTypeId,Value=CT_ID
```

Note

You can use any `CreateRfc` parameters with any RFC whether or not they are part of the schema for the change type. For example, to get notifications when the RFC status changes, add this line, `--notification "{\"Email\": {\"EmailRecipients\" : [\"email@example.com\"]}}"` to the RFC parameters part of the request (not the execution parameters). For a list of all CreateRfc parameters, see the AMS Change Management API Reference.

If you don't write down the RFC ID, and need to find it later, you can use the AMS change management (CM) system to search for it and narrow the results with a filter or query.

1. The CM API ListRfcSummaries operation has filters. You can Filter results based on an `Attribute` and `Value` combined in a logical AND operation, or based on an `Attribute`, a `Condition`, and `Values`.

   RFC filtering

Attribute	Valid values	Valid conditions	Default condition	Notes
ActualEndTime	Any string representing an ISO8601 datetime (for example, "20170101T000000Z")	Before, After, Between	None	The Before or After condition only accepts one value in the Values field. The Between condition must have exactly two values in the Values field, where the first value should represent a date that happens

Attribute	Valid values	Valid conditions	Default condition	Notes
				before the second value
ActualStartTime	Any string representing an ISO8601 datetime (for example, “20170101T000000Z”)	Before, After, Between	None	The Before or After condition only accepts one value in the Values field. The Between condition must have exactly two values in the Values field, where the first value should represent a date that happens before the second value
AutomationStatusId	Manual, Automated	Equals	Equals	There are only two automation statuses
ChangeTypeId	Any valid change type ID; for example, ct-123h45t6uz7jl	Equals	Equals	Finding a Change Type or CSIO
ChangeTypeVersion	Any valid change type ID; for example, 1.0	Equals	Equals	Finding a Change Type or CSIO
CreatedBy	Any string (maximum allowed length is 2048 characters)	Contains	Contains	The CreatedBy field of the RFC contains the ARN of the user who created it

Attribute	Valid values	Valid conditions	Default condition	Notes
CreatedTime	Any string representing an ISO8601 datetime (for example, “20170101T000000Z”)	Before, After, Between	None	The Before or After condition only accepts one value in the Values field. The Between condition must have exactly two values in the Values field, where the first value should represent a date that happens before the second value
LastModifiedTime	Any string representing an ISO8601 datetime (for example, “20170101T000000Z”)	Before, After, Between	None	The Before or After condition only accepts one value in the Values field. The Between condition must have exactly two values in the Values field, where the first value should represent a date that happens before the second value

Attribute	Valid values	Valid conditions	Default condition	Notes
LastSubmittedTime	Any string representing an ISO8601 datetime (for example, “20170101T000000Z”)	Before, After, Between	None	The Before or After condition only accepts one value in the Values field. The Between condition must have exactly two values in the Values field, where the first value should represent a date that happens before the second value
RequestedEndTime	Any string representing an ISO8601 datetime (for example, “20170101T000000Z”)	Before, After, Between	None	The Before or After condition only accepts one value in the Values field. The Between condition must have exactly two values in the Values field, where the first value should represent a date that happens before the second value

Attribute	Valid values	Valid conditions	Default condition	Notes
RequestedStartTime	Any string representing an ISO8601 datetime (for example, “20170101T000000Z”)	Before, After, Between	None	The Before or After condition only accepts one value in the Values field. The Between condition must have exactly two values in the Values field, where the first value should represent a date that happens before the second value
RfcStatusId	Canceled, Editing, Failure, InProgress, PendingApproval, Rejected, Scheduled, Success	Equals	Equals	Refresh the RFC list in the AMS console or run GetRfc
Title	Any valid RFC title	Contains	Contains	Regular expressions in each individual field are not supported. Case insensitive search

Examples:

To find the IDs of all the RFCs related to SQS (where SQS is contained in the Item portion of the CT), you can use this command:

```
list-rfc-summaries --query 'RfcSummaries[?contains(Item.Name,`SQS`)].
[Category.Id,Subcategory.Id,Type.Id,Item.Id,RfcId]' --output table
```

Which returns something like this:

```
-------------------------------------------------------------------------
|                          ListRfcSummaries                             |
+----------+---------------------------------+-------+-------+---------------+
|Deployment| Advanced Stack Components      |SQS    |Create |ct-123h45t6uz7jl|
|Management| Monitoring & Notification  |SQS    |Update |ct-123h45t6uz7jl|
+----------+---------------------------------+-------+-------+---------------+
```

Another filter available for `list-rfc-summaries` is `AutomationStatusId`, to look for RFCs that are automated or manual:

```
aws amscm list-rfc-summaries --filter Attribute=AutomationStatusId,Value=Automated
```

Another filter available for `list-rfc-summaries` is `Title` (**Subject** in the console):

```
 Attribute=Title,Value=RFC-TITLE
```

Example of the new request structure in JSON that returns RFCs where:

- (Title CONTAINS the phrase "Windows 2012" OR "Amazon Linux") AND
- (RfcStatusId EQUALS "Success" OR "InProgress") AND
- (20170101T000000Z <= RequestedStartTime <= 20170103T000000Z) AND (ActualEndTime <= 20170103T000000Z)

```
{
  "Filters": [
    {
      "Attribute": "Title",
      "Values": ["Windows 2012", "Amazon Linux"],
      "Condition": "Contains"
    },
    {
      "Attribute": "RfcStatusId",
      "Values": ["Success", "InProgress"],
      "Condition": "Equals"
    },
    {
```

```
      "Attribute": "RequestedStartTime",
      "Values": ["20170101T000000Z", "20170103T000000Z"],
      "Condition": "Between"
    },
    {
      "Attribute": "ActualEndTime",
      "Values": ["20170103T000000Z"],
      "Condition": "Before"
    }
  ]
}
```

> **Note**
>
> With more advanced `Filters`, AMS intends to deprecate the following fields in an upcoming release:
>
> - Value: The Value field is part of the Filters field. Use the Values field that supports more advanced functionality.
> - RequestedEndTimeRange: Use the RequestedEndTime inside the Filters field that supports more advanced functionality
> - RequestedStartTimeRange: Use the RequestedStartTime inside the Filters field that supports more advanced functionality.

For information about using CLI queries, see How to Filter the Output with the --query Option and the query language reference, JMESPath Specification.

2. If you're using the AMS console:

   Go to the **RFCs** list page. If needed, you can filter on the RFC **Subject**, which is what you entered as the RFC `Title` when you created it.

Tips

> **Note**
>
> This procedure applies only to scheduled RFCs, that is, RFCs that did not use the **ASAP** option.

Canceling RFCs

You can cancel an RFC using the Console or the AMS API/CLI.

To cancel an RFC with the console, find the RFC in your RFC list, open it, click **Cancel**.

Required Data:

- `Reason`: Why you are canceling the RFC.
- `RfcId`: The RFC you are canceling.

1. Typically you would cancel an RFC right after submitting it (so the RFC ID should be handy); otherwise, you would not be able to cancel it unless you scheduled it and it's before the specified start time. If you need to find the RFC ID, you can use this command (you can substitute the `Value` with `PendingApproval` for an RFC that is manually approved):

   ```
   aws amscm list-rfc-summaries --filter Attribute=RfcStatusId,Value=Scheduled
   ```

2. Example command to cancel an RFC:

   ```
   aws amscm cancel-rfc --reason "Bad Stack ID" --rfc-id "RFC_ID" --profile saml --
   region us-east-1
   ```

RFCs and the AMS console

The AMS console offers many features to help you succeed with creating and submitting RFCs.

RFC List page (Console)

The AMS console **RFCs** list page provides you with the following options:

- Advanced RFC search through a **Filter**. For information, see Finding RFCs.
- Finding the last time the RFC was **Modified**. This value represents that last time that the RFC status was changed.
- Viewing RFC details with the RFC **Subject**. Choosing this link opens the details page for that RFC.
- Viewing RFC status. For information, see RFC status codes

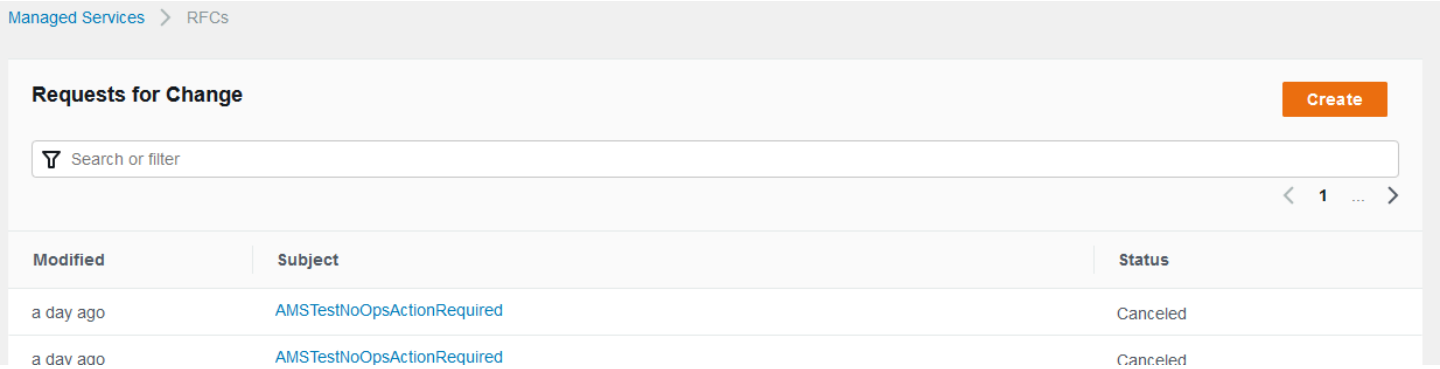

RFC quick create (console)

Use the RFC quick create cards, or list table, or choose change types for RFCs by classification.

To learn more, see Creating an RFC.

RFC correspondence and attachment (console)

You can add correspondence to an RFC after it has been submitted and before it is approved; for example, while it's in the state of "PendingApproval". After an RFC is approved (in a state of "Scheduled" or "InProgress"), correspondence cannot be added, because it could be construed as a change to the request. After an RFC is completed (in a state of "Canceled", "Rejected", "Success", or "Failure"), correspondence is once again enabled, though correspondence is disabled once an RFC is closed for more than 30 days.

> **Note**
>
> Each correspondence is limited to 5,000 characters.

Limitations for attachments:

- Only three attachments per correspondence.
- Limit fifty attachments per RFC.
- Each attachment must be less than 5 MB in size.
- Only text files are accepted such as plaintext (`.txt`), comma-separated values (`.csv`), JSON (`.json`), or YAML (`.yaml`). In the case of YAML format, the file must be attached using file extension `.yaml`.

> **Note**
>
> Text files that have XML content are prohibited. If you have XML content to share with AMS, use a service request.

- File names are limited to 255 characters, with only numbers, letters, spaces, dashes (-), underscores (_), and dots (.).
- Updating and deleting attachments on an RFC is not currently supported.

To add correspondence and attachments to an RFC, follow these steps:

1. In the AMS console, on the RFC details page for an RFC, find the **Correspondence** section at the bottom of the page.

   Before any correspondence:

   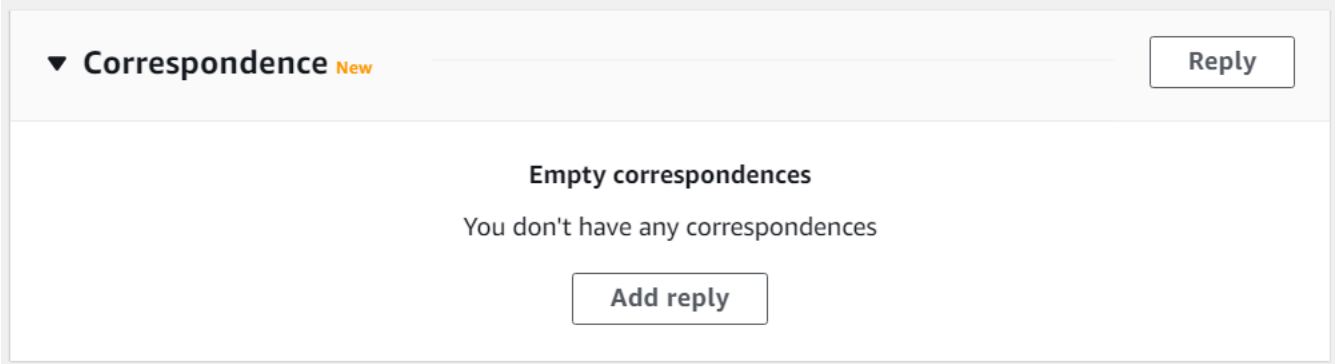

   After some correspondence:

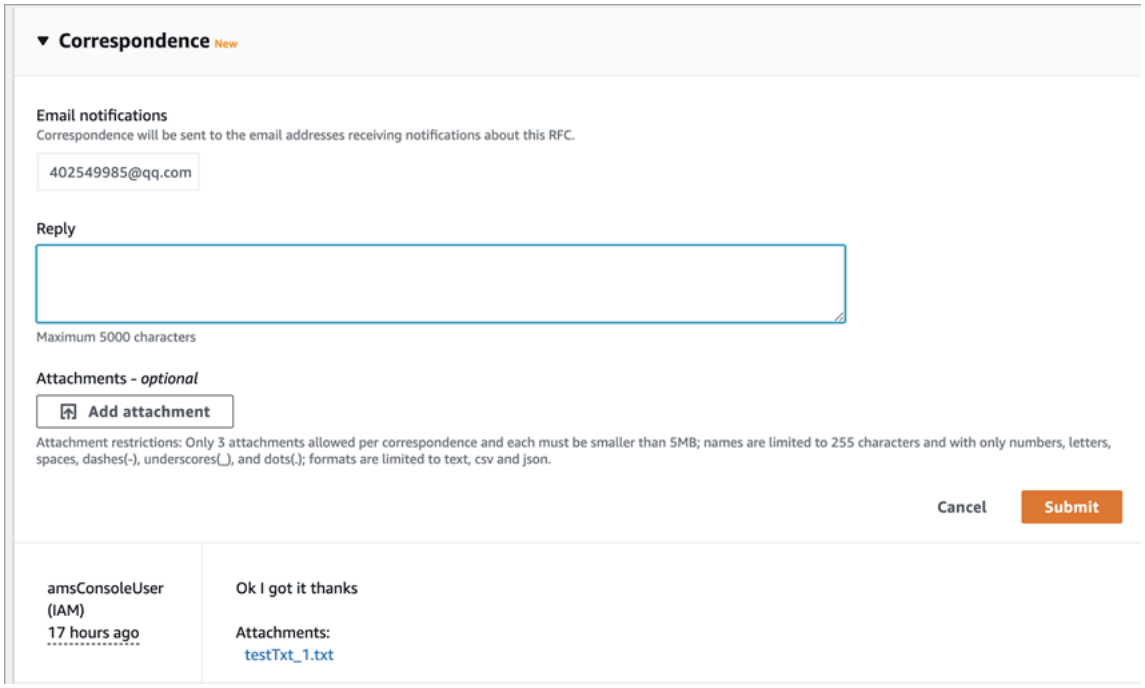

2. To add a new correspondence, type your message in the **Reply** text box. To attach files related to the correspondence, choose **Add Attachment**, and then choose the files you want.

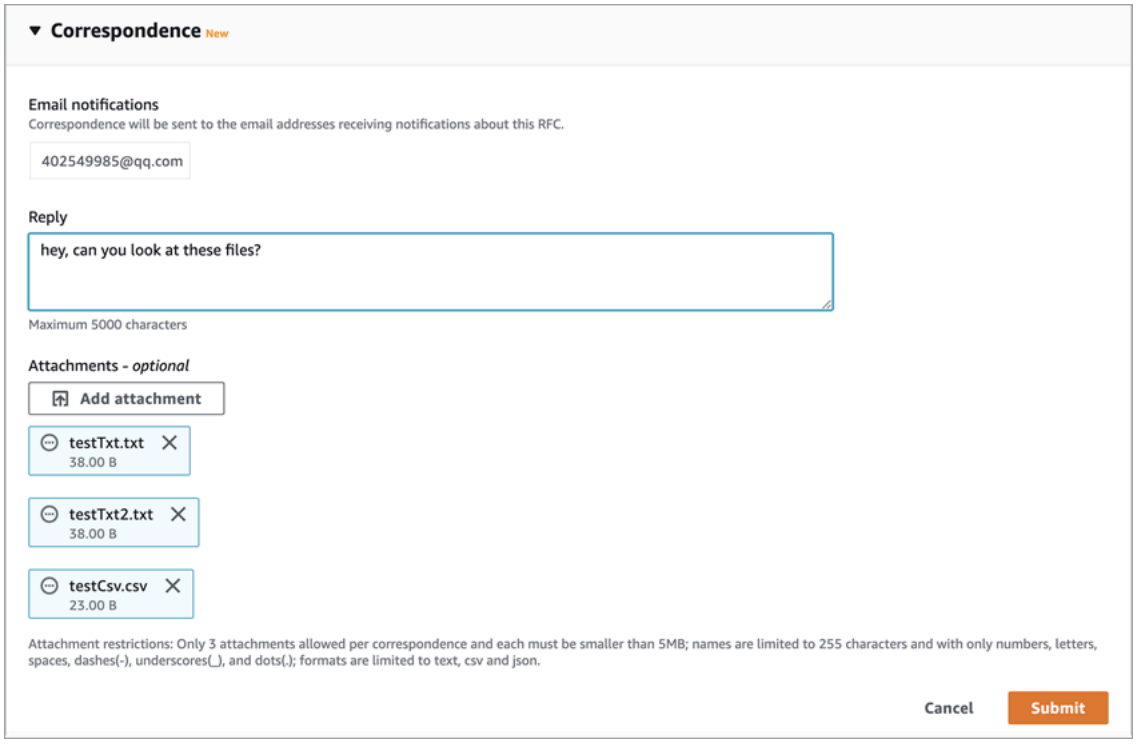

3. When you're finished, choose **Submit**.

   The new correspondence, along with links to the attached files, appear in the correspondence list on the RFC details page.

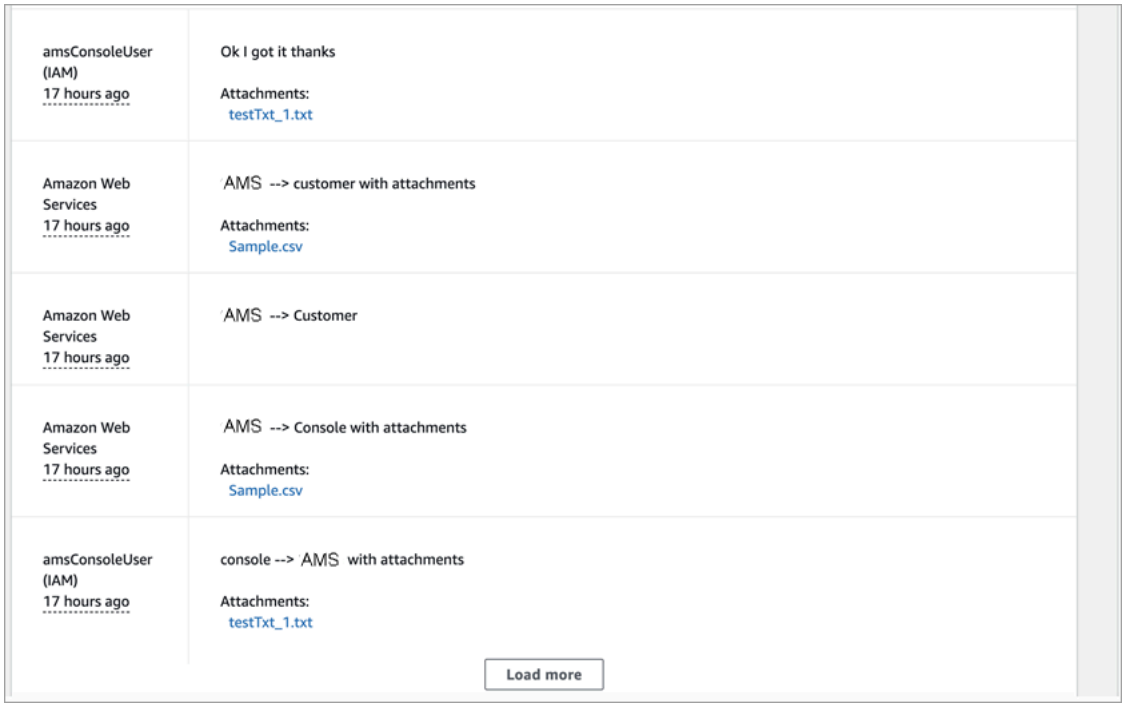

RFC email notifications (console)

The AMS console **Requests for Change** create page provides you with an option to add email addresses to receive notifications of RFC state changes:

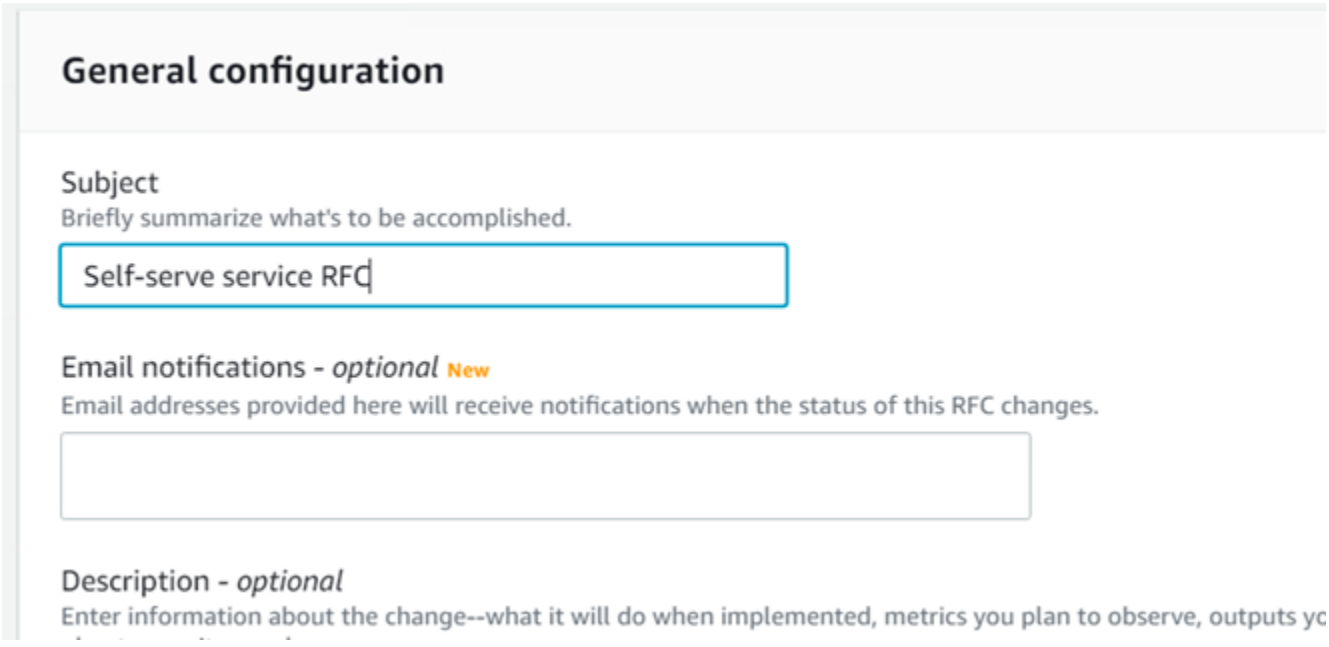

Additionally, you can add email addresses for notifications to any change type, for example:

```
aws amscm create-rfc --change-type-id ct-1e1xtak34nx76
```

```
                    --change-type-version 1.0 --title "TITLE"
                    --notification "{\"Email\": {\"EmailRecipients\" :
[\"email@example.com\"]}}"
```

Add a similar line (`--notification "{\"Email\": {\"EmailRecipients\" : [\"email@example.com\"]}}"`) to any change type inline or template request in the RFC parameters part of the request, not the parameters part.

RFC common parameters

RFC parameters required to submit, and common parameters:

- Change type information: ChangeTypeId and ChangeTypeVersion. Ror a list of change type IDs and version numbers, see Change Type Reference.

  Run `list-change-type-classification-summaries` in the CLI with the `query` argument to narrow the results. For example, narrow results to change types that contain "Access" in the `Item` name.

  ```
  aws amscm list-change-type-classification-summaries --query
   "ChangeTypeClassificationSummaries [?contains (Item, 'access')].
  [Category,Subcategory,Item,Operation,ChangeTypeId]" --output table
  ```

  Run `get-change-type-version` and specify the change type ID. The following command gets the CT version for ct-2tylseo8rxfsc.

  ```
  aws amscm get-change-type-version --change-type-id ct-2tylseo8rxfsc
  ```

- Title: A name for the RFC; this becomes the **Subject** of the RFC in the AMS console RFC list and you can search on it with the `GetRfc` command and a filter on `Title`
- Scheduling: If you want a scheduled RFC, you must include the `RequestedStartTime` and `RequestedEndTime` parameters, or use the **Schedule this change** console option. For an **ASAP** RFC (that runs as soon as it's approved), when using the CLI, leave `RequestedStartTime` and `RequestedEndTime` null. When using the console, accept the **ASAP** option.

  If the `RequestedStartTime` is missed, the RFC is rejected.
- Provisioning CTs: The execution parameters, or `Parameters` are the specific settings that are required to provision the resource. They vary widely depending on the CT.

- Non-provisioning CTs: CTs that do not provision a resource, such as access CTs or Other | Other, or delete stack, have minimal execution parameters and no `Parameters` block.
- Some RFCs also require that you specify a `TimeoutInMinutes`, or how many minutes are allowed for the creation of the stack before the RFC is failed. Valid values are 60 (minutes) up to 360, for long-running UserData. If the execution can't be completed before the `TimeoutInMinutes` is exceeded, the RFC fails. However, this setting doesn't delay the execution of the RFC.
- RFCs that create instances, such as an S3 bucket or an ELB, generally provide a schema that allows you to add up to seven tags (key/value pairs). You can add more tags to your S3 bucket by submitting a service request or a Management | Other | Other | Update CT. EC2, EFS, RDS, and the multi-tiered (HA Two-Tiered and HA One-Tiered) schemas allow up to fifty tags. Tags are specified in the `ExecutionParameters` part of the schema. Providing tags can be of great value. For more information, see Tagging Your Amazon EC2 Resources.

  When using the AMS console, you must open the **Additional configuration** area in order to add tags.

  > **Tip**
  >
  > Many CT schemas have a `Description` and `Name` field near the top of the schema. Those fields are used to name the stack or stack component, they don't name the resource you're creating. Some schemas offer a parameter to name the resource you're creating, and some do not. For example, the CT schema for Create EC2 stack doesn't offer a parameter to name the EC2 instance. In order to do so, you must create a tag with the key "Name" and the value of what you want the name to be. If you do not create such a tag, your EC2 instance displays in the EC2 console without a name attribute.

RFC AWS Region option

The AMS API/CLI (`amscm` and `amsskms`) endpoints are in `us-east-1`. If you are federating with Security Assertion Markup Language (SAML), you will have scripts provided to you at onboarding that set your Region to us-east-1 and, in that case, you do not need to specify the `--region` option when issuing a command. However, if your account is not in that AWS Region, and your SAML is configured to use us-east-1, when you issue other AWS commands (for example, `aws s3`), you will need to specify your account-onboarded Region.

> **Note**
>
> Most of the command examples provided in this guide do not include the `--region` option.

RFC daily email

The RFC digest feature allows you to sign up for a daily email summarizing the RFC activity in your account over the last 24 hours. The RFC digest feature is a streamlined process that can dramatically reduce the number of email notifications you receive regarding your account's RFCs. Thus, the RFC digest can reduce the likelihood of missing actions that are pending your response.

If you would like to enable the RFC digest feature, contact your AMS CSDM and they can subscribe you. You can request up to 20 email addresses (or aliases) be included on your RFC digest email list. The current email schedule is fixed at 09:00 UTC-8.

To disable the RFC digest feature, contact your CSDM with your request.

If you choose not to set up RFC digest and want notifications regarding your RFCs, or if you want more detailed information on your RFCs than what the RFC digest provides, you must use the Change Management System to set up CloudWatch Events notifications or email notifications for every individual RFC that you want information on. For information on setting up RFC notifications, see RFC State Change Notifications.

The topics you can expect to see in the digest, which are also shown in the below screenshot, include:

- Pending Customer Approval: Lists RFCs that are in **PendingApproval** status, awaiting your approval
- Pending Customer Reply: Lists RFCs that are awaiting your reply on RFC correspondence
- Pending AWS Approval or Reply: Lists RFCs that are waiting on AMS for reply or approval
- Completed: Lists RFCs in **Success**, **Failure**, **Cancelled** and **Rejected** status

```
ACCOUNT ID: 123456789012

Total RFCs in this digest: 10
(Some RFCs may be included in more than one category below)

----

PENDING CUSTOMER REPLY - 1

RfcId:12345678-1234-5678-0912-123456789012
Title: Title of the First RFC
Activity State: Pending Customer Reply
Status: Pending AWS Approval
CreationTime: 2020-10-23T15:41:39Z

----

PENDING CUSTOMER APPROVAL - 1

RfcId:12345678-1234-5678-0912-123456789012
Title: Title of the First RFC
Activity State: Pending AWS Reply
Status: Pending Customer Approval
CreationTime: 2020-10-23T15:41:39Z

----

PENDING AWS REPLY OR APPROVAL - 2

RfcId:12345678-1234-5678-0912-123456789012
Title: Title of the First RFC
Activity State: Pending Customer Reply
Status: Pending AWS Approval
CreationTime: 2020-10-23T15:41:39Z

RfcId:12345678-1234-5678-0912-123456789012
Title: Title of the Second RFC
Activity State: Pending AWS Reply
Status: Pending Customer Approval
CreationTime: 2020-10-23T15:41:39Z

----

COMPLETED - 8

RfcId:12345678-1234-5678-0912-123456789012
Title: Title of the First RFC
Activity State: NoActionPending
Status: Success
LastUpdatedTime: 2020-10-23T15:41:39Z

RfcId:98765432-1098-7654-3210-987654321098
```

What are change types?

Change type refers to the action that an AWS Managed Services (AMS) request for change (RFC) performs and encompasses the change action itself, and the type of change – manual vs automated. AMS has a large collection of change types not used by other Amazon web services. You use these change types when submitting a request for change (RFC) to deploy, or manage, or gain access to, resources.

Topics

- Automated and manual CTs
- CT approval requirements
- Change type versions
- Create change types
- Update change types
- Internal-only change types
- Change type schemas
- Managing permissions for change types
- Redacting sensitive information from change types
- Finding a change type, using the query option

Automated and manual CTs

A constraint on change types is whether they are automated or manual, this is the change type `AutomationStatusId` attribute, called the **Execution mode** in the AMS console.

Automated change types have expected results and execution times and run through the AMS automated system, generally within an hour or less (this largely depends on what resources the CT is provisioning). Manual change types are uncommon, but they are treated differently because they require that an AMS operator act on the RFC before it can be run. That sometimes means communicating with the RFC submitter, so, manual change types require varying lengths of time to complete.

For all scheduled RFCs, an unspecified end time is written to be the time of the specified `RequestedStartTime` plus the `ExpectedExecutionDurationInMinutes` attribute of the submitted change type. For example, if the `ExpectedExecutionDurationInMinutes` is "60" (minutes), and the specified `RequestedStartTime` is `2016-12-05T14:20:00Z` (December

5, 2016 at 4:20 AM), the actual end time would be set to December 5, 2016 at 5:20 AM. To find the `ExpectedExecutionDurationInMinutes` for a specific change type, run this command:

```
aws amscm --profile saml get-change-type-version --change-type-id CHANGE_TYPE_ID --
query "ChangeTypeVersion.{ExpectedDuration:ExpectedExecutionDurationInMinutes}"
```

Note

Scheduled RFCs with **Execution mode**= Manual, in the Console, must be set to run at least 24 hours in the future. This caveat does not apply to the AMS API/CLI, but it is still important to schedule manual RFCs at least 8 hours ahead.

Note

When using "review required" CTs, AMS recommends that you use the ASAP **Scheduling** option (choose **ASAP** in the console, leave start and end time blank in the API/CLI) as these CTs require an AMS operator to examine the RFC, and possibly communicate with you before it can be approved and run. If you schedule these RFCs, be sure to allow at least 24 hours. If approval does not happen before the scheduled start time, the RFC is rejected automatically.

AMS aims to respond to a manual CT within four hours, and will correspond as soon as possible, but it could take much longer for the RFC to actually be run.

For a list of the CTs that are Manual and require AMS review, see the Change Type CSV file, available on the **Developer's Resources** page of the Console.

YouTube Video: How can I find automated change types for AMS RFCs?

To find the **Execution mode** for a CT in the AMS console, you must use the **Browse change types** search option. The results show the execution mode of the matching change type or change types.

To find the `AutomationStatus` for a specific change type by using the AMS CLI, run this command:

```
aws amscm --profile saml get-change-type-version --change-type-id CHANGE_TYPE_ID --
query "ChangeTypeVersion.{AutomationStatus:AutomationStatus.Name}"
```

You can also look up change types in the AMS Change Type Reference, which provides information about all AMS change types.

> **Note**
>
> The AMS API/CLI are not currently part of the AWS API/CLI. To access the AMS API/CLI, you download the AMS SDK through the AMS console.

CT approval requirements

AMS CTs always have two approval conditions, **AwsApprovalId** and **CustomerApprovalId** that indicate whether the RFC requires AMS or you, or anyone, to approve the execution.

The approval condition is somewhat related to the execution mode; for details, see Automated and manual CTs.

To find out the approval condition for a CT, you can look in the AMS Change Type Reference, or run GetChangeTypeVersion. Both will also give you the CT `AutomationStatusId` or **Execution mode**.

You can approve RFCs by using the AMS console or with the following command:

```
aws amscm approve-rfc --rfc-id RFC_ID
```

CT approval condition

If the CT approval condition is	It requires approval from	And
`AwsApprovalId: Required`	The AMS change type system,	No action is required. This condition is typical for automated CTs.
`AwsApprovalId: NotRequiredIfSubmitter`	The AMS change type system and no one else, if the submitted RFC is for the	No action is required. This condition is typical for manual CTs because they will always be reviewed by AMS operators.

If the CT approval condition is	It requires approval from	And
	account it was submitted against,	
`CustomerApprovalId: NotRequired`	The AMS change type system,	If the RFC passes syntax and parameter checks, it is auto approved.
`CustomerApprovalId: Required`	The AMS change type system and you,	A notification is sent to you, and you must explicitly approve the RFC, either by responding to the notice, or running the ApproveRfc operation.
`CustomerApprovalId: NotRequiredIfSubmitter`	The AMS change type system and no one else, if you submitted the RFC.	If the RFC passes syntax and parameter checks, it is auto approved.
Urgent Security Incident or Patch	AMS	Is auto approved and implemented.

Change type versions

Change types are versioned and the version changes when a major update is made to the change type.

After selecting a change type using the AMS console, you have the option of opening the **Additional configuration** area and selecting a change type version. You can also specify a change type version at the API/CLI command line. You might want to do this for various reasons, including:

- You know that the version of the **Update** change type that you want must match the version of the **Create** change type that you used to create the resource that you now want to update. For example, you might have an Elastic Load Balancer (ELB) instance that you created with ELB Create change type version 1. To update it, choose ELB Update version 1.

- You want to use a change type version that has different options in it than the most recent change type. We don't recommend this because AMS updates change types mainly for security reasons and we recommend that you always choose the most recent version.

Create change types

Create change types are matched version-to-version with the Update change types. That is, the change type version that you use to provision a resource must match the version of the Update change type that you would use later to modify that resource. For example, if you create an S3 bucket with the Create S3 Bucket change type version 2.0, and later want to submit an RFC to modify that S3 bucket, you must use the Update S3 Bucket change type version 2.0 as well, even if there is an Update S3 Bucket change type with version 3.0.

We recommend keeping a record of the change type ID and version that you use when provisioning a resource with a Create change type in case you later want to use an Update change type to modify it.

Update change types

AMS provides Update change types to update resources that were created with Create change types. The Update change types must be matched version-to-version with the Create change type originally used to provision the resource.

We recommend keeping a record of the change type ID and version that you use when provisioning a resource to make it easy to update it.

YouTube Video: How do I use update CTs to change resources in an AWS Managed Services (AMS) account?

Internal-only change types

You can see change types that are for internal use only. This is so you know what actions AMS can, or does, take. If there is an internal-only change type that you would like to have available for your use, submit a service request.

For example, there is a Management | Monitoring and notification | CloudWatch alarm suppression | Update CT that is internal-only. AMS uses it to deploy infrastructure updates (such as patching) to turn off alarm notifications that the updates might erroneously trigger. When this CT is submitted, you will notice the RFC for the CT in your RFC list. Any internal-only CT that is deployed in an RFC appears in your RFC list.

Change type schemas

All change types provide a JSON schema for your input in the creation, modification, or access, of resources. The schema provides the parameters, and their descriptions, for you to create a request for change (RFC).

The successful execution of an RFC results in execution output. For provisioning RFCs, the execution output includes a "stack_id" that represents the stack in CloudFormation and can be searched in the CloudFormation console. The execution output sometimes includes output of the ID of the instance created and that ID can be used to search for the instance in the corresponding AWS console. For example, the Create ELB CT execution output includes a "stack_id" that is searchable in CloudFormation and outputs a key=ELB value=<stack-xxxx> that is searchable in the Amazon EC2 console for Elastic Load Balancing.

Let's examine a CT schema. This is the schema for CodeDeploy Application Create, a fairly small schema. Some schemas have very large `Parameter` areas.

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "name": "Create CodeDeploy application",
  "description": "Use to create an AWS CodeDeploy applicati
on
  resource with the specified name.",
  "type": "object",
  "properties": {
    "Description": {
      "description": "The reason for the request.",
      "type": "string",
      "minLength": 1,
      "maxLength": 500
    },
    "VpcId": {
      "description": "ID of the vpc to use, in the form
 vpc-0123abcd or vpc-01234567890abcdef.",
      "type": "string",
      "pattern": "^vpc-[a-z0-9]{8}$"
    },
    "StackTemplateId": {
      "description": "Must be stm-sft6rv00000000000",
      "type": "string",
      "enum": ["stm-sft6rv00000000000"]
```

The first part of the schema provides information to AMS about the requested change type.

```
    },
    "Name":{
      "description": "A name for the stack or stack component
;
      this becomes the Stack Name.",
      "type": "string",
      "minLength": 1,
      "maxLength": 255
    },
    "Tags": {
      "description": "Up to seven tags (key/value pairs) to
      categorize the resource.",
      "type": "array",
      "items": {
        "type": "object",
        "properties": {
          "Key": {
            "type": "string",
            "minLength": 1,
            "maxLength": 127
          },
          "Value": {
            "type": "string",
            "minLength": 1,
            "maxLength": 255
          }
        },
        "additionalProperties": false,
        "required": [
          "Key",
          "Value"
        ]
      },
      "minItems": 1,
      "maxItems": 7
    },
    "TimeoutInMinutes": {
      "description": "The maximum amount of time, in minutes,
 to
      allow for execution of the change. This will not prolong
 execution,
      but the RFC fails if the change is not completed in the
 specified time.
```

The TimeoutIn Minutes parameter allows you to indicate a boundary time for running the change type. Valid values are 60 up to 360, for long-running UserData.

The Parameters section is where you specify settings for the resource you are creating, or the action you are requesting.

```
      Valid values are 60 up to 360, for long-running
 UserData.",
      "type": "number",
      "minimum": 0,
      "maximum": 60
    },
    "Parameters": {
      "description": "Specifications for the stack.",
      "type": "object",
      "properties": {
        "CodeDeployApplicationName": {
          "description": "The name of an AWS CodeDeploy
 application.",
          "type": "string",
          "minLength": 1,
          "maxLength": 100,
          "pattern": "^[a-zA-Z0-9._+=,@-]{1,100}$"
        }
      },
      "additionalProperties": false,
      "required": [
        "CodeDeployApplicationName"
      ]
    }
  },
  "additionalProperties": false,
  "required": [
    "Description",
    "VpcId",
    "StackTemplateId",
    "Name",
    "TimeoutInMinutes",
    "Parameters"
  ]
}
```

The "additional properties" sections let you know what parameters are required and which are optional.

Note

This schema allows up to seven tags; however, EC2, EFS, RDS, and the multi-tier create schemas allow up to 50 tags.

Managing permissions for change types

You can use a custom policy to restrict which change types (CTs) are available to different groups or users.

To learn more about doing this, see the AMS User Guide section Setting Permissions.

Redacting sensitive information from change types

AMS change type schemas offer a parameter attribute, `"metadata":"ams:sensitive":"true"` that is used for parameters that would contain sensitive information, such as a password. When this attribute is set, the input provided is obscured. Note that you cannot set this parameter attribute; however, if you are working with AMS to create a change type and have a parameter that you would like obscured at input, you can request this.

Finding a change type, using the query option

This example demonstrates how to use the AMS Console to find the appropriate change type for the RFC that you want to submit.

You can use the console or the API/CLI to find a change type ID (CT) or version. There are two methods, either a search or choosing the classification. For both selection types, You can sort the search by choosing either **Most frequently used**, **Most recently used**, or **Alphabetical**.

YouTube Video: How do I create an RFC using the AWS Managed Services CLI and where can I find the CT Schema?

In the AMS console, on the **RFCs** -> **Create RFC** page:

- With **Browse by change type** selected (the default), either:
  - Use the **Quick create** area to select from AMS's most popular CTs. Click on a label and the **Run RFC** page opens with the **Subject** option auto-filled for you. Complete the remaining options as needed and click **Run** to submit the RFC.
  - Or, scroll down to the **All change types** area and start typing a CT name in the option box, you don't have to have the exact or full change type name. You can also search for a CT by change type ID, classification, or execution mode (automated or manual) by entering the relevant words.

With the default **Cards** view selected, matching CT cards appear as you type, select a card and click **Create RFC**. With the **Table** view selected, choose the relevant CT and click **Create RFC**. Both methods open the **Run RFC** page.

- Alternatively, and to explore change type choices, click **Choose by category** at the top of the page to open a series of drop-down option boxes.
- Choose **Category**, a **Subcategory**, an **Item**, and an **Operation**. The information box for that change type appears a panel appears at the bottom of the page.
- When you're ready, press **Enter**, and a list of matching change types appears.
- Choose a change type from the list. The information box for that change type appears at the bottom of the page.
- After you have the correct change type, choose **Create RFC**.

> **Note**
>
> The AMS CLI must be installed for these commands to work. To install the AMS API or CLI, go to the AMS console **Developers Resources** page. For reference material on the AMS CM API or AMS SKMS API, see the AMS Information Resources section in the User Guide. You may need to add a `--profile` option for authentication; for example, `aws amsskms` *`ams-cli-command`* `--profile SAML`. You may also need to add the `--region` option as all AMS commands run out of us-east-1; for example `aws amscm` *`ams-cli-command`* `--region=us-east-1`.

> **Note**
>
> The AMS API/CLI (amscm and amsskms) endpoints are in the AWS N. Virginia Region, `us-east-1`. Depending on how your authentication is set, and what AWS Region your account and resources are in, you may need to add `--region us-east-1` when issuing commands. You may also need to add `--profile saml`, if that is your authentication method.

To search for a change type using the AMS CM API (see ListChangeTypeClassificationSummaries) or CLI:

You can use a filter or query to search. The ListChangeTypeClassificationSummaries operation has Filters options for `Category`, `Subcategory`, `Item`, and `Operation`, but the values must match the existing values exactly. For more flexible results when using the CLI, you can use the `--query` option.

Change type filtering with the AMS CM API/CLI

Attribute	Valid values	Valid/Default condition	Notes
ChangeTypeId	Any string representing a ChangeTypeId (For ex: ct-abc123xyz7890)	Equals	For change type IDs, see the Change Type Reference. For change type IDs, see Finding a Change Type or CSIO.
Category	Any free-form text	Contains	Regular expressions in each individual field are not supported. Case insensitive search
Subcategory			
Item			
Operation			

1. Here are some examples of listing change type classifications:

   The following command lists all change type categories.

   ```
   aws amscm list-change-type-categories
   ```

   The following command lists the subcategories belonging to a specified category.

   ```
   aws amscm list-change-type-subcategories --category CATEGORY
   ```

   The following command lists the items belonging to a specified category and subcategory.

   ```
   aws amscm list-change-type-items --category CATEGORY --subcategory SUBCATEGORY
   ```

2. Here are some examples of searching for change types with CLI queries:

   The following command searches CT classification summaries for those that contain "S3" in the Item name and creates output of the category, subcategory, item, operation, and change type ID in table form.

```
aws amscm list-change-type-classification-summaries --query
 "ChangeTypeClassificationSummaries [?contains(Item, 'S3')].
[Category,Subcategory,Item,Operation,ChangeTypeId]" --output table
```

```
+-------------------------------------------------------------+
|               ListChangeTypeClassificationSummaries         |
+----------+-------------------------+--+------+--------------+
|Deployment|Advanced Stack Components|S3|Create|ct-1a68ck03fn98r|
+----------+-------------------------+--+------+--------------+
```

3. You can then use the change type ID to get the CT schema and examine the parameters. The following command outputs the schema to a JSON file named CreateS3Params.schema.json.

```
aws amscm get-change-type-version --change-type-id "ct-1a68ck03fn98r"
 --query "ChangeTypeVersion.ExecutionInputSchema" --output text >
 CreateS3Params.schema.json
```

   For information about using CLI queries, see How to Filter the Output with the --query Option and the query language reference, JMESPath Specification.

4. After you have the change type ID, we recommend verifying the version for the change type to make sure it's the latest version. Use this command to find the version for a specified change type:

```
aws amscm list-change-type-version-summaries --filter
 Attribute=ChangeTypeId,Value=CHANGE_TYPE_ID
```

   To find the `AutomationStatus` for a specific change type, run this command:

```
aws amscm --profile saml get-change-type-version --change-type-id CHANGE_TYPE_ID --
query "ChangeTypeVersion.{AutomationStatus:AutomationStatus.Name}"
```

   To find the `ExpectedExecutionDurationInMinutes` for a specific change type, run this command:

```
aws amscm --profile saml get-change-type-version --change-type-id ct-14027q0sjyt1h
 --query "ChangeTypeVersion.{ExpectedDuration:ExpectedExecutionDurationInMinutes}"
```

Troubleshooting RFC errors in AMS

Many AMS provisioning RFC failures can be investigated through the CloudFormation documentation. See Troubleshooting AWS CloudFormation: Troubleshooting Errors

Additional troubleshooting suggestions are provided in the following sections.

"Management" RFC errors in AMS

AMS "Management" Category change types (CTs) allow you to request access to resources as well as manage existing resources. This section describes some common issues.

RFC access errors

- Make sure the Username and FQDN you specified in the RFC are correct and exist in the domain. For help finding your FQDN, see Finding your FQDN.
- Make sure the stack ID you specified for access is an EC2-related stack. Stacks such as ELB and Amazon Simple Storage Service (S3) are not candidates for access RFCs, instead, use your read only access role to get access these stacks resources. For help finding a stack ID, see Finding stack IDs
- Make sure the stack ID you provided is correct and belongs to the relevant account.

For help with other access RFC failures, see Access management.

YouTube Video: How do I raise a Request for Change (RFC) properly to avoid rejections and failures?

RFC (manual) CT scheduling errors

Most change types are ExecutionMode=Automated, but some are ExecutionMode=Manual and that affects how you should schedule them to avoid RFC failure.

Scheduled RFCs with ExecutionMode=Manual, must be set to execute at least 24 hours in the future if you are using the AMS Console to create the RFC. This caveat does not apply to the AMS API/CLI, but it is still important to schedule Manual RFCs at least 8 hours ahead.

AMS aims to respond to a manual CT within four hours, and will correspond as soon as possible, but it could take much longer for the RFC to actually be executed.

Using RFCs with manual update CTs

AMS Operations reject Management | Other | Other RFCs for updates to stacks, when there is an Update change type for the type of stack that you want to update.

RFC delete stack errors

RFC delete stack failures: If you use the Management | Standard stacks | Stack | Delete CT, you will see the detailed events in the AWS CloudFormation Console for the stack with the AMS stack name. You can identify your stack by checking it against the name it has in the AMS Console. The AWS CloudFormation Console provides more details about failure causes.

Before deleting a stack, you should consider how the stack was created. If you created the stack using an AMS CT and did not add or edit the stack resources, then you can expect to delete it without issue. However, it is a good idea for you remove any manually-added resources from a stack before submitting a delete stack RFC against it. For example, if you create a stack using the full stack CT (HA Two Tier), it includes a security group - SG1. If you then use AMS to create another security group - SG2, and reference the new SG2 in the SG1 created as part of the full stack, and then use the delete stack CT to delete the stack, the SG1 will not delete as it is referenced by SG2.

Important

Deleting stacks can have unwanted and unanticipated consequences. AMS prefers to *not* delete stacks or stack resources on behalf of customers for this reason. Note, that AMS will only delete resources on your behalf (through a submitted Mangement | Other | Other | Update change type) that are not possible to delete using the appropriate, automated, change type to delete. Additional considerations:

- If the resources are enabled for 'delete protection', AMS can help to unblock this if you submit a Management | Other | Other | Update change type and, after the deletion protection is removed, you can use the automated CT to delete that resource.
- If there are multiple resources in a stack, and you want to delete only a subset of the stack resources, use the CloudFormation Update change type (see CloudFormation Ingest Stack: Updating). You can also submit a Management | Other | Other | Update change type and AMS engineers can help you craft the changeset, if needed.

- If there are issues using the CloudFormation Update CT due to drifts, AMS can help if you submit a Management | Other | Other | Update to resolve the drift (as far as supported by the AWS CloudFormation Service) and provide a ChangeSet that you can then validate and execute using the automated CT, Management/Custom Stack/Stack From CloudFormation Template/Approve Changeset and Update.

AMS maintains the above restrictions to help ensure there are no unexpected or unanticipated resource deletions.

For more information, see Troubleshooting AWS CloudFormation: delete stack fails.

RFC update DNS errors

Multiple RFCs to update a DNS hosted zone can fail, some without reason. Creating multiple RFCs at the same time to update DNS hosted zones (private or public) can cause some RFCs to fail because they are trying to update the same stack at the same time. AMS change management rejects or fails RFCs that are not able to update a stack because the stack is already being updated by another RFC. AMS recommends that you create one RFC at a time and wait for the RFC to succeed before raising a new one for the same stack.

RFC IAM entities errors

AMS provisions a number of default IAM roles and profiles into AMS accounts that are designed to meet your needs. However, you may need to request additional IAM resources occasionally.

The process for submitting RFCs requesting custom IAM resources follows the standard workflow for manual RFCs, but the approval process also includes a security review to ensure appropriate security controls are in place. Therefore, the process typically takes longer than other manual RFCs. To reduce the cycle time on these RFCs, please follow the following guidelines.

For information on what we mean by an IAM review and how it maps to our Technical Standards and Risk Acceptance process, see RFC security reviews.

Common IAM resources requests:

- If you are asking for a policy pertaining to a major cloud-compatible application, such as CloudEndure, see the AMS pre-approved IAM CloudEndure policy: Unpack the WIGs Cloud Endure Landing Zone Example file and open the `customer_cloud_endure_policy.json`

> **Note**
>
> If you want a more permissive policy, discuss your needs with your CloudArchitect/CSDM and obtain, if needed, an AMS Security Review and Signoff before submitting an RFC implementing the policy.

- If you want to modify a resource deployed by AMS in your account by default, we recommend that you ask for a modified copy of that resource instead of changes to the existing one.
- If you are requesting permissions for a human user (instead of attaching the permissions to the user) attach the permissions to a role, and then grant the user permission to assume that role. For details on doing this, see Temporary AMS Advanced console access.
- If you require exceptional permissions for a temporary migration or workflow, provide an end date for those permissions in your request.
- If you've already discussed the subject of your request with your security team, provide evidence of their approval to your CSDM with as much detail as possible.

If AMS rejects an IAM RFC we provide a clear reason for the rejection. For example, we might reject an IAM policy create request and explain what about the policy is inappropriate. In that case, you can make the identified changes and resubmit the request. If additional clarity on the status of a request is required, submit a service request, or contact your CSDM.

The following list describes the typical risks that AMS tries to mitigate as we review your IAM RFCs. If your IAM RFC has any of these risks, it may result in the rejection of the RFC. In cases where you require an exception, AMS asks for approvals from your security team. To seek such an exception, coordinate with your CSDM.

> **Note**
>
> AMS may, for any reason, decline any change to IAM resources inside of an account. For concerns regarding any RFC rejection, reach out to AMS Operations via a service request, or contact your CSDM.

- Privilege escalation, such as permissions that allow you to modify your own permissions, or to modify the permissions of other resources inside the account. Examples:
  - The use of `iam:PassRole` with another, more privileged role.

  - Permission to attach/detach IAM policies from a role or user.
  - The modification of IAM policies in the account.
  - The ability to make API calls in the context of management infrastructure.
- Permissions to modify resources or applications that are required to provide AMS services to you. Examples:
  - Modification of AMS infrastructure like the bastions, management host, or EPS infrastructure.
  - Deletion of log management AWS Lambda functions, or log streams.
  - The deletion or modification of the default CloudTrail monitoring application.
  - The modification of the Directory Services Active Directory (AD).
  - Disabling CloudWatch (CW) alarms.
  - The modification of the principals, policies, and namespaces deployed in the account as a part of the landing zone.
- Deployment of infrastructure outside of best practices, such as permissions that allow the creation of infrastructure in a state that endangers your information security. Examples:
  - The creation of public, or unencrypted, S3 buckets or public sharing of EBS volumes.
  - The provisioning of public IP addresses.
  - The modification of security groups to allow broad access.
- Overly broad permissions capable of causing application impact, such as permissions that can result in data loss, integrity loss, inappropriate configuration, or interruptions of service for your infrastructure and the applications inside the account. Examples:
  - Disabling, or redirecting, network traffic through APIs like `ModifyNetworkInterfaceAttribute` or `UpdateRouteTable`.
  - The disabling of managed infrastructure by detaching volumes from managed hosts.
- Permissions for services not a part of the AMS service description and not supported by AMS.

  Services not listed in the AMS Service description cannot be used in AMS accounts. To request support for a feature or service, please reach out to your CSDM.
- Permissions that do not meet your stated goal as they are either too generous, or too conservative, or are applied to the wrong resources. Examples:
  - A request for `s3:PutObject` permissions to an S3 bucket that has mandatory KMS encryption, without `KMS:Encrypt` permissions to the relevant key.
  - Permissions that pertain to resources that don't exist in the account.

- IAM RFCs where the description of the RFC does not seem to match the request.

"Deployment" RFC errors

AMS "Deployment" Category change types (CTs) allow you to request various AMS-supported resources be added to your account.

Most AMS CTs that create a resource are based on AWS CloudFormation templates. As a customer you have read-only access to all AWS services including AWS CloudFormation, you can quickly identify the AWS CloudFormation stack that represents your stack based on the stack description using the AWS CloudFormation Console. The failed stack will likely be in a state of DELETE_COMPLETE. Once you have identified the AWS CloudFormation stack, the events will show you the specific resource that failed to create, and why.

Using CloudFormation documentation to troubleshoot

Most AMS provisioning RFCs use a CloudFormation template and that documentation can be helpful for troubleshooting. See documentation for that AWS CloudFormation template:

- Create application load balancer failure: AWS::ElasticLoadBalancingV2::LoadBalancer (Application Load Balancer)
- Create Auto scaling group: AWS::AutoScaling::AutoScalingGroup (Auto Scaling Group)
- Create memcached cache: AWS::ElastiCache::CacheCluster (Cache Cluster)
- Create Redis cache: AWS::ElastiCache::CacheCluster (Cache Cluster)
- Create DNS Hosted Zone (used with Create DNS private/public): AWS::Route53::HostedZone (R53 Hosted Zone)
- Create DNS Record Set (used with Create DNS private/public): AWS::Route53::RecordSet (Resource Record Sets)
- Create EC2 stack: AWS::EC2::Instance (Elastic Compute Cloud)
- Create Elastic File System (EFS): AWS::EFS::FileSystem (Elastic File System)
- Create Load balancer: AWS::ElasticLoadBalancing::LoadBalancer (Elastic Load Balancer)
- Create RDS DB: AWS::RDS::DBInstance (Relational Database)
- Create Amazon S3: AWS::S3::Bucket (Simple Storage Service)
- Create Queue: AWS::SQS::Queue (Simple Queue Service)

RFC creating AMIs errors

An Amazon Machine Image (AMI) is a template that contains a software configuration (for example, an operating system, an application server, and applications). From an AMI, you launch an instance, which is a copy of the AMI running as a virtual server in the cloud. AMIs are very useful, and required to create EC2 instances or Auto Scaling groups; however, you must observe some requirements:

- The instance you specify for `Ec2InstanceId` must be in a stopped state for the RFC to succeed. Do not use Auto Scaling group (ASG) instances for this parameter because the ASG will terminate a stopped instance.
- To create an AMS Amazon Machine Image (AMI), you must start with an AMS instance. Before you can use the instance to create the AMI, you must prepare it by ensuring that it is stopped and dis-joined from its domain. For details, see Create a Standard Amazon Machine Image Using Sysprep
- The name you specify for the new AMI must be unique within the account or the RFC fails. How to do this is described in AMI | Create, and for more details, see and AWS AMI Design.

> **Note**
>
> For additional information for prepping for AMI creation, see AMI | Create.

RFCs creating EC2s or ASGs errors

For EC2 or ASG failures with timeouts, AMS recommends that you confirm if the AMI used is customized. If it is, please refer to the AMI creation steps included in this guide (see AMI | Create) to ensure that the AMI was created correctly. A common mistake when creating a custom AMI is not following the steps in the guide to rename or invoke Sysprep.

RFCs creating RDS errors

Amazon Relational Database Service (RDS) failures can occur for many different reasons because you can use many different engines when you create the RDS, and each engine has its own requirements and limitations. Before attempting to create an AMS RDS stack, carefully review AWS RDS parameter values, see CreateDBInstance.

To learn more about Amazon RDS in general, including size recommendations, see Amazon Relational Database Service Documentation.

RFCs creating Amazon S3s errors

One common error when creating an S3 storage bucket is not using a unique name for the bucket. If you submitted an S3 bucket Create CT with the same name as one previously submitted, it would fail because an S3 bucket would already exist with that BucketName. This would be detailed in the AWS CloudFormation Console, where you will see that the stack event shows that the bucket name is already in use.

RFC validation versus execution errors

RFC failures and related messages differ in the output messages on the AMS console RFC details page for a selected RFC:

- Validation Failures reasons are available in Status Field only
- Execution Failures reasons are available in Execution Output as well as Status Fields.

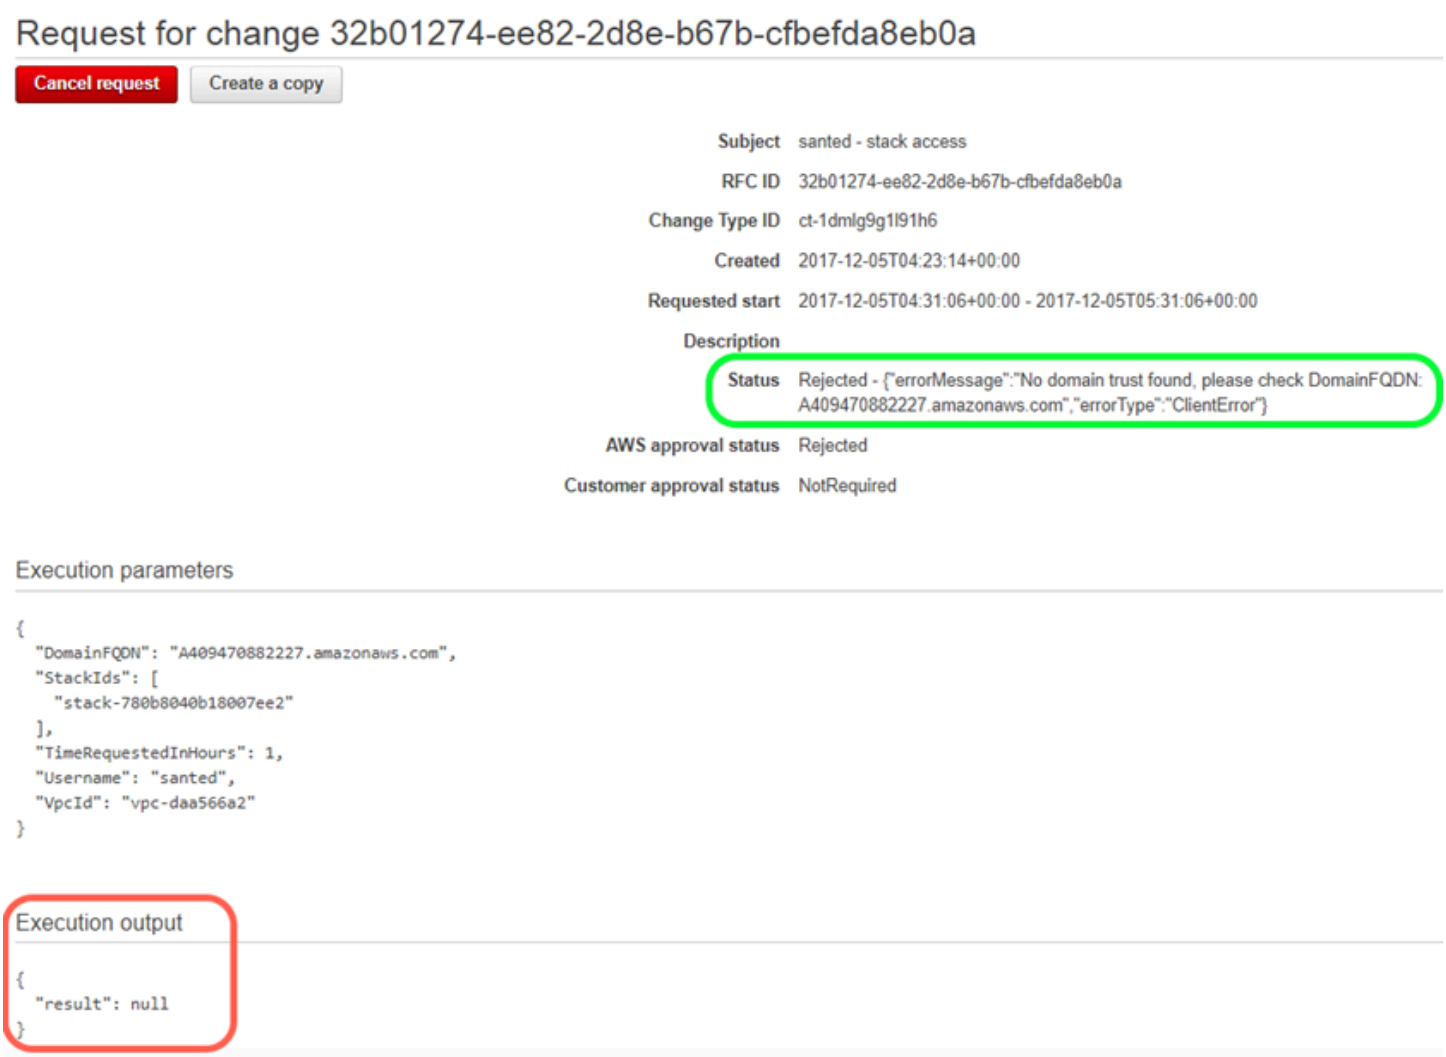

RFC error messages

When you come across the following error for the listed change types (CTs), you can use these solutions to help you find the source of the problems and fix them.

```
{"errorMessage":"An error has occurred during RFC execution. We are investigating the issue.","errorType":"InternalError"}
```

If you require further assistance after referring to the following troubleshooting options, then engage AMS via RFC correspondence or create a service request. See RFC Correspondence and Attachment (Console) and Creating a Service Request in AMS for more details.

Workload ingestion (WIGS) errors

> **Note**
>
> Validation tools for both Windows and Linux can be downloaded and run directly on your on-premises servers, as well as EC2 instances in AWS. These can be found through the *AMS Advanced Application Developer's Guide* Migrating workloads: Linux pre-ingestion validation and Migrating workloads: Windows pre-ingestion validation.

- Make sure EC2 instance exists in target AMS account. For example, if you have shared your AMI from a non-AMS account to an AMS account, you'll have to create an EC2 instance in your AMS account with the shared AMI before you can submit a Workload Ingest RFC.
- Check to see if the security groups attached to the instance have egress traffic allowed. The SSM Agent needs to be able to connect to its public endpoint.
- Check to see if the instance has the right permissions to connect with the SSM agent. These permissions come with the `customer-mc-ec2-instance-profile`, you can check for this in the EC2 console:

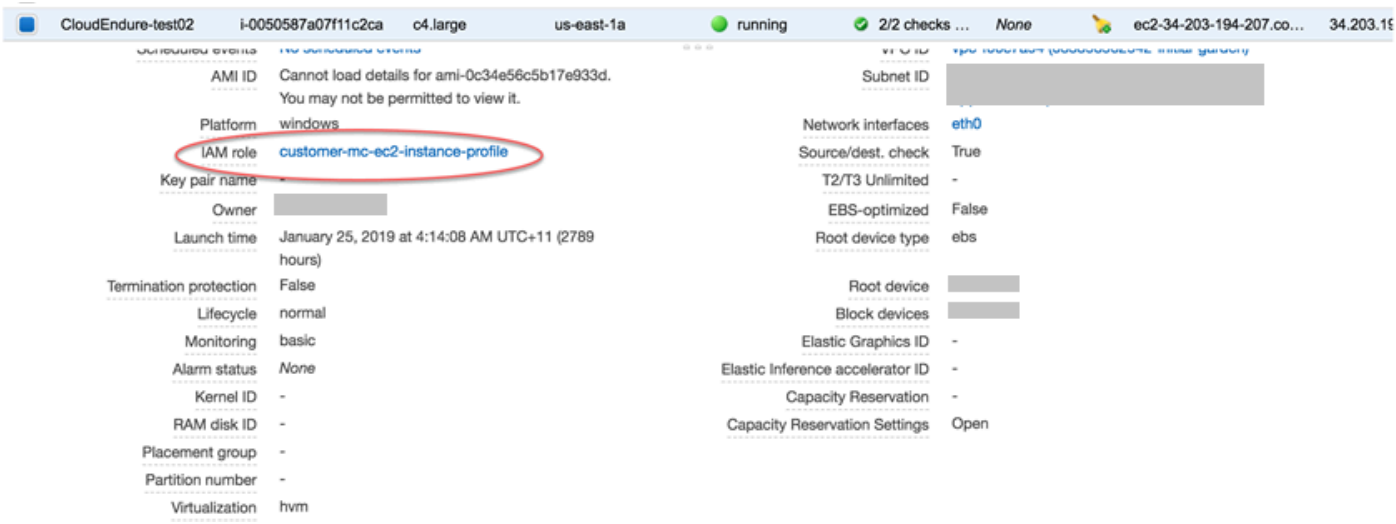

EC2 instance stack stop errors

- Check to see if the instance is already in a stopped or terminated state.
- If the EC2 instance is online and you see the `InternalError` error, then submit a service request for AMS to investigate.

- Note that you can't use the change type Management | Advanced stack components | EC2 instance stack | Stop ct-3mvvt2zkyveqj to stop an Auto Scaling group (ASG) instance. If you need to stop an ASG instance, then submit a service request.

EC2 instance stack create errors

The `InternalError` message is from CloudFormation; a CREATION_FAILED status reason. You can find details on the stack failure in CloudWatch stack events by following these steps:

- In the AWS Management console, you can view a list of stack events while your stack is being created, updated, or deleted. From this list, find the failure event and then view the status reason for that event.

  The status reason might contain an error message from AWS CloudFormation or from a particular service that can help you understand the problem.
- For more information about viewing stack events, see Viewing AWS CloudFormation Stack Data and Resources on the AWS Management Console.

EC2 instance volume restore errors

AMS creates an internal troubleshooting RFC when EC2 instance volume restore fails. This is done because EC2 instance volume restore is an important part of disaster recovery (DR) and AMS creates this internal troubleshooting RFC for you automatically.

When the internal troubleshooting RFC is created, a banner is displayed providing you with links to the RFC. This internal troubleshooting RFC provides your with more visibility into RFC failures and, rather than submitting retry RFCs leading to the same errors, or making you manually reach out to AMS for this failure, you can keep track of your changes and know that the failure is being worked on by AMS. This also reduces the time-to-recovery (TTR) metric for their change as AMS Operators proactively work on the RFC failure instead of waiting for your request.

How to get help with an RFC

You can reach out to AMS to identify the root cause of your failure. AMS business hours are 24 hours a day, 7 days a week, 365 days a year.

AMS provides several avenues for you to ask for help or make service requests.

- To ask for information or advice, or for access to an AMS-managed IT service, or to request an additional service from AMS, use the AMS console and submit a service request. For details, see Creating a Service Request. For general information about AMS service requests, see Service Request Management.
- To report an AWS or AMS service performance issue that impacts your managed environment, use the AMS console and submit an incident report. For details, see Reporting an Incident. For general information about AMS incident management, see Incident response.
- For specific questions about how you or your resources or applications are working with AMS, or to escalate an incident, email one or more of the following:
  1. First, if you are unsatisfied with the service request or incident report response, email your CSDM: ams-csdm@amazon.com
  2. Next, if escalation is required, you can email the AMS Operations Manager (but your CSDM will probably do this): ams-opsmanager@amazon.com
  3. Further escalation would be to the AMS Director: ams-director@amazon.com
  4. Finally, you are always able to reach the AMS VP: ams-vp@amazon.com

Direct Change mode in AMS

Topics

- Getting Started with Direct Change mode
- Security and compliance
- Change management in Direct Change mode
- Creating stacks using Direct Change mode
- Direct Change Mode use cases

AWS Managed Services (AMS) Direct Change mode (DCM) extends AMS Advanced change management by providing native AWS access to AMS Advanced Plus and Premium accounts to provision and update AWS resources. With DCM, you have the option to use native AWS API (console or CLI/SDK) or AMS Advanced change management requests for change (RFCs), and in either case the resources and changes to them are fully supported by AMS, including monitoring, patch, backup, incident response management. Resources provisioned through DCM are registered in the AMS service knowledge management system (SKMS), joined to the AMS managed Active Directory domain (when applicable), and run AMS management agents. Use existing tooling

(for example, CloudFormation, AWS SDK, and CDK) to develop and deploy AMS-managed CloudFormation stacks.

> **Note**
>
> Direct Change mode does not remove AMS change management RFCs. You have full access to AMS RFCs with DCM.

Watch Akash's video to learn more (6:30)

Getting Started with Direct Change mode

Begin by checking prerequisites and then submitting a request for change (RFC) in your eligible AMS Advanced account.

1. Confirm that the account that you want to use with DCM meets the requirements:

   - The account is AMS Advanced Plus or Premium.
   - The account doesn't have Service Catalog enabled. We currently don't support onboarding accounts to both DCM and Service Catalog at the same time. If you are already onboarded to Service Catalog but are interested in DCM, discuss your needs with your cloud service delivery manager (CSDM). If you decide to switch from Service Catalog to DCM, offboard Service Catalog, to do that, include the ask in the request for change below. For details about Service Catalog in AMS, see AMS and Service Catalog.

2. Submit a request for change (RFC) using the Management | Managed account | Direct Change mode | Enable change type (ct-3rd4781c2nnhp). For an example walkthrough, see Direct Change mode | Enable.

   After the CT is processed, the predefined IAM roles, `AWSManagedServicesCloudFormationAdminRole` and `AWSManagedServicesUpdateRole` are provisioned in the specified account.

3. Assign the appropriate role to the users that require DCM access using your internal federation process.

Note

You can specify any number of SAMLIdentityProviders, AWS Services, and IAM Entities (Roles, Users etc) to assume the roles. You must provide at least one: `SAMLIdentityProviderARNs`, `IAMEntityARNs`, or `AWSServicePrincipals`. For more information, consult with your company's IAM department or with your AMS cloud architect (CA).

Direct Change mode IAM roles and policies

When Direct Change mode is enabled in an account, these new IAM entities are deployed:

`AWSManagedServicesCloudFormationAdminRole`: This role grants access to the CloudFormation console, create and update CloudFormation stacks, view drift reports, and create and execute CloudFormation ChangeSets. Access to this role is managed through the your SAML provider.

Managed policies that are deployed and attached to the role `AWSManagedServicesCloudFormationAdminRole` are:

- AMS Advanced multi-account landing zone (MALZ) Application account
  - AWSManagedServices_CloudFormationAdminPolicy1
  - AWSManagedServices_CloudFormationAdminPolicy2
    - This policy represents the permissions granted to the `AWSManagedServicesCloudFormationAdminRole`. You and partners use this policy to grant access to an existing role in the account and allow that role to launch and update CloudFormation stacks in the account. This might require additional AMS service control policy (SCP) updates to allow other IAM entities to launch CloudFormation stacks.
- AMS Advanced single-account landing zone (SALZ) account
  - AWSManagedServices_CloudFormationAdminPolicy1
  - AWSManagedServices_CloudFormationAdminPolicy2
  - cdk-legacy-mode-s3-access [in-line policy]
  - AWS ReadOnlyAccess policy

`AWSManagedServicesUpdateRole`: This role grants restricted access to downstream AWS service APIs. The role is deployed with managed policies that provide mutating and non-mutating API operations, but in general restricts mutating operations (such as Create/Delete/PUT), against certain services such as IAM, KMS,GuardDuty, VPC, AMS infrastructure resources and configuration, and so forth. Access to this role is managed through the your SAML provider.

Managed policies that are deployed and attached to the role `AWSManagedServicesUpdateRole` are:

- AMS Advanced multi-account landing zone Application account
  - AWSManagedServicesUpdateBasePolicy
  - AWSManagedServicesUpdateDenyPolicy
  - AWSManagedServicesUpdateDenyProvisioningPolicy
  - AWSManagedServicesUpdateEC2AndRDSPolicy
  - AWSManagedServicesUpdateDenyActionsOnAMSInfraPolicy
- AMS Advanced single-account landing zone account
  - AWSManagedServicesUpdateBasePolicy
  - AWSManagedServicesUpdateDenyProvisioningPolicy
  - AWSManagedServicesUpdateEC2AndRDSPolicy
  - AWSManagedServicesUpdateDenyActionsOnAMSInfraPolicy1
  - AWSManagedServicesUpdateDenyActionsOnAMSInfraPolicy2

Besides these, the managed policy `AWSManagedServicesUpdateRole` role also has the AWS managed policy `ViewOnlyAccess` attached to it.

Security and compliance

Security and compliance is a shared responsibility between AMS Advanced and you, as our customer. AMS Advanced Direct Change mode does not change this shared responsibility.

Security in Direct Change mode

AMS Advanced offers additional value with a prescriptive landing zone, a change management system, and access management. When using Direct Change mode, this responsibility model does not change. However, you should be aware of additional risks.

The Direct Change Mode "Update" role (see Direct Change mode IAM roles and policies) provides elevated permissions allowing the entity with access to it, to make changes to infrastructure resources of AMS-supported services within your account. With elevated permissions, varied risks exist depending on the resource, service, and actions, especially in situations where an incorrect change is made due to oversight, mistake, or lack of adherence to your internal process and control framework.

As per AMS Technical Standards, the following risks have been identified and recommendations are made as follows. Detailed information about AMS Technical Standards is available through AWS Artifact. To access AWS Artifact, contact your CSDM for instructions or go to Getting Started with AWS Artifact.

AMS-STD-001: Tagging

Standards	Does it break	Risks	Recommendations
All the AMS owned resources must have following key-value pair All the AMS-owned tags other than those listed above must have prefixes like `AMS*` or `MC*` in upper/lower/mix case.	Yes. Breaks for CloudFormation,CloudTrail, EFS, OpenSearch, CloudWatch Logs, SQS, SSM, Tagging api - as these services do not support the `aws:TagsKey` condition to restrict tagging for the AMS namespace. Standard given in table **AMS-STD-003**, following, states that you can change AppId, Environment and AppName, but not for AMS-owned resources. Not	Incorrect tagging of AMS resources may adversly impact the reporting, alerting and patching operations of your resources, on the AMS side.	Access must be restrticted to make any changes on the AMS default tagging requirements for anyone other than AMS teams.

Standards	Does it break	Risks	Recommendations
	achievable through IAM permissions.		
Any tag on AMS-owned stacks must not be deleted based on your change requests.	Yes. CloudFormation does not support the `aws:TagsKey` condition to restrict tags for the AMS namespace.		
You are not permitted to use AMS tag naming convention in your infrastructure as mentioned in table **AMS-STD-002**, next.	Yes. Breaks for CloudFormation, CloudTrail, Amazon Elastic File System (EFS), OpenSearch, CloudWatch Logs, Amazon Simple Queue Service (SQS), Amazon EC2 Systems Manager (SSM), Tagging API; these services do not support the `aws:TagsKey` condition to restrict tagging for the AMS namespace.		

AMS-STD-002: Identity and Access Management (IAM)

Standards	Does it break	Risks	Recommendations
4.7 Actions, which bypass the change management	Yes. The purpose of self service actions allow you to perform	The secure access model is a core technical facet of	The IAM user should be time-bounded and granted permissio

Standards	Does it break	Risks	Recommendations
process (RFC), must not be permitted such as starting or stopping of an instance, creation of S3 buckets or RDS instances, and so forth. Developer mode accounts and Self-Service Provisioned mode services (SSPS) are exempted as long as actions are performed within the boundaries of the assigned role.	actions bypassing the AMS RFC system.	AMS and an IAM user for console or programmatic access circumvents this access control. The IAM users access is not monitored by AMS change management. Access is logged in CloudTrail only.	ns based on least-privilege and need-to-know.

AMS-STD-003: Network Security

Standards	Does it break	Risks	Recommendations
S2. Elastic IP on EC2 instances must be used only with a formal risk acceptance agreement, or with a valid use case by internal teams.	Yes. Self service actions allow you to associate and disassociate elastic IP addresses (EIP).	Adding an elastic IP to an instance exposes it to the Internet. This increases the risk of information disclosure and unauthorized activity.	Block any unnecessary traffic to that instance through security groups, and verify that your security groups are attached with the instance to ensure that it allows the traffic only as needed for business reasons.
S14. VPC Peering and endpoint connectio	Yes. Not possible through IAM policy.	Traffic leaving your AMS account is not	We recommend peering only with

Standards	Does it break	Risks	Recommendations
ns between accounts that belong to the same customer can be permitted.		monitored once egressing the account boundary.	AMS accounts that you own. If your use case requires this, use security groups and route tables to limit what traffic ranges, resources, and types can egress through the relevant connection.
AMS base AMIs can be shared between AMS-managed and unmanaged accounts as long as we can verify that they are owned by the same AWS organization.		AMIs may contain sensitive data and it may be exposed to unintended accounts.	Share AMIs with only the account owned by your organizat ion or validate the use-case and account information before sharing outside the organization.

AMS-STD-007: Logging

Standards	Does it break	Risks	Recommendations
19. Any log can be forwarded from one AMS account to another AMS account of the same customer.	Yes. Potential insecurity for customer logs as verification of the customer accounts being in the same organization can not be achieved through IAM policy.	Logs may contain sensitive data and it may be exposed to unintended accounts.	Share logs with only accounts managed by your AWS Org, or validate the use-case and account information before sharing outside of your organization. We can verify this via multiple ways,
20. Any log can be forwarded from AMS to a non-AMS account only if the			

Standards	Does it break	Risks	Recommendations
non-AMS account is owned by the same AMS customer (by confirming that they are under the same AWS Organizations account or by matching the email domain with the customer's company name and PAYER linked account) using internal tools.			check with your cloud service delivery manager (CSDM).

Work with your internal authorization and authentication team to control the permissions to the Direct Change mode roles accordingly.

Compliance in Direct Change mode

Direct Change mode is compatible with both production and non-production workloads. It's your responsibility to ensure adherence to any compliance standards (for example, PHI, HIPAA, PCI), and to ensure that the use of Direct Change mode complies with your internal control frameworks and standards.

Change management in Direct Change mode

Change management is the process that AMS Advanced uses to implement requests for change. A request for change (RFC) is a request created by either you, or AMS Advanced through the AMS Advanced interface to make a change to your managed environment and includes an AMS Advanced change type (CT) ID for a particular operation. For more information, see Change management.

> **Note**
>
> Direct Change mode does not remove AMS change management RFCs, you still have full access to AMS RFCs with DCM.

AMS Direct Change mode (DCM) extends AMS Advanced change management by providing native AWS access to AMS Advanced Plus and Premium accounts to provision and update AWS resources. Users who have been granted Direct Change mode permission through the IAM roles, can use native AWS API access to provision and make changes to resources in their AMS Advanced accounts. The users can still use AMS Advanced change management RFCs using the same IAM roles. In both cases the resources and changes to them are fully supported by AMS, including monitoring, patch, backup, incident response management. Users who do not have the appropriate role in these accounts must use the AMS Advanced change management RFC process to make changes.

Change management use cases

For security reasons, some changes in AMS Advanced can only be done through the change management request for change (RFC) process. The `AWSManagedServicesCloudFormationAdminRole` is restricted to actions taken through CloudFormation (CFN). For more about how to create stacks through DCM, see Creating stacks using Direct Change mode. The `AWSManagedServicesUpdateRole` is restricted for the following actions.

For example walkthroughs for each change type, including the Management | Managed account | Direct Change mode | Enable (ct-3rd4781c2nnhp) change type, see the "Additional Information" section for the relevant change type in the *AMS Advanced Change Type Reference* Change Types by Classification section.

Service	Action
AWS Key Management Service (AWS KMS)	Update
AWS Certificate Manager	Create
AWS Identity and Access Management (IAM)	Any
AWS VPN	Any

Service	Action
AMS Resource Scheduler	
AWS Backup	Create backup plan
AMS Workload Ingestion (WIGs)	Any
AMS Egress Filtering (Managed Palo Alto)	
AMS Advanced MALZ account changes	
Amazon GuardDuty	
AMS Advanced Stack Access	Any
Amazon Elastic Block Store (EBS) volume	Delete
Amazon Elastic Block Store (EBS) default encryption	Enable default encryption
Amazon Elastic Compute Cloud (Amazon EC2)	Change hostname
Amazon Machine Images (AMI)	Delete, share
Amazon EC2 Security Group	Any
AMS Advanced SSPS	
AWS Managed Microsoft AD	
AMS Advanced developer mode	
Amazon Simple Storage Service (Amazon S3)	Create S3 bucket policies
AWS Systems Manager	Create

Creating stacks using Direct Change mode

There are two requirements when launching stacks in CloudFormation using the `AWSManagedServicesCloudFormationAdminRole`, in order for the stack to be managed by AMS:

- The template must contain an `AmsStackTransform`.
- The stack name must start with the prefix `stack-` followed by a 17 character alphanumeric string.

Note

To successfully use the `AmsStackTransform`, you must acknowledge that your stack template contains the `CAPABILITY_AUTO_EXPAND` capability in order for AWS CloudFormation (CFN) to create or update the stack. You do this by passing the `CAPABILITY_AUTO_EXPAND` as part of your create-stack request. CFN rejects the request if this capability is not acknowledged when the `AmsStackTransform` is included in the template. The CFN console ensures that you pass this capability if you have the transform in your template, but this can be missed when you are interacting with CFN via their APIs. You must pass this capability whenever you use the following CFN API calls:

- CreateChangeSet
- CreateStack
- UpdateStack

When creating or updating a stack with DCM, the same validations and augmentations of CFN Ingest and Stack Update CTs are performed on the stack, for more information see AWS CloudFormation Ingest Guidelines, Best Practices, and Limitations. The exception to this is that the AMS default security groups (SGs) will not be attached to any stand-alone EC2 instances or EC2 instances in Auto Scaling groups (ASGs). When you create your CloudFormation template, with stand-alone EC2 instances or ASGs, you can attach the default SGs.

> **Note**
>
> IAM roles can now be created and managed with the `AWSManagedServicesCloudFormationAdminRole`.

The AMS default SGs have ingress and egress rules that allow the instances to launch successfully and to be accessed later through SSH or RDP by AMS operations and you. If the you find that the AMS default security groups are too permissive, you can create your own SGs with more restrictive rules and attach them to your instance, as long as it still allows you and AMS operations to access the instance during incidents.

The AMS default security groups are the following:

- SentinelDefaultSecurityGroupPrivateOnly: Can be accessed in the CFN template through this SSM parameter `/ams/${VpcId}/SentinelDefaultSecurityGroupPrivateOnly`
- SentinelDefaultSecurityGroupPrivateOnlyEgressAll: Can be accessed in the CFN template through this SSM parameter `/ams/${VpcId}/SentinelDefaultSecurityGroupPrivateOnlyEgressAll`

AMS Transform

Add a `Transform` statement to your CloudFormation template. This adds a CloudFormation macro that validates and registers the stack with AMS at launch time.

JSON Example

```
"Transform": {
    "Name": "AmsStackTransform",
    "Parameters": {
      "StackId": {"Ref" : "AWS::StackId"}
    }
  }
```

YAML Example

```
Transform:
  Name: AmsStackTransform
  Parameters:
```

```
    StackId: !Ref 'AWS::StackId'
```

Also add the `Transform` statement when updating the template of an existing stack.

JSON Example

```
{
  "AWSTemplateFormatVersion": "2010-09-09",
  "Description" : "Create an SNS Topic",
    "Transform": {
      "Name": "AmsStackTransform",
      "Parameters": {
        "StackId": {"Ref" : "AWS::StackId"}
     }
  },
  "Parameters": {
    "TopicName": {
      "Type": "String",
      "Default": "HelloWorldTopic"
    }
  },
  "Resources": {
    "SnsTopic": {
      "Type": "AWS::SNS::Topic",
      "Properties": {
        "TopicName": {"Ref": "TopicName"}
      }
    }
  }
}
```

YAML Example

```
AWSTemplateFormatVersion: '2010-09-09'
Description: Create an SNS Topic
Transform:
  Name: AmsStackTransform
  Parameters:
    StackId: !Ref 'AWS::StackId'
Parameters:
  TopicName:
    Type: String
    Default: HelloWorldTopic
```

```
Resources:
  SnsTopic:
    Type: AWS::SNS::Topic
    Properties:
      TopicName: !Ref TopicName
```

Stack name

The stack name must start with the prefix `stack-` followed by a 17 character alphanumeric string. This is to maintain compatibility with other AMS systems that operate on AMS stack IDs.

The following are examples of ways to generate compatible stack IDs:

Bash:

```
echo "stack-$(env LC_CTYPE=C tr -dc 'a-z0-9' < /dev/urandom | head -c 17)"
```

Python:

```
import string
import random

'stack-' + ''.join(random.choices(string.ascii_lowercase + string.digits, k=17))
```

Powershell:

```
"stack-" + ( -join ((0x30..0x39) + ( 0x61..0x7A) | Get-Random -Count 17  | %
 {[char]$_}) )
```

Direct Change Mode use cases

The following are uses cases for Direct Change Mode:

Resource provision and management through AWS CloudFormation

- Integrate existing CloudFormation-based tooling and processes.

Ongoing resource management and updates

- Small atomic changes with low risk.
- Changes that would otherwise run through a manual or automated RFC.

- Tooling that requires native AWS API access.
- The DCM role can be utilized if you're in the migration stage. Migration teams leverage the permissions on the DCM to create or modify stacks.
- DCM roles can be used in the CI/CD pipeline to build new AMIs, create Amazon ECS tasks, and so on.

AMS Advanced Developer mode

Topics

- Getting started with AMS Advanced Developer mode
- Security and compliance in Developer mode
- Change management in Developer mode
- Provisioning infrastructure in AMS Developer mode
- Detective controls in AMS Developer mode
- Logging, monitoring, and event management in AMS Developer mode
- Incident management in AMS Developer mode
- Patch management in AMS Developer mode
- Continuity management in AMS Developer mode
- Security and access management in AMS Developer mode

AWS Managed Services (AMS) Developer mode uses elevated permissions in AMS Advanced Plus and Premium accounts to provision and update AWS resources outside of the AMS Advanced change management process. AMS Advanced Developer mode does this by leveraging native AWS API calls within the AMS Advanced Virtual Private Cloud (VPC), enabling you to design and implement infrastructure and applications in your managed environment.

When using an account that has Developer mode enabled, continuity management, patch management, and change management are provided for resources provisioned through the AMS Advanced change management process or by using an AMS Amazon Machine Image (AMI). However, these AMS management features are not offered for resources provisioned through native AWS APIs.

You are responsible for monitoring infrastructure resources that are provisioned outside of the AMS Advanced change management process. Developer mode is compatible with both production

and non-production workloads. With elevated permissions, you have an increased responsibility to ensure adherence to internal controls.

> **Important**
>
> Resources that you create using Developer mode can be managed by AMS Advanced only if they are created using AMS Advanced change management processes.

Developer mode is one of the AMS Advanced modes you can employ. For more information, see Modes overview.

Getting started with AMS Advanced Developer mode

Learn the various AMS Advanced accounts with AMS Advanced Developer mode and how to successfully implement Developer mode.

Topics

- Before you begin with AMS Developer mode
- Prerequisites for AMS Developer mode
- How to implement AMS Advanced Developer mode
- AMS Advanced Developer mode permissions

Before you begin with AMS Developer mode

Before implementing Developer mode, there are a few things you should know.

AMS Advanced cannot manage existing stacks or resources in a DevMode account that were created outside of the AMS Advanced change management process through requests for change (RFCs). However, while the account is in DevMode, AMS Advanced continues to manage resources provisioned through the AMS Advanced change management process with RFCs.

You cannot start with a DevMode account and later covert it to an AMS Advanced-managed application account.

Prerequisites for AMS Developer mode

The following are the prerequisites for implementing Developer mode:

- You must be an AMS Advanced customer with at least one onboarded AMS Advanced Plus or Premium account.
- Any account you use must be an AMS Advanced Plus or Premium account.
- **Multi-Account Landing Zone (MALZ)**: You must use the `AWSManagedServicesDevelopmentRole` predefined AWS Identity and Access Management (IAM) role. You request this role. The next section describes how to acquire Developer mode permissions.
- **Single-Account Landing Zone (SALZ)**: You must use the `customer_developer_role` predefined AWS Identity and Access Management (IAM) role. You request this role. The next section describes how to acquire Developer mode permissions.

How to implement AMS Advanced Developer mode

You implement Developer mode by requesting that your eligible AMS Advanced account be provisioned with the predefined IAM role:

- **MALZ**: `AWSManagedServicesDevelopmentRole`
- **SALZ**: `customer_developer_role`

You then assign the role to the relevant users in your federated network.

AMS Advanced recommends that you ensure that your use of Developer mode complies with your internal control frameworks and standards as Developer mode creates two vectors of change: AMS Advanced change management for AMS Advanced-managed resources and customer-managed role federation for resources that you, as our customer, manage. While AMS Advanced processes remain compliant with our declarations, customer processes and control frameworks might need to be updated.

To implement Developer mode in your AMS Advanced account

1. Confirm the account that you want to use with Developer mode meets the requirements listed in Prerequisites for AMS Developer mode.
2. Submit a request for change (RFC) using the change type (CT) Management | Managed account | Developer mode | Enable (review required). For an example of how to use this CT, see Developer Mode | Enable (Review Required).

After the CT is processed, the predefined IAM role, (`AWSManagedServicesDevelopmentRole` for **MALZ**, `customer_developer_role` for **SALZ**), is provisioned in the requested account.

3. Assign the appropriate role to the users that require Developer mode access using your internal federation process.

   AMS Advanced recommends that you limit access to prevent unwanted or unapproved provisioning of, or changes to, resources.

AMS Advanced Developer mode permissions

The predefined role (`AWSManagedServicesDevelopmentRole` for **MALZ**, `customer_developer_role` for **SALZ**), grants permission to create application infrastructure resources within the AMS Advanced VPC, including IAM roles, while restricting access to *shared service* components that are operated by AMS Advanced (for example, management hosts, domain controllers, Trend Micro EPS, bastions, and unsupported AWS services). The role also restricts access to the following AWS services: Amazon GuardDuty, AWS Organizations, AWS Directory Service APIs, and AMS Advanced logs.

While the role allows you to create additional IAM roles, the same permissions boundaries included in Developer mode access are enforced on any IAM role created by the `AWSManagedServicesDevelopmentRole`.

Security and compliance in Developer mode

Security and compliance is a shared responsibility between AMS Advanced and you as our customer. AMS Advanced Developer mode shifts the shared responsibility to you for resources provisioned outside of the change management process or provisioned through change management but updated with Developer mode permissions. For more information about shared responsibility, see AWS Managed Services.

Cautions:

- DevMode allows you and your authorized team to bypass the deny-by-default principles at the core of AMS security. The advantages, self-service, less time waiting for AMS must be weighed against the disadvantages, anyone can perform unexpected and destructive actions without the knowledge of their security team. Automated change types to enable Dev mode and Direct

Change mode are exposed, and any authorized person in your org can run these CTs and enable these modes.

- You are responsible for managing the permissions of CT execution from your user base.
- AMS doesn't manage CT execution permissions

Recommendations:

- **Protect**
  - Customers can prevent access to this CT via permissioning, see Restrict permissions with IAM role policy statements
  - Prevent access to this CT by implementing a proxy such as an ITSM system
  - Utilize service control policies (SCPs) that prevent policies and behaviors as needed, see AMS Preventative and Detective Controls Library
- **Detect**
  - Monitor your RFC's for these CTs (Enable developer mode ct-1opjmhuddw194 and Direct change mode, Enable ct-3rd4781c2nnhp) being executed and respond accordingly
  - Review and/or audit your accounts for the presence of the IAM resources to identify those accounts where Developer mode or Direct Change mode have been deployed
- **Respond**
  - Remove accounts in Developer mode as needed

Security in Developer mode

AMS Advanced offers additional value with a prescriptive landing zone, a change management system, and access management. When using Developer mode the security value of AMS Advanced is persisted by using the same account configuration of standard AMS Advanced accounts that establishes the baseline AMS Advanced security hardened network. The network is protected by the permissions boundary enforced in the role (`AWSManagedServicesDevelopmentRole` for **MALZ**, `customer_developer_role` for **SALZ**), which restricts the user from breaking down the parameter protections established when the account is set up.

For example, users with the role can access Amazon Route 53 but AMS Advanced internal hosted zone is restricted. The same permissions boundaries are enforced on an IAM role created by the `AWSManagedServicesDevelopmentRole`, enforcing permissions boundaries on the

`AWSManagedServicesDevelopmentRole` that restricts the user from breaking down the parameter protections established when the account is onboarded to AMS Advanced.

Compliance in Developer mode

Developer mode is compatible with both production and non-production workloads. It's your responsibility to ensure adherence to any compliance standards (for example, PHI, HIPAA, PCI), and to ensure that the use of Developer mode complies with your internal control frameworks and standards.

Change management in Developer mode

Change management is the process the AMS Advanced service uses to implement requests for change. A request for change (RFC) is a request created by either you or AMS Advanced through the AMS Advanced interface to make a change to your managed environment and includes a change type (CT) ID for a particular operation. For more information, see Change management modes.

Change management is not enforced in AMS Advanced accounts where Developer mode permissions are granted. Users who have been granted Developer mode permission with the IAM role (`AWSManagedServicesDevelopmentRole` for **MALZ**, `customer_developer_role` for **SALZ**), can use native AWS API access to provision and make changes to resources in their AMS Advanced accounts. Users who do not have the appropriate role in these accounts must use the AMS Advanced change management process to make changes.

> **Important**
>
> Resources that you create using Developer mode can be managed by AMS Advanced only if they are created using AMS Advanced change management processes. Requests for changes submitted to AMS Advanced for resources created outside of the AMS Advanced change management process are rejected by AMS Advanced because they must be handled by you.

Self-service provisioning services API restrictions

All AMS Advanced self-provisioned services are supported with Developer mode. Access to self-provisioned services are subject to the limitations outlined in the respective user guide sections for each. If a self-provisioned service is not available with your Developer mode role, you can request an updated role through the Developer mode change type.

The following services do not provide full access to service APIs:

Self-Provisioned Services Restricted in Developer mode

Service	Notes
Amazon API Gateway	All Gateway APIs calls are allowed except `SetWebACL` .
Application Auto Scaling	Can only register or deregister scalable targets, and put or delete a scaling policy.
AWS CloudFormation	Can't access or modify CloudFormation stacks that have a name prefixed with `mc-`.
AWS CloudTrail	Can't access or modify CloudTrail resources that have a name prefixed with `ams-` and/or `mc-`.
Amazon Cognito (User Pools)	Can't associate software tokens. Can't create user pools, user import jobs, resource servers, or identity providers.
AWS Directory Service	Only the following AWS Directory Service actions are required by `Connect` and `WorkSpaces` services. All other Directory Service actions are denied by the Developer mode permission boundary policy: • `ds:AuthorizeApplication` • `ds:CreateAlias` • `ds:CreateIdentityPoolDirectory` • `ds:DeleteDirectory` • `ds:DescribeDirectories` • `ds:GetAuthorizedApplication Details` • `ds:ListAuthorizedApplications`

Service	Notes
	• `ds:UnauthorizeApplication` In single-account landing zone accounts, the boundary policy explicitly denies access to the AMS Advanced managed directory used by AMS Advanced for maintaining access to dev-mode enabled accounts.
Amazon Elastic Compute Cloud	Can't access Amazon EC2 APIs that contain the string: `DhcpOptions` , `Gateway`, `Subnet`, `VPC`, and `VPN`. Can't access or modify Amazon EC2 resources that have a tag prefixed with `AMS`, `mc`, `ManagementHostASG` , and/or `sentinel`.
Amazon EC2 (Reports)	Only view access is granted (cannot modify). Note: Amazon EC2 Reports is moving. The **Reports** menu item will be removed from the Amazon EC2 console navigation menu. To view your Amazon EC2 usage reports after it has been removed, use the AWS Billing and Cost Management console.

Service	Notes
AWS Identity and Access Management (IAM)	Can't delete existing permission boundaries, or modify IAM user password policies. Can't create or modify IAM resources unless you are using the correct IAM role (`AWSManagedServicesDevelopmentRole` for **MALZ**, `customer_developer_role` for **SALZ**)). Can't modify IAM resources that are prefixed with: `ams`, `mc`, `customer_deny_policy` , and/or `sentinel`. When creating a new IAM resource (role, user, or group), the permission boundary (**MALZ**: `AWSManagedServicesDevelopmentRolePermissionsBoundary` , **SALZ**: `ams-app-infra-permissions-boundary`) must be attached.
AWS Key Management Service (AWS KMS)	Can't access or modify AMS Advanced-managed KMS keys.
AWS Lambda	Can't access or modify AWS Lambda functions that are prefixed with `AMS`.
CloudWatch Logs	Can't access CloudWatch log streams that a name prefixed with: `mc`, `aws`, `lambda`, and/or `AMS`.
Amazon Relational Database Service (Amazon RDS)	Can't access or modify Amazon Relational Database Service (Amazon RDS) databases (DBs) that have a name prefixed with: `mc-`.
AWS Resource Groups	Can only access `Get`, `List`, and `Search` Resource Group API actions.

Service	Notes
Amazon Route 53	Can't access or modify Route53 AMS Advanced-maintained resources.
Amazon S3	Can't access Amazon S3 buckets that have a name prefixed with: `ams-*`, `ams`, `ms-a`, or `mc-a`.
AWS Security Token Service	The only security token service API allowed is `DecodeAuthorizationMessage` .
Amazon SNS	Can't access SNS topics that have a name prefixed with: `AMS-`, `Energon-Topic` , or `MMS-Topic` .
AWS Systems Manager Manager (SSM)	Can't modify SSM parameters that are prefixed with `ams`, `mc`, or `svc`. Can't use the SSM API `SendCommand` against Amazon EC2 instances that have a tag prefixed with `ams` or `mc`.
AWS Tagging	You only have access to AWS Tagging API actions that are prefixed with `Get`.

Service	Notes
AWS Lake Formation	The following AWS Lake Formation API actions are denied: • `lakeformation:DescribeResource` • `lakeformation:GetDataLakeSettings` • `lakeformation:DeregisterResource` • `lakeformation:RegisterResource` • `lakeformation:UpdateResource` • `lakeformation:PutDataLakeSettings`
Amazon Elastic Inference	You can only call the Elastic Inference API action `elastic-inference:Connect` . This permission is included in the `customer_sagemaker_admin_policy` that is attached to the `customer_sagemaker_admin_role` . This action gives you access to the Elastic Inference accelerator.
AWS Shield	No access to any of this services APIs or console.
Amazon Simple Workflow Service	No access to any of this services APIs or console.

Provisioning infrastructure in AMS Developer mode

Users that don't have the Developer mode IAM role, `AWSManagedServicesDevelopmentRole`, in accounts where Developer mode is enabled, are required to follow the AMS Advanced change management process that leverages AMS Advanced AMIs. Users with correct role (**MALZ**: `AWSManagedServicesDevelopmentRole`, **SALZ**: `customer_developer_role`) can use the AMS Advanced change management system and AMS Advanced AMIs but are not required to.

> **Note**
>
> An AWS AMI, that has not been processed through AMS Advanced workload ingestion, or created in an AMS Advanced account, will not include AMS Advanced-required configurations.

Detective controls in AMS Developer mode

This section has been redacted because it contains sensitive AMS security-related information. This information is available through the AMS console **Documentation**. To access AWS Artifact, you can contact your CSDM for instructions or go to Getting Started with AWS Artifact.

Logging, monitoring, and event management in AMS Developer mode

Logging, monitoring, and event management aren't available for resources provisioned outside of the AMS Advanced change management process, or for resources provisioned through change management and then altered by an account using Developer mode permissions.

Incident management in AMS Developer mode

No change to incident response times. Incident resolution is a best effort for resources provisioned outside the change management process, or resources provisioned through change management and then altered by an account using Developer mode permissions.

> **Note**
>
> AMS service level agreement (SLA) does not apply for resources created or updated outside of the AMS change management system (requests for change or RFCs), Developer mode included; therefore, resources updated or created in Developer mode are automatically degraded to a P3 and AMS support is best effort.

Patch management in AMS Developer mode

Patch management is not available for resources provisioned outside of the AMS Advanced change management process, or for resources provisioned through change management and then altered by an account using Developer mode permissions. Patching times:

- For a critical security update: Within 10 business days of release by the vendor for resources provisioned through change management and then altered by an account using Developer mode permissions.
- For an important update: Within 2 months of release by the vendor for resources provisioned through change management and then altered by an account using Developer mode permissions.

Continuity management in AMS Developer mode

Continuity management is not available for resources provisioned outside of the AMS Advanced change management process, or for resources provisioned through change management and then altered by an account using Developer mode permissions.

Environment recovery initiation time can take up to 12 hours for resources provisioned outside of the AMS Advanced change management process, or for resources provisioned through change management and then altered by an account using Developer mode permissions.

Security and access management in AMS Developer mode

Anti-malware protection is your responsibility for resources provisioned outside of the AMS Advanced change management process, or for resources provisioned through change management and then altered by an account using Developer mode permissions. Access to Amazon Elastic Compute Cloud (Amazon EC2) instances not provisioned through AMS Advanced change management might be controlled by key pairs instead of providing federated access.

Self-Service Provisioning mode

AWS Managed Services (AMS) Self-Service Provisioning (SSP) mode provides full access to native AWS service and API Capabilities in AMS managed accounts. You access services through standardized, scoped down, IAM roles. AMS provides service requests and incident management. Alerting, monitoring, logging, patch, back up, and change management are your responsibility. In many cases, Self-Service Provisioning services (SSPS) are self-managed, or serverless, and don’t require management of certain operational tasks like patching. You benefit from using these services within the environment boundary defined by AMS guardrails and any IAM changes (including service linked roles, service roles, cross-account roles, or policy updates) need to be approved by AMS Operations to maintain the baseline security of the platform. You can leverage

CloudFormation templates to automate deployment of these services but this is not supported for all SSP services currently.

> **Important**
>
> Use SSP mode in your AWS Managed Services (AMS) accounts to access and employ AWS services, with restrictions as noted.

There are some AWS services that you can use without AMS management, in your AMS account. The Self-Service Provisioning mode services, or SSPS for short, how to add them into your AMS account and FAQs for each, are described here.

Self-service provisioning services are offered as is, and you're responsible for managing them. AMS provides no alerts, monitoring, logging, or patching for the resources associated with those services. AMS provides IAM roles that enable you to use the service in your AMS account safely. AMS SLAs do not apply.

For resources that you provision through self-service, AMS provides incident management, detective controls and guardrails, reporting, designated resources (Cloud Service Delivery Manager and Cloud Architect), security and access, and technical support through service requests. Additionally, where applicable, you assume responsibility for continuity management, patch management, infrastructure monitoring, and change management for resources provisioned or configured outside of the AMS change management system.

Getting started with SSP mode

Self-service provisioning is one of the AMS modes for multi-account landing zone (MALZ) that you can employ. For more information, see Modes overview.

To provide self-service provisioning capabilities, AMS has created elevated IAM roles with permission boundaries to limit unintended changes from direct AWS service access. The roles do not prevent all changes and you are responsible to adhere to your internal controls and compliance policies, and to validate that all AWS services being used meet the required certifications. We call this the self-service provisioning mode. For details on AWS compliance requirements, see AWS Compliance.

To add a self-service provisioning service to your multi-account landing zone Application account, use the **Management | AWS service | Self-provisioned service | Add** change type (CT), either the review-required CT or automated CT, as instructed per service.

Currently, these are the self-service provisioning service (SSPS) options, you can choose from those listed.

> **Note**
>
> To request that AMS provide an additional self-service provisioning service, file a service request.

Amazon API Gateway (AMS SSPS)

Amazon API Gateway is a fully managed service that makes it easy for developers to create, publish, maintain, monitor, and secure APIs at any scale. Using the AWS Management Console you can create REST and WebSocket APIs that act as a front door for applications to access data, business logic, or functionality from your back-end services, such as workloads running on Amazon Elastic Compute Cloud (Amazon EC2), code running on AWS Lambda, any web application, or real-time communication applications.

API Gateway handles all the tasks involved in accepting and processing up-to hundreds of thousands of concurrent API calls, including traffic management, authorization and access control, monitoring, and API version management. API Gateway has no minimum fees or startup costs. You pay only for the API calls you receive and the amount of data transferred out and, with the API Gateway tiered pricing model, you can reduce your cost as your API usage scales. To learn more, see Amazon API Gateway.

FAQs: API Gateway in AMS

Q: How do I request access to Amazon API Gateway in my AMS account?

Request access to API Gateway by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM roles to your account: `customer_apigateway_author_role` and `customer_apigateway_cloudwatch_role`. After provisioned in your account, you must onboard the roles in your federation solution.

Q: What are the restrictions to using Amazon API Gateway in my AMS account?

- API Gateway configuration is limited to resources without `AMS-` or `MC-` prefixes to prevent any modifications to AMS infrastructure.

- CREATE privileges for VPCLink are disabled in order to prevent unregulated creation of Elastic Load Balancers. If VPCLinks are required, see Application Load Balancer | Create.

Q: What are the prerequisites or dependencies to using Amazon API Gateway in my AMS account?

It depends on the type of API Gateway you want to deploy. It can be a standalone service, but it can also request access to existing services (for instance, network loadbalancer).

Alexa for Business (AMS SSPS)

Alexa for Business is a service that enables your organization and employees to use Alexa to get more work done. With Alexa for Business, you can use Alexa as your intelligent assistant to be more productive in meeting rooms, at your desk, and even with the Alexa devices you already use at home or on the go. IT and facilities managers can use Alexa for Business to measure and increase the utilization of the existing meeting rooms in their workplace.

To learn more, see Alexa for Business.

Alexa for Business in AWS Managed Services FAQs

Q: How do I request access to Alexa for Business in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_alexa_console_role`. A `customer_alexa_device_setup_user` is also created for the Device Setup Tool provided by Alexa for Business; this Device Setup Tool can then be used to set up your devices. Once provisioned in your account, you must onboard the roles in your federation solution.

The Alexa for Business gateway enables you to connect Alexa for Business to your Cisco Webex and Poly Group Series endpoints to control meetings with your voice. The gateway software runs on your on-premises hardware and securely proxies conferencing directives from Alexa for Business to your Cisco endpoint. The gateway needs two pairs of AWS credentials to communicate with Alexa for Business. We provide two limited-access IAM users: `customer_alexa_gateway_installer_user` and `customer_alexa_gateway_execution_user` for your Alexa for Business gateways, one for installing the gateway and one for operating the gateway; these can be requested by submitting an RFC with the Management | Other | Other change type.

> **Note**
>
> To generate usage reports and send them to Amazon S3, specify the Amazon S3 bucket name in the self-provisioned service RFC.

Q: What are the restrictions to using Alexa for Business in my AMS account?

There are no restrictions. Full functionality of Alexa for Business is available with the Alexa for Business self-provisioned service role.

Q: What are the prerequisites or dependencies to using Alexa for Business in my AMS account?

If you intend to use WPA2 Enterprise Wi-Fi to set up your shared devices, please specify this network security type in the Device Setup Tool, for which an AWS Private Certificate Authority is required.

What Alexa for Business functionality requires separate RFCs?

To register an Alexa Voice Service (AVS) device with Alexa for Business, provide access to the Alexa built-in device maker. To do this, an IAM role needs to be created in the Alexa for Business console that can be deployed using the Management | Other | Other change type. This allows the AVS device maker to register and manage devices with Alexa for Business on your behalf.

Amazon AppStream 2.0 (AMS SSPS)

Amazon AppStream 2.0 (AppStream 2.0) lets you move your desktop applications to AWS, without rewriting them. You can install your applications on AppStream 2.0, set launch configurations, and make your applications available to users. AppStream 2.0 offers a wide selection of virtual machine options so that you can select the instance type that best matches your application requirements, and set the auto-scale parameters so that you can easily meet the needs of your end users. AppStream 2.0 enables you to launch applications in your own network, which means your applications can interact with your existing AWS resources.

Amazon AppStream 2.0 enables you to quickly and easily install, test, and update your applications using the image builder. Any application that runs on Microsoft Windows Server 2012 R2, Windows Server 2016, or Windows Server 2019 is supported, and you don’t need to make any modifications. When your testing is complete, you can set application launch configurations, default user settings, and publish your image for users to access.

To learn more, see AppStream 2.0.

AppStream 2.0 in AWS Managed Services FAQs

Q: How do I request access to AppStream 2.0 in my AMS account?

Request access to AppStream 2.0 by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer_appstream_console_role`.

A `customer_appstream_stream_role` is also deployed to stream applications that require users to be authenticated using their Active Directory login credentials.

Once provisioned in your account, you must onboard the roles in your federation solution.

Q: What are the restrictions to using AppStream 2.0 in my AMS account?

- The following functionality must be configured by the AMS Support team, and requires specific RFCs. Instruction on requesting additional functionality can be found in section 4.
  - Creating and Streaming from Interface VPC Endpoints.
  - Support for Amazon S3 endpoints for home folders and application setting persistence on a private network.
  - Creating and choosing the IAM role that will be available on all fleet streaming instances.
  - Joining AppStream 2.0 fleets and image builders Microsoft Active Directory domains.
  - Creating AppStream 2.0 Custom Usage Reports.
  - Custom branding is currently not supported.

Q: What are the prerequisites or dependencies to using AppStream 2.0 in my AMS account?

There are no prerequisites to using AppStream 2.0 in your AMS account.

Q: What AppStream 2.0 functionality requires separate RFCs?

- In order to choose an interface VPC endpoint for AppStream 2.0, submit a Management | Other | Other | Update change type RFC to create a VPC endpoint in your account. For steps to create custom endpoints for AppStream 2.0, see Creating and Streaming from Interface VPC Endpoints in the AppStream 2.0 user guide.
- Support for Amazon S3 endpoints for home folders and application setting persistence on a private network can be configured by requesting Amazon S3 VPC endpoints with a Management

| Other | Other | Create change type RFC. The RFC must include the target Amazon S3 bucket hosting the home folder contents, or application settings Amazon S3 buckets, respectively. This RFC will provide AppStream 2.0 the permissions it needs to access Amazon S3 VPC endpoints. For steps to create custom endpoints for streams, see Using Amazon S3 VPC Endpoints for Home Folders and Application Settings Persistence in the AppStream 2.0 user guide.

- In order to create and choose an IAM role that will be available on all fleet streaming instances, submit a Management | Other | Other | Create change type RFC requesting the IAM role with the required policy. The IAM role name should always start with prefix : "customer_appstream".
- Amazon AppStream 2.0 fleets and image builders can be joined to domains in Microsoft Active Directory by submitting a Management | Other | Other | Update change type RFC for the Service Account creation in Active Directory (AD). Minimal permissions required to join Microsoft Active Directory are defined in the AppStream 2.0 documentation at Granting Permissions to Create and Manage Active Directory Computer Objects.
- In order to create custom AppStream 2.0 Usage Reports, submit a Management | Other | Other | Create change type RFC requesting following:
  - "AppStreamUsageReports" CFN stack creation
  - "customer_appstream_usagereports_role" be provisioned in the account
  - Also, provide the following details:
    - Provide CRON expression to schedule Crawler run. By default it is 23:00 UTC everyday.
    - Amazon S3 bucket ARN to be used for athena query results. This bucket should have prefix: "aws-athena-query-results"
    - Amazon S3 bucket ARN for AppStream Usage Reports Logs.

  After the role is provisioned, onboard the role into your federation solution and login, then access AWS GlueAWS Glue and Athena for generating custom reports using the usage report role. For details about using AppStream 2.0 Usage Reports see Create Custom Reports and Analyze AppStream 2.0 Usage Data, in the AppStream 2.0 documentation.

Amazon Athena (AMS SSPS)

Amazon Athena (Athena) is an interactive query service that helps you to analyze data in Amazon S3 using standard SQL. Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run. You point to your data in Amazon S3, define the schema, and start querying using standard SQL. Most results are delivered within seconds. With Athena, there's no need for complex exact-trandform-load (ETL) jobs to prepare your data for analysis. This makes

it straight-forward for anyone with SQL skills to quickly analyze large-scale datasets. To learn more, see Amazon Athena.

FAQs: Athena in AMS

Q: How do I request access to Amazon Athena in my AMS account?

Request access to Athena by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer_athena_console_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon Athena in my AMS account?

There are no restrictions. Full functionality of Amazon Athena is available in your AMS account.

Q: What are the prerequisites or dependencies to using Amazon Athena in my AMS account?

Athena has a major dependency on the AWS Glue service, as it uses the data catalog/metastore created with AWS Glue. Therefore, AWS Glue permissions are included in the successful Athena RFC.

The role `customer_athena_console_role` has a prerequisite for an Amazon S3 bucket. To create a new bucket, use the automated CT `ct-1a68ck03fn98r` (Deployment | Advanced stack components | S3 storage | Create). When you use this automated CT to create an S3 bucket for Athena, the bucket name must begin with prefix `athena-query-results-*`.

Amazon CloudSearch (AMS SSPS)

Amazon CloudSearch is a managed service in the AWS Cloud that you use to cost-effective to set up, manage, and scale a search solution for your website or application. Amazon CloudSearch supports 34 languages and popular search features such as highlighting, autocomplete, and geospatial search. To learn more, see Amazon CloudSearch.

Amazon CloudSearch in AWS Managed Services FAQs

Q: How do I request access to Amazon CloudSearch in my AMS account?

Request access to Amazon CloudSearch by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM roles to your account: `customer_csearch_admin_role` and

`customer_csearch_dev_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon CloudSearch in my AMS account?

Full functionality of Amazon CloudSearch is available in your AMS account. All AMS-supported database solutions are currently supported on Amazon CloudSearch. Note that, currently, DynamoDB is the only managed AWS database solution that can’t be indexed.

Q: What are the prerequisites or dependencies to using Amazon CloudSearch in my AMS account?

Amazon CloudSearch depends on Amazon S3 working with Identity Providers to automatically analyze input data and determine the table fields. Access to Amazon S3 is not provided with this RFC, and must be requested separately in a service request.

Amazon CloudWatch Synthetics (AMS SSPS)

You can use Amazon CloudWatch Synthetics to create 'canaries' to monitor your endpoints and APIs.

Canaries are configurable scripts, written in Node.js or Python, that run on a schedule. They create Lambda functions in your account that use Node.js or Python as a framework. Canaries work over both HTTP and HTTPS protocols. Canaries check the availability and latency of your endpoints and can store load time data and UI screenshots. They monitor your REST APIs, URLs, and website content, and they can check for unauthorized changes from phishing, code injection and cross-site scripting.

Canaries follow the same routes and perform the same actions as a customer, making it possible for you to continually verify your customer experience even when you don't have any customer traffic on your applications. By using canaries, you can discover issues before your customers do. To learn more, see Amazon CloudWatch: Using synthetic monitoring.

Amazon CloudWatch Synthetics in AWS Managed Services FAQs

Q: How do I request access to Amazon CloudWatch Synthetics in my AMS account?

Request access to Amazon CloudWatch Synthetics by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: 'customer_cw_synthetics_console_role' and

'customer_cw_synthetics_canary_lambda_role'. Once provisioned in your account, you must onboard the 'customer_cw_synthetics_console_role' role in your federation solution.

Q: What are the restrictions to using Amazon CloudWatch Synthetics in my AMS account?

There are no restrictions for the use of Amazon CloudWatch Synthetics in your AMS account. Creating roles for canaries outside of the AMS-provided service role 'customer_cw_synthetics_canary_lambda_role' is prohibited.

Q: What are the prerequisites or dependencies to using Amazon CloudWatch Synthetics in my AMS account?

Canaries create and use a default Amazon CloudWatch Synthetics S3 bucket: "cw-syn-results-*${accountnumber}*-*${default-region}*"

Amazon Cognito (User Pools) (AMS SSPS)

Amazon Cognito user pools provide a secure user directory that scales to hundreds of millions of users. As a fully managed service, Amazon Cognito user pools can be set up without any worries about standing up server infrastructure. This service enables you to manage a pool of final users that you can use to integrate with your internal applications. This service provides you an alternative to a customized database or a directory of final users for web or mobile applications. At the same time, Amazon Cognito user pools provides the full set of functionalities of a directory service like passwords policies, multi factor authentication, password recovery and self-sign up into services. It also allows the application to federate the access in other popular public services like OpenID, Facebook, Amazon or Google.

Amazon Cognito is divided into two main products. Amazon Cognito user pools and Amazon Cognito Identity Provider. This section focuses on Amazon Cognito user pools, which provide access to other AWS services like Amazon S3 or DynamoDB. The service allows you to use Amazon Cognito user pools, or a third party identity provider, to provide access to AWS services. It also provides access to AWS services using anonymous guest access. Because of the powerful nature of Amazon Cognito user pools, it would be managed manually on a case-by-case basis as an operation manual service, in order to avoid potential security breaks into the account. To learn more, see Amazon Cognito User Pools.

Amazon Cognito user pools in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to Amazon Cognito user pools in my AMS account?

Implementation of Amazon Cognito user pools in AMS is a 2 step process:

1. Submit a Management | Other | Other | Create (ct-1e1xtak34nx76) change type and request the creation of the Amazon Cognito user pools in your AMS Account. Include the following information:
   - AWS Region.
   - Name for the Cognito User Pool.
   - If the you want to use the Amazon Simple Email Service (Amazon SES) to send messages and notifications instead of the default internal Cognito mail service, then the customer should provide an already validated email address for the Amazon SES Service in the account. This address will be used for the "From" and "REPLY-TO" fields of the message. They must also indicate the Region where Amazon SES was activated (us-east-1, eu-west-1 or us-west-2).
   - If the you want to use SMS messages for one-time passwords and verification, then the customer should indicate so.
2. Request user access by submitting a Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM roles to your account: `customer_cognito_admin_role` and `customer_cognito_importjob_role`. After it's provisioned in your account, you must onboard the role in your federation solution. These roles allow you to manage the Cognito User Pool, manage your users and groups in the pool, create importjobs for users, modify the notification and subscription messages, associate applications to the user pool, add federation services to the pool in a self-managed fashion, and delete already created pools.

Q: What are the restrictions to using Amazon Cognito user pools in my AMS account?

You won't be able to create the Amazon Cognito user pools. That action requires the creation of IAM roles to leverage services used by Amazon Cognito, like Amazon SES and Amazon Simple Notification Service (Amazon SNS).

Q: What are the prerequisites or dependencies to using Amazon Cognito user pools in my AMS account?

If you want to use Amazon SES to send messages and notifications by email to your user pools, they should already activate the Amazon SES service in the account, and already validate the email address that should be used in the "FROM" and "REPLY-TO" fields of the sent emails. For more information about validating email address using Amazon SES, see Verifying Email Addresses in Amazon SES.

Amazon Comprehend (AMS SSPS)

Amazon Comprehend is a natural language processing (NLP) service that uses machine learning to find insights and relationships in text, no machine learning experience is required. Amazon Comprehend uses machine learning to help you uncover the insights and relationships in your unstructured data. The service identifies the language of the text; extracts key phrases, places, people, brands, or events; understands how positive or negative the text is; analyzes text using tokenization and parts of speech; and automatically organizes a collection of text files by topic. You can also use AutoML capabilities in Amazon Comprehend to build a custom set of entities or text classification models that are tailored uniquely to your organization's needs. To learn more, see Amazon Comprehend.

Amazon Comprehend in AWS Managed Services FAQs

Q: How do I request access to Amazon Comprehend in my AMS account?

Amazon Comprehend console and data access roles can be requested through the submission of two AMS Service RFCs:

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_comprehend_console_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon Comprehend in my AMS account?

Create New IAM Role functionality through the Amazon Comprehend console is restricted. Otherwise, full functionality of Amazon Comprehend is available in your AMS account.

Q: What are the prerequisites or dependencies to using Amazon Comprehend in my AMS account?

Amazon S3 and AWS Key Management Service (AWS KMS) are required in order to use Amazon Comprehend, if Amazon S3 buckets are encrypted with AWS KMS keys.

Amazon Connect (AMS SSPS)

Amazon Connect is an omnichannel cloud contact center that helps companies provide superior customer service at a lower cost. Amazon Connect provides a seamless experience across voice and chat for customers and agents. This includes one set of tools for skills-based routing, powerful

real-time and historical analytics, and easy-to-use intuitive management tools – all with pay-as-you-go pricing.

You can create one or more instances of the virtual contact center instances in either AMS multi-account landing zone or single-account landing zone accounts. You can use existing SAML 2.0 identity providers for agent access or use Amazon Connect native support for user life cycle management.

Additionally, you can claim toll free/direct dial phone numbers for each Amazon Connect instance from the Amazon Connect console. You can create rich contact flows to achieve the desired customer experience and routing using an easy-to-use graphical user interface. The contact flows can leverage AWS Lambda functions to integrate with on-premises data stores and API's. You can also enable data streaming using Kinesis Streams and Firehose.

The call recordings, chat transcripts, and reports, are stored in an Amazon S3 bucket encrypted using an AWS KMS key. The contact flow logs can be saved to CloudWatch log groups.

To learn more, see Amazon Connect.

Amazon Connect in AWS Managed Services FAQs

Q: How do I request access to Amazon Connect in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM roles to your account: `customer_connect_console_role` and `customer_connect_user_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon Connect in my AMS account?

There are no restrictions. Full functionality of Amazon Connect is available in your AMS account.

Q: What are the prerequisites or dependencies to using Amazon Connect in my AMS account?

- You must create an AWS KMS Key and an Amazon S3 bucket using standard AMS RFCs; the Amazon S3 bucket is required for storing call recordings and chat transcripts.
- If you want to integrate with Active Directory (AD), an AD Connector is required for integration between AMS-hosted Amazon Connect instances and your on-premises directory services. AD Connector can be configured in your account by requesting a 'Management | Other | Other' RFC.
- You can enable the following optional self-provisioned services based on your contact flow requirements.

- **AWS Lambda**: You can use Lambda functions to extend the contact flows to leverage existing on-premises data stores or APIs. You can use the Lambda self-provisioned service to create the Lambda functions.
- **Amazon Kinesis Data Streams**: You can create data streams to enable Data streaming to external applications. You can stream contact trace records or Agent Events.
- **Amazon Kinesis Data Firehose**: You can create Data Firehose to stream high volume contact trace records to external applications.
- **Amazon Lex**: You can leverage Amazon Lex Chatbots to create smart contact flows leveraging Amazon Alexa services for rich customer experience and automation.
- **Q: How do I request to add list of countries for outbound or inbound calls?**

  To add a list of countries for outbound or inbound calls, submit a service request to AMS.

Amazon DevOps Guru (AMS SSPS)

Amazon DevOps Guru is a fully managed operations service that makes it easy for developers and operators to improve the performance and availability of their applications. DevOps Guru lets you offload the administrative tasks associated with identifying operational issues so that you can quickly implement recommendations to improve your application. DevOps Guru creates reactive insights you can use to improve your application now. It also creates proactive insights to help you avoid operational issues that might affect your application in the future. DevOps Guru applies machine learning to analyze your operational data and application metrics and events to identify behaviors that deviate from normal operating patterns. You are notified when DevOps Guru detects an operational issue or risk. For each issue, DevOps Guru presents intelligent recommendations to address current and predicted future operational issues.

To learn more, see What is Amazon DevOps Guru.

Amazon DevOps Guru in AWS Managed Services FAQs

Q: How do I request access to Amazon DevOps Guru in my AMS account?

To request access, submit a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_devopsguru_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon DevOps Guru in my AMS account?

There are no restrictions. Full functionality of Amazon DevOps Guru is available in your AMS account.

Q: What are the prerequisites or dependencies to using Amazon DevOps Guru in my AMS account?

There are no prerequisites. DevOps Guru leverages the following AWS services: Amazon CloudWatch Logs, RDS Insights, AWS X-Ray, AWS Lambda, and AWS CloudTrail.

Amazon DocumentDB (with MongoDB compatibility) (AMS SSPS)

Amazon DocumentDB (with MongoDB compatibility) is a fast, scalable, highly available, and fully managed document database service that supports MongoDB workloads. Amazon DocumentDB gives you the performance, scalability, and availability you need when operating mission-critical MongoDB workloads at scale. Amazon DocumentDB implements the Apache 2.0 open source MongoDB 3.6 API by emulating the responses that a MongoDB client expects from a MongoDB server, allowing you to use your existing MongoDB drivers and tools with Amazon DocumentDB. In Amazon DocumentDB, the storage and compute are decoupled, allowing each to scale independently, and you can increase the read capacity to millions of requests per second by adding up to 15 low latency read replicas, regardless of the size of your data. Amazon DocumentDB is designed for 99.99% availability and replicates six copies of your data across three AWS Availability Zones (AZs). You can use AWS Database Migration Service (DMS) for free (for six months) to migrate your on-premises or Amazon Elastic Compute Cloud (Amazon EC2) MongoDB databases to Amazon DocumentDB with virtually no downtime. To learn more, see Amazon DocumentDB (with MongoDB compatibility).

Amazon DocumentDB in AWS Managed Services FAQs

Q: How do I request access to Amazon DocumentDB in my AMS account?

Amazon DocumentDB console and data access roles can be requested through the submission of two AMS RFCs, console access and data access:

Request access to Amazon DocumentDB by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer_documentdb_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon DocumentDB in my AMS account?

Amazon DocumentDB requires Amazon RDS-specific permissions. Because AMS fully manages Amazon RDS, the IAM role for Amazon DocumentDB includes some restrictions to actions on Amazon RDS. The following restrictions apply:

- Access to the `DeleteDBInstance` and `DeleteDBCluster` APIs have been restricted. To use those deletion APIs, submit an RFC with the Management | Other | Other | Create (ct-1e1xtak34nx76) change type.
- You can't add or remove tags from Amazon RDS instances.
- You can't make your Amazon DocumentDB instance public.

Q: What are the prerequisites or dependencies to using Amazon DocumentDB in my AMS account?

Amazon S3 and AWS KMS are required in order to use Amazon DocumentDB, if Amazon S3 buckets are encrypted with AWS KMS keys.

Amazon DynamoDB (AMS SSPS)

Amazon DynamoDB (DynamoDB) is a key value and document database that delivers single-digit millisecond performance at any scale. It's a fully managed, multi-region, multi-active database with built-in security, backup and restore, and in-memory caching for internet scale applications. To learn more, see Amazon DynamoDB.

Amazon DynamoDB Accelerator (DAX) is a write-through caching service that is designed to simplify the process of adding a cache to DynamoDB tables. DAX is intended for applications that require high-performance reads.

DynamoDB in AWS Managed Services FAQs

Q: How do I request access to DynamoDB and DAX in my AMS account?

Request access to DynamoDB and DAX by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM roles and policies to your account:

- DynamoDB role name: `customer_dynamodb_role`

  DAX service role name: `customer_dax_service_role`
- DynamoDB policy name: `customer_dynamodb_policy`

DAX service policy: `customer_dax_service_policy`

Once provisioned in your account, you must onboard the `customer_dynamodb_role` in your federation solution.

Q: What are the restrictions to using DynamoDB in my AMS account?

All DynamoDB functionality are supported including DynamoDB Accelerator (DAX).

When creating alarms for any given table, the alarm name must be prefixed with "customer*"; for example, `customer-employee-table-high-put-latency`.

When creating an Amazon SNS topic for DynamoDB, it must be named: `dynamodb`.

To delete the Amazon SNS topic created by DynamoDB, submit a Management | Other | Other | Update change type RFC.

Q: What are the prerequisites or dependencies to using DynamoDB in my AMS account?

There are no prerequisites or dependencies to use DynamoDB in your AMS account.

Amazon Elastic Container Registry (AMS SSPS)

Amazon Elastic Container Registry (Amazon ECR) is a fully-managed Docker container registry that makes it easy for developers to store, manage, and deploy Docker container images. Amazon ECR is integrated with Amazon Elastic Container Service (Amazon ECS), simplifying your development to production workflow. Amazon ECR eliminates the need to operate your own container repositories or worry about scaling the underlying infrastructure. Amazon ECS hosts your images in a highly available and scalable architecture, allowing you to reliably deploy containers for your applications. Integration with AWS Identity and Access Management (IAM) provides resource-level control of each repository. With Amazon ECR, there are no upfront fees or commitments. You pay only for the amount of data you store in your repositories and data transferred to the Internet.

To learn more, see Amazon Elastic Container Registry.

Amazon Elastic Container Registry in AWS Managed Services FAQs

Q: How do I request access to Amazon ECR in my AMS account?

Request access to Amazon ECR by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM

role to your account: `customer_ecr_console_role`. Once provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon ECR in my AMS account?

There are restrictions around AMS namespaces for the use of Amazon ECR in your AMS account. Container images may not be prefixed with "AMS-" or "Sentinel-".

Q: What are the prerequisites or dependencies to using Amazon ECR in my AMS account?

There are no prerequisites or dependencies to use Amazon ECR in your AMS account.

Q: Is it possible to have an instance profile with Amazon ECR power user permissions?

Yes, use change type Management | Applications | IAM instance profile | Create (ct-0ixp4ch2tiu04).

EC2 Image Builder (AMS SSPS)

EC2 Image Builder is a fully managed AWS service that makes it easier to automate the creation, management, and deployment of customized, secure, and up-to-date "golden" server images that are pre-installed and pre-configured with software and settings to meet specific IT standards.

You can use the AWS Management Console, AWS CLI, or APIs to create custom images in your AWS account. When you use the AWS Management Console, the Amazon EC2 Image Builder wizard guides you through steps to:

- Provide starting artifacts
- Add and remove software
- Customize settings and scripts
- Run selected tests
- Distribute images to AWS Regions

The images you build are created in your account and can be configured for operating system patches on an ongoing basis. To learn more, see EC2 Image Builder.

EC2 Image Builder in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to EC2 Image Builder in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. Through this RFC, the following IAM role will be provisioned in your account: `customer_ec2_imagebuilder_role`. Once provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions for EC2 Image Builder?

AMS does not support the use of Service Defaults for infrastructure configuration. You can create a new infrastructure configuration or use an existing one.

AMS does not currently support the creation of container recipes.

Q: What are the prerequisites or dependencies to enable EC2 Image Builder?

- EC2 Image Builder service-linked role: You don't need to manually create a service-linked role. When you create your first Image Builder resource in the AWS Management Console, the AWS CLI, or the AWS API, Image Builder creates the service-linked role for you.
- Instances used to build images and run tests using Image Builder must have access to the Systems Manager service. The SSM Agent will be installed on the source image if it is not already present, and it will be removed before the image is created.
- AWS IAM: The IAM role that you associate with your instance profile must have permissions to run the build and test components included in your image. The following IAM role policies must be attached to the IAM role that is associated with the instance profiles: `EC2InstanceProfileForImageBuilder` and `AmazonSSMManagedInstanceCore`. The IAM role name should contain the `*imagebuilder*` keyword.
- If you configure logging, the instance profile specified in your infrastructure configuration must have `s3:PutObject` permissions for the target bucket (`arn:aws:s3:::`*`{bucket-name}`*`/*`). For example:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "s3:PutObject"
            ],
```

```
            "Resource": "arn:aws:s3:::{bucket-name}/*"
        }
    ]
}
```

- Create an SNS topic with name 'imagebuilder' to receive any alerts and notification from EC2 Image Builder.

Amazon ECS on AWS Fargate (AMS SSPS)

AWS Fargate is a technology that you can use with Amazon ECS to run containers (see Containers on AWS) without having to manage servers or clusters of Amazon EC2 instances. With AWS Fargate, you no longer have to provision, configure, or scale, clusters of virtual machines to run containers. This removes the need to choose server types, decide when to scale your clusters, or optimize cluster packing.

To learn more, see Amazon ECS on AWS Fargate.

Amazon ECS on Fargate in AWS Managed Services FAQs

Q: How do I request access to Amazon ECS on Fargate in my AMS account?

Request access to Amazon ECS on Fargate by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM roles to your account: `customer_ecs_fargate_console_role` (if no existing IAM role is provided to associate the ECS policy to), `customer_ecs_fargate_events_service_role`, `customer_ecs_task_execution_service_role`, `customer_ecs_codedeploy_service_role`, and `AWSServiceRoleForApplicationAutoScaling_ECSService`. Once provisioned in your account, you must onboard the roles in your federation solution.

Q: What are the restrictions to using Amazon ECS on Fargate in my AMS account?

- Amazon ECS task monitoring and logging are considered your responsibility since container level activities occur above the hypervisor, and logging capabilities are limited by Amazon ECS on Fargate. As a user of Amazon ECS on Fargate, we recommend that you take the necessary steps to enable logging on your Amazon ECS tasks. For more information, see Enabling the awslogs Log Driver for Your Containers.

- Security and malware protection at the container level are also considered to be your responsibility. Amazon ECS on Fargate will not include Trend Micro or preconfigured network security components.
- This service is available for both multi-account landing zone and single-account landing zone AMS accounts.
- Amazon ECS Service Discovery is restricted by default in the self-provisioned role since elevated permissions are required to create Route 53 private hosted zones. To enable Service Discovery on a service, submit a Management | Other | Other | Update change type. To provide the information required to enable Service Discovery for your Amazon ECS Service, see the Service Discovery manual.
- AMS does not currently manage or restrict images used to deploy to containers onto Amazon ECS Fargate. You will be able to deploy images from Amazon ECR, Docker Hub, or any other private image repository. Therefore, we advised that public or any unsecured images not be deployed, since they may result in malicious activity on the account.

Q: What are the prerequisites or dependencies to using Amazon ECS on Fargate in my AMS account?

- The following are dependencies of Amazon ECS on Fargate; however, no additional action is required to enable these services with your self-provisioned role:
  - CloudWatch logs
  - CloudWatch events
  - CloudWatch alarms
  - CodeDeploy
  - App Mesh
  - Cloud Map
  - Route 53
- Depending on your use case, the following are resources that Amazon ECS relies on, and may require prior to using Amazon ECS on Fargate in your account:
  - Security group to be used with the Amazon ECS service. You can use the Deployment | Advanced stack components | Security Group | Create (auto) (ct-3pc215bnwb6p7), or, if your security group requires special rules, use Deployment | Advanced stack components | Security Group | Create (review required) (ct-1oxx2g2d7hc90). Note: The security group your select with Amazon ECS has to be created specifically for Amazon ECS where the Amazon ECS service or

cluster reside. You can learn more in the **Security Group** section at Setting Up with Amazon ECS and Security in Amazon Elastic Container Service.

- Application load balancer (ALB), network load balancer (NLB), classic load balancer (ELB) for load balancing between tasks.
- Target Groups for ALBs.
- App mesh resources (for instance, Virtual Routers, Virtual Services, Virtual Nodes) to integrate with your Amazon ECS Cluster.

- Currently, there is no way for AMS to automatically mitigate risk associated with supporting security groups' permissions when created outside of the standard AMS change types. We recommend that you request a specific security group for use with your Fargate cluster to limit the possibility of using a security group not designated for the use with Amazon ECS.

Amazon EKS on AWS Fargate (AMS SSPS)

AWS Fargate is a technology that provides on-demand, right-sized compute capacity for containers (to understand containers, see What are Containers?). With AWS Fargate, you no longer have to provision, configure, or scale groups of virtual machines to run containers. This removes the need to choose server types, decide when to scale your node groups, or optimize cluster packing.

Amazon Elastic Kubernetes Service (Amazon EKS) integrates Kubernetes with AWS Fargate by using controllers that are built by AWS using the upstream, extensible model provided by Kubernetes. These controllers run as part of the Amazon EKS-managed Kubernetes control plane and are responsible for scheduling native Kubernetes pods onto Fargate. The Fargate controllers include a new scheduler that runs alongside the default Kubernetes scheduler in addition to several mutating and validating admission controllers. When you start a pod that meets the criteria for running on Fargate, the Fargate controllers running in the cluster recognize, update, and schedule the pod onto Fargate.

To learn more, see Amazon EKS on AWS Fargate Now Generally Available, also Amazon EKS Best Practices Guide for Security (includes "Recommendations" such as "Review and revoke unnecessary anonymous access" and more).

Tip

AMS has a change type, Deployment | Advanced stack components | Identity and Access Managment (IAM) | Create OpenID Connect provider (ct-30ecvfi3tq4k3), that you can use

with Amazon EKS. For an example, see Identity and Access Management (IAM) | Create OpenID Connect Provider.

Amazon EKS on AWS Fargate in AWS Managed Services FAQs

Q: How do I request access to Amazon EKS on Fargate in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account.

- `customer_eks_fargate_console_role`.

  After it's provisioned in your account, you must onboard the role in your federation solution.
- These service roles give Amazon EKS on Fargate permission to call other AWS services on your behalf:
  - `customer_eks_pod_execution_role`
  - `customer_eks_cluster_service_role`

Q: What are the restrictions to using Amazon EKS on Fargate in my AMS account?

- Creating managed or self-managed EC2 nodegroups is not supported in AMS. If you have a requirement for using EC2 worker nodes, reach out to your AMS Cloud Service Delivery Manager(CSDM) or Cloud Architect(CA).
- AMS does not include Trend Micro or preconfigured network security components for container images. You are expected to manage your own image scanning services to detect malicious container images prior to deployment.
- EKSCTL is not supported due to CFN interdependencies.
- During cluster creation, you will have permissions to disable cluster control plane logging. For more information, see Amazon EKS control plane logging. We advise that you enable all important API, Authentication, and Audit logging on cluster creation.
- During cluster creation, cluster endpoint access for Amazon EKS clusters are defaulted to public; for more information, see Amazon EKS cluster endpoint access control. We recommend that Amazon EKS endpoints be set to private. If endpoints are required for public access, we recommend that they be set to public only for specific CIDR rangers.

- AMS does not have any method to force and restrict images used to deploy to containers on Amazon EKS Fargate. You're enabled to deploy images from Amazon ECR, Docker Hub, or any other private image repository. Therefore, there is a risk of deploying a public image that may perform malicious activity on the account.
- Deploying EKS clusters through the cloud development kit (CDK) or CloudFormation Ingest is not supported in AMS.

Q: What are the prerequisites or dependencies to using Amazon EKS on Fargate in my AMS account?

In order to use the service, the following dependencies must be configured:

- For authenticating against the service, both KUBECTL and aws-iam-authenticator must be installed; for more information, see Managing cluster authentication.
- Kubernetes rely on a concept called "service accounts." In order to utilize the service accounts functionality inside of a kubernetes cluster on EKS, a Management | Other | Other | Update RFC is required with the following inputs:
  - [Required] Amazon EKS Cluster name
  - [Required] Amazon EKS Cluster namespace where service account (SA) will be deployed.
  - [Required] Amazon EKS Cluster SA name.
  - [Required] IAM Policy name and permissions/document to be associated.
  - [Required] IAM Role name being requested.
  - [Optional] OpenID Connect provider URL. For more information, see
    - Enabling IAM roles for service accounts on your cluster
    - Introducing fine-grained IAM roles for service accounts
- We recommend that Config rules be configured and monitored for
  - Public cluster endpoints
  - Disabled API logging

  Monitoring and remediating these Config rules will be your responsibility.

If you want to deploy an ALB Ingress controller, submit a Management | Other | Other Update RFC to provision the necessary IAM role to be used with the ALB Ingress Controller pod. The following

inputs are required for creating IAM resources to be associated with ALB Ingress Controller (include these with your RFC):

- [Required] Amazon EKS Cluster name
- [Optional] OpenID Connect provider URL
- [Optional] Amazon EKS Cluster namespace where the application load balancer (ALB) ingress controller service will be deployed. [default: kube-system]
- [Optional] Amazon EKS Cluster service account (SA) name. [default: aws-load-balancer-controller]

If you want to enable envelope secrets encryption in your cluster (which we recommend), provide the KMS key IDs you intend to use, in the description field of the RFC to add the service (Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct). To learn more about envelope encryption, see Amazon EKS adds envelope encryption for secrets with AWS KMS.

Amazon EMR (AMS SSPS)

Amazon EMR is the industry-leading cloud big data platform for processing vast amounts of data using open source tools such as Apache Spark, Apache Hive, Apache HBase, Apache Flink, Apache Hudi, and Presto. With Amazon EMR you can run Petabyte-scale analysis at less than half of the cost of traditional on-premises solutions and over 3x faster than standard Apache Spark. For short-running jobs, you can spin up and spin down clusters and pay per second for the instances used. For long-running workloads, you can create highly available clusters that automatically scale to meet demand.

You can create one or more instances of the Amazon EMR clusters in either AMS multi-account landing zone or single-account landing zone accounts to support both transient and persistent Amazon EMR clusters. You can also enable Kerberos authentication to enable authenticate users from on-premises Active Directory domain.

You can leverage multiple data stores with the Amazon EMR clusters to support use-case specific Hadoop tools and libraries. The Amazon EMR clusters can be created using OnDemand or Spot instances and configure autoscaling to manage capacity and reduce the cost.

The cluster log files can be archived to an Amazon S3 bucket for logging and debugging. You can also access the web interfaces hosted in the Amazon EMR cluster to support hadoop administration requirements or note book experiences for customers.

To learn more, see Amazon EMR.

Amazon EMR in AWS Managed Services FAQs

Q: How do I request access to Amazon EMR in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM roles to your account:

- `customer_emr_cluster_instance_profile`
- `customer_emr_cluster_autoscaling_role`
- `customer_emr_console_role`
- `customer_emr_cluster_service_role`

After it's provisioned in your account, you must onboard the customer_emr_console_role in your federation solution.

Q: What are the restrictions to using Amazon EMR in my AMS account?

While creating Amazon EMR on an EC2 cluster from the AWS console, we advise you to use the **Create Cluster – Advanced** option. Amazon EMR clusters must be created by adding the tag with the Key **"for-use-with-amazon-emr-managed-policies"** with Value **"true"**. Select the following configurations in the **Security** options:

- Select custom roles for your cluster:
  - EMR Role : customer_emr_cluster_service_role
  - EC2 Instance Profile : customer_emr_cluster_instance_profile
  - Auto Scaling Role : customer_emr_cluster_autoscaling_role
- EC2 Security groups:
  - Master : ams-emr-master-security-group
  - Core & Task : ams-emr-worker-security-group
  - Service Access : ams-emr-serviceaccess-security-group

Q: What are the prerequisites or dependencies to using Amazon EMR in my AMS account?

AMS creates default security groups for the Amazon EMR master, worker, and services nodes.

The launch templates and security groups to be used with Amazon EMR clusters must have the tag key **"for-use-with-amazon-emr-managed-policies"** with value **"true"**.

The default Amazon EMR cluster instance profile enables access to the resources such as s3 buckets and dynamodb tables with their names containing "emr". You can request additional IAM policies to use any additional resources to be used with Amazon EMR. The following resource ARN's can be used with Amazon EMR jobs using the **customer_emr_cluster_instance_profile**:

- arn:aws:dynamodb:*:*:table/*emr*
- arn:aws:kinesis:*:*:stream/*emr*
- arn:aws:sns:*:*:*emr*arn:aws:sqs:*:*:*emr*
- arn:aws:sqs:*:*:*emr*
- arn:aws:sqs:*:*:AWS-ElasticMapReduce-*
- arn:aws:sdb:*:*:domain:*emr*
- arn:aws:s3:::*emr*

If kerberos authentication is required for the Amazon EMR cluster:

- Provide the realm name to be used for each kerberized Amazon EMR cluster and the on-premise Active Directory IP addresses.
- Infrastructure requirements:

  Multi-Account Landing Zone (MALZ): Submit an RFC to create a new Managed application account or a new VPC in an existing application account.

  Single-Account Landing Zone (SALZ): Submit an RFC to create a new subnet in your VPC.
- Configure the incoming trust for the cluster’s realm on the on-premise Active Directory.
- Submit an RFC to configure DNS zones for the realm in the Managed AD.
- Realm configuration:

  MALZ: Submit a Management | Other | Other | Update (ct-0xdawir96cy7k) RFC to update the VPC DHCP option set to use the realm name for domain name suffix.

  SALZ: Submit a Management | Other | Other | Update (ct-0xdawir96cy7k) RFC to generate a new Amazon EMR AMI to use the specific realm for domain name suffix.

To deploy Amazon EMR studio, the role `customer_emr_cluster_service_role` has a prerequisite for an Amazon Simple Storage Service bucket. To create the bucket, use the automated CT `ct-1a68ck03fn98r` (Deployment | Advanced stack components | S3 storage | Create). When you use this automated CT to create an Amazon S3 bucket for Amazon EMR, the bucket name must begin with the prefix `customer-emr-*`. And, you must create the bucket in the same AWS Region as the Amazon EMR cluster.

Amazon EventBridge (AMS SSPS)

Amazon EventBridge is a serverless event bus service that makes it easy to connect your applications with data from a variety of sources. EventBridge delivers a stream of real-time data from your own applications, Software-as-a-Service (SaaS) applications, and AWS services and routes that data to targets such as AWS Lambda. You can set up routing rules to determine where to send your data to build application architectures that react in real time to all of your data sources. EventBridge allows you to build event driven architectures, which are loosely coupled and distributed.

To learn more, see Amazon EventBridge.

EventBridge in AWS Managed Services FAQs

Q: How do I request access to EventBridge in my AMS account?

Request access to EventBridge by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer_eventbridge_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using EventBridge in my AMS account?

You must submit AMS RFCs and create the following resources: Service roles to trigger the batch job, SQS queue, CodeBuild, CodePipeline, and SSM commands.

Q: What are the prerequisites or dependencies to using EventBridge in my AMS account?

You must request an EventBridge service role with an RFC using the Management | Other | Other | Create change type prior to using EventBridge to trigger other AWS resources, such as AWS Batch, Lambda, Amazon SNS, Amazon SQS, or Amazon CloudWatch Logs resources. Please specify which services you wish to invoke when requesting your service role. To learn about permissions

required to invoke targets, see the EventBridge documentation, Using Resource-Based Policies for EventBridge.

EventBridge is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in EventBridge. CloudTrail should be enabled and allowed to store the log files to S3 buckets. Note: All AMS Accounts have CloudTrail enabled, so no action is needed.

Amazon Forecast (AMS SSPS)

Amazon Forecast (Forecast) is a fully managed service that uses machine learning to deliver highly accurate forecasts.

Based on the same technology used at Amazon.com, Forecast uses machine learning to combine time series data with additional variables to build forecasts. Forecast requires no machine learning experience to get started. You only need to provide historical data, plus any additional data that you believe may impact your forecasts. For example, the demand for a particular color of a shirt may change with the seasons and store location. This complex relationship is hard to determine on its own, but machine learning is ideally suited to recognize it. Once you provide your data, Forecast will automatically examine it, identify what is meaningful, and produce a forecasting model capable of making predictions that are up to 50% more accurate than looking at time series data alone.

To learn more, see Amazon Forecast.

Amazon Forecast in AWS Managed Services FAQs

Q: How do I request access to Forecast in my AMS account?

Request access to AWS Firewall Manager by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer_forecast_admin_role`. Once provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Forecast in my AMS account?

The default S3 bucket access only allows you to access buckets with the naming pattern 'customer-forecast-*'. If you have your own naming convention for data buckets, discuss bucket naming and related access setup with your Cloud Architect (CA). For example:

- You could define your specific Amazon Forecast service role with naming like 'AmazonForecast-ExecutionRole-*' and associated proper S3 bucket access. See the Service role - AmazonForecast-

ExecutionRole-Admin and IAM policy - customer_forecast_default_s3_access_policy, in the IAM console.

- You may need to associate related S3 buckets access to IAM federation role. See the IAM policy - customer_forecast_default_s3_access_policy, in the IAM console.

Q: What are the prerequisites or dependencies to using Forecast in my AMS account?

- Proper Amazon S3 bucket(s) must be created before using Forecast. Especially, the default S3 buckets access is with naming pattern ‘customer-forecast-*’
- If you want to use naming patterns on S3 buckets other than 'customer-forecast-*', you must create a new service role with S3 access permissions on the buckets:
  1. A new service role to be created with naming 'AmazonForecast-ExecutionRole-{suffix}'.
  2. A new IAM policy to be created which is similar to customer_forecast_default_s3_access_policy and to be associated with the new service role and related federation admin role (e.g. 'customer_forecast_admin_role')

Q: How can I enhance data security while using Amazon Forecast?

- For data encryption at rest, you can use AWS KMS to provision a customer-managed CMK to protect data storage on Amazon S3 service:
  - Enable default encryption on the bucket with the provision key and set up bucket policy to accept AWS KMS data encryption while putting data.
  - Enable the Amazon Forecast service role 'AmazonForecast-ExecutionRole-*' and federation admin role (e.g. 'customer_forecast_admin_role') as the AWS KMS key user.
- For data encryption in transit, you can set up the HTTPS protocol, which is required while transferring objects on Amazon S3 bucket policy.
- Further restrictions on access control, enable a bucket policy for approved access for the Amazon Forecast service role 'AmazonForecast-ExecutionRole-*' and admin role (e.g. 'customer_forecast_admin_role').

Q: What are the best practices while using Amazon Forecast?

- You should have a good understanding of your data classification practices and map out the related data security needs while using S3 buckets with Amazon Forecast.

- For Amazon S3 bucket configuration, we strongly advise you to enable HTTPS enforcement in your S3 bucket policy.
- You must be aware of the admin role 'customer_forecast_admin_role' support permissive access (Get/Delete/Put S3 objects) on Amazon S3 buckets with naming of 'customer-forecast-*'. NOTE: If you require fine-grained access control for multiple teams, follow these practices:
  - Define your team-based access IAM identity (role/user) with least-privilege access to related Amazon S3 buckets.
  - Create team/project based AWS KMS CMKs grant proper access to corresponding IAM identities. (user access and 'AmazonForecast-ExecutionRole-{team/project}'.
  - Setup S3 bucket default encryption with the created AWS KMS CMKs.
  - Enforce S3 API traffics with HTTPS protocol on S3 bucket policy.
  - Enforce S3 bucket configuration for approved access for related IAM identities (user access and 'AmazonForecast-ExecutionRole-{team/project}' to the buckets.
- If you want to use the 'customer_forecast_admin_role' for general purpose, consider points listed previously to protect S3 buckets.

Q: Where is compliance information about Amazon Forecast?

See the AWS services Compliance Program.

Amazon FSx (AMS SSPS)

Amazon FSx provides fully managed third-party file systems. Amazon FSx provides you with the native compatibility of third-party file systems with feature sets for workloads such as Windows-based storage, high-performance computing (HPC), machine learning, and electronic design automation (EDA). Amazon FSx automates the time-consuming administration tasks such as hardware provisioning, software configuration, patching, and backups. Amazon FSx integrates the file systems with cloud-native AWS services, making them even more useful for a broader set of workloads.

Amazon FSx provides you with two file systems to choose from: Amazon FSx for Windows File Server for Windows-based applications and Amazon FSx for Lustre for compute-intensive workloads. To learn more, see Amazon FSx.

Amazon FSx in AWS Managed Services FAQs

Q: How do I request access to Amazon FSx in my AMS account?

Request access to Amazon FSx by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer_fsx_admin_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon FSx in my AMS account?

There are no restrictions. Full functionality of the service is available.

Q: What are the prerequisites or dependencies to using Amazon FSx in my AMS account?

There are no prerequisites. However, for advance configurations like Multi-AZ, you must install and manage the DFS Replication and DFS Namespaces services. For more information, see Deploying Multi-AZ File Systems.

Q: How do I integrate my Amazon FSx file system with my multi-account landing zone Managed AD?

When creating an Amazon FSx file system, you can specify your MALZ Managed AD as the 'AWS Managed Microsoft Active Directory' for Windows Authentication. For more information see, Using Amazon FSx with AWS Directory Service for Microsoft Active Directory

You must also share the Managed AD to the application account first. Do this by submitting an RFC with the Management | Other | Other | Create (ct-1e1xtak34nx76) change type.

Q: Which users belong in the AWS Delegated FSx Administrators group?

Only IT file server administrators. This group has **Full Access** privileges across all file shares.

Q: Should I use the default file share, share, which is created when the FSx system is provisioned?

No, we don't recommend using the the default file share, **share**, as provisioned. It grants **Full Access** to **Everyone**, which which violates the principle of least privilege. Instead, create smaller, custom file shares that match your business needs.

Q: How can I create custom file shares for specific organizations in my business?

See File Shares for instructions on creating custom file shares. Restrict access on each file share using the principle of least privilege.

Amazon FSx OpenZFS (AMS SSPS)

Amazon FSx for OpenZFS is a fully managed file storage service that makes it easy to move data residing in on-premises ZFS or other Linux-based file servers to AWS without changing your application code or how you manage data. It offers highly reliable, scalable, performant, and feature-rich file storage built on the open-source OpenZFS file system, providing the familiar features and capabilities of OpenZFS file systems with the agility, scalability, and simplicity of a fully managed AWS service. For developers building cloud-native applications, it offers simple, high-performance storage with rich capabilities for working with data.

Amazon FSx for OpenZFS file systems are broadly accessible from Linux, Windows, and macOS compute instances and containers using the industry-standard NFS protocol (v3, v4.0, v4.1, v4.2). Powered by AWS Graviton processors and the latest AWS disk and networking technologies (including AWS Scalable Reliable Datagram networking and the AWS Nitro system), Amazon FSx for OpenZFS delivers up to 1 million IOPS with latencies of hundreds of microseconds. With complete support for OpenZFS features like instant point-in-time snapshots and data cloning, FSx for OpenZFS makes it easy for you to replace your on-premises file servers with AWS storage that provides familiar file system capabilities and eliminates the need to perform lengthy qualifications and change or re-architect existing applications or tools. And, by combining the power of OpenZFS data management capabilities with the high performance and cost efficiency of the latest AWS technologies, FSx for OpenZFS enables you to build and run high-performance, data-intensive applications.

As a fully managed service, Amazon FSx for OpenZFS makes it easy to launch, run, and scale fully managed file systems on AWS that replace the file servers you run on premises while helping to provide better agility and lower costs. With Amazon FSx for OpenZFS, you no longer have to worry about setting up and provisioning file servers and storage volumes, replicating data, installing and patching file server software, detecting and addressing hardware failures, and manually performing backups. It also provides rich integration with other AWS services, such as AWS Identity and Access Management (IAM), AWS Key Management Service (AWS KMS), Amazon CloudWatch, and AWS CloudTrail.

Amazon FSx provides you with two file systems to choose from: Amazon FSx for Windows File Server for Windows-based applications and Amazon FSx for Lustre for compute-intensive workloads. To learn more, see Amazon FSx.

Amazon FSx OpenZFS in AWS Managed Services FAQs

Q: How do I request access to Amazon FSx OpenZFS in my AMS account?

Request access to Amazon FSx OpenZFS by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer_fsx_ontap_admin_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon FSx OpenZFS in my AMS account?

Replacing the security group on the Amazon FSx elastic network interfaces (ENIs) requires you to submit Management | Other | Other | Update RFCs since security groups are a critical perimeter for the AMS environment. That is the only restriction.

Q: What are the prerequisites or dependencies to using Amazon FSx OpenZFS in my AMS account?

There are no prerequisites. However, you must have Amazon FSx (AMS SSPS) installed.

Amazon FSx NetApp ONTAP (AMS SSPS)

Amazon FSx for NetApp ONTAP is a fully managed service that provides highly reliable, scalable, performant, and feature-rich file storage built on NetApp's popular ONTAP file system. It provides the familiar features, performance, capabilities, and APIs of NetApp file systems with the agility, scalability, and simplicity of a fully managed AWS service.

Amazon FSx for NetApp ONTAP provides feature-rich, fast, and flexible shared file storage that’s broadly accessible from Linux, Windows, and macOS compute instances running in AWS or on premises. FSx for ONTAP offers high-performance SSD storage with sub-millisecond latencies, and makes it quick and easy to manage your data by enabling you to snapshot, clone, and replicate your files with the click of a button. It also automatically tiers your data to lower-cost, elastic storage, eliminating the need to provision or manage capacity and allowing you to achieve SSD levels of performance for your workload while only paying for SSD storage for a small fraction of your data. It provides highly available and durable storage with fully managed backups and support for cross-region disaster recovery, and supports popular data security and anti-virus applications that make it even easier to protect and secure your data. For customers who use NetApp ONTAP on-premises, FSx for ONTAP is an ideal solution to migrate, back up, or burst your file-based applications from on-premises to AWS without the need to change your application code or how you manage your data.

As a fully managed service, Amazon FSx for NetApp ONTAP makes it simple to launch and scale reliable, performant, and secure shared file storage in the cloud. With Amazon FSx for NetApp

ONTAP, you no longer have to worry about setting up and provisioning file servers and storage volumes, replicating data, installing and patching file server software, detecting and addressing hardware failures, managing failover and failback, and manually performing backups. It also provides rich integration with other AWS services, such as AWS Identity and Access Management (IAM), Amazon WorkSpaces, AWS Key Management Service (AWS KMS), and AWS CloudTrail.

Amazon FSx provides you with two file systems to choose from: Amazon FSx for Windows File Server for Windows-based applications and Amazon FSx for Lustre for compute-intensive workloads. To learn more, see Amazon FSx.

Amazon FSx NetApp ONTAP in AWS Managed Services FAQs

Q: How do I request access to Amazon FSx NetApp ONTAP in my AMS account?

Request access to Amazon FSx NetApp ONTAP by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer_fsx_ontap_admin_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon FSx NetApp ONTAP in my AMS account?

Replacing the security group on the Amazon FSx NetApp ONTAP elastic network interfaces (ENIs) requires you to submit Management | Other | Other | Update RFCs since security groups are a critical perimeter for the AMS environment. That is the only restriction.

Q: What are the prerequisites or dependencies to using Amazon FSx NetApp ONTAP in my AMS account?

There are no prerequisites. However, you must have Amazon FSx (AMS SSPS) installed.

Amazon Inspector Classic (AMS SSPS)

Amazon Inspector Classic is an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. Amazon Inspector Classic automatically assesses applications for exposure, vulnerabilities, and deviations from best practices. After performing an assessment, Amazon Inspector Classic produces a detailed list of security findings prioritized by level of severity. These findings can be reviewed directly or as part of detailed assessment reports, which are available via the Amazon Inspector Classic console or API. To learn more, see Amazon Inspector Classic.

Amazon Inspector in AWS Managed Services FAQs

Q: How do I request access to Amazon Inspector Classic in my AMS account?

Request access to Amazon Inspector Classic by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the `customer_inspector_admin_role` IAM role to your account. The role includes the AWS-managed AmazonInspectorFullAccess policy. Once provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon Inspector Classic in my AMS account?

There are no restrictions. Full functionality of Amazon Inspector Classic is available in your AMS account.

Q: What are the prerequisites or dependencies to using Amazon Inspector Classic in my AMS account?

There are no prerequisites or dependencies to use Amazon Inspector Classic in your AMS account.

Use the new Amazon Inspector in AMS

You can now use the new Amazon Inspector in your AMS account.

For Amazon Inspector Classic, the `customer-inspector-admin-role-ssm-inspector-agent-policy` and `AmazonInspectorFullAccess` were required. However, there has been an update to the SSPS role `customer-inspector-admin-role`, which now includes an additional `policyAmazonInspector2FullAccess`. This new policy allows API permissions for the new version of Amazon Inspector.

Amazon Managed Service for Apache Flink (AMS SSPS)

Managed Service for Apache Flink (KDA) Amazon Managed Service for Apache Flink is the easiest way to analyze streaming data, gain actionable insights, and respond to your business and customer needs in real time. Amazon Managed Service for Apache Flink reduces the complexity of building, managing, and integrating streaming applications with other AWS services. SQL users can easily query streaming data or build entire streaming applications using templates and an interactive SQL editor. Java developers can quickly build sophisticated streaming applications using open source Java libraries and AWS integrations to transform and analyze data in real time. Amazon Managed Service for Apache Flink takes care of everything required to run your real-time

applications continuously and scales automatically to match the volume and throughput of your incoming data. With Amazon Managed Service for Apache Flink, you only pay for the resources your streaming applications consume. There is no minimum fee or setup cost. To learn more, see Amazon Managed Service for Apache Flink.

Managed Service for Apache Flink in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to Amazon Managed Service for Apache Flink in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_kinesis_analytics_application_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon Managed Service for Apache Flink in my AMS account?

- Configurations are limited to resources without 'AMS-' or 'MC-' prefixes to prevent any modifications to AMS infrastructure.
- Permission to delete or create new Kinesis Data Streams or Firehose has been removed from the policy. We have another policy that allows that.

Q: What are the prerequisites or dependencies to using Amazon Kinesis Data Streams in my AMS account?

There are a few dependencies:

- Amazon Managed Service for Apache Flink requires that Kinesis Data Streams or Firehose must be created prior to configuring an application with Managed Service for Apache Flink.
- The resource-based policy permissions should indicate a particular input data source.

Amazon Managed Service for Prometheus (AMS SSPS)

Amazon Managed Service for Prometheus (AMP) is a serverless, Prometheus-compatible monitoring service for container metrics that makes it easier to securely monitor container environments at scale. With Amazon Managed Service for Prometheus, you can use the same

open-source Prometheus data model and query language that you use today to monitor the performance of your containerized workloads, and also enjoy improved scalability, availability, and security without having to manage the underlying infrastructure.

Amazon Managed Service for Prometheusautomatically scales the ingestion, storage, and querying of operational metrics as workloads scale up and down. It integrates with AWS security services to enable fast and secure access to data. For more information, see What is Amazon Managed Service for Prometheus?

Amazon Managed Service for Prometheus in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to Amazon Managed Service for Prometheus in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer-prometheus-console-role`. After it's provisioned in your account, you must onboard the `customer-prometheus-console-role` role in your federation solution.

Q: What are the restrictions to using Amazon Managed Service for Prometheus in my AMS account?

All features are supported.

Q: What are the prerequisites or dependencies to using Amazon Managed Service for Prometheus in my AMS account?

There are no prerequisites or dependencies to get started with Amazon Managed Service for Prometheus. However, depending on your specific use case, you might require access to other AWS services.

Amazon Data Firehose (AMS SSPS)

Firehose (KDF) is the easiest way to reliably load streaming data into data lakes, data stores, and analytics tools. It can capture, transform, and load streaming data into Amazon S3, Amazon Redshift, Amazon OpenSearch Service, and Splunk, enabling near real-time analytics with existing business intelligence tools and dashboards you're already using today. It is a fully managed service that automatically scales to match the throughput of your data and requires no ongoing administration. It can also batch, compress, transform, and encrypt the data before loading it,

minimizing the amount of storage used at the destination and increasing security. To learn more, see What Is Amazon Data Firehose?

Firehose in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to Amazon Data Firehose in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_kinesis_firehose_user_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Firehose in my AMS account?

There are no restrictions. Full functionality of Amazon Data Firehose is available in your AMS account.

Q: What are the prerequisites or dependencies to using Firehose in my AMS account?

New service-linked IAM roles must be requested for each delivery stream. You can also re-use a single service-linked role for all streams by updating the role policy with the required resource permissions (including S3 buckets/ KMS Keys / Lambda Functions / Kinesis streams).

After you have submitted the RFC to add Firehose, an AMS Operations engineer will reach out to you through a Service Request for the ARNs of resources that you would like to connect with Data Firehose (for example, AWS KMS, S3, Lambda, and Kinesis Streams).

Amazon Kinesis Data Streams (AMS SSPS)

Amazon Kinesis Data Streams (KDS) is a highly scalable, and durable, real-time data streaming service. KDS can continuously capture gigabytes of data per second from hundreds of thousands of sources such as website clickstreams, database event streams, financial transactions, social media feeds, IT logs, and location-tracking events. The data collected is available in milliseconds to enable real-time analytics use cases such as real-time dashboards, real-time anomaly detection, dynamic pricing, and more. To learn more, see Amazon Kinesis Data Streams.

Kinesis Data Streams in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to Amazon Kinesis Data Streams in my AMS account?

Request access to Amazon Kinesis Data Streams by submitting an RFC with the Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM role to your account: `customer_kinesis_data_streaming_user_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon Kinesis Data Streams in my AMS account?

There are no restrictions. Full functionality of Amazon Kinesis Data Streams is available in your AMS account.

Q: What are the prerequisites or dependencies to using Amazon Kinesis Data Streams in my AMS account?

There are no prerequisites or dependencies to use Amazon Kinesis Data Streams in your AMS account.

Amazon Kinesis Video Streams (AMS SSPS)

Amazon Kinesis Video Streams (KVS) helps you to securely stream video from connected devices to AWS for analytics, machine learning (ML), playback, and other processing. Kinesis Video Streams automatically provisions, and elastically scales, all the infrastructure needed to ingest streaming video data from millions of devices. It also durably stores, encrypts, and indexes video data in your streams, and allows you to access your data through easy-to-use APIs. Kinesis Video Streams enables you to playback video for live and on-demand viewing, and quickly build applications that take advantage of computer vision and video analytics through integration with Amazon Rekognition Video, and libraries for ML frameworks such as Apache MxNet, TensorFlow, and OpenCV. To learn more, see Amazon Kinesis Video Streams.

Amazon Kinesis Video Streams in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to Amazon Kinesis Video Streams in my AMS account?

Request access to Amazon Kinesis Video Streams by submitting an RFC with the Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM role to your account: `customer_kinesis_video_streaming_user_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon Kinesis Video Streams in my AMS account?

There are no restrictions. Full functionality of Amazon Kinesis Video Streams is available in your AMS account.

Q: What are the prerequisites or dependencies to using Amazon Kinesis Video Streams in my AMS account?

There are no prerequisites or dependencies to use Amazon Kinesis Video Streams in your AMS account.

Amazon Lex (AMS SSPS)

Amazon Lex is a service for building conversational interfaces into any application using voice and text. Amazon Lex provides the advanced deep learning functionalities of automatic speech recognition (ASR) for converting speech to text, and natural language understanding (NLU) to recognize the intent of the text, to enable you to build applications with highly engaging user experiences and lifelike conversational interactions. With Amazon Lex, the same deep learning technologies that power Amazon Alexa are now available to any developer, enabling you to quickly and easily build sophisticated, natural language, conversational bots or chatbots. To learn more, see Amazon Lex.

Amazon Lex in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to Amazon Lex in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM role to your account: `customer_lex_author_role`. Once provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon Lex in my AMS account?

Amazon Lex integration with Lambda is limited to Lambda functions without an "AMS-" prefix, in order to prevent any modifications to AMS infrastructure.

Q: What are the prerequisites or dependencies to using Amazon Lex in my AMS account?

There are no prerequisites or dependencies to use Amazon Lex in your AMS account.

Amazon MQ (AMS SSPS)

Amazon MQ is a managed message broker service for Apache ActiveMQ that helps you to set up and operate message brokers in the cloud. Message brokers allow different software systems, often using different programming languages and on different platforms, to communicate and exchange information. Amazon MQ reduces your operational load by managing the provisioning, setup, and maintenance of ActiveMQ, a popular open-source message broker. Connecting your current applications to Amazon MQ uses industry standard APIs and protocols for messaging, including JMS, NMS, AMQP, STOMP, MQTT, and WebSocket. Using standards means that, in most cases, there's no need to rewrite any messaging code when you migrate to AWS. To learn more, see What Is Amazon MQ?

Amazon MQ in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to Amazon MQ in my AMS account?

Utilization of Amazon MQ in your AMS account is a two-step process:

1. Provision the Amazon MQ Broker. To do this, submit a CFN Template, with the Amazon MQ Broker included, through an RFC with the Deployment | Ingestion | Stack from CloudFormation Template | Create change type (ct-36cn2avfrrj9v), or submit an RFC with the Management | Other | Other | Create change type (ct-1e1xtak34nx76) change type requesting that Amazon MQ Broker be provisioned in your account.
2. Access the Amazon MQ console. After the Amazon MQ Broker is provisioned, obtain access to the Amazon MQ console by submitting an RFC with the Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM role to your account: `customer_mq_console_role`.

After the role is provisioned in your account, you must onboard it in your federation solution.

Q: What are the restrictions to using Amazon MQ in my AMS account?

Full functionality of Amazon MQ is available in your AMS account; however, provisioning Amazon MQ Broker is not available through the policy due to the elevated permission required. See above for details on how to provision Amazon MQ broker in your accounts.

Q: What are the prerequisites or dependencies to using Amazon MQ in my AMS account?

There are no prerequisites or dependencies to use Amazon MQ in your AMS account.

Amazon Managed Streaming for Apache Kafka (AMS SSPS)

Amazon Managed Streaming for Apache Kafka (Amazon MSK) is a fully managed AWS streaming data service makes it easy for you to build and run applications that use Apache Kafka to process streaming data without needing to become an expert in operating Apache Kafka clusters. Amazon MSK manages the provisioning, configuration, and maintenance of Apache Kafka clusters and Apache ZooKeeper nodes for you. Amazon MSK also shows key Apache Kafka performance metrics in the AWS Console.

Amazon MSK provides multiple levels of security for your Apache Kafka clusters, including VPC network isolation, AWS IAM for control-plane API authorization, encryption at rest, TLS encryption in-transit, TLS based certificate authentication, SASL/SCRAM authentication secured by AWS Secrets Manager. To learn more, see Amazon MSK.

Amazon MSK in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to Amazon MSK in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM policies and role to your account:

- `customer-msk-admin-policy.json`
- `AmazonMSKFullAccess`
- `customer-msk-admin-role.json`

Once provisioned in your account you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon MSK?

For Amazon MSK to deliver broker logs to the destinations that you configure, ensure that the `AmazonMSKFullAccess` policy is attached to your IAM role. So full access permissions are already in place.

Q: What are the prerequisites or dependencies to using Amazon MSK?

Before creating your MSK cluster, you must have a VPC and subnets within that VPC. By default, AMS has this covered as part of default AMS VPC creation.

To learn about the limitation of Amazon MSK, refer to Amazon MSK Limits.

Amazon Personalize (AMS SSPS)

Amazon Personalize is a machine learning service that makes it easy for developers to create individualized recommendations for customers using their applications.

Machine learning is being increasingly used to improve customer engagement by powering personalized product and content recommendations, tailored search results, and targeted marketing promotions. However, developing the machine-learning capabilities necessary to produce these sophisticated recommendation systems has been beyond the reach of most organizations today due to the complexity. Amazon Personalize allows developers with no prior machine learning experience to easily build sophisticated personalization capabilities into their applications, using machine learning technology perfected from years of use on Amazon.com.

With Amazon Personalize, you provide an activity stream from your application – clicks, page views, signups, purchases, and so forth – as well as an inventory of the items you want to recommend, such as articles, products, videos, or music. You can also choose to provide Amazon Personalize with additional demographic information from your users such as age, or geographic location. Amazon Personalize will process and examine the data, identify what is meaningful, select the right algorithms, and train and optimize a personalization model that is customized for your data. All data analyzed by Amazon Personalize is kept private and secure, and only used for your customized recommendations. You can start serving personalized recommendations via a simple API call. You pay only for what you use, and there are no minimum fees and no upfront commitments.

To learn more, see Amazon Personalize.

Amazon Personalize in AWS Managed Services FAQs

Q: How do I request access to Amazon Personalize in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type, and you need to specify which S3 bucket contains the data to be used by AWS personalize to generate the recommendations. This RFC provisions the following IAM roles to your account: `customer_personalize_console_role` and `customer_personalize_service_role`.

- Once the `customer_personalize_console_role` is provisioned in your account, you must onboard the role in your federation solution. You can also attach the `customer_personalize_console_policy` to another existing role other than `Customer_ReadOnly_Role`.
- After the `customer_personalize_service_role` is provided to your account, then you can refer its ARN when creating a new dataset group.

At this time, AMS Operations will also deploy this service role in your account: `aws_code_pipeline_service_role_policy`.

Q: What are the restrictions to using Amazon Personalize in my AMS account?

Amazon Personalize configuration is limited to resources without 'ams-' or 'mc-' prefixes, to prevent any modifications to AMS infrastructure.

Q: What are the prerequisites or dependencies to using Amazon Personalize in my AMS account?

- If the S3 bucket where data is stored is encrypted, the KMS key ID must be provided, so we can allow the role used by Amazon Personalize to decrypt the bucket.

  Amazon Personalize does not support the default KMS S3 key. If required to use KMS, create a custom key and add the following policy to it by opening an RFC with change type KMS Key | Create (Review Required):

```
{
    "Version": "2012-10-17",
    "Id": "key-consolepolicy-3",
    "Statement": [
        {
            "Sid": "Enable IAM User Permissions",
            "Effect": "Allow",
            "Principal": {
                "Service": "personalize.amazonaws.com"
            },
            "Action": "kms:*",
            "Resource": "*"
        }
    ]
}
```

- An S3 bucket must be created with the following bucket policy. Do this by submitting an RFC with change type S3 Storage | Create Policy. This policy allows Amazon Personalize to access data; that bucket will contain the data to be used by Amazon Personalize.

```
{
"Version": "2012-10-17",
"Id": "PersonalizeS3BucketAccessPolicy",
"Statement": [
{
"Sid": "PersonalizeS3BucketAccessPolicy",
"Effect": "Allow",
"Principal": {
"Service": "personalize.amazonaws.com"
},
"Action": [
"s3:GetObject",
"s3:ListBucket"
],
"Resource": [
"arn:aws:s3:::bucket-name",
"arn:aws:s3:::bucket-name/*"
]
}
]
}
```

Amazon QuickSight (AMS SSPS)

Amazon QuickSight is a fast, cloud-powered business intelligence service that delivers insights to everyone in your organization. As a fully managed service, Amazon QuickSight lets you easily create and publish interactive dashboards that include machine learning (ML) insights. To learn more, see Amazon QuickSight.

Amazon QuickSight in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to Amazon QuickSight in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM role to your account:

`customer_quicksight_console_admin_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon QuickSight in my AMS account?

- AWS resource settings on Amazon QuickSight won't be accessible to you because of the IAM policy dependency. However, the AMS team enables each resource for you in response to your request to enable the service.
- Resource access for individual users and groups are not supported in this model because this feature enables users to alter IAM permissions that could compromise AMS infrastructure.
- The ability to invite IAM identities from within QuickSight is not supported due to the risk involved altering IAM objects.
- Amazon QuickSight service offers two editions: Enterprise and Standard. Both provide a single sign-on (SSO) option that is supported on AMS. However, the Enterprise Edition has an option to integrate Amazon QuickSight with Active Directory (AD). Amazon QuickSight on AMS does not support integration with AD due to incompatibilities between AMS account structure and the Amazon QuickSight trust requirements.

Q: What are the prerequisites or dependencies to using Amazon QuickSight in my AMS account?

- When AMS receives this RFC to add Amazon QuickSight, you are sent a service request for additional information; provide them the following:
  - Amazon QuickSight account name (for example, *`CustomerName`*`-quicksight`
  - Amazon QuickSight Edition (Standard versus Enterprise)
  - The AWS Region in which to enable the Amazon QuickSight service (defaults to your AMS AWS Region).
  - A notification email address for Amazon QuickSight account.
  - (Optional) The S3 bucket where data files to be analyzed are located.
  - The VPC and subnet IDs that connect to Amazon QuickSight support a feature to add a VPC connection, which enables private connectivity between Amazon QuickSight and resources inside the account.

An AMS operator performs the sign up process on your behalf and configures two QuickSight functionalities:

- Auto discovery to data sources.
- VPC connections.

> **Note**
>
> These actions need to be performed by an AMS operator because elevated IAM and VPC permissions are required during the sign-in process.

Amazon Rekognition (AMS SSPS)

Amazon Rekognition makes it easy to add image and video analysis to your applications using proven, highly scalable, deep learning technology that requires no machine learning expertise to use. With Amazon Rekognition, you can identify objects, people, text, scenes, and activities in images and videos, as well as detect any inappropriate content. Amazon Rekognition also provides highly accurate facial analysis and facial search capabilities that you can use to detect, analyze, and compare faces for a wide variety of user verification, people counting, and public safety use cases.

With Amazon Rekognition Custom Labels, you can identify objects and scenes in images that are specific to your business needs. For example, you can build a model to classify specific machine parts on your assembly line or to detect unhealthy plants. Amazon Rekognition Custom Labels takes care of the model development heavy lifting for you, so no machine learning experience is required. You simply need to supply images of objects or scenes you want to identify, and the service handles the rest.

To learn more, see Amazon Rekognition.

Amazon Rekognition in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to Amazon Rekognition in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_rekognition_console_role & customer_rekognition_service_role`. Once provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon Rekognition in my AMS account?

Full functionality of Amazon Rekognition is available with the Amazon Rekognition self-provisioned service role.

Q: What are the prerequisites or dependencies to using Amazon Rekognition in my AMS account?

If you use Kinesis Video Streams that provide the source streaming video for an Amazon Rekognition Video stream processor or a data stream as a destination to write data to Kinesis Data Streams, kindly provide AMS with a `kinesisStreamName` when creating the RFC.

Amazon SageMaker (AMS SSPS)

SageMaker provides every developer and data scientist with the ability to build, train, and deploy machine learning models quickly. Amazon SageMaker is a fully-managed service that covers the entire machine learning workflow to label and prepare your data, choose an algorithm, train the model, tune and optimize it for deployment, make predictions, and take action. Your models get to production faster with much less effort and lower cost. To learn more, see Amazon SageMaker.

SageMaker in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to SageMaker in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM roles to your account: `customer_sagemaker_admin_role` and service role `AmazonSageMaker-ExecutionRole-Admin`. After SageMaker is provisioned in your account, you must onboard the `customer_sagemaker_admin_role` role in your federation solution. The service role cannot be accessed by you directly; the SageMaker service uses it while doing various actions as described here: Passing Roles.

Q: What are the restrictions to using SageMaker in my AMS account?

- The following use cases are not supported by the AMS Amazon SageMaker IAM role:
  - SageMaker Studio is not supported at this time.
  - SageMaker Ground Truth to manage private workforces is not supported since this feature requires overly permissive access to Amazon Cognito resources. If managing a private workforce is required, you can request a custom IAM role with combined SageMaker and

Amazon Cognito permissions. Otherwise, we recommend using public workforce (backed by Amazon Mechanical Turk), or AWS Marketplace service providers, for data labeling.

- Creating VPC Endpoints to support API calls to SageMaker services (aws.sagemaker.{region}.notebook, com.amazonaws.{region}.sagemaker.api & com.amazonaws.{region}.sagemaker.runtime) is not supported as permissions can't be scoped down to SageMaker related services only. To support this use case, submit a Management | Other | Other RFC to create related VPC endpoints.
- SageMaker endpoint auto scaling is not supported as SageMaker requires `DeleteAlarm` permissions on any ("*") resource. To support endpoint auto scaling, submit a Management | Other | Other RFC to setup auto scaling for a SageMaker endpoint.

Q: What are the prerequisites or dependencies to using SageMaker in my AMS account?

- The following use cases require special configuration prior to use:
  - If an S3 bucket will be used to store model artifacts and data, then you must request an S3 bucket named with the required keywords ("SageMaker", "Sagemaker", "sagemaker" or "aws-glue") with a Deployment | Advanced stack components | S3 storage | Create RFC.
  - If Elastic File Store (EFS) will be used, then EFS storage must be configured in the same subnet, and allowed by security groups.
  - If other resources require direct access to SageMaker services (notebooks, API, runtime, and so on), then configuration must be requested by:
    - Submitting an RFC to create a security group for the endpoint (Deployment | Advanced stack components | Security group | Create (auto)).
    - Submitting a Management | Other | Other | Create RFC to set up related VPC endpoints.

Q: What are the supported naming conventions for resources that the `customer_sagemaker_admin_role` can access directly? (The following are for update and delete permissions; if you require additional supported naming conventions for your resources, reach out to an AMS Cloud Architect for consultation.)

- Resource: Passing `AmazonSageMaker-ExecutionRole-*` role
  - Permissions: The SageMaker self-provisioned service role supports your use of the SageMaker service role (`AmazonSageMaker-ExecutionRole-*`) with AWS Glue, AWS RoboMaker, and AWS Step Functions.
- Resource: Secrets on AWS Secrets Manager

  - Permissions: Describe, Create, Get, Update secrets with a `AmazonSageMaker-*` prefix.
  - Permissions: Describe, Get secrets when the `SageMaker` resource tag is set to `true`.
- Resource: Repositories on AWS CodeCommit
  - Permissions: Create/ delete repositories with a `AmazonSageMaker-*` prefix.
  - Permissions: Git Pull/Push on repositories with following prefixes, `*sagemaker*`, `*SageMaker*`, and `*Sagemaker*`.
- Resource: Amazon ECR (Amazon Elastic Container Registry) Repositories
  - Permissions: Permissions: Set, delete repository policies, and upload container images, when the following resource naming convention is used, `*sagemaker*`.
- Resource: Amazon S3 buckets
  - Permissions: Get, Put, Delete object, abort multipart upload S3 objects when resources have the following prefixes: `*SageMaker*`, `*Sagemaker*`, `*sagemaker*` and `aws-glue`.
  - Permissions: Get S3 objects when the `SageMaker` tag is set to `true`.
- Resource: Amazon CloudWatch Log Group
  - Permissions: Create Log Group or Stream, Put Log Event, List, Update, Create , Delete log delivery with following prefix: `/aws/sagemaker/*`.
- Resource: Amazon CloudWatch Metric
  - Permissions: Put metric data when the following prefixes are used: `AWS/SageMaker`, `AWS/SageMaker/`, `aws/SageMaker`, `aws/SageMaker/`, `aws/sagemaker`, `aws/sagemaker/`, and `/aws/sagemaker/.`.
- Resource: Amazon CloudWatch Dashboard
  - Permissions: Create/Delete dashboards when the following prefixes are used: `customer_*`.
- Resource: Amazon SNS (Simple Notification Service) topic
  - Permissions: Subscribe/Create topic when following prefixes are used: `*sagemaker*`, `*SageMaker*`, and `*Sagemaker*`.

Q: What's the difference between `AmazonSageMakerFullAccess` and `customer_sagemaker_admin_role`?

The `customer_sagemaker_admin_role` with the `customer_sagemaker_admin_policy` provides almost the same permissions as AmazonSageMakerFullAccess except:

- Permission to connect with RoboMaker, Cognito, and Glue resources.

- Sagemaker endpoint autoscaling. You must submit a Management | Other | Other | Update RFC to elevate to autoscaling permissions temporarily, or permanently, as autoscaling requires permissive access on CloudWatch service.

Q: How do I adopt KMS CMKs in data encryption at rest?

You must ensure that the key policy has been set up properly on the CMKs so that related IAM users or roles can use the keys. For more information, see the AWS KMS Key Policy document.

Amazon Simple Email Service (AMS SSPS)

Amazon Simple Email Service (Amazon SES) is a cloud-based email sending service designed to help digital marketers and application developers, send marketing, notification, and transactional emails.

You can use the SMTP interface or one of the AWS SDKs to integrate Amazon SES directly into your existing applications. You can also integrate the email sending capabilities of Amazon SES into the software you already use, such as ticketing systems and email clients.

To learn more, see Amazon Simple Email Service.

Amazon SES in AWS Managed Services FAQs

Q: How do I request access to Amazon SES in my AMS account?

Request access to Amazon SES by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer_ses_admin_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the prerequisites or dependencies to using Amazon SES in my AMS account?

- You must configure an S3 bucket policy to allow Amazon SES to publish events to the bucket.
- You must use a default (AWS SES), or configure, a CMK key to allow Amazon SES to encrypt emails and push events to other service resources like S3, SNS, Lambda, and Firehose, belonging to the account.

Q: What are the restrictions to using Amazon SES in my AMS account?

You must raise RFCs to create the following resources:

- An SMTP user and IAM service role with PutEvents permission, to a Kinesis Firehose stream.
- You must create new AWS resources such as S3 bucket, Firehose stream, SNS topic by using AMS change types in order for your Amazon SES rules and configuration sets' destinations to work with those resources.
- SMTP credentials. To request new SMTP credentials, use the Change Type (Management | Other | Other | Create). AMS creates the credentials and adds them to Secrets Manager for you.

Amazon Simple Workflow Service (AMS SSPS)

Amazon Simple Workflow Service (Amazon SWF) helps developers build, run, and scale background jobs that have parallel or sequential steps. You can think of Amazon SWF as a fully-managed state tracker and task coordinator in the Cloud. If your application's steps take more than 500 milliseconds to complete, you need to track the state of processing, or you need to recover or retry if a task fails, Amazon SWF can help you. To learn more, see Amazon Simple Workflow Service.

Amazon SWF in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to Amazon SWF in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM role to your account: `customer_swf_role`. Once provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon SWF in my AMS account?

The Lambda `InvokeFunction` permissions have been included in this service however, the AMS `customer_deny_policy` that is added to all AMS customer roles explicitly denies access to AMS Lambda functions and AMS-owned resources. In order to tag or untag resources within Amazon SWF, submit a Management | Other | Other Change Type.

Q: What are the prerequisites or dependencies to using Amazon SWF in my AMS account?

Amazon SWF is dependent on the AWS Lambda service, therefore, permissions to invoke Lambda have been provided as a part of this role and no additional permissions are required to invoke Lambda from Amazon SWF. Otherwise, there are no prerequisites to using Amazon SWF.

Amazon Textract (AMS SSPS)

Amazon Textract is a fully managed machine learning service that automatically extracts printed text, handwriting, and other data from scanned documents that goes beyond simple optical character recognition (OCR) to identify, understand, and extract data from forms and tables. To learn more, see Amazon Textract.

Amazon Textract in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request Amazon Textract to be set up in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM roles to your account: `customer_textract_console_role`, `customer_textract_human_review_execution_role`, and `customer_ec2_textract_instance_profile`. Once provisioned in your account, you must onboard the role `customer_textract_console_role` in your federation solution.

Q: What are the restrictions to using Amazon Textract in my AMS account?

There are no restrictions for the use of Amazon Textract in your AMS account.

Q: What are the prerequisites or dependencies to using Amazon Textract in my AMS account?

You must request the creation of an S3 bucket by submitting an RFC Deployment | Advanced stack components |S3 storage | Create (ct-1a68ck03fn98r).

Amazon Transcribe (AMS SSPS)

Powered by deep learning technologies, Amazon Transcribe is a fully managed and continuously trained automatic speech recognition service that automatically generates time-stamped text transcripts from audio files. Amazon Transcribe makes it easy for developers to add speech-to-text capabilities to their applications. Audio data is virtually impossible for computers to search and analyze. Therefore, recorded speech needs to be converted to text before it can be used in applications. Historically, customers had to work with transcription providers that required them to sign expensive contracts and were hard to integrate into their technology stacks to accomplish this task. Many of these providers use outdated technology that does not adapt well to different scenarios, like low-fidelity phone audio common in contact centers, which results in poor accuracy.

Amazon Transcribe uses a deep learning process called automatic speech recognition (ASR) to convert speech into text, quickly and accurately. Amazon Transcribe can be used to transcribe customer service calls, automate closed captioning and subtitling, and generate metadata for media assets to create a fully searchable archive. You can use Amazon Transcribe Medical to add medical speech-to-text capabilities to clinical documentation applications. To learn more, see Amazon Transcribe.

Amazon Transcribe in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request Amazon Transcribe to be set up in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_transcribe_role`. Once provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon Transcribe in my AMS account?

You must use 'customer-transcribe*' as the prefix for your buckets when working with transcribe, unless RA and specified otherwise.

You are not able to create an IAM role within Amazon transcribe.

You cannot use a service-managed S3 bucket for output data in default SSPS (if this is needed, please reach out to your account CA).

You must submit Risk Acceptance if you want to use customer-managed KMS Keys that do not fall under the AMS namespace.

Q: What are the prerequisites or dependencies to using Amazon Transcribe in my AMS account?

S3 must have access to the buckets with the name 'customer-transcribe*'. KMS is required in order to use Amazon Transcribe if your S3 buckets are encrypted with KMS keys. If a bucket doesn't need to be encrypted "KMStranscribeAllow" can be removed.

Amazon WorkDocs (AMS SSPS)

Amazon WorkDocs is a fully-managed, secure content creation, storage, and collaboration service. With Amazon WorkDocs, you can easily create, edit, and share content, and because it's stored centrally on AWS, access it from anywhere on any device. Amazon WorkDocs helps

you to collaborate with others, and lets you easily share content, provide rich feedback, and collaboratively edit documents. You can use Amazon WorkDocs to retire your legacy file share infrastructure by moving file shares to the cloud. Amazon WorkDocs lets you integrate with your existing systems, and offers a rich API so that you can develop your own content-rich applications. Amazon WorkDocs is built on AWS, where your content is secured on the world's largest cloud infrastructure. To learn more, see Amazon WorkDocs.

Amazon WorkDocs in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to Amazon WorkDocs in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM role to your account: `customer_workdocs_console_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Amazon WorkDocs in my AMS account?

Full functionality of Amazon WorkDocs is available in your AMS account. However, you can't delete an Amazon WorkDocs site. The permissions required to de-register an Amazon WorkDocs site require modification to the AWS Managed Microsoft AD directory. To do this, submit an AMS service request for the deletion of an Amazon WorkDocs site.

Q: What are the prerequisites or dependencies to using Amazon WorkDocs in my AMS account?

Amazon WorkDocs has a dependency on AWS Directory Service for Microsoft Active Directory (MAD). AMS has MAD already implemented in AMS accounts; however, it is limited to a one-way trust. You must submit a service request to AMS to have an AD Connector set up to proxy your on-premises domain.

Amazon WorkSpaces (AMS SSPS)

WorkSpaces enables you to provision virtual, cloud-based Microsoft Windows or Amazon Linux desktops for your users, known as WorkSpaces. WorkSpaces eliminates the need to procure and deploy hardware or install complex software. You can quickly add or remove users as your needs change. Users access their WorkSpaces by using a client application from a supported device or, for Windows WorkSpaces, a web browser, and they log in by using their existing on-premises Active Directory (AD) credentials.

To learn more, see Amazon WorkSpaces.

WorkSpaces in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to WorkSpaces in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_workspaces_console_role`. Once provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using WorkSpaces in my AMS account?

Full functionality of Workspaces is available with the Amazon WorkSpaces self-provisioned service role.

Q: What are the prerequisites or dependencies to using WorkSpaces in my AMS account?

- WorkSpaces are limited by AWS Region; therefore, the AD Connector must be configured in the same AWS Region where the WorkSpaces instances are hosted.

  Customers can connect WorkSpaces to customer AD using one of the following two methods:

  1. Using AD connector to proxy authentication to on-premises Active Directory service (preferred):

     Configure Active Directory (AD) Connector in your AMS account prior to integrating your WorkSpaces instance with your on-premises directory service. The AD Connector acts as a proxy for your existing AD users (from your domain) to connect to WorkSpaces using existing on-premises AD credentials. This is preferred because WorkSpaces are directly joined to the customer's on-prem domain, which acts as both Resource and User forest, leading to more control on the customer side.

     For more information, see Best Practices for Deploying Amazon WorkSpaces (Scenario 1).
  2. Using AD Connector with AWS Microsoft AD, Shared Services VPC, and a one-way trust to on-premises:

     You can also authenticate users with your on-premises directory by first establishing a one-way outgoing trust from AMS-managed AD to your on-premises AD. WorkSpaces will join

AMS-managed AD using an AD Connector. WorkSpaces access permissions will then be delegated to the WorkSpaces instances through the AMS-managed AD, without the need to establish a two-way trust with your on-premises environment. In this scenario, the User forest will be in the customer AD and the Resource forest will be in the AMS-managed AD (changes to AMS-managed AD can be requested via RFC). Note that the connectivity between WorkSpaces VPC and the MALZ Shared Services VPC running AMS-managed AD is established via Transit Gateway.

For more information, see Best Practices for Deploying Amazon WorkSpaces (Scenario 6).

> **Note**
>
> The AD Connector can be configured by submitting a Management | Other | Other | Create change type RFC with the prerequisite AD configuration details; for more information, see Create an AD Connector. If method 2 is used to create a Resource forest in AMS-managed AD, submit another Management | Other | Other | Create change type RFC in AMS shared-services account by running the AMS-managed AD.

AMS Code services (AMS SSPS)

AMS Code services is a proprietary bundling of AWS code management services as detailed next. You can choose to deploy all of the services in AMS with AMS Code services, or you can deploy them in AMS individually.

AMS Code services includes the following services:

- AWS CodeCommit: A fully managed source control service that hosts secure Git-based repositories. It makes it so teams can collaborate on code in a secure and highly scalable ecosystem. CodeCommit eliminates the need to operate your own source control system or worry about scaling its infrastructure. You can use CodeCommit to securely store anything from source code to binaries, and it works seamlessly with your existing Git tools. To learn more, see AWS CodeCommit

  To deploy this in your AMS account independently of AMS Code services, see AWS CodeCommit (AMS SSPS).
- AWS CodeBuild: A fully managed continuous integration service that compiles source code, runs tests, and produces software packages that are ready to deploy. With CodeBuild, you don't

need to provision, manage, and scale your own build servers. CodeBuild scales continuously and processes multiple builds concurrently, so your builds are not left waiting in a queue. You can get started quickly by using prepackaged build environments, or you can create custom build environments that use your own build tools. With CodeBuild, you are charged by the minute for the compute resources you use. To learn more, see AWS CodeBuild

To deploy this in your AMS account independently of AMS Code services, see AWS CodeBuild (AMS SSPS).

- AWS CodeDeploy: A fully managed deployment service that automates software deployments to a variety of compute services such as Amazon EC2 and your on-premises servers. AWS CodeDeploy helps you to rapidly release new features, helps you avoid downtime during application deployment, and handles the complexity of updating your applications. You can use AWS CodeDeploy to automate software deployments, eliminating the need for error-prone manual operations. The service scales to match your deployment needs. To learn more, see AWS CodeDeploy

  To deploy this in your AMS account independently of AMS Code services, see AWS CodeDeploy (AMS SSPS).
- AWS CodePipeline: A fully managed continuous delivery service that helps you automate your release pipelines for fast and reliable application and infrastructure updates. CodePipeline automates the build, test, and deploy phases of your release process every time there is a code change, based on the release model you define. This enables you to rapidly and reliably deliver features and updates. You can easily integrate AWS CodePipeline with third-party services such as GitHub or with your own custom plugin. With AWS CodePipeline, you only pay for what you use. There are no upfront fees or long-term commitments. To learn more, see AWS CodePipeline

  To deploy this in your AMS account independently of AMS Code services, see AWS CodePipeline (AMS SSPS).

AMS Code services in AWS Managed Services FAQs

Q: How do I request access to AMS Code services in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_code_suite_console_role`. After provisioned in your account, you must onboard the role in your federation solution. At this time AMS Operations will also

deploy the `customer_codebuild_service_role`, `customer_codedeploy_service_role`, `aws_code_pipeline_service_role` service roles in your account for CodeBuild, CodeDeploy and CodePipeline services. If additional IAM permissions for the are required for the `customer_codebuild_service_role` are needed, submit an AMS service request.

> **Note**
>
> You can also add these services separately; for information, see AWS CodeBuild (AMS SSPS), AWS CodeDeploy (AMS SSPS), and AWS CodePipeline (AMS SSPS), respectively.

Q: What are the restrictions to using AMS Code services in my AMS account?

- AWS CodeCommit: The triggers feature on CodeCommit is disabled given the associated rights to create SNS topics. Directly authenticating against CodeCommit is restricted; users should authenticate with Credential Helper. Some KMS commands are also restricted: kms:Encrypt, kms:Decrypt, kms:ReEncrypt, kms:GenereteDataKey, kms:GenerateDataKeyWithoutPlaintext, and kms:DescribeKey.
- CodeBuild: For AWS CodeBuild console admin access, permissions are limited at the resource level; for example, CloudWatch actions are limited on specific resources and the `iam:PassRole` permission is controlled.
- CodeDeploy: Currently CodeDeploy supports deployments on Amazon EC2/On-premises only. Deployments on ECS and Lambda through CodeDeploy is not supported.
- CodePipeline: CodePipeline features, stages, and providers are limited to the following:
  - Deploy Stage: Amazon S3 and AWS CodeDeploy
  - Source Stage: Amazon S3, AWS CodeCommit, Bit Bucket, and GitHub
  - Build Stage: AWS CodeBuild and Jenkins
  - Approval Stage: Amazon SNS
  - Test Stage: AWS CodeBuild, Jenkins, BlazeMeter, Ghost Inspector UI Testing, Micro Focus StormRunner Load, Runscope API Monitoring
  - Invoke Stage: Step Functions and Lambda

> **Note**
>
> AMS Operations will deploy the `customer_code_pipeline_lambda_policy` in your account; it must be attached with the Lambda execution role for Lambda invoke

stage. Please provide the Lambda service/execution role name that you want this policy added with. If there is no custom Lambda service/execution role, AMS will create a new role named `customer_code_pipeline_lambda_execution_role`, which will be a copy of `customer_lambda_basic_execution_role` along with `customer_code_pipeline_lambda_policy`.

Q: What are the prerequisites or dependencies to using AMS Code services in my AMS account?

- CodeCommit: If S3 buckets are encrypted with AWS KMS keys, S3 and AWS KMS are required to use AWS CodeCommit.
- CodeBuild: If additional IAM permissions are required for the defined AWS CodeBuild service role, request them through an AMS service request.
- CodeDeploy: None.
- CodePipeline: None. AWS supported services—AWS CodeCommit, AWS CodeBuild, AWS CodeDeploy—must be launched prior to, or along with, the launch of CodePipeline. However this will be taken care by an AMS engineer; you don't need to do anything.

AWS Amplify (AMS SSPS)

The AWS Amplify is a complete solution that allows frontend web and mobile developers to easily build, connect, and host fullstack applications. Amplify provides flexibility to leverage the breadth of AWS services as your use cases evolve. Amplify provides products to build fullstack iOS, Android, Flutter, Web, and React Native apps. To learn more, see AWS Amplify.

AWS Amplify in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request AWS Amplify to be set up in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_amplify_console_role`. After provisioned to your account, you must onboard the role in your federation solution.

Additionally, you must provide a Risk Acceptance because AWS Amplify has infrastructure-mutating permissions. To do this, work with your Cloud Service Delivery Manager (CSDM).

Q: What are the restrictions to using AWS Amplify in my AMS account?

You must use `'amplify*'` as the prefix for your buckets when working with Amplify, unless RA and specified otherwise.

Q: What are the prerequisites or dependencies to using AWS Amplify in my AMS account?

There are no prerequisites for the use of AWS Amplify in your AMS account.

AWS AppSync (AMS SSPS)

AWS AppSync simplifies application development by letting you create a flexible API to securely access, manipulate, and combine data from one or more data sources. AWS AppSync is a managed service that uses GraphQL to make it easy for applications to get exactly the data they need.

With AWS AppSync, you can build scalable applications, including those requiring real-time updates, on a range of data sources such as NoSQL data stores, relational databases, HTTP APIs, and your custom data sources with AWS Lambda. For mobile and web apps, AWS AppSync additionally provides local data access when devices go offline, and data synchronization with customizable conflict resolution, when they are back online. To learn more, see AWS AppSync.

AWS AppSync in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access AWS AppSync in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM roles to your account: `customer_appsync_service_role` and `customer_appsync_author_role`. Once provisioned in your account, you must onboard the `customer_appsync_author_role` in your federation solution.

Q: What are the restrictions to using the AWS AppSync?

- When creating a Data Source on AppSync the customer need to specify the previously created service role, creation of a new role is not allowed and therefore will return an access denied
- AppSync roles are configured to restrict permissions to resources containing 'AMS-' or 'MC-' prefixes to prevent any modifications to AMS infrastructure.

Q: What are the prerequisites or dependencies to using AWS AppSync?

The service allows multiple other services to be used as a data source, The basic permissions to use them as such is included in the service role (`customer_appsync_service_role`), but you must manually select the service role when using the service.

AWS App Mesh (AMS SSPS)

AWS App Mesh provides application level networking to make it easy for your services to communicate with each other across multiple types of compute infrastructure. App Mesh standardizes how your services communicate, giving you end-to-end visibility and ensuring high-availability for your applications.

AWS App Mesh makes it easy to run services by providing consistent visibility and network traffic controls for services built across multiple types of compute infrastructure. App Mesh removes the need to update application code to change how monitoring data is collected or traffic is routed between services. App Mesh configures each service to export monitoring data and implements consistent communications control logic across your application. This makes it easy to quickly pinpoint the exact location of errors and automatically re-route network traffic when there are failures or when code changes need to be deployed. To learn more, see AWS App Mesh.

AWS App Mesh in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access AWS App Mesh in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM role to your account: `customer_app_mesh_console_role`. After it is provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using the AWS App Mesh?

Full functionality of AWS App Mesh is available in your AMS account.

Q: What are the prerequisites or dependencies to using AWS App Mesh?

There are no prerequisites or dependencies to use AWS App Mesh in your AMS account.

AWS Audit Manager (AMS SSPS)

Audit Manager helps you continuously audit your AWS usage to simplify how you assess risk and compliance with regulations and industry standards. Audit Manager automates evidence collection to make it easier to assess if your policies, procedures, and activities are operating effectively. When it is time for an audit, Audit Manager helps you manage stakeholder reviews of your controls and helps you build audit-ready reports with significantly less manual effort. To learn more, see Audit Manager.

AWS Audit Manager in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to AWS Audit Manager in my AMS account?

You can request access through the submission of the AWS Services RFC (Management | AWS service | Compatible Service). Through this RFC the following IAM Role will be provisioned in your account: `customer-audit-manager-admin-Role`. Once provisioned in your account you must onboard the role in your federation solution.

Q: What are the restrictions to using AWS Audit Manager?

There are no restrictions for the use of AWS Audit Manager in your AMS account. Full functionality for AWS Audit Manager will be provided.

Q: What are the prerequisites or dependencies to using AWS Audit Manager?

1. You need to provide AMS with the s3 bucket where you want reports/assessments to reside.
2. If you want to have encryption with the service, you need to provide AMS with the KMS CMK ARN to use.
3. If you want to send an SNS notifications to a Topic, you must provide the name of the topic or arn.
4. **(Optional)** There is an additional prerequisite if you want to enable Organizations as part of your multi-account landing zone in Audit Manager and you want a delegated administrator account: In the description field for RFC (Management | AWS service | Compatible Service| Add), mention that you want to use the delegated administrator account as part of Audit Manager Setup and provide the below details:
   - KMS CMK ARN (used to set up Audit Manager, initially)

- Delegated administrator account ID for Audit Manager to use as part of this multi-account landing zone (can be a MALZ application account)

AWS Batch (AMS SSPS)

AWS Batch enables developers, scientists, and engineers to easily and efficiently run hundreds of thousands of batch computing jobs on AWS. AWS Batch dynamically provisions the optimal quantity and type of compute resources (such as CPU or memory optimized instances) based on the volume and specific resource requirements of the batch jobs submitted. With AWS Batch, there is no need to install and manage batch computing software or server clusters that you use to run your jobs, allowing you to focus on analyzing results and solving problems. To learn more, see AWS Batch.

AWS Batch in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to AWS Batch in my AMS account?

1. You can request access through the submission of the AWS Services RFC (Management | AWS service | Compatible Service). Through this RFC, the following IAM roles and policies will be provisioned in your account:

IAM roles:

- `customer_batch_console_role`
- `customer_batch_ecs_instance_role`
- `customer_batch_events_service_role`
- `customer_batch_service_role`
- `customer_batch_ecs_task_role`

Policies:

- `customer_batch_console_role_policy`
- `customer_batch_service_role_policy`
- `customer_batch_events_service_role_policy`

2. Once provisioned in your account, you must onboard the role `customer_batch_console_role` in your federation solution.

Q: What are the restrictions to using AWS Batch?

When creating the Compute Environment, you should tag EC2 instances as "customer_batch" or "customer-batch". If the instances are not tagged, instances will not be terminated by batch when the job completes.

Q: What are the prerequisites or dependencies to using AWS Batch?

There are no prerequisites or dependencies to use AWS Batch in your AMS account.

AWS Certificate Manager (AMS SSPS)

AWS Certificate Manager (ACM) is a service that lets you provision, manage, and deploy public and private Secure Sockets Layer/Transport Layer Security (SSL/TLS) certificates for use with AWS services and your internal connected resources. SSL/TLS certificates are used to secure network communications and establish the identity of websites over the internet as well as resources on private networks. AWS Certificate Manager removes the time-consuming manual process of purchasing, uploading, and renewing SSL/TLS certificates.

With AWS Certificate Manager, you can request a certificate, deploy it on ACM-integrated AWS resources, such as Elastic Load Balancers, Amazon CloudFront distributions, and APIs on API Gateway, and let AWS Certificate Manager handle certificate renewals. It also enables you to create private certificates for your internal resources and manage the certificate lifecycle centrally. Public and private certificates provisioned through AWS Certificate Manager for use with ACM-integrated services are free. You pay only for the AWS resources you create to run your application. With AWS Private Certificate Authority, you pay monthly for the operation of the AWS Private CA and for the private certificates you issue. To learn more, see AWS Certificate Manager - AWS Documentation.

ACM in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to AWS Certificate Manager in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM role to your account: `customer_acm_create_role`. You can use this role to create and manage ACM certificates. After it's provisioned in your account, you must onboard the role in your federation solution.

ACM certificates can be created using the following change types, even if you haven't added the `customer_acm_create_role` IAM role:

- ACM | Create Public Certificate
- ACM | Create Private Certificate
- ACM Certificate with additional SANs | Create

Q: What are the restrictions to using the AWS Certificate Manager?

You must submit a Request for Change (RFC) to AMS to delete or modify existing certificates, as those actions require full admin access (use the Management | Other | Other | Update change type (ct-0xdawir96cy7k). Note that the IAM policy can't exclude rights based on tag names (mc*, ams*, etc). Certificates do not incur a cost, so deleting unused certificates is not time sensitive.

Q: What are the prerequisites or dependencies to using Certificate Manager?

Existing public DNS name, and access to create DNS CNAME records, but those do not need to be hosted in the managed account.

AWS Private Certificate Authority (AMS SSPS)

Private certificates are used for identifying and securing communication between connected resources on private networks, such as servers, mobile, and IoT devices and applications. AWS Private CA is a managed private CA service that helps you easily and securely manage the lifecycle of your private certificates. AWS Private CA provides you a highly-available private CA service without the upfront investment and ongoing maintenance costs of operating your own private CA. AWS Private CA extends ACM's certificate management capabilities to private certificates, enabling you to create and manage public and private certificates centrally. You can easily create and deploy private certificates for your AWS resources using the AWS Management Console or the ACM API. For EC2 instances, containers, IoT devices, and on-premises resources, you can easily create and track private certificates and use your own client-side automation code to deploy them. You also have the flexibility to create private certificates and manage them yourself for applications that require custom certificate lifetimes, key algorithms, or resource names To learn more, see AWS Private CA.

AWS Private CA in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access AWS Private CA in my AMS account?

Request access through the submission of the AWS Services RFC (Management | AWS service | Compatible Service). Through this RFC the following IAM role will be provisioned in your account: `customer_acm_pca_role`. Once provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using the AWS Private CA?

Currently, AWS Resource Access Manager (AWS RAM) cannot be used to share your AWS Private CA cross-account.

Q: What are the prerequisites or dependencies to using AWS Private CA?

1. If you plan to create a CRL, you need an S3 bucket to store it in. AWS Private CA automatically deposits the CRL in the Amazon S3 bucket you designate and updates it periodically. It is a pre requisite that the S3 bucket has the below bucket policy before you can set-up a CRL. In order to proceed with this request; create a RFC with ct-0fpjlxa808sh2 (Management | Advanced stack components | S3 storage | Update policy) as follows:

- Provide the S3 bucket name or ARN.
- Copy the below policy onto RFC and replace `bucket-name` with your desired S3 bucket name.

```
{
   "Version":"2012-10-17",
   "Statement":[
      {
         "Effect":"Allow",
         "Principal":{
            "Service":"acm-pca.amazonaws.com"
         },
         "Action":[
            "s3:PutObject",
            "s3:PutObjectAcl",
            "s3:GetBucketAcl",
            "s3:GetBucketLocation"
         ],
         "Resource":[
            "arn:aws:s3:::bucket-name/*",
            "arn:aws:s3:::bucket-name"
         ]
```

```
      }
   ]
}
```

2. If the above S3 bucket is encrypted, then the Service Principal acm-pca.amazonaws.com requires permissions to decrypt. In order to proceed with this request; create a RFC with ct-3ovo7px2vsa6n (Management | Advanced stack components | KMS key | Update) as follows:

- Provide the KMS Key ARN on which the policy must be updated.
- Copy the below policy onto RFC and replace `bucket-name` with your desired S3 bucket name.

```
{
   "Sid":"Allow ACM-PCA use of the key",
   "Effect":"Allow",
   "Principal":{
      "Service":"acm-pca.amazonaws.com"
   },
   "Action":[
      "kms:GenerateDataKey",
      "kms:Decrypt"
   ],
   "Resource":"*",
   "Condition":{
      "StringLike":{
         "kms:EncryptionContext:aws:s3:arn":[
            "arn:aws:s3:::bucket_name/acm-pca-permission-test-key",
            "arn:aws:s3:::bucket_name/acm-pca-permission-test-key-private",
            "arn:aws:s3:::bucket_name/audit-report/*",
            "arn:aws:s3:::bucket_name/crl/*"
         ]
      }
   }
}
```

3. AWS Private CA CRLs don't support the S3 setting "Block public access to buckets and objects granted through new access control lists (ACLs)". You must disable this setting with the S3 account and bucket in order to allow the AWS Private CA to write CRLs as mentioned in How to securely create and store your CRL for ACM Private CA If you would like to disable, create a new RFC with ct-0xdawir96cy7k (Management | Other | Other | Update) and attach a Risk Acceptance. If you have any questions on risk acceptance, reach out to your Cloud Architect.

AWS CloudEndure (AMS SSPS)

AWS CloudEndure migration simplifies, expedites, and automates large-scale migrations from physical, virtual, and cloud-based infrastructure to AWS. CloudEndure Disaster Recovery (DR) protects against downtime and data loss from any threat, including ransomware and server corruption.

AWS CloudEndure in AWS Managed Services FAQs

Q: How do I request access to CloudEndure in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM User to your account: `customer_cloud_endure_user`. After it's provisioned in your account, the access key and secret key for the user is shared in AWS Secrets Manager.

These policies are provisioned to the account as well: `customer_cloud_endure_policy` and `customer_cloud_endure_deny_policy`.

Additionally, you must provide a Risk Acceptance as the CloudEndure DR solution for application integration has infrastructure-mutating permissions. To do this, work with your cloud service delivery manager (CSDM).

Q: What are the restrictions to using CloudEndure in my AMS account?

The cloud endure replication and conversion instances can be launched only in the subnet you indicate.

Q: What are the prerequisites or dependencies to using CloudEndure in my AMS account? Share the following via RFC bidirectional correspondence:

- VPC Subnet details for Replication and Conversion instances to be launched.
- The KMS Key Amazon Resource Name (ARN) if the EBS volumes are encrypted.

AWS CloudHSM (AMS SSPS)

The AWS CloudHSM service helps you meet corporate, contractual, and regulatory compliance requirements for data security by using dedicated Hardware Security Module (HSM) instances within the AWS cloud. AWS, and AWS Marketplace partners, offer a variety of solutions for protecting sensitive data within the AWS platform, but for some applications and data subject

to contractual or regulatory mandates for managing cryptographic keys, additional protection may be necessary. AWS CloudHSM complements existing data protection solutions and allows you to protect your encryption keys within HSMs that are designed and validated to government standards for secure key management. AWS CloudHSM allows you to securely generate, store, and manage cryptographic keys used for data encryption in a way that keys are accessible only by you. To learn more, see AWS CloudHSM.

AWS CloudHSM in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to AWS CloudHSM in my AMS account?

Utilization of in your AMS account is a two-step process:

1. Request an AWS CloudHSM cluster. Do this by submitting an RFC with the Management | Other | Other | Create (ct-1e1xtak34nx76) change type. Include the following details:
   - AWS Region.
   - VPC ID/ARN. Provide a VPC ID/VPC ARN that is in the same account as the RFC that you submit.
   - Specify at least two Availability Zones for the cluster.
   - Amazon EC2 instance ID that will connect to the HSM cluster.
2. Access the AWS CloudHSM console. Do this by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer_cloudhsm_console_role`.

After the role is provisioned in your account, you must onboard it in your federation solution.

Q: What are the restrictions to using AWS CloudHSM in my AMS account?

Access to the AWS CloudHSM console doesn't provide you with the ability to create, terminate or restore your cluster. To do those things, submit a Management | Other | Other | Create change type (ct-1e1xtak34nx76) change type.

Q: What are the prerequisites or dependencies to using AWS CloudHSM in my AMS account?

You must allow TCP traffic using port 2225 through a client Amazon EC2 instance within a VPC, or use Direct Connect VPN for on-premise servers that want access to the HSM cluster. AWS CloudHSM is dependent on Amazon EC2 for security groups and network interfaces. For log

monitoring or auditing, HSM relies on CloudTrail (AWS API operations) and CloudWatch Logs for all local HSM device activity.

Q: Who will apply updates to the AWS CloudHSM client and related software libraries?

You are responsible for applying the library and client updates. You'll want to monitor the CloudHSM version history page for releases, and then apply updates using the CloudHSM client upgrade.

> **Note**
>
> Software patches for the HSM appliance are always automatically applied by the AWS CloudHSM service.

AWS CodeBuild (AMS SSPS)

AWS CodeBuild is a fully managed continuous integration service that compiles source code, runs tests, and produces software packages that are ready to deploy. With CodeBuild, you don’t need to provision, manage, and scale your own build servers. CodeBuild scales continuously and processes multiple builds concurrently, so your builds are not left waiting in a queue. You can get started quickly by using prepackaged build environments, or you can create custom build environments that use your own build tools. With CodeBuild, you are charged by the minute for the compute resources you use. To learn more, see AWS CodeBuild.

> **Note**
>
> To onboard CodeCommit, CodeBuild, CodeDeploy, and CodePipeline with a single RFC, submit the Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type and request the three services: CodeBuild, CodeDeploy and CodePipeline. Then, all three roles, `customer_codebuild_service_role`, `customer_codedeploy_service_role`, and `aws_code_pipeline_service_role` are provisioned in your account. After provisioning in your account, you must onboard the role in your federation solution.

CodeBuild in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to AWS CodeBuild in my AMS account?

Utilization of AWS CodeBuild in your AMS account is a two-step process:

1. Provision the `CodeBuild Service Role` for build process to coordinate with AWS S3 buckets, Amazon CloudWatch and Log groups
2. Request access to the CodeBuild console

You can request that both be set up in your AMS account by submitting an RFC with the Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using AWS CodeBuild in my AMS account?

For AWS CodeBuild console administrator access, permissions are limited at resource level; for example, CloudWatch actions are limited on specific resources and the `iam:PassRole` permission is controlled.

Q: What are the prerequisites or dependencies to using CodeBuild in my AMS account?

If additional IAM permissions are required for the defined AWS CodeBuild service role, request them through an AMS service request.

AWS CodeCommit (AMS SSPS)

AWS CodeCommit is a fully managed source control service that hosts secure Git-based repositories. It helps teams to collaborate on code in a secure and highly scalable ecosystem. CodeCommit eliminates the need to operate your own source control system or worry about scaling its infrastructure. You can use CodeCommit to securely store anything from source code to binaries, and it works seamlessly with your existing Git tools. To learn more, see AWS CodeCommit.

> **Note**
>
> To onboard CodeCommit, CodeBuild, CodeDeploy, and CodePipeline with a single RFC, submit the Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type and request the three services: CodeBuild, CodeDeploy and CodePipeline. Then, all three roles, `customer_codebuild_service_role`, `customer_codedeploy_service_role`, and `aws_code_pipeline_service_role` are

provisioned in your account. After provisioning in your account, you must onboard the role in your federation solution.

CodeCommit in AWS Managed Services FAQs

Q: How do I request access to CodeCommit in my AMS account?

AWS CodeCommit console and data access roles can be requested through the submission of two AWS Service RFCs, console access, and data access:

- Request access to AWS CodeCommit by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer_codecommit_console_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

  Data access (such as Training and Entity Lists) require separate CTs for each data source specifying the S3 data source (mandatory), output bucket (mandatory) and KMS (optional). There are no limitations to AWS CodeCommit job creation as long as all data sources have been granted access roles. To request data access, submit an RFC with the Management | Other | Other | Create (ct-1e1xtak34nx76).

Q: What are the restrictions to using AWS CodeCommit in my AMS account?

Triggers feature on CodeCommit are disabled given the associated rights to create SNS topics. Directly authenticating against CodeCommit is restricted, users should authenticate with Credential Helper. Some KMS commands are also restricted: `kms:Encrypt`, `kms:Decrypt`, `kms:ReEncrypt`, `kms:GenereteDataKey`, `kms:GenerateDataKeyWithoutPlaintext`, and `kms:DescribeKey`.

Q: What are the prerequisites or dependencies to using AWS CodeCommit in my AMS account?

If S3 buckets are encrypted with KMS keys, S3 and KMS are required to use AWS CodeCommit.

AWS CodeDeploy (AMS SSPS)

AWS CodeDeploy is a fully managed deployment service that automates software deployments to a variety of compute services such as Amazon EC2, AWS Fargate, AWS Lambda, and your on-premises servers. AWS CodeDeploy helps you to rapidly release new features, helps you avoid downtime during application deployment, and handles the complexity of updating your

applications. You can use AWS CodeDeploy to automate software deployments, eliminating the need for error-prone manual operations. The service scales to match your deployment needs. To learn more, see AWS CodeDeploy.

> **Note**
>
> To onboard CodeCommit, CodeBuild, CodeDeploy, and CodePipeline with a single RFC, submit the Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type and request the three services: CodeBuild, CodeDeploy and CodePipeline. Then, all three roles, `customer_codebuild_service_role`, `customer_codedeploy_service_role`, and `aws_code_pipeline_service_role` are provisioned in your account. After provisioning in your account, you must onboard the role in your federation solution.

CodeDeploy in AWS Managed Services FAQs

Q: How do I request access to CodeDeploy in my AMS account?

Request access to CodeDeploy by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM roles to your account: `customer_codedeploy_console_role` and `customer_codedeploy_service_role`. After it's provisioned in your account, you must onboard the `customer_codedeploy_console_role` role in your federation solution.

Q: What are the restrictions to using CodeDeploy in my AMS account?

Currently we are only supporting Compute Platform as — Amazon EC2/On-premises. Blue/Green Deployments are not supported.

Q: What are the prerequisites or dependencies to using CodeDeploy in my AMS account?

There are no prerequisites or dependencies to use CodeDeploy in your AMS account.

AWS CodePipeline (AMS SSPS)

AWS CodePipeline is a fully managed continuous delivery service that helps you automate your release pipelines for fast and reliable application and infrastructure updates. CodePipeline automates the build, test, and deploy phases of your release process every time there is a code

change, based on the release model you define. This enables you to rapidly and reliably deliver features and updates. You can easily integrate AWS CodePipeline with third-party services such as GitHub or with your own custom plugin. With AWS CodePipeline, you only pay for what you use. There are no upfront fees or long-term commitments. To learn more, see AWS CodePipeline.

> **Note**
>
> To onboard CodeCommit, CodeBuild, CodeDeploy, and CodePipeline with a single RFC, submit the Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type and request the three services: CodeBuild, CodeDeploy and CodePipeline. Then, all three roles, `customer_codebuild_service_role`, `customer_codedeploy_service_role`, and `aws_code_pipeline_service_role` are provisioned in your account. After provisioning in your account, you must onboard the role in your federation solution.
> CodePipeline in AMS does not support "Amazon CloudWatch Events" for Source Stage because it needs elevated permissions to create the service role and policy, which bypasses the least-privileges model and AMS change management process.

CodePipeline in AWS Managed Services FAQs

Q: How do I request access to CodePipeline in my AMS account?

Request access to CodePipeline by submitting a service request for the `customer_code_pipeline_console_role` in the relevant account. After it's provisioned in your account, you must onboard the role in your federation solution.

At this time, AMS Operations will also deploy this service role in your account: `aws_code_pipeline_service_role_policy`.

Q: What are the restrictions to using CodePipeline in my AMS account?

Yes. CodePipeline features, stages, and providers are limited to the following:

1. Deploy Stage: Limited to Amazon S3, and AWS CodeDeploy
2. Source Stage: Limited to Amazon S3, AWS CodeCommit, BitBucket, and GitHub
3. Build Stage: Limited to AWS CodeBuild, and Jenkins
4. Approval Stage: Limited to Amazon SNS

5. Test Stage: Limited to AWS CodeBuild, Jenkins, BlazeMeter, Ghost Inspector UI Testing, Micro Focus StormRunner Load, and Runscope API Monitoring
6. Invoke Stage: Limited to Step Functions, and Lambda

> **Note**
>
> AMS Operations will deploy `customer_code_pipeline_lambda_policy` in your account; it must be attached with the Lambda execution role for Lambda invoke stage. Please provide the Lambda service/execution role name that you want this policy added with. If there is no custom Lambda service/execution role, AMS will create a new role named `customer_code_pipeline_lambda_execution_role`, which will be a copy of `customer_lambda_basic_execution_role` along with `customer_code_pipeline_lambda_policy`.

Q: What are the prerequisites or dependencies to using CodePipeline in my AMS account?

AWS supported services AWS CodeCommit, AWS CodeBuild, AWS CodeDeploy must be launched prior to, or along with, the launch of CodePipeline.

AWS Compute Optimizer (AMS SSPS)

AWS Compute Optimizer recommends optimal AWS Compute resources for your workloads to reduce costs and improve performance by using machine learning to analyze historical utilization metrics. Over-provisioning compute (Amazon EC2 and ASGs) can lead to unnecessary infrastructure cost and under-provisioning compute can lead to poor application performance. Compute Optimizer helps you choose the optimal Amazon EC2 instance types, including those that are part of an Amazon EC2 Auto Scaling group, based on your utilization data. To learn more, see AWS Compute Optimizer.

Compute Optimizer in AWS Managed Services FAQs

Q: How do I request access to Compute Optimizer in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_compute_optimizer_readonly_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Compute Optimizer in my AMS account?

There are no restrictions. Full functionality of AWS Compute Optimizer is available in your AMS account.

Q: What are the prerequisites or dependencies to using Compute Optimizer in my AMS account?

- You must submit an RFC (Management | Other | Other | Update) authorizing AMS Ops to enable the service in the account. During deployment, a service linked role (SLR) is created to allow metrics gathering and report generation. The SLR is labeled "AWSServiceRoleForComputeOptimizer". For more information, see Using Service-Linked Roles for AWS Compute Optimizer
- CloudWatch metrics must be enabled for the following metrics:
  - **CPU utilization**: The percentage of allocated Amazon EC2 compute units that are in use on the instance. This metric identifies the processing power required to run an application upon a selected instance.
  - **Memory utilization**: The amount of memory that has been used in some way during the sample period. This metric identifies the memory required to run an application upon a selected instance. Memory utilization is analyzed only for resources that have the unified CloudWatch agent installed on them. For more information, see Enabling Memory Utilization with the CloudWatch Agent (p. 10).
  - **Network in**: The number of bytes received on all network interfaces by the instance. This metric identifies the volume of incoming network traffic to a single instance.
  - **Network out**: The number of bytes sent out on all network interfaces by the instance. This metric identifies the volume of outgoing network traffic from a single instance.
  - **Local disk input/output (I/O)**: The number of input/output operations for the local disk. This metric identifies the performance of the root volume of an instance

AWS DataSync (AMS SSPS)

AWS DataSync moves large amounts of data online between on-premises storage and Amazon S3, Amazon Elastic File System (Amazon Elastic File System) or Amazon FSx. Manual tasks related to data transfers can slow down migrations and burden IT operations. DataSync eliminates or automatically handles many of these tasks, including scripting copy jobs, scheduling and monitoring transfers, validating data, and optimizing network utilization. The DataSync software agent connects to your Network File System (NFS) and Server Message Block (SMB) storage, so you

don't have to modify your applications. DataSync can transfer hundreds of terabytes and millions of files at speeds up to 10 times faster than open-source tools, over the internet or AWS Direct Connect links. You can use DataSync to migrate active data sets or archives to AWS, transfer data to the cloud for timely analysis and processing, or replicate data to AWS for business continuity.

To learn more, see AWS DataSync.

DataSync in AWS Managed Services FAQs

Q: How do I request access to DataSync in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_datasync_console_role`.

Once provisioned in your account, you must onboard the roles in your federation solution.

The CloudWatch log group to use in order to stream task logs is "/aws/datasync".

Q: What are the restrictions to using DataSync in my AMS account?

Full functionality of AWS DataSync is available in your AMS account.

Q: What are the prerequisites or dependencies to using DataSync in my AMS account?

- S3 ARNs (Amazon Resource Names) are required for all S3 buckets associated with DataSync tasks that will be performed using the DataSync service role. All bucket names must have the prefix `datasync-`, as these are the only buckets the DataSync service-linked role (`customer_datasync_service_role`) will be able to access.
- VPC Endpoints and security groups for DataSync agents must be requested with an RFC with the Management | Other | Other | Create (ct-1e1xtak34nx76) change type prior to using VPC Endpoints.
- AWS DataSync agents run in AMS as an appliance. The AWS DataSync agent is patched and updated by the service; for details, see AWS DataSync FAQs.
- To launch an AWS DataSync agent, submit an RFC with the Management | Other | Other | Create (ct-1e1xtak34nx76) change type, requesting the agent be deployed. Provide the AWS DataSync Amazon EC2 AMI ID, instance type, subnet, security group; and either reference an existing Amazon EC2 keypair or request the creation of a new keypair.

> **Note**
>
> AMS will provision the AWS DataSync agent manually on behalf of customer, and it does not require the WIGS ingestion process on the AWS DataSync Amazon EC2 AMI.

AWS Device Farm (AMS SSPS)

AWS Device Farm is an application testing service that lets you improve the quality of your web and mobile apps by testing them across an extensive range of desktop browsers and real mobile devices; without having to provision and manage any testing infrastructure. The service enables you to run your tests concurrently on multiple desktop browsers or real devices to speed up the execution of your test suite, and generates videos and logs to help you quickly identify issues with your app.

To learn more, see AWS Device Farm.

AWS Device Farm in AWS Managed Services FAQs

Q: How do I request access to AWS Device Farm in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_devicefarm_role`.

Once provisioned in your account, you must onboard the roles in your federation solution.

Q: What are the restrictions to using AWS Device Farm in my AMS account?

Full access to the AWS Device Farm service is provided with the exception of using the AMS namespace in the 'Name' tag.

Q: What are the prerequisites or dependencies to using AWS Device Farm in my AMS account?

None.

AWS Elastic Disaster Recovery (AMS SSPS)

AWS Elastic Disaster Recovery minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery. You can increase IT resilience when you use AWS Elastic Disaster Recovery

to replicate on-premises or cloud-based applications running on supported operating systems. Use the AWS Management Console to configure replication and launch settings, monitor data replication, and launch instances for drills or recovery.

To learn more, see AWS Elastic Disaster Recovery.

AWS Elastic Disaster Recovery in AWS Managed Services FAQs

Q: How do I request access to AWS Elastic Disaster Recovery in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_drs_console_role`.

After its provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using AWS Elastic Disaster Recovery in my AMS account?

There are no restrictions to use AWS Elastic Disaster Recovery in your AMS account.

Q: What are the prerequisites or dependencies to using AWS Elastic Disaster Recovery in my AMS account?

- After you have access to the console role, you must initialize the Elastic Disaster Recovery service to create the needed IAM roles within the account.
  - You must submit a Management | Other | Other RFC to create a clone of the `customer-mc-ec2-instance-profile` instance profile and attach the `AWSElasticDisasterRecoveryEc2InstancePolicy` policy. You must specify which machines to attach the new policy to.
  - If the instance isn't using the default instance profile, then AMS can attach `AWSElasticDisasterRecoveryEc2InstancePolicy` through automation.
- You must use a customer-owned KMS key for cross-account recovery. The source account's KMS key must be updated following the policy to allow target account access. For more information, see Share the EBS encryption key with the target account.
- The KMS key policy must be updated to allow the allow `customer_drs_console_role` to view the policy if you don't want to switch roles to view.
- For cross-account, cross-Region disaster recovery, AMS must set up the source and target account as Trusted Accounts and deploy the Failback and in-AWS right-sizing roles through AWS CloudFormation.

AWS Elemental MediaConvert (AMS SSPS)

AWS Elemental MediaConvert is a file-based video transcoding service with broadcast-grade features. It enables you to create video-on-demand (VOD) content for broadcast and multiscreen delivery at scale. The service combines advanced video and audio capabilities with a simple web services interface and pay-as-you-go pricing. With AWS Elemental MediaConvert, you can focus on delivering compelling media experiences without having to worry about the complexity of building and operating your own video processing infrastructure.

To learn more, see AWS Elemental MediaConvert.

MediaConvert in AWS Managed Services FAQs

Q: How do I request access to MediaConvert in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_mediaconvert_author_role`. Once provisioned in your account, you must onboard the role in your federation solution.

A second role will be provided, `customer_MediaConvert_Default_Role`, that is used by MediaConvert in order to read from the source S3 bucket and write the output to the destination S3 bucket, and also to invoke the API gateway in case you need digital rights management (DRM).

Q: What are the restrictions to using MediaConvert in my AMS account?

There are no restrictions for the use of MediaConvert in AMS.

Q: What are the prerequisites or dependencies to using MediaConvert in my AMS account?

There are no prerequisites or dependencies to use MediaConvert in your AMS account.

AWS Elemental MediaLive (AMS SSPS)

AWS Elemental MediaLive is a broadcast-grade live video processing service. It enables you to create high-quality video streams for delivery to broadcast televisions and internet-connected multiscreen devices, like connected TVs, tablets, smartphones, and set-top boxes. The service works by encoding your live video streams in real-time, taking a larger-sized live video source and compressing it into smaller versions for distribution to your viewers. With AWS Elemental MediaLive, you can easily set up streams for both live events and 24x7 channels with advanced

broadcasting features, high availability, and pay-as-you-go pricing. AWS Elemental MediaLive lets you focus on creating compelling live video experiences for your viewers without the complexity of building and operating broadcast-grade video processing infrastructure.

To learn more, see AWS Elemental MediaLive.

MediaLive in AWS Managed Services FAQs

Q: How do I request access to MediaLive in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_medialive_author_role`.

As a part of this RFC, a second role is deployed into your account; `customer_medialive_service_role` role, this role can be assigned to your Media Live channels and inputs to interact with other services such as Amazon S3, MediaStore, and CloudWatch Logs.

After the roles are provisioned in your account, you must onboard the roles in your federation solution.

Q: What are the restrictions to using MediaLive in my AMS account?

There are no restrictions for the use of MediaLive in AMS.

Q: What are the prerequisites or dependencies to using MediaLive in my AMS account?

There are no prerequisites or dependencies to use MediaLive in your AMS account.

AWS Elemental MediaPackage (AMS SSPS)

AWS Elemental MediaPackage reliably prepares and protects your video for delivery over the internet. From a single video input, AWS Elemental MediaPackage creates video streams formatted to play on connected TVs, mobile phones, computers, tablets, and game consoles. It makes it easy to implement popular video features for viewers (start-over, pause, rewind, and so on.), like those commonly found on DVRs. AWS Elemental MediaPackage can also protect your content using Digital Rights Management (DRM). AWS Elemental MediaPackage scales automatically in response to load, so your viewers will always get a great experience without you having to accurately predict in advance the capacity you'll need.

To learn more, see AWS Elemental MediaPackage.

MediaPackage in AWS Managed Services FAQs

Q: How do I request access to AWS Elemental MediaPackage in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_mediapackage_author_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

A second role will be provided, `customer_mediapackage_service_role`, that can be assigned to your Media Live channels and inputs to interact with other services such as S3 and Secrets Manager.

Q: What are the restrictions to using MediaPackage in my AMS account?

There are no restrictions for the use of MediaPackage in AMS.

Q: What are the prerequisites or dependencies to using MediaPackage in my AMS account?

There are no prerequisites or dependencies to use MediaPackage in your AMS account.

AWS Elemental MediaStore (AMS SSPS)

AWS Elemental MediaStore is an AWS storage service optimized for media. It gives you the performance, consistency, and low latency required to deliver live streaming video content. AWS Elemental MediaStore acts as the origin store in your video workflow. Its high performance capabilities meet the needs of the most demanding media delivery workloads, combined with long-term, cost-effective storage. To learn more, see AWS Elemental MediaStore.

MediaStore in AWS Managed Services FAQs

Q: How do I request access to MediaStore in my AMS account?

Request access to MediaStore by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer_mediastore_author_role`. As a part of this RFC, a second role is deployed into your account; `MediaStoreAccessLogs` role, which is used by the MediaStore

service to log activity in CloudWatch, if you choose to enable that feature. After it's provisioned in your account, you must onboard the roles in your federation solution.

At this time, AMS Operations will also deploy this service role in your account: `aws_code_pipeline_service_role_policy`.

Q: What are the restrictions to using MediaStore in my AMS account?

There are no restrictions for the use of MediaStore in AMS.

Q: What are the prerequisites or dependencies to using MediaStore in my AMS account?

There are no prerequisites or dependencies to use MediaStore in your AMS account.

AWS Elemental MediaTailor (AMS SSPS)

AWS Elemental MediaTailor lets video providers insert individually targeted advertising into their video streams without sacrificing broadcast-level quality-of-service. With AWS Elemental MediaTailor, viewers of your live or on-demand video each receive a stream that combines your content with ads personalized to them. But unlike other personalized ad solutions, with AWS Elemental MediaTailor your entire stream – video and ads – is delivered with broadcast-grade video quality to improve the experience for your viewers. AWS Elemental MediaTailor delivers automated reporting based on both client and server-side ad delivery metrics, to accurately measure advertising impressions and viewer behavior. You can easily monetize unexpected high-demand viewing events with no up-front costs using AWS Elemental MediaTailor. It also improves ad delivery rates, helping you make more money from every video, and it works with a wider variety of content delivery networks, ad decision servers, and client devices.

To learn more, see AWS Elemental MediaTailor.

MediaTailor in AWS Managed Services FAQs

Q: How do I request access to MediaTailor in my AMS account?

Request access to MediaTailor by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer-mediatailor-role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using MediaTailor in my AMS account?

There are no restrictions for the use of MediaTailor in AMS.

Q: What are the prerequisites or dependencies to using MediaTailor in my AMS account?

There are no prerequisites or dependencies to use MediaTailor in your AMS account.

AWS Global Accelerator (AMS SSPS)

Global Accelerator is a network layer service in which you create accelerators to improve availability and performance for internet applications used by a global audience. To learn more, see Global Accelerator.

Global Accelerator in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request Global Accelerator to be set up in my AMS account?

Request access through the submission of the AWS Services RFC (Management | AWS service | Self-provisioned Service). Through this RFC, the following IAM roles will be provisioned in your account: `customer_global_accelerator_console_role`. Once provisioned in your account you must onboard the console role in your federation solution.

Q: What are the restrictions to using Global Accelerator in my AMS account?

Global Accelerator is a global service that supports endpoints in multiple AWS Regions, which are listed in the AWS Region Table.

Q: What are the prerequisites or dependencies to using Global Accelerator in my AMS account?

When you set up your accelerator with Global Accelerator, you associate the static IP addresses to regional endpoints in one or more AWS Regions. For standard accelerators, the endpoints are Network Load Balancers, Application Load Balancers, Amazon EC2 instances, or Elastic IP addresses. For custom routing accelerators, endpoints are virtual private cloud (VPC) subnets with one or more EC2 instances.

AWS Glue (AMS SSPS)

AWS Glue is a fully managed extract, transform, and load (ETL) service that helps you to prepare and load your data for analytics. You can create and run an ETL job with a few clicks in the AWS

Management Console. You point AWS Glue to your data stored on AWS, and AWS Glue discovers your data and stores the associated metadata (e.g. table definition and schema) in the AWS Glue Data Catalog. Once cataloged, your data is immediately searchable, queryable, and available for ETL actions. To learn more, see AWS Glue.

AWS Glue in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request AWS Glue to be set up in my AMS account?

Request access to AWS Glue by submitting an RFC with the Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM role to your account: `customer_glue_console_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

For access to Crawlers, Jobs, and Development endpoints (roles needed for specific use cases), submit an RFC with the Deployment | Advanced stack components | Identity and Access Management (IAM) | Create entity or policy (ct-3dpd8mdd9jn1r).

Q: What are the restrictions to using AWS Glue in my AMS account?

There are no restrictions. Full functionality of AWS Glue is available in your AMS account with the exception of Notebooks. AWS Glue Notebooks are a non-managed resource that launches Amazon EC2 instances in an account. AMS recommends that you launch your own Amazon EC2 instances and install the software necessary to support a notebook environment and development. For more information, see Tutorial: Set Up a Local Apache Zeppelin Notebook to Test and Debug ETL Scripts and Using Development Endpoints for Developing Scripts.

Q: What are the prerequisites or dependencies to using AWS Glue in my AMS account?

AWS Glue has a dependency on Amazon S3, CloudWatch, and CloudWatch Logs. Transitive dependencies vary based on data sources, and other AWS Glue service features may be interacting with (example: Amazon Redshift, Amazon RDS, Athena).

AWS Lake Formation (AMS SSPS)

AWS Lake Formation is a service that makes it easy to set up a secure data lake in days. A data lake is a centralized, curated, and secured repository that stores all your data, both in its original form

and prepared for analysis. A data lake enables you to break down data silos and combine different types of analytics to gain insights and guide better business decisions.

Creating a data lake with Lake Formation is as simple as defining data sources and what data access and security policies you want to apply. Lake Formation then helps you collect and catalog data from databases and object storage, move the data into your new Amazon S3 data lake, clean and classify your data using machine learning algorithms, and secure access to your sensitive data. Your users can access a centralized data catalog (for details, see AWS Glue FAQs) that describes available data sets and their appropriate usage. Your users then leverage these data sets with their choice of analytics and machine learning services, like Amazon Redshift, Amazon Athena, and (in beta) Amazon EMR for Apache Spark. Lake Formation builds on the capabilities available in AWS Glue.

To learn more, see AWS Lake Formation.

Lake Formation in AWS Managed Services FAQs

Q: How do I request access to AWS Lake Formation in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer_lakeformation_data_analyst_role`. After it's provisioned in your account, you must onboard the roles in your federation solution.

Q: What are the restrictions to using AWS Lake Formation in my AMS account?

Full functionality of Lake Formation is available in AMS.

Q: What are the prerequisites or dependencies to using AWS Lake Formation in my AMS account?

Lake Formation integrates with the AWS Glue service, therefore AWS Glue users can access only the databases and tables on which they have Lake Formation permissions. Additionally AWS Athena and Amazon Redshift users can only query the AWS Glue databases and tables on which they have Lake Formation permissions.

AWS Lambda (AMS SSPS)

AWS Lambda lets you run code without provisioning or managing servers. You pay only for the compute time you consume, there is no charge when your code is not running. With Lambda, you can run code for virtually any type of application or back-end service, all with zero administration.

upload your code and Lambda takes care of everything required to run and scale your code with high availability. You can set up your code to automatically trigger from other AWS services, or call it directly from any Web or mobile app. To learn more, see AWS Lambda.

Lambda in AWS Managed Services FAQs

Q: How do I request access to AWS Lambda in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM roles to your account: `customer_lambda_admin_role` and `customer_lambda_basic_execution_role`. After it's provisioned in your account, you must onboard the roles in your federation solution.

Q: What are the restrictions to using AWS Lambda in my AMS account?

- A Lambda function is designed to be invoked by event sources. For a list of services that can be used as a Lambda event source, see Using AWS Lambda with Other Services. Not all of these services are currently available in AMS accounts. If you require a service that is not available, work with your AMS CSDM to file an exception.
- By default AMS will provide you with a basic Lambda initiation role containing the `AWSLambdaBasicExecutionRole` and `AWSXrayWriteOnlyAccess` permissions; for information, see AWS Lambda Initiation Role. If you require additional permissions, such as the ability to provision Lambda functions within your AMS VPC, submit an RFC with the Management | Other | Other | Create (ct-1e1xtak34nx76).

Q: What are the prerequisites or dependencies to using AWS Lambda in my AMS account?

There are no prerequisites or dependencies to get started with AWS Lambda; however, depending on your specific use case, you may require access to other AWS services in order to create event sources, or additional permissions for your function to perform various actions. If additional permissions are needed, submit an RFC with the Management | Other | Other | Create (ct-1e1xtak34nx76).

AWS License Manager (AMS SSPS)

AWS License Manager integrates with AWS services to simplify the management of licenses across multiple AWS accounts, IT catalogs, and on-premises, through a single AWS account. AWS License Manager lets administrators create customized licensing rules that emulate the terms of

their licensing agreements, and then enforces these rules when an instance of Amazon EC2 gets launched. The rules in AWS License Manager enable you to limit a licensing breach by physically stopping the instance from launching or by notifying administrators about the infringement. To learn more, see AWS License Manager.

License Manager in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request AWS License Manager to be set up in my AMS account?

Request access to AWS License Manager by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer_license_manager_role`. Once the License Manager IAM role is provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using AWS License Manager in my AMS account?

You're able to associate AWS License Manager rules to the AMIs you own (filtered under "Owned by me"). If you choose to enforce a limit association to an AMI (example: can only support 100 vCPU of this AMI) and exhaust the limit, future launches with that AMI are blocked and return an error stating "No licenses available." This is the intended behavior of this service (not allowing license exhaustion). In the event you exhaust the limit but need to launch the AMI again, you must modify the rule configured in AWS License Manager.

Q: What are the prerequisites or dependencies to using AWS License Manager in my AMS account?

There are no prerequisites or dependencies to use AWS License Manager in your AMS account.

AWS Migration Hub (AMS SSPS)

AWS Migration Hub provides a single location where you can track the progress of application migrations across multiple AWS and partner solutions. Using Migration Hub allows you to choose the AWS and partner migration tools that best fit your needs, while providing visibility into the status of migrations across your application portfolio. Migration Hub also provides key metrics and progress for individual applications, regardless of which tools are being used to migrate them. This allows you to quickly get progress updates across all of your migrations, easily identify and troubleshoot any issues, and reduce the overall time and effort spent on your migration projects. To learn more, see AWS Migration Hub.

Migration Hub in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to Migration Hub in my AMS account?

Request access to Migration Hub by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer_migrationhub_author_role`. Once provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions for Migration Hub?

None.

Q: What are the prerequisites to enable Migration Hub?

There are no prerequisites to start using Migration Hub in your AMS account. However, permissions outside Migration Hub might be required during the management of the service, such as writing permissions to Amazon S3 to upload server information.

AWS Outposts (AMS SSPS)

AWS Outposts is a fully managed service that extends AWS infrastructure, AWS services, APIs, and tools to virtually any datacenter, co-location space, or on-premises facility for a consistent hybrid experience. AWS Outposts is good for workloads that require low latency access to on-premises systems, local data processing, or local data storage. To learn more, see AWS Outposts.

AWS Outposts in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request AWS Outposts to be set up in my AMS account?

Request access to AWS Outposts by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer_outposts_role`. Once the role is provisioned in your account, you must onboard it in your federation solution.

Q: What are the restrictions to using AWS Outposts in my AMS account?

There are no restrictions for the use of AWS Outposts in your AMS account.

Q: What are the prerequisites or dependencies to using AWS Outposts in my AMS account?

There are no prerequisites or dependencies to use AWS Outposts in your AMS account.

AWS Resilience Hub (AMS SSPS)

AWS Resilience Hub helps you proactively prepare and protect your AWS applications from disruptions. The Resilience Hub offers resiliency assessment and validation that integrate into your software development lifecycle to uncover resiliency weaknesses. Resilience Hub helps you estimate whether or not your applications can meet the recovery time objective (RTO) and recovery point objective (RPO) targets, and helps resolve issues before they are released into production. After you deploy an AWS application into production, you can use Resilience Hub to continue tracking the resiliency posture of your application. If an outage occurs, Resilience Hub sends a notification to the operator to launch the associated recovery process.

AWS Resilience Hub in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to AWS Resilience Hub in my AMS account?

Request access to Resilience Hub by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM roles and policies to your account:

IAM roles

- `customer_resiliencehub_console_role`
- `customer_resiliencehub_service_role`

Policies

- `customer_resiliencehub_console_policy`
- `customer_resiliencehub_service_policy`

After the role is provisioned in your account, you must onboard the role `customer_resiliencehub_console_role` in your federation solution.

Q: What are the restrictions to using AWS Resilience Hub in my AMS account?

There are no restrictions. Full functionality of Resilience Hub is available in your AMS acount.

Q: What are the prerequisites or dependencies to using AWS Resilience Hub in my AMS account?

There are no prerequisites or dependencies to use Resilience Hub in your AMS account.

AWS Secrets Manager (AMS SSPS)

AWS Secrets Manager helps you protect secrets needed to access your applications, services, and IT resources. The service enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Users and applications retrieve secrets with a call to the Secrets Manager APIs, eliminating the need to hardcode sensitive information in plain text. Secrets Manager offers secret rotation with built-in integration for Amazon RDS, Amazon Redshift, and Amazon DocumentDB. Also, the service is extensible to other types of secrets, including API keys and OAuth tokens. To learn more, see AWS Secrets Manager.

> **Note**
>
> By default, AMS operators can access secrets in AWS Secrets Manager that are encrypted using the account's default AWS KMS key (CMK). If you want your secrets to be inaccessible to AMS Operations, use a custom CMK, with an AWS Key Management Service (AWS KMS) key policy that defines permissions appropriate to the data stored in the secret.

Secrets Manager in AWS Managed Services FAQs

Q: How do I request access to AWS Secrets Manager in my AMS account?

Request access to Secrets Manager by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM roles to your account: `customer_secrets_manager_console_role` and `customer-rotate-secrets-lambda-role`. The `customer_secrets_manager_console_role` is used as an Admin role to provision and manage the secrets, and `customer-rotate-secrets-lambda-role` is used as the Lambda execution role for the Lambda functions that rotate the secrets. After it's provisioned in your account, you must onboard the `customer_secrets_manager_console_role` role in your federation solution.

Q: What are the restrictions to using AWS Secrets Manager in my AMS account?

Full functionality of AWS Secrets Manager is available in your AMS account, along with automatic rotation functionality of secrets. However, note that setting up your rotation using 'Create a new Lambda function to perform rotation' is not supported because it requires elevated permissions to create the AWS CloudFormation stack (IAM Role and Lambda function creation), which bypasses the Change Management process. AMS Advanced only supports 'Use an existing Lambda function to perform rotation' where you manage your Lambda functions to rotate secrets using the AWS Lambda SSPS Admin role. AMS Advanced doesn't create or manage Lambda to rotate the secrets.

Q: What are the prerequisites or dependencies to using AWS Secrets Manager in my AMS account?

The following namespaces are reserved for use by AMS and are unavailable as part of direct access to AWS Secrets Manager:

- arn:aws:secretsmanager:*:*:secret:ams-shared/*
- arn:aws:secretsmanager:*:*:secret:customer-shared/*
- arn:aws:secretsmanager:*:*:secret:ams/*

Sharing keys using Secrets Manager (AMS SSPS)

Sharing secrets with AMS in the plain text of an RFC, service request, or incident report, results in an information disclosure incident and AMS redacts that information from the case and requests that you regenerate the keys.

You can use AWS Secrets Manager (Secrets Manager) under this namespace, `customer-shared`.

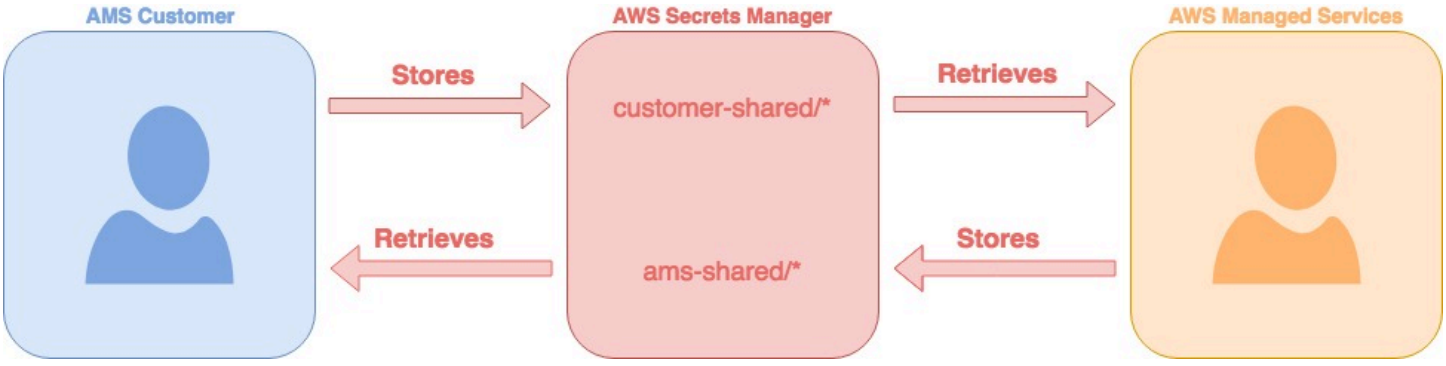

Sharing Keys using Secrets Manager FAQs

Q: What type of secrets must be shared using Secrets Manager?

A few examples are pre-shared keys for VPN creation, confidential keys such as Authentication keys (IAM, SSH), License keys and Passwords.

Q: How can I share the keys with AMS using Secrets Manager?

1. Login to the AWS Management console using your federated access and the appropriate role:

   for SALZ, the `Customer_ReadOnly_Role`

   for MALZ, `AWSManagedServicesChangeManagementRole`.
2. Navigate to the AWS Secrets Manager console and click **Store a new secret**.
3. Select **Other type of secrets**.
4. Enter the secret value as a plain-text and use the default KMS encryption. Click **Next**.
5. Enter the secret name and description, the name always starts with **customer-shared/**. For example **customer-shared/mykey2022**. Click **Next**.
6. Leave automatic rotation disabled, Click **Next**.
7. Review and click **Store** to save the secret.
8. Reply to us with the secret name through the Service request, RFC, or incident report, so we can identify and retrieve the secret.

Q: What permissions are required for sharing the keys using Secrets Manager?

SALZ: Look for the `customer_secrets_manager_shared_policy` managed IAM policy and verify that the policy document is the same as the one attached in the creation steps below. Confirm that the policy is attached to the following IAM Roles: `Customer_ReadOnly_Role`.

MALZ: Validate that the `AMSSecretsManagerSharedPolicy`, is attached to the `AWSManagedServicesChangeManagementRole` role that allows you the `GetSecretValue` action in the `ams-shared`namespace.

Example:

```
{
 "Action": "secretsmanager:*",
 "Resource": [
 "arn:aws:secretsmanager:*:*:secret:ams-shared/*",
 "arn:aws:secretsmanager:*:*:secret:customer-shared/*"
 ],
```

```
"Effect": "Allow",
"Sid": "AllowAccessToSharedNameSpaces"
}
```

> **Note**
>
> The requisite permissions are granted when you add AWS Secrets Manager as a self-service provisioned service.

AWS Security Hub (AMS SSPS)

AWS Security Hub provides you with a comprehensive view of your security state within AWS and your compliance with security industry standards and best practices. Security Hub centralizes and prioritizes security and compliance findings from across AWS accounts, services, and supported third-party partners to help you analyze your security trends and identify the highest priority security issues. To learn more, see AWS Security Hub.

Security Hub in AWS Managed Services FAQs

Q: How do I request access to AWS Security Hub in my AMS account?

Request access to Security Hub by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer_securityhub_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using Security Hub in my AMS account?

Archiving functionality has been noted as a potential security and operational risk and has been restricted as a part of the self-provisioned service Security role.

Q: What are the prerequisites or dependencies to using AWS Security Hub in my AMS account?

There are no prerequisites or dependencies to use AWS Security Hub in your AMS account.

AWS Service Catalog AppRegistry

AppRegistry enables application search, reporting, and management actions from a central location. Builders seldom create applications in a single AWS account. They typically separate

application resources by lifecycle phases, such as development, test, and production. AppRegistry allows you to group and view all your resource collections across the AWS accounts that you define.

With AppRegistry, you can store your AWS applications, the collection of resources that are associated with your applications, and application attribute groups. To learn more, see What is AppRegistry.

FAQs: AWS Service Catalog AppRegistry in AMS

Q: How do I request access to AWS Service Catalog AppRegistry in my AMS account?

Request access to AppRegistry by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type. This RFC provisions the following IAM role to your account: `customer-appregistry-console-role`. After provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using AWS Service Catalog AppRegistry in my AMS account?

Full access to the AppRegistry service is provided with the exception of using the AMS namespace in the `'Name'` tag.

Q: What are the prerequisites or dependencies to using AWS Service Catalog AppRegistry in my AMS account?

There are no prerequisites or dependencies to use AppRegistry in your AMS account.

AWS Shield

AWS Shield Advanced is a managed Distributed Denial of Service (DDoS) protection service that safeguards applications running on AWS. Shield Advanced provides always-on detection and automatic inline mitigations that minimize application downtime and latency, so there is no need to engage AWS Support to benefit from DDoS protection. There are two tiers of AWS Shield - Standard and Advanced; AMS offers Shield Advanced. To learn more, see Shield Advanced.

All AWS customers benefit from the automatic protections of AWS Shield Standard, at no additional charge. AWS Shield Standard defends against most common, frequently occurring, network and transport layer DDoS attacks that target your website or applications. When you use AWS Shield Standard with Amazon CloudFront and Amazon Route 53, you receive comprehensive availability protection against all known infrastructure (Layer 3 and 4) attacks.

For higher levels of protection against attacks targeting your applications running on Amazon Elastic Compute Cloud (Amazon EC2), Elastic Load Balancing (ELB), Amazon CloudFront, AWS Global Accelerator, and Amazon Route 53 resources, you can subscribe to AWS Shield Advanced.

In addition to the network and transport layer protections that come with AWS Shield Standard, AWS Shield Advanced provides additional detection and mitigation against large and sophisticated DDoS attacks, near real-time visibility into attacks, and integration with AWS WAF, a web application firewall. AWS Shield Advanced also gives you 24x7 access to the AWS Shield Response Team (SRT) and protection against DDoS related spikes in your Amazon Elastic Compute Cloud (Amazon EC2), Elastic Load Balancing (Elastic Load Balancing), Amazon CloudFront, AWS Global Accelerator, and Amazon Route 53 charges.

Shield Advanced in AWS Managed Services FAQs

Q: How do I request access to Shield Advanced in my AMS account?

Request access to Shield Advanced by submitting an RFC with the Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM roles to your account: `customer_shield_role` and `aws_drt_shield_role`. Once provisioned in your account, you must onboard the roles in your federation solution.

After the roles are deployed into your account, you can use the `customer_shield_role` to confirm your subscription to AWS Shield Advanced in your account.

> **Note**
>
> Note that there is a monthly fee and a one-year commitment associated with the use of AWS Shield Advanced. Additionally, using AWS Shield Advanced in AMS authorizes AMS to escalate to the AWS Shield (SRT), who may make changes to your web application firewall (AWS WAF) rules during escalated distributed denial of service (DDoS) incidents. These changes will be made in coordination with AMS.

Q: What are the restrictions to using Shield Advanced in my AMS account?

Although not a restriction, you should understand that using Shield Advanced deploys the `aws_drt_shield_role`, which allows AWS Shield teams (SRT) to make emergency changes to AWS WAF rules inside of AMS accounts during escalated DDoS incidents. This is recommended by AMS for the fastest remediation of DDoS attacks, and would occur after an AMS escalation to the SRT.

Q: What are the prerequisites or dependencies to using Shield Advanced in my AMS account?

There are no prerequisites or dependencies to use Shield Advanced in your AMS account.

AWS Snowball (AMS SSPS)

Snowball is a petabyte-scale data transport solution that uses devices designed to be secure, to transfer large amounts of data into and out of the AWS Cloud. Snowball addresses common challenges with large-scale data transfers including high network costs, long transfer times, and security concerns. You can use Snowball to migrate analytics data, genomics data, video libraries, image repositories, backups, and to archive part of data center shutdowns, tape replacement or application migration projects. Transferring data with Snowball is simple, fast, more secure, and can be as little as one-fifth the cost of transferring data by way of high-speed Internet.

With Snowball, you don't need to write any code or purchase any hardware to transfer your data. Start by using the AWS Management Console to Create an Import Job for Snowball, and a Snowball device will be automatically shipped to you. Once it arrives, attach the device to your local network, download and run the Snowball Client ("Client") to establish a connection, and then use the Client to select the file directories that you want to transfer to the device. The Client then encrypts and transfers the files to the device at high speed. Once the transfer is complete and the device is ready to be returned, the E Ink shipping label automatically updates and you can track the job status with Amazon Simple Notification Service (Amazon SNS), text messages, or directly in the Console. To learn more, see AWS Snowball.

Snowball in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to AWS Snowball in my AMS account?

Implementation of Snowball in AMS is a two-step process:

1. Submit a Management | Other | Other | Create (ct-1e1xtak34nx76) change type and request a service role for Snowball for your AMS Account.
2. Request user access by submitting a Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM roles to your account: `customer_snowball_console_role`, `customer_snowball_export_role`, and `customer_snowball_import_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using AWS Snowball in my AMS account?

Full functionality of the AWS Snowball is available in your AMS account.

Q: What are the prerequisites or dependencies to using AWS Snowball in my AMS account?

You must have the service role account as noted above.

AWS Step Functions (AMS SSPS)

AWS Step Functions is a Web service that enables you to coordinate the components of distributed applications and microservices by using visual workflows. You build applications from individual components that each perform a discrete function, or task, allowing you to scale and change applications quickly. Step Functions provides a reliable way to coordinate components and step through the functions of your application. Step Functions offers a graphical console to visualize the components of your application as a series of steps. It automatically triggers and tracks each step, and retries when there are errors, so your application runs in order and as expected, every time. Step Functions logs the state of each step, so when things do go wrong, you can diagnose and debug problems quickly. To learn more, see AWS Step Functions.

Step Functions in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to AWS Step Functions in my AMS account?

Request access to AWS Step Functions by submitting an RFC with the Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM role to your account: `customer_step_functions_role`. Once provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using AWS Step Functions in my AMS account?

Full functionality of the AWS Step Functions is available in your AMS account.

Q: What are the prerequisites or dependencies to using AWS Step Functions in my AMS account?

At runtime, the role used by Step Functions must have access to the services used by the step function. For example, a step function could depend on Lambda functions. Someone authoring a

step function is likely to be creating Lambda functions at the same time and would have to request access to that service as well.

AWS Systems Manager Parameter Store (AMS SSPS)

AWS Systems Manager Parameter Store provides secure, hierarchical storage for configuration data management and secrets management. You can store data such as passwords, database strings, and license codes as parameter values. You can store values as plain text or encrypted data. You can then reference values by using the unique name that you specified when you created the parameter. Highly scalable, available, and durable, Parameter Store is backed by the AWS Cloud. To learn more, see AWS Systems Manager Parameter Store.

> **Note**
>
> If you want a dedicated secrets store with lifecycle management, use AWS Secrets Manager (AMS SSPS) instead of Parameter Store. Secrets Manager helps you meet your security and compliance requirements by enabling you to rotate secrets automatically. Secrets Manager offers built-in integration for MySQL, PostgreSQL, and Amazon Aurora on Amazon RDS, that's extensible to other types of secrets by customizing Lambda functions.

AWS Systems Manager Parameter Store in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to Systems Manager Parameter Store in my AMS account?

Request access to AWS Systems Manager Parameter Store by submitting an RFC with the Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM role to your account: `customer_systemsmanager_parameterstore_console_role`. Once provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using AWS Systems Manager Parameter Store in my AMS account?

You are required to use AWS Managed keys; access is restricted from creating custom KMS keys. However, if a custom key is required, submit an RFC to create a customer-managed key (CMK) using the Deployment | Advanced Stack Components | KMS Key | Create change type (ct-1d84keiri1jhg)

with this IAM role, `customer_systemsmanager_parameterstore_console_role` as the value for the `IAMPrincipalsRequiringDecryptPermissions` and `IAMPrincipalsRequiringEncryptPermissionsPrincipal` parameters. After the KMS Key is created, you can create a Secure String using it.

Q: What are the prerequisites or dependencies to using AWS Systems Manager Parameter Store in my AMS account?

There are no prerequisites; however, SSM Parameter Store is dependent on KMS to create a Secure String so you can encrypt and decrypt their Values stored in Parameter Store.

AWS Systems Manager Automation (AMS SSPS)

AWS Systems Manager Automation simplifies common maintenance and deployment tasks of Amazon Elastic Compute Cloud instances and other AWS resources using runbooks, actions and service quotas. It enables you to build, execute and monitor automations at scale. A Systems Manager Automation is a type of Systems Manager document that defines the actions that Systems Manager performs on your managed instances. A runbook you use to perform common maintenance and deployment tasks such as running commands or automation scripts within your managed instances. Systems Manager includes features that help you target large groups of instances by using Amazon Elastic Compute Cloud tags, and velocity controls that help you roll out changes according to the limits you define. The runbooks are written using JavaScript Object Notation (JSON) or YAML. Using the Document Builder in the Systems Manager Automation console, however, you can create a runbook without having to author in native JSON or YAML. Alternatively you can use Systems Manager-provided runbooks with pre-defined steps that suits your needs. To learn more, see Working with runbooks in AWS Systems Manager documentation.

> **Note**
>
> Although Systems Manager Automation supports 20 action types that can be used in the runbook, a limited number of actions you can use while authoring runbook to be used in your AMS Advanced account. Similarly, a limited number of Systems Manager-provided runbook can be used either directly or from within your own runbook. For details, see the restrictions in the following FAQ.

AWS Systems Manager Automation in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to Systems Manager Automation in my AMS account?

Request access to AWS Systems Manager Automation by submitting an RFC with the Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM role to your account: `customer_systemsmanager_automation_console_role`. Once provisioned in your account, you must onboard the role in your federation solution.

Q: What are the limitations to using AWS Systems Manager Automation in my AMS account?

You are required to author your runbook, with limited set of Systems Manager supported actions for automation, only to run commands and/or scripts within your managed instances. The actions that are available to you along with any restrictions are outlined as below.

AWS Systems Manager Automation Limitations

Action	Description	Limitation
aws:assertAwsResourceProperty –	Assert an AWS resource state or event state	Only EC2 instances
aws:aws:branch –	Run conditional automation steps	No limitation
aws:createTags –	Create tags for AWS resources	Only to SSM automation runbooks that you author
aws:executeAutomation –	Run another automation	Only the automation runbook that you author
aws:executeScript –	Run a script	Only script that does not make any API call to any services
aws:pause –	Pause an automation	No limitation
aws:runCommand –	Run a command on a managed instance	Only using System Manager provided document - AWS-RunShellScript and AWS-RunPowerShellScript

Action	Description	Limitation
aws:sleep –	Delay an automation	No limitation
aws:waitForAwsResourceProperty –	Wait on an AWS resource property	Only EC2 instances

You can also chose to run command or script directly with Systems Manager provided runbook AWS-RunShellScript and AWS-RunPowerShellScript using the 'Run Command' feature from within the Systems Manager console. You can also nest these runbooks within your runbook that caters for additional pre and/or post validation or any complex automation logic.

The role adheres to least privilege principle and only provides permission required to author, execute and retrieve execution details of runbooks aimed to executing command and/or scripts within your managed instances. It does not provide permission for any other capabilities that AWS Systems Manager service provides. While the feature allows you to author automation runbooks, execution of the runbooks can not be targeted for AMS owned resources.

Q: What are the prerequisites or dependencies to using AWS Systems Manager Automation in my AMS account?

There are no prerequisites; however, you must ensure your internal process and/or compliance controls are adhered to while authoring runbooks. We also recommend to thoroughly test runbooks before executing them against production resources.

Q: Can the Systems Manager policy `customer_systemsmanager_automation_policy` be attached to other IAM roles?

No, unlike other self-provision enabled services, this policy can only be assigned to the provisioned default role `customer_systemsmanager_automation_console_role`.

Unlike the policies of other SSPS roles, this SSM SSPS policy cannot be shared with other custom IAM roles, because this AMS service is only for running commands or automation scripts within your managed instances. If these permissions were allowed to be attached to other custom IAM roles, potentially with permissions on other services, the scope of allowed actions could extend to managed services, and potentially lower the security posture of your account.

To evaluate any requests for change (RFCs) against our AMS technical standards, work with your respective Cloud Architect or Service Delivery Manager, see RFC security reviews.

Note

AWS Systems Manager allows you to use runbooks that are shared with your account. We recommend you exercise caution and perform a due-diligence check when using shared runbooks and make sure to review the content to understand the command/scripts they run before executing the runbooks. For details refer to Best practices for shared SSM documents.

AWS Transfer Family (AMS SSPS)

AWS Transfer Family (Transfer Family) is a fully managed AWS service that enables you to transfer files over Secure File Transfer Protocol (SFTP), into and out of Amazon Simple Storage Service (Amazon S3) storage. SFTP is also known as Secure Shell (SSH) File Transfer Protocol. SFTP is used in data exchange workflows across different industries such as financial services, healthcare, advertising, and retail, among others.

With AWS SFTP, you get access to an SFTP server in AWS without the need to run any server infrastructure. You can use this service to migrate your SFTP-based workflows to AWS while maintaining your end users' clients and configurations as is. You first associate your hostname with the SFTP server endpoint, then add your users and provision them with the right level of access. After you do, your users' transfer requests are serviced directly out of your AWS SFTP server endpoint. To learn more, see AWS Transfer for SFTP, also Create an SFTP-enabled server.

AWS Transfer for SFTP in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to AWS Transfer for SFTP in my AMS account?

Request access to AWS Transfer for SFTP by submitting an RFC with the Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). Through this RFC the following IAM roles, and a policy, are provisioned in your account:

- `customer_transfer_author_role`. This role is designed for you to manage the SFTP service through the console.
- `customer_transfer_sftp_server_logging_role`. This role is designed to be attached on the SFTP Server. It allows the SFTP server to pull logs into CloudWatch.

- `customer_transfer_sftp_user_role`. This role is designed to be attached on the SFTP users. It allows the SFTP Users to interact with the S3 bucket.
- `policy customer_transfer_scope_down_policy`. This policy is a scope-down policy that can be applied to the SFTP User to limit their access on the S3 bucket to their home folders.
- `customer_transfer_sftp_efs_user_role`. This role is designed to be attached on the SFTP users. It allows the SFTP Users to interact with the EFS file system.

After it's provisioned in your account, you must onboard the roles in your federation solution.

Q: What are the restrictions to using AWS Transfer for SFTP in my AMS account?

AWS Transfer for SFTP configuration is limited to resources without "AMS-" or "MC-" prefixes to prevent any modifications to AMS infrastructure.

Q: What are the prerequisites or dependencies to using AWS Transfer for SFTP in my AMS account?

- You must have an S3 bucket before creating the AWS Transfer for SFTP server and users.
- To use a "Customer Identify Provider," you must deploy the API Gateway, Lambda function, and your user repository (AD, Secrets Manager, and so on). For more information, see Enable password authentication for AWS Transfer for SFTP using AWS Secrets Manager and Working with Identity Providers

AWS Transit Gateway (AMS SSPS)

AWS Transit Gateway is a service that enables you to connect your Amazon Virtual Private Clouds (VPCs) and your on-premises networks to a single gateway. As you grow the number of workloads running on AWS, you need to be able to scale your networks across multiple accounts and Amazon VPCs to keep up with the growth. Today, you can connect pairs of Amazon VPCs using peering. However, managing point-to-point connectivity across many Amazon VPCs, without the ability to centrally manage the connectivity policies, can be operationally costly and cumbersome. For on-premises connectivity, you need to attach your AWS VPN to each individual Amazon VPC. This solution can be time consuming to build and hard to manage when the number of VPCs grows into the hundreds. To learn more, see AWS Transit Gateway.

Transit Gateway in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to AWS Transit Gateway in my AMS account?

Request access to AWS Transit Gateway by submitting an RFC with the Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM role to your account: `customer_tgw_console_role`. Once provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using AWS Transit Gateway in my AMS account?

Full functionality of AWS Transit Gateway is available in your AMS single-account landing zone account for the exception of route table modifications for TransitGateway routing. Request route table changes by submitting a Management | Other | Other | Create change type (ct-1e1xtak34nx76).

> **Note**
>
> This service is only supported for single-account landing zone (SALZ), not multi-account landing zone (MALZ).

Q: What are the prerequisites or dependencies to using AWS Transit Gateway in my AMS account?

There are no prerequisites or dependencies to use AWS Transit Gateway in your AMS account.

AWS WAF - Web Application Firewall (AMS SSPS)

AWS WAF is a web application firewall (AWS WAF) that helps protect your web applications from common web exploits that could affect application availability, compromise security, or consume excessive resources. AWS WAF gives you control over which traffic to allow, or block, to your web applications by defining customizable web security rules. You can use AWS WAF to create custom rules that block common attack patterns, such as SQL injection or cross-site scripting; and rules that are designed for your specific application.

To learn more, see AWS WAF - Web Application Firewall.

AMS doesn't support monitoring (CloudWatch alarms / events / MMS alerts) for AWS WAF. Due to the nature of AWS WAF, you must create custom rules for your applications; AMS can't quantify and create alarms for you, without context of your application. To learn more, see AWS WAF - Web Application Firewall.

AWS WAF in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request AWS WAF to be set up in my AMS account?

Request access to AWS WAF by submitting an RFC with the Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM role to your account: `customer_waf_role`. After the AWS WAF IAM role is provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using AWS WAF?

After permissions are provisioned, you have the full functionality of AWS WAF.

Q: What are the prerequisites or dependencies to using AWS WAF?

There are no prerequisites or dependencies to use AWS WAF in your AMS account.

AWS Well-Architected Tool (AMS SSPS)

The AWS Well-Architected Tool helps you review the state of your workloads and compares them to the latest AWS architectural best practices. The tool is based on the AWS Well-Architected Framework, developed to help cloud architects build secure, high-performing, resilient, and efficient application infrastructure. This framework provides a consistent approach for you to evaluate architectures, has been used in tens of thousands of workload reviews conducted by the AWS solutions architecture team, and provides guidance to help implement designs that scale with application needs over time. To learn more, see AWS Well-Architected Tool.

AWS WA Tool in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to AWS Well-Architected Tool in my AMS account?

Request access to AWS Well-Architected Tool by submitting an RFC with the Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM role to your account: `customer_well_architected_tool_console_admin_role`. After it's provisioned in your account, you must onboard the role in your federation solution.

Q: What are the restrictions to using AWS Well-Architected Tool in my AMS account?

Full functionality of the AWS Well-Architected Tool is available in your AMS account.

Q: What are the prerequisites or dependencies to using AWS Well-Architected Tool in my AMS account?

There are no prerequisites or dependencies to use AWS Well-Architected Tool in your AMS account.

AWS X-Ray (AMS SSPS)

AWS X-Ray (X-Ray) helps developers analyze and debug production, distributed applications, such as those built using a microservices architecture. With X-Ray, you can understand how your application and its underlying services are performing, to identify and troubleshoot the root cause of performance issues and errors. X-Ray provides an end-to-end view of requests as they travel through your application, and shows a map of your application's underlying components. You can use X-Ray to analyze both applications in development and in production, from simple three-tier applications, to complex microservices applications consisting of thousands of services. To learn more, see AWS X-Ray.

X-Ray in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to AWS X-Ray in my AMS account?

Request access by submitting a Management | AWS service | Self-provisioned service | Add (ct-1w8z66n899dct) change type. This RFC provisions the following IAM role to your account: `customer_xray_console_role`. After it's provisioned in your account, you must onboard the role in your federation solution. Additionally, you must have the `customer_xray_daemon_write_instance_profile` to push data from your Amazon EC2 instances to X-Ray. This instance profile is created when you receive the `customer_xray_console_role`.

You can submit a service request to AMS Operations to assign the `customer_xray_daemon_write_policy` to the existing instance profile, or you can use the instance profile that is created when AMS Operations enables X-Ray for you.

Q: What are the restrictions to using AWS X-Ray in my AMS account?

Full functionality of AWS X-Ray is available in your AMS account except for encryption with AWS KMS key (KMS key). AWS X-Ray encrypts all trace data by default. By default, X-Ray encrypts traces and related data at rest. If you need to encrypt data at rest with a key, you can choose either AWS-

managed KMS key (aws/xray) or KMS Customer-Managed key. For KMS Customer-Managed key for X-Ray encryption, submit a Management | Other | Other | Create change type (ct-1e1xtak34nx76).

Q: What are the prerequisites or dependencies to using AWS X-Ray in my AMS account?

AWS X-Ray has a dependency on Amazon S3, CloudWatch, and CloudWatch Logs, which are already implemented in AMS accounts. Transitive dependencies vary based on data sources and other AWS service AWS X-Ray that features may be interacting with (for example, Amazon Redshift, Amazon RDS, Athena).

VM Import/Export (AMS SSPS)

VM Import/Export enables you to easily import virtual machine images from your existing environment to Amazon EC2 instances and export them back to your on-premises environment. This offering allows you to leverage your existing investments in the virtual machines that you have built to meet your IT security, configuration management, and compliance requirements by bringing those virtual machines into Amazon EC2 as ready-to-use instances. You can also export imported instances back to your on-premises virtualization infrastructure, allowing you to deploy workloads across your IT infrastructure. To learn more, see VM Import/Export.

VM Import/Export in AWS Managed Services FAQs

Common questions and answers:

Q: How do I request access to VM Import/Export in my AMS account?

Request access to VM Import/Export by submitting an RFC with the Management | AWS service | Self-provisioned service | Add change type (ct-1w8z66n899dct). This RFC provisions the following IAM policy to your account: `customer_vmimport_policy`. After it's provisioned in your account, you must onboard the role in your federation solution.

An additional role, the **VM Import/Export Service** role, is required for the service to perform actions in your account.

Q: What are the restrictions to using VM Import/Export in my AMS account?

- Functionality to import custom machine images and data volumes is both available in AMS VM Import/Export. However, permissions to S3 have been scoped down to limit actions to buckets matching the name `customer-vmimport-*` in order to limit access to information within the account.

- Image and snapshot import is supported in AMS VM Import/Export. However, instance import and instance export functionality is not available due to security measures.
- Additionally, export functionality has been disabled to mitigate the risk of exporting restricted and sensitive data.

Q: What are the prerequisites or dependencies to using VM Import/Export in my AMS account?

- You must provide a supported disk image to import into the AWS environment. For information, see VM Import/Export Requirements.
- Note: VM Import/Export is not accessible through the AWS console. The service can only be accessed through the AWS CLI, AWS Tools for PowerShell, and the AWS SDKs. A **VM Import/Export enabled** role must be requested by an AMS RFC (Management | Other | Other | Create), and then you have to access the service directly with the previously mentioned tools. Alternatively, you can request an instance profile by request for change (RFC, ct-19jq3ulr3g9zg) through which the tools can perform commands from an instance.

Customer Managed mode

AWS Managed Services (AMS) Customer Managed mode provides a governance model that is flexible and can be adapted to your requirements. This can be considered a fallback option for services and applications that AMS is unable to operate for you. AMS does not operate infrastructure hosted in accounts created under this mode. However, you can leverage centralized multi-account management in this mode. The following Multi-Account Landing Zone features can be leveraged in this mode:

- Automated Account deployment
- Connectivity through Transit Gateway in networking account
- AMS Config Rules library
- Store copies of logs in logging account
- Aggregation of customer managed Guard Duty alerts to Security account
- Consolidated Billing
- Enablement of custom Service Control Policies.

For example: If you want to run workloads on Ubuntu Pro, which is not an Operating System managed by AMS, you could use a customer managed account for hosting it. You can also

consolidate workloads through customer managed accounts, to take advantage of the bulk discount on Reserved Instances/Sharing Plans available through sharing across an AWS organization.

Getting started with Customer Managed mode

The AMS Customer Managed mode is available through a special multi-account landing zone Application account.

For details, including how to create a Customer Managed Application account, see Customer Managed application accounts.

AMS and AWS Service Catalog

Service Catalog in AWS Managed Services (AMS) allows organizations to create and manage catalogs of AWS information technology (IT) services and enables IT administrators to create, manage, and distribute catalogs of approved products to end users in their accounts, who can then access the products they need in a personalized portal of services. Administrators can control which users have access to each product to enforce compliance with organizational business policies. Administrators can also set up roles so that end users only require IAM access to Service Catalog in order to deploy approved resources. Service Catalog allows your organization to benefit from increased agility and reduced costs because end users can find and launch only the products they need from a catalog that you control.

Service Catalog provides you with an alternative to the AMS request for change (RFC) process for provisioning and updating resources in your AMS managed account(s). AMS manages all of the infrastructure operations tasks needed to run AWS at scale for all infrastructure resources provisioned through Service Catalog including security, compliance, provisioning, availability, patch, monitoring, alerting, reporting, incident response, and cost optimization. Utilizing Service Catalog in your AMS managed account provides you with a mechanism to centrally manage commonly deployed IT services and helps you achieve consistent governance while enabling users to quickly deploy only the approved IT services they need into their managed environments.

Getting started with Service Catalog

To get started with Service Catalog in AMS, submit a service request through the AMS console to request access to Service Catalog. Upon submission of the request, three IAM roles will be deployed into your account(s) along with an AMS managed stack containing the CloudFormation macro that invokes the AMS `Transform` (described previously) so we can register the products

in our systems, and to perform operations against the infrastructure provisioned through Service Catalog. The three IAM roles deployed include a role for IT admins to manage products as Service Catalog admins; a role for application owners and end-users to configure, launch, and manage products; and a role that will be used as a launch constraint, that defines the permissions that Service Catalog will use while launching or updating the your product.

Service Catalog in AMS before you begin

Does Service Catalog replace the existing AMS request for change (RFC) process?

In accounts where Service Catalog is enabled, it will act as the change management system in which you provision and update IT services in your AMS account through your predefined product catalog; AMS will provide a default portfolio/product catalog, and your IT admins can create and configure your own. Service Catalog will only acknowledge stacks provisioned through Service Catalog. Likewise, services provisioned through Service Catalog will not be modifiable through the AMS RFC process as modification outside of Service Catalog will drift the stack from the approved product configuration.

Can I see stacks provisioned through service catalog in the AMS Console?

Yes. You can view all stacks provisioned through service catalog in the AMS console. Stacks provisioned through service catalog are easily identifiable by the stack ID of "SC-". Although stacks are viewable in the AMS console you will not be able to update through the AMS RFC process. Access to the AMS change management system (RFCs) is limited to access request, patch orchestration and back-up RFCs only.

If I provision and/or update a stack through Service Catalog will there be a corresponding RFC in the AMS Console?

The only RFC that will show in the AMS console is an RFC to register the stack with AMS when a stack is initially provisioned. This RFC is filed automatically by the AMS validation process that is triggered when a stack is launched through Service Catalog. All other provisioning and changes are tracked directly in Service Catalog and are viewable in the Service Catalog console. Furthermore, you can use the **Provisioned Product Plan** feature in Service Catalog to view the list of changes that will be made to the resources in advance of provisioning or updating the product.

Do I have to do anything specific for provisioning products in my AMS managed account?

Yes. All Service Catalog products provisioned in AMS accounts must contain this line of JSON in the CFN template that defines that product:

```
"Transform":{"Name":"AmsStackTransform","Parameters":{"StackId":
{"Ref":"AWS::StackId"}}}
```

This snippet of CloudFormation code triggers the AMS validations required before the resource can be provisioned in your AMS managed account. It is your responsibility to include this line of code as part of the product definition. If it is not included, provisioning will fail and the following error message will be displayed: "Failed to create product. This account is managed by AMS. All products in AMS accounts must have the AMS `Transform` code in the template."

Is there any Service Catalog functionality not available and/or limited for AMS customers at launch?

Yes, the following SC features are not available for AMS customers at initial launch:

- Account Creation through Service Catalog
- Ability to launch all AWS Services through Service Catalog into an AMS-managed account. AWS Service availability is limited to AMS supported services (managed and self-provisioned). For more information on AMS-supported services, see the AMS service description.
- Service Catalog IT service manager (ITSM) connectors will not communicate with AMS incident reports, and service requests.
- Ability to leverage Service Catalog quick starts and reference architectures without modification. Remember that Service Catalog products for AMS accounts must contain this line of JSON code:

  ```
  "Transform":{"Name":"AmsStackTransform","Parameters":{"StackId":
  {"Ref":"AWS::StackId"}}}
  ```

  in the CNF template. Note that this line is *not* part of a typical AWS CloudFormation template and must be explicitly added.
- Terraform is not currently supported by AMS for provisioning Service Catalog products.
- AWS CFN stacksets are not supported in AMS.
- You cannot create custom IAM roles.
- Service Actions are limited to:
  - AWS-RebootRdsInstance
  - AWS-RestartEC2Instance
  - AWS-StartEC2Instance

- AWS-StartRdsInstance
- AWS-StopEC2Instance
- AWS-StopRdsInstance
- AWS-CreateImage
- AWS-CreateRdsSnapshot
- AWS-CreateSnapshot

> **Note**
>
> When creating service actions, you can configure the execution role to be the end user's permissions, the launch role, or a custom IAM role of your choosing. The selected execution role must have sufficient permissions to perform the service action, and have a TrustPolicy that allows it to be assumed by Service Catalog, otherwise that service action will fail at execution time. We recommend using the AWSManagedServicesServiceCatalogLaunchRole, which has the correct permissions and trust policy to be used as a service action.

What will I still need to use the AMS RFC system for?

At general availability (GA) you will still need to use RFCS to run the following actions:

- Configuring Patch Orchestrator
- Configuring Back up policies
- Requesting instance access
- Creating and assigning security groups that fall outside AMS guidelines.
- Performing workload ingest (WIGS)
- Creating IAM roles

Can I use the Service Catalog CLI to access Service Catalog in my AMS managed account?

Yes, Service Catalog APIs are available and enabled through the CLI. Actions from the management of Service Catalog artifacts through the provisioning and terminating of those artifacts, are available. For more information, see AWS Service Catalog Resources, or download the latest AWS SDK or CLI.

Who creates, manages, and distributes customers' catalogs of approved products?

The customer's catalog administrator and/or IT administrator, or assigned resource, is responsible for the management of your Service Catalog catalogs and approved products.

Can I use AMS AMIs?

AMS AMIs vended after March 2020 can be deployed through AWS Service Catalog.

How do I migrate to AMS using Service Catalog?

To migrate your workload to AMS using Service Catalog you begin by following the Workload Ingest (WIGs) process to create an AMI in AMS. You use the AMI produced by WIGS to create a product in Service Catalog. How to do this is detailed in AWS Service Catalog - Getting Started.

Finding the data you need (SKMS), AMS

Finding the data you need when using your AWS Managed Services (AMS) accounts calls on the AMS service knowledge management, or SKMS, system. AMS.

SKMS stands for service knowledge management system and refers to all information related to the AWS Managed Services (AMS) service for a customer. AMS has an SKMS API for finding data.

Topics

- What Is service knowledge management?
- Finding VPC IDs in AMS
- Finding subnet IDs in AMS
- Find AMI IDs, AMS
- Find security group (SG) IDs, AMS
- Find IAM entities in AMS
- Find stack IDs in AMS
- Find instance IDs or IP addresses in AMS
- Find Amazon Resource Names (ARNs) in AMS
- Find resources by ARN in AMS
- Find AMS account settings

What Is service knowledge management?

Service knowledge management is the store of all information on your AMS account. Information about the following is obtained from the AMS service knowledge management system (SKMS), through the AMS SKMS API or through the AMS Console:

- VPCs
- Managed subnets
- Stacks and stack components, including Amazon EC2 instances and other resources
- Amazon Machine Images (AMIs)

You can use information from the SKMS to understand the infrastructure under management and as input to change management and service requests to create, change, or remove infrastructure.

> **Note**
>
> All AMS SKMS API calls are recorded in AWS CloudTrail.

Access the SKMS through the AMS SKMS API, which provides operations for discovering information about an environment (VPCs and subnets) and the application resources (stacks, Amazon EC2 instances, and instance images or AMIs) that can be deployed there.

VPCs and instance images are set up in an account, with the necessary access permissions, during onboarding. After they have been established, you can use the change management system to populate the VPCs with working stacks.

Finding VPC IDs in AMS

A virtual private cloud (VPC) has one or more subnets. In AMS your VPC is in an AWS Region and you have private and public subnets.

See also Finding subnet IDs in AMS.

Some CTs require the VpcId. To find a VPC ID, you can use either the AMS console or API/CLI.

AMS Console:

In the navigation pane, select **VPCs** and the relevant VPC. The VPC details page for the selected VPC opens with information including the VPC ID.

AMS SKMS API ListVpcSummaries or CLI:

> **Note**
>
> The AMS CLI must be installed for these commands to work. To install the AMS API or CLI, go to the AMS console **Developers Resources** page. For reference material on the AMS CM API or AMS SKMS API, see the AMS Information Resources section in the User Guide. You may need to add a `--profile` option for authentication; for example, `aws amsskms` *`ams-cli-command`* `--profile SAML`. You may also need to add the `--region` option as

all AMS commands run out of us-east-1; for example `aws amscm ams-cli-command --region=us-east-1`.

Note

The AMS API/CLI (amscm and amsskms) endpoints are in the AWS N. Virginia Region, `us-east-1`. Depending on how your authentication is set, and what AWS Region your account and resources are in, you may need to add `--region us-east-1` when issuing commands. You may also need to add `--profile saml`, if that is your authentication method.

1. In the following examples, the first command requests a list of summaries for all VPCs in the account. The second command requests the list of VPCs, with a query filter to list only those VPCs created in 2016, and output the CreatedTime, VpcId, and Name.

   Note

   You can obtain the AMS SKMS CLI through the **Developer's Resources** page in the AMS console.

   ```
   aws amsskms list-vpc-summaries --output table
   ```

   ```
   -----------------------------------------------------
   |                 ListVPCSummaries                  |
   +---------------------------------------------------+
   |                  VPCSummaries                     |
   |+------------------+------------------------------|
   |   CreatedTime     |   2016-01-15T18:50:11Z       |
   |   VpcId           |   vpc-01234567890abcdef      |
   |   LastModifiedTime|   2016-01-15T18:50:11Z       |
   |   Name            |   952444781316-initial-vpc   |
   |+------------------+------------------------------|
   |                  Visibility                       |
   |+------------------+------------------------------|
   |   Id              |  PrivateAndPublic            |
   |   Name            |  PrivateAndPublic            |
   ```

```
|+------------------+-------------------------------|
```

2. This time with a query:

```
aws amsskms list-VPC-summaries --query "VPCSummaries[?
starts_with(@.CreatedTime,to_string(`2016`))].[CreatedTime, VpcId, Name]" --output
 table
```

```
-------------------------------------------------------------------------
|                              ListVPCSummaries                         |
+---------------------+-----------------------+--------------------------+
|2016-01-15T18:50:11Z | vpc-01234567890abcdef | 952444781316-initial-VPC |
+---------------------+-----------------------+--------------------------+
```

Finding subnet IDs in AMS

Several resources require that you specify a subnet, or list of subnets, at configuration time. To find subnets, you can use either the AMS console or AMS SKMS API/CLI. Note that the AMS SKMS API/CLI is private and must be installed before you can use it.

AMS Console:

1. In the navigation pane, select **VPCs** and the relevant VPC. The VPC details page for the selected VPC opens with a table of subnets, click a subnet ID to open the details page and find the ID.

AMS SKMS API ListSubnetSummaries or CLI:

> **Note**
>
> The AMS CLI must be installed for these commands to work. To install the AMS API or CLI, go to the AMS console **Developers Resources** page. For reference material on the AMS CM API or AMS SKMS API, see the AMS Information Resources section in the User Guide. You may need to add a `--profile` option for authentication; for example, `aws amsskms` *`ams-cli-command`* `--profile SAML`. You may also need to add the `--region` option as all AMS commands run out of us-east-1; for example `aws amscm` *`ams-cli-command`* `--region=us-east-1`.

> **Note**
>
> The AMS API/CLI (amscm and amsskms) endpoints are in the AWS N. Virginia Region, `us-east-1`. Depending on how your authentication is set, and what AWS Region your account and resources are in, you may need to add `--region us-east-1` when issuing commands. You may also need to add `--profile saml`, if that is your authentication method.

To find the subnets for your VPC, you can search with the `list-subnet-summaries` command as shown.

> **Note**
>
> If you're looking for subnets that are not in an AMS account, you can try `aws ec2 describe-subnets --region us-west-2`.

1. The SKMS API/CLI ListSubnetSummaries operation:

   A simple list:

   ```
   aws amsskms list-subnet-summaries
   ```

   Output to a table:

   ```
   aws amsskms list-subnet-summaries --output table
   ```

2. The SKMS API ListSubnetSummaries operation has parameters to narrow the results based on visibility. In addition, you can Filter results based on name. If you're using the CLI, you can also use the `--query` option to narrow the output or search on a portion of a value. For example, to find all of the subnets for a particular VPC, you can use this command:

   ```
   aws amsskms list-subnet-summaries --query
    "SubnetSummaries.sort_by(@,&Visibility.Name)[].[Visibility.Name,SubnetId,Name]" --
   output table
   ```

   Which returns something like this:

```
-------------------------------------------------------------------
|                   ListSubnetSummaries                           |
+---------+------------         -------+---------------------------+
|  Private|  subnet-01234567890abcdef  |  Demo Deployment Zone #1  |
|  Private|  subnet-01234567890abcdef  |  Demo Deployment Zone #1  |
|  Public |  subnet-01234567890abcdef  |  Demo DMZ #1              |
|  Public |  subnet-01234567890abcdef  |  Demo DMZ #1              |
+---------+----------         ---------+---------------------------+
```

For information about using CLI queries, see How to Filter the Output with the --query Option and the query language reference, JMESPath Specification.

3. If you have multiple VPCs, include a VPC filter in the command, and then run the command for each VPC. For example:

```
list-subnet-summaries --filter Attribute=VpcId,Value=vpc-xxxxxxxx --query
 "SubnetSummaries.sort_by(@,&Visibility.Name)[].[Visibility.Name,SubnetId,Name]" --
output table
```

4. In AWS, use describe-subnets.

For information about using CLI queries, see How to Filter the Output with the --query Option and the query language reference, JMESPath Specification..

Subnet names

Your AMS subnets are created automatically after input is gathered from you and added to the system. AMS uses a formula to create your subnet names: A*ACCOUNT_ID*-*SUBNET-TYPE*-*AZ-IDENTIFIER*. The subnet type would be either `dmz`, `shared-services`, or `customer-application`. Should you have more than one customer-application subnet, an optional identifier may be added to the subnet name, after the account ID, to indicated that the subnet is an "additional" or "reserved" subnet.

Find AMI IDs, AMS

An Amazon Machine Image, or AMI, is a template for Amazon EC2 instances, created from an Amazon EC2 instance. AWS provides updated AMIs (with patches, for example) every month; however, AWS Managed Services (AMS) requires AMIs that have been modified for AMS use. AMS releases new AMIs that you can use shortly after Patch Tuesday every month.

Amazon Machine Images (AMIs) are instance configuration templates that are used to create EC2 instances in AWS. AMS requires that specific AMIs be used for AMS-managed resources. The change types for creating EC2 instances and EC2 Auto Scaling groups require that you specify an AMI for AMS to use as the basis for the instances that the change type creates. AMS recommends that you always select the most recent AMI available to you.

To learn more about AWS AMIs, see AWS AMI Design.

When creating an Amazon EC2 stack or Amazon EC2 Auto Scaling group for your AMS account, you must specify an AMI by **AmiId**. You're limited to AMIs that begin with "customer-" and we recommend that you always choose the most recent AMI.

To find the most recent AMI for your account, you can search with an AMS SKMS CLI command or use the AMS console details page for relevant VPC:

- Use the AMS console: Available AMIs are listed on the **AMI** page in the AMS console. Select from AMIs with names that begin with "customer-".
- Use the AMS SKMS API/CLI ListAmis operation.

> **Note**
>
> The AMS CLI must be installed for these commands to work. To install the AMS API or CLI, go to the AMS console **Developers Resources** page. For reference material on the AMS CM API or AMS SKMS API, see the AMS Information Resources section in the User Guide. You may need to add a `--profile` option for authentication; for example, `aws amsskms` *`ams-cli-command`* `--profile SAML`. You may also need to add the `--region` option as all AMS commands run out of us-east-1; for example `aws amscm` *`ams-cli-command`* `--region=us-east-1`.

Here is a CLI example with a `query` option that restricts the results to customer AMIs:

```
aws amsskms list-amis --vpc-id VPC_ID --query "Amis.sort_by(@,&Name)[?
starts_with(Name,'customer')].[Name,AmiId]" --output table
```

This example uses the `filter` option with the `query` option to find Windows AMIs that start with "customer":

```
aws amsskms list-amis --vpc-id VPC_ID --query "Amis.sort_by(@,&Name)[?
starts_with(Name,'customer')].[Name,AmiId]" --filter Attribute=Platform,Value=windows
 --output table
```

- For information about using CLI queries, see How to Filter the Output with the --query Option and the query language reference, JMESPath Specification.

Find security group (SG) IDs, AMS

Amazon EC2 create and OpenSearch create domain CTs require a security group ID. This will be in the form `sg-02ce123456e7893c7`. Your account has at least two default security groups; see Security groups. Additionally, you may have security groups that you created for specific purposes. To discover your security groups:

- AWS Console: Use the EC2 or VPC console to view all security groups for the selected VPC.
- API/CLI (when logged into your AMS account):

  List your security groups:

```
aws ec2 describe-security-groups
```

Find IAM entities in AMS

Your account has default IAM Roles and Policies; see IAM user role in AMS and default IAM instance profiles; see EC2 IAM instance profile with default policies. To discover your IAM roles and policies:

- Console: Use the IAM console to view all IAM policies and roles for your account.
- API/CLI (when logged into your AMS account):

> **Note**
>
> The AMS CLI must be installed for these commands to work. To install the AMS API or CLI, go to the AMS console **Developers Resources** page. For reference material on the AMS CM API or AMS SKMS API, see the AMS Information Resources section in the User Guide. You may need to add a `--profile` option for authentication; for example, `aws amsskms` *`ams-cli-command`* `--profile SAML`. You may also need to add the `--`

`region` option as all AMS commands run out of us-east-1; for example `aws amscm ams-cli-command --region=us-east-1`.

List your roles:

```
aws --profile saml iam list-roles
```

List your policies:

```
aws --profile saml iam list-role-policies --role-name ROLE_NAME
```

Find stack IDs in AMS

To find a Stack ID, you can use either the Amazon EC2 console, AMS console, or the AMS SKMS API/CLI.

AMS Console:

- In the navigation pane, select **RFCs**, and then click the RFC that created the stack. Use the filter option at the top to reduce the list. The RFC details page opens and includes the run output with the stack ID.
- Alternatively, you can select **Stacks in the navigation pane** to open the stacks list page, and then page through the stack list to the stack you're interested in. This method is more useful if you know the subject of the stack you are looking for.

Amazon EC2 Console:

In the navigation pane, select **Instances** or **Load Balancers** or **Auto Scaling Groups**.

AMS SKMS API ListStackSummaries or CLI:

Note

The AMS CLI must be installed for these commands to work. To install the AMS API or CLI, go to the AMS console **Developers Resources** page. For reference material on the AMS CM API or AMS SKMS API, see the AMS Information Resources section in the User Guide.

You may need to add a `--profile` option for authentication; for example, `aws amsskms` *`ams-cli-command`* `--profile SAML`. You may also need to add the `--region` option as all AMS commands run out of us-east-1; for example `aws amscm` *`ams-cli-command`* `--region=us-east-1`.

Note

The AMS API/CLI (amscm and amsskms) endpoints are in the AWS N. Virginia Region, `us-east-1`. Depending on how your authentication is set, and what AWS Region your account and resources are in, you may need to add `--region us-east-1` when issuing commands. You may also need to add `--profile saml`, if that is your authentication method.

To view a list of stacks in the current account, run the ListStackSummaries operation of the SKMS API (CLI: `list-stack-summaries`). To get complete information about a particular stack instance, by StackId, run GetStack.

- In the following examples, the first command requests a list of summaries for all stack instances in the account. The second command requests the list of stack instances, with a query filter to list only those of a specific stack template, and output the VpcId, Name, and StackId.

```
aws amsskms list-stack-summaries --output table
```

```
-------------------------------------------------------------------
|                                ListStackSummaries               |
|                                 StackSummaries                  |
+------------+----------------------+--------------------+--------+
|   VpcId    |         StackId      |  StackTemplateId   |  Name  |
+------------+----------------------+--------------------+--------+
|vpc-0123abcd|stack-1fb7fe2212345678|stm-sdhopvbb123456789|Test ELB|
|vpc-0123abcd|stack-8323cc0e12345678|stm-s2b72beb123456789|S3 store|
|vpc-0123abcd|stack-2309fa0712345678|stm-sdhopvbb123456789|ELB     |
|vpc-0123abcd|stack-5e61a70512345678|stm-sdpabqbb123456789|PatchSim|
|vpc-0123abcd|stack-bd0e080d12345678|stm-s2b72beb123456789|CLI demo|
+------------+----------------------+--------------------+--------+
```

For information about using CLI queries, see How to Filter the Output with the --query Option and the query language reference, JMESPath Specification..

Find instance IDs or IP addresses in AMS

- To request access to an instance, to log in to an instance, or to create an AMI, you must have the instance ID. For an EC2 instance (either a standalone instance or a part of a stack), or a database instance, you can find the ID in a few different ways:
  - The AMS Console for an instance in an ASG stack: Look on the RFC detail page for the RFC that created the stack. In the Execution Output section, you will find the stack ID for the ASG stack and you can then go to the EC2 Console **Auto Scaling Groups** page and search for that stack ID and find instances for it. When you find the instance, select it and an area opens at the bottom of the page with details, including the IP address.
  - The AMS Console for a standalone EC2 or database (DB) instance: Look on the RFC detail page for the RFC that created the EC2 stack or DB instance. In the Execution Output section, you will find the Instance ID and IP address.
  - AWS EC2 Console:
    1. In the navigation pane, select **Instances**. The **Instances** page opens.
    2. Click the instance that you want the ID for. The instance details page opens and displays the ID and IP address.
  - AWS Database Console:
    1. On the Home page, select **DB Instances**. The **Instances** page opens.
    2. Filter for the DB instance that you want the ID for. The instance details page opens and displays the ID.
  - AMS CLI/API.

    Note

    The AMS CLI must be installed for these commands to work. To install the AMS API or CLI, go to the AMS console **Developers Resources** page. For reference material on the AMS CM API or AMS SKMS API, see the AMS Information Resources section in the User Guide. You may need to add a `--profile` option for authentication; for example, `aws amsskms` *`ams-cli-command`* `--profile SAML`. You may also

need to add the `--region` option as all AMS commands run out of us-east-1; for example `aws amscm ams-cli-command --region=us-east-1`.

Note

The AMS API/CLI (amscm and amsskms) endpoints are in the AWS N. Virginia Region, `us-east-1`. Depending on how your authentication is set, and what AWS Region your account and resources are in, you may need to add `--region us-east-1` when issuing commands. You may also need to add `--profile saml`, if that is your authentication method.

Run the following command to get stack execution output details:

```
aws amsskms get-stack --stack-id STACK_ID
```

The output looks similar to this with the InstanceId appearing near the bottom, under `Outputs` (values shown are examples):

```
{
    "Stack": {
        "StackId": "stack-7fa52bd5eb8240123",
        "Status": {
            "Id": "CreateCompleted",
            "Name": "CreateCompleted"
        },
        "VpcId": "vpc-01234567890abcdef",
        "Description": "Amazon",
        "Parameters": [
            {
                "Value": "sg-01234567890abcdef,sg-01234567890abcdef",
                "Key": "SecurityGroups"
            },
            {
                "Value": "subnet-01234567890abcdef",
                "Key": "InstanceSubnetId"
            },
            {
                "Value": "t2.large",
```

```
                "Key": "InstanceType"
            },
            {
                "Value": "ami-01234567890abcdef",
                "Key": "InstanceAmiId"
            }
        ],
        "Tags": [],
        "Outputs": [
            {
                "Value": "i-0b22a22eec53b9321",
                "Key": "InstanceId"
            },
            {
                "Value": "10.0.5.000",
                "Key": "InstancePrivateIP"
            }
        ],
        "StackTemplateId": "stm-s6xvs000000000000",
        "CreatedTime": "1486584508416",
        "Name": "Amazon"
    }
}
```

Find Amazon Resource Names (ARNs) in AMS

An Amazon Resource Name (ARN) is a string that uniquely identifies an AWS resource, such as EC2 instances, S3 buckets, accounts, Lambda functions, and so forth. AWS requires an ARN when you want to specify a resource unambiguously across all of AWS, such as in IAM policies, Amazon Relational Database Service (Amazon RDS) tags, and API calls. ARNs are constructed from identifiers that specify the service, Region, account, and other information. There are three ARN formats:

```
arn:aws:service:region:account-id:resource-id
arn:aws:service:region:account-id:resource-type/resource-id
arn:aws:service:region:account-id:resource-type:resource-id
```

> **Note**
>
> The exact format of an ARN depends on the service and resource type. To learn more about ARNs, see Amazon Resource Names (ARNs) and AWS Service Namespaces and ARN Formats. For ARN format examples by resource, see the AWS *Service Authorization Reference* resource types table.

Finding the ARN of an AWS object can be difficult. Here are three ways to try:

- AWS service console: Go to the relevant AWS service console, locate the resource and find the ARN in the details for the resource.
- AWS API/CLI (you must first install the AWS CLI): Look for the relevant service in the AWS CLI Command Reference, then, depending on the AWS service, look for the relevant operation, such as `describe`, or `get`, and so forth. For example, for all IAM roles, policies and users, you can get the ARN in the output from the CLI with:

  ```
  aws iam get-role --role-name EMR_DefaultRole
  ```

- Construct the ARN based on the relevant format: Find the ARN format for the resource, by looking at the Actions, resources, and condition keys for AWS services page, finding the relevant service, and then the relevant action, and drilling down to the resource ARN format. Once you have the format, replace the variables with the relevant settings.

You can construct the ARN yourself by following the appropriate format (the formats change per service and resource type) and filling in the information. Here are some ARN examples:

- An AWS account ARN has the following syntax:

  ```
  arn:aws:iam::ACCOUNT-ID:root
  ```

- An S3 ARN has a flat hierarchy of buckets and associated objects:

  ```
  arn:aws:s3:::ams-bucket
  ```

- An EC2 ARN has sub resource-types like image, security groups, instance, and so forth. This example includes the instance ID at the end:

```
arn:aws:ec2:us-east-1:123456789012:instance/i-012abcd34efghi56
```

- A Lambda ARN has the function name for the resource-id part, and you may need to include the version number at the end, as shown in this example:

```
arn:aws:lambda:us-east-1:123456789012:function:api-function:1
```

The AWS Key Management Service service provides this information: Finding the key ID and key ARN.

To find the ARN of a DynamoDB table, use the DynamoDB describe-table CLI.

For an outsider's look at finding AWS ARNs, see AWS ARN Explained: Amazon Resource Name Guide.

Find resources by ARN in AMS

Amazon Resource Names (ARNs) uniquely identify AWS resources. To learn about ARNs and ARN formats, see Amazon Resource Names (ARNs) and AWS Service Namespaces and ARN Formats.

> **Note**
>
> In order to obtain details about a resource from its ARN, *you must have access to the account that created the resource*.

There is no direct path in AWS to look up all resource details from the resource ARN, because services have multiple resource types with various related information. If you have the ARN for a resource, you can determine:

- The related AWS service (the third ARN segment) tells you what AWS console to look at to find the resource
- The resource ID (the sixth or seventh ARN segment) confirms that you've found the right resource

Or you can look for the AWS CLI commands available for that service in the *AWS CLI Command Reference* for information about obtaining details about the resource.

For example, from the following ARN, you can determine that the service is `lambda`, the account is `123456789012`, the resource type is `function`, and the name of the function is `TestFunction`.

```
arn:aws:lambda:us-east-1:123456789012:function:TestFunction
```

From this, you can review the AWS CLI documentation for the Lambda service to learn how more details can be retrieved with various commands, such as `get-function` and `get-function-configuration`.

For example, you can use the following commands to get more information about a Lambda function if you have its name or ARN:

```
aws lambda get-function-configuration --function-name TestFunction
```

```
aws lambda get-function-configuration --function-name arn:aws:lambda:us-
east-1:123456789012:function:TestFunction
```

Find AMS account settings

Account settings that are used to create AMS RFCs, set schedules, and determine who receives notifications.

Some settings are created during onboarding and require a service request to change. You should make a note of these account details because you will use them when communicating with AMS:

- **Credentials**: If you need to retrieve your AMS user name or password, contact your local IT administrator--AMS uses your corporate Active Directory.
- **Cloud Service Delivery Manager (CSDM)**: This person is your liaison with AMS and is available to answer service questions. You are given this person's contact information at onboarding and should keep it available to all in your organization who interact with AMS. You can expect to receive monthly reports on your AMS service from this person.
- **Console access**: You access the AMS console at a URL set up specifically for your account. You can get the URL from your CSDM.
- **AMS CLI**: You can obtain the AMS CLI through the AMS console **Developer's resources** page, or the distributables package that you get from your CSDM. After you have the distributables package, follow the steps outlined in Installing or upgrading the AMS CLI.

- **Maintenance window**: Your maintenance window determines when patching happens for your EC2 instances. The AWS Managed Services Maintenance Window (or Maintenance Window) performs maintenance activities for AWS Managed Services (AMS) and recurs the second Thursday of every month from 3 PM to 4 PM Pacific Time. AMS may change the maintenance window with 48 hours notice. You may have chosen a different window at onboarding--keep a record of your chosen maintenance window.
- **Monitoring**: AMS provides a set of CloudWatch metrics by default, but you can also request additional metrics. If you do, keep record of those.
- **Logs**: By default, your logs are stored at ams-a-*ACCOUNT_ID*-log-management-*REGION* where *REGION* is the region where the log was generated.
- **Mitigation**: At onboarding, AMS records the mitigation action of your choice in case a malware attack against your resources is identified. For example, contact certain people. Keep this information available to all in your organization who interact with AMS.
- **Region**: You can look at the VPC details page in the AMS console. You can also run this command after you have installed the AMS SKMS CLI (this command uses a SAML profile, remove if your authentication method is different):

```
aws --profile saml amsskms get-vpc --vpc-id VPC_ID
```

Important

Note

The AMS API/CLI (amscm and amsskms) endpoints are in the AWS N. Virginia Region, `us-east-1`. Depending on how your authentication is set, and what AWS Region your account and resources are in, you may need to add `--region us-east-1` when issuing commands. You may also need to add `--profile saml`, if that is your authentication method.

Find FQDNs in AMS

AWS Managed Services (AMS) access change types (CTs) require the fully qualified domain name, or FQDN, of your AMS-trusted domain, in the form of `C844273800838.amazonaws.com`. To discover your AWS FQDN:

- AWS Console: Look in the AWS Directory Service console (under Security, Identity, and Compliance category) in the **Directory name** column
- CLI: Use these commands while logged into your domain:

  Windows (returns user and FQDN):

  ```
  whoami /upn
  ```

  or (DC+DC+DC=FQDN)

  ```
  whoami /fqdn
  ```

  Linux:

  ```
  hostname --fqdn
  ```

> **Note**
>
> The AMS API/CLI (amscm and amsskms) endpoints are in the AWS N. Virginia Region, `us-east-1`. Depending on how your authentication is set, and what AWS Region your account and resources are in, you may need to add `--region us-east-1` when issuing commands. You may also need to add `--profile saml`, if that is your authentication method.

Find availability zones (AZs) in AMS

Availability Zone: All accounts have at least two availability zones. To accurately find your availability zone names, you must first know the associated subnet ID.

- AMS Console: In the navigation pane click **VPCs**, and then click the relevant VPC, if necessary. On the VPCs details page, select the relevant subnet in the table of subnets to open the subnet details page with the name of the associated availability zone.
- AMS SKMS API/CLI:

```
aws amsskms list-subnet-summaries --output table
```

```
aws amsskms get-subnet --subnet-id SUBNET_ID
```

> **Note**
>
> The AMS API/CLI (amscm and amsskms) endpoints are in the AWS N. Virginia Region, `us-east-1`. Depending on how your authentication is set, and what AWS Region your account and resources are in, you may need to add `--region us-east-1` when issuing commands. You may also need to add `--profile saml`, if that is your authentication method.

Find SNS topics in AMS

Your SNS topics determine who is notified under various circumstances. AMS provides SNS topics for AMI notifications (see AMS AMI notifications with SNS), CloudWatch alarms and EC2 resources (see Receiving alerts generated by AMS) and more. To discover your existing SNS topics:

- AWS Console: Use the SNS console to view all topics, applications, and subscriptions, and a graph of messages. Also create, delete, subscribe to, and publish to topics.
- API/CLI (when logged into your AMS account, requires the AWS CLI):

  List your SNS topics:

```
aws sns list-topics
```

  List your SNS subscriptions:

```
aws sns list-subscriptions
```

Note

The AMS API/CLI (amscm and amsskms) endpoints are in the AWS N. Virginia Region, `us-east-1`. Depending on how your authentication is set, and what AWS Region your account and resources are in, you may need to add `--region us-east-1` when issuing commands. You may also need to add `--profile saml`, if that is your authentication method.

Find backup settings in AMS

Backups and snapshots are managed by AMS through the native AWS Backup service.

The configuration is managed through AWS Backup plans. You can have multiple AWS Backup plans that associate tagged resources with backup schedules and retention policies. To find your AMS account AWS Backup settings, use the https://console.aws.amazon.com/backup console, or the *AWS CLI Command Reference* for backup commands.

For more information about AMS and AWS Backup, see Continuity Management.

Access management in AMS

Learn how to access resources by using SSH, or remote desktop protocol (RDP), and how to use bastions.

The AWS Managed Services (AMS) access management system is configured during onboarding. Only users with the AMS IAM user role, federated through AMS, can access AMS resources in the account.

In addition to the federated trust, described next, AMS security groups are an important element in private and public application access. For information about AMS security groups and how to change them, see Security groups.

Topics

- What is Access Management?
- How and when to use the root user account in AMS
- AMS Advanced console and Amazon EC2 access
- Accessing the AWS Management console and the AMS console
- Accessing instances using bastions

What is Access Management?

Access management is how AMS protects your resources by allowing only authorized and authenticated access. AMS uses a default IAM user role and instance profile, as well as multi-factor authentication, security groups, DNS-friendly bastion names, and more to keep your resources protected.

AMS focuses on three types of access that require management:

- Console access: Leveraging federation, users in the account's Active Directory can access the console using single sign-on (SSO). If you have multi-factor authentication configured for these accounts, you can continue to require MFA to gain access to the console.
- Instance access with RDP or SSH: Leveraging an Active Directory trust, users in the account's existing Active Directory can request access to an instance, and then successfully authenticate to a bastion and the instance by using their existing corporate credentials. If you have multi-factor

authentication configured for those accounts, you can continue to require MFA to request access to an instance. AMS uses an MFA solution of its own to restrict AMS engineer access to instances.

- Application access: Varies by use case.

Topics

- Why and when AMS accesses your account

Why and when AMS accesses your account

AWS Managed Services (AMS) manages your AWS infrastructure and sometimes, for specific reasons, AMS operators and administrators access your account. These access events are documented in your AWS CloudTrail (CloudTrail) logs.

Why, when, and how AMS accesses your account is explained in the following topics.

AMS customer account access triggers

AMS customer account access activity is driven by triggers. The triggers today are the AWS tickets created in our issues management system in response to Amazon CloudWatch (CloudWatch) alarms and events, and incident reports or service requests that you submit. Multiple service calls and host-level activities might be performed for each access.

Access justification, the triggers, and the initiator of the trigger are listed in the following table.

Access Triggers

Access	Initiator	Trigger
Patching	AMS	Patch issue
Infrastructure deployments	AMS	Deployment issue
Internal problem investigation	AMS	Problem issue (an issue that has been identified as systemic)
Alert investigation and remediation	AMS	AWS Systems Manager operational work items (SSM OpsItems)

Access	Initiator	Trigger
Manual RFC execution	You	Request for Change (RFC) issue. (Non-automated RFCs may require AMS access to your resources)
Incident investigation and remediation	You	Inbound support case (an incident or service request you submit)
Inbound service request fulfillment	You	

AMS customer account access IAM roles

When triggered, AMS accesses customer accounts using AWS Identity and Access Management (IAM) roles. Like all activity in your account, the roles and their usage are logged in CloudTrail.

> **Important**
>
> Do not modify or delete these roles.

IAM roles for AMS access to customer accounts

Role Name	Account Type (SALZ, MALZ Management, MALZ Application, etc.)	Description
ams-service-admin	SALZ, MALZ	AMS Service automation access and automated infrastructure deployments e.g Patch, Backup, Automated Remediation.
ams-application-infra-read-only	SALZ, MALZ Application, MALZ Tools-Application	Operator read only access
ams-application-infra-operations		Operator access for incidents/service requests

Role Name	Account Type (SALZ, MALZ Management, MALZ Application, etc.)	Description
ams-application-infra-admin		AD Admin access
ams-primary-read-only	MALZ Management	Operator read only access
ams-primary-operations		Operator access for incidents/service requests
ams-primary-admin		AD Admin access
ams-logging-read-only	MALZ Logging	Operator read only access
ams-logging-operations		Operator access for incidents/service requests
ams-logging-admin		AD Admin access
ams-networking-read-only	MALZ Networking	Operator read only access
ams-networking-operations		Operator access for incidents/service requests
ams-networking-admin		AD Admin access
ams-shared-services-read-only	MALZ Shared Services	Operator read only access
ams-shared-services-operations		Operator access for incidents/service requests
ams-shared-services-admin		AD Admin access
ams-security-read-only	MALZ Security	Operator read only access
ams-security-operations		Operator access for incidents/service requests

Role Name	Account Type (SALZ, MALZ Management, MALZ Application, etc.)	Description
ams-security-admin		AD Admin access
ams-access-security-analyst	SALZ, MALZ Application, MALZ Tools-Application, MALZ Core	AMS Security access
ams-access-security-analyst-read-only		AMS Security, read only access
Sentinel_AdminUser_Role_PXHazRQadu0PVcCDcMbHE	SALZ	[BreakGlassRole]Used to breakGlass into the customer accounts
Sentinel_PowerUser_Role_wZuPuS0ROOl0IazDbRI9	SALZ, MALZ	Poweruser access to customer accounts for RFC execution
Sentinel_ReadOnlyUser_Role_Pd4L6Rw9RD0lnLkD5JOo		ReadOnly access to customer accounts for RFC execution
ams_admin_role		Admin access to customer accounts for RFC execution
AWSManagedServices_Provisioning_CustomerStacksRole		Used to launch and update CFN stacks on behalf of customers through CloudFormation Ingest
customer_ssm_automation_role		Role passed by CT executions to SSM Automation for runbook execution
ams_ssm_automation_role	SALZ, MALZ Application, MALZ Core	Role passed by AMS services to SSM Automation for runbook execution

Role Name	Account Type (SALZ, MALZ Management, MALZ Application, etc.)	Description
ams_ssm_iam_deployment_role	MALZ Application	Role used by IAM catalog
ams_ssm_shared_svcs_intermediary_role	MALZ Shared Services	Role used by application ams_ssm_automation_role to execute specific SSM Documents in Shared Services account
AmsOpsCenterRole	SALZ, MALZ	Used to create and update OpsItems in customer accounts
AMSOpsItemAutoExecutionRole		Used to get SSM Documents, describe resource tags, update OpsItems, and start automation
customer-mc-ec2-instance-profile		Default customer EC2 instance profile (role)

Requesting instance access

To access a resource, you must first submit a request for change (RFC) for that access. There are two types of access that you can request: admin (read/write permissions) and read-only (standard user access). Access lasts for eight hours, by default. This information is required:

- Stack ID, or set of stack IDs, for the instance or instances you want to access.
- The fully qualified domain name of your AMS-trusted domain.
- The Active Directory username of the person who wants access.
- The ID of the VPC where the stacks are that you want access to.

Once you've been granted access, you can update the request as needed.

For examples of how to request access, see Stack Admin Access | Grant or Stack Read-only Access | Grant.

How and when to use the root user account in AMS

The root user is the superuser within your AWS account. AMS monitors root usage. We recommend that you use root only for the few tasks that require it, for example: changing your account settings, activating AWS Identity and Access Management (IAM) access to billing and cost management, changing your root password, and enabling multi-factor authentication (MFA). See Tasks that require root user credentials.

> **Note**
>
> MFA is enabled during AMS Advanced onboarding to specifically disallow root user access. Root access in AMS-managed accounts is different from other AWS accounts, and is critical to the security of your entire AMS-managed environment. The MFA configured is a virtual MFA and is performed using an AMS-owned device. After the virtual MFA is configured with AMS' assistance, the virtual token is immediately deleted. This ensures that neither you nor AMS retains the ability to log in to the account as the root user. Root login can only be re-enabled on special requests (explained next) and AMS expects such accesses to be used only when absolutely necessary. For information about MFA, see Secure New Account with Multi-Factor Authentication.

Root access always triggers an AMS Security and Operations team response. AMS monitors API calls for root access, and alarms are triggered if such access is detected.

Requesting root access is slightly different between AMS account types.

Root access with AMS Advanced single-account landing zone:

If you have a single-account landing zone, contact your cloud service deliver manager (CSDM) and cloud architects (CAs) to advise them of the root access work that you require. It is best to give twenty-four hours notice before the proposed activity.

Root access with AMS Advanced multi-account landing zone:

For multi-account landing zone Application, Shared Services, Security, or Networking accounts, use the Management | Other | Other (ct-1e1xtak34nx76) change type. Include the date, time, and

the purpose of using the root user credentials and schedule the RFC to be sure to give twenty-four hours notice before the proposed activity. Use your multi-account landing zone Management account to submit the RFC.

Additionally, contact your CSDM and CAs twenty-four hours in advance, to advise them of the root access work you require.

AMS operations and security response to root usage:

AMS receives an alarm when the root user account is used. If the root credentials usage is unscheduled, they contact the AMS Security team, and your account team, to verify if this is expected activity. If it is not expected activity, AMS works with your Security team to investigate the issue.

AMS Advanced console and Amazon EC2 access

Accessing the AMS Advanced console.

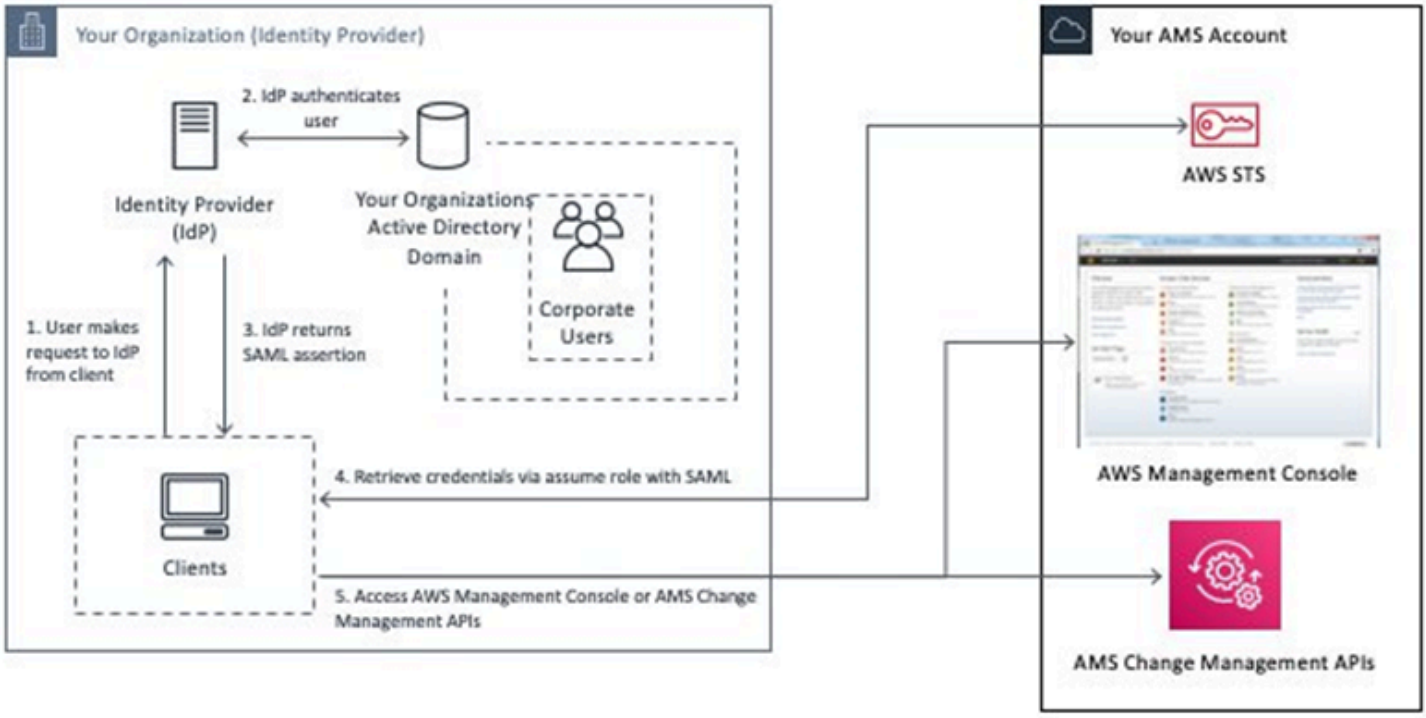

Accessing your Amazon EC2 instances.

Submit access request:

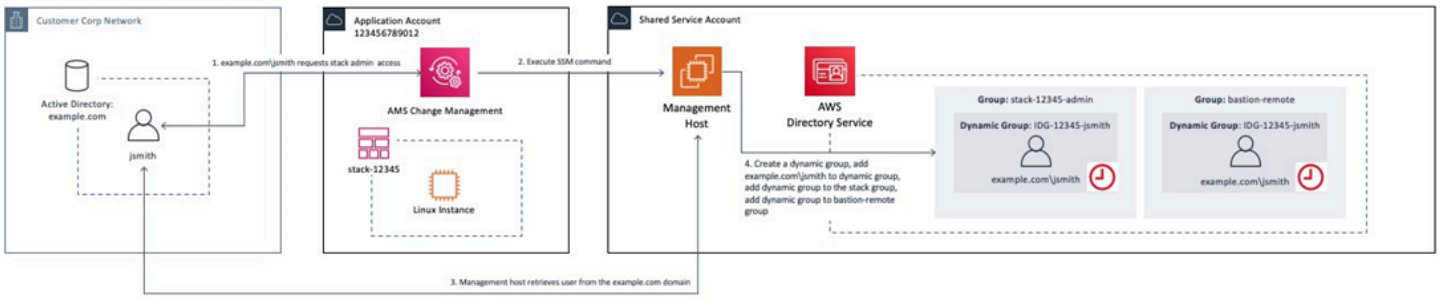

Logging into your AMS Amazon EC2 instances:

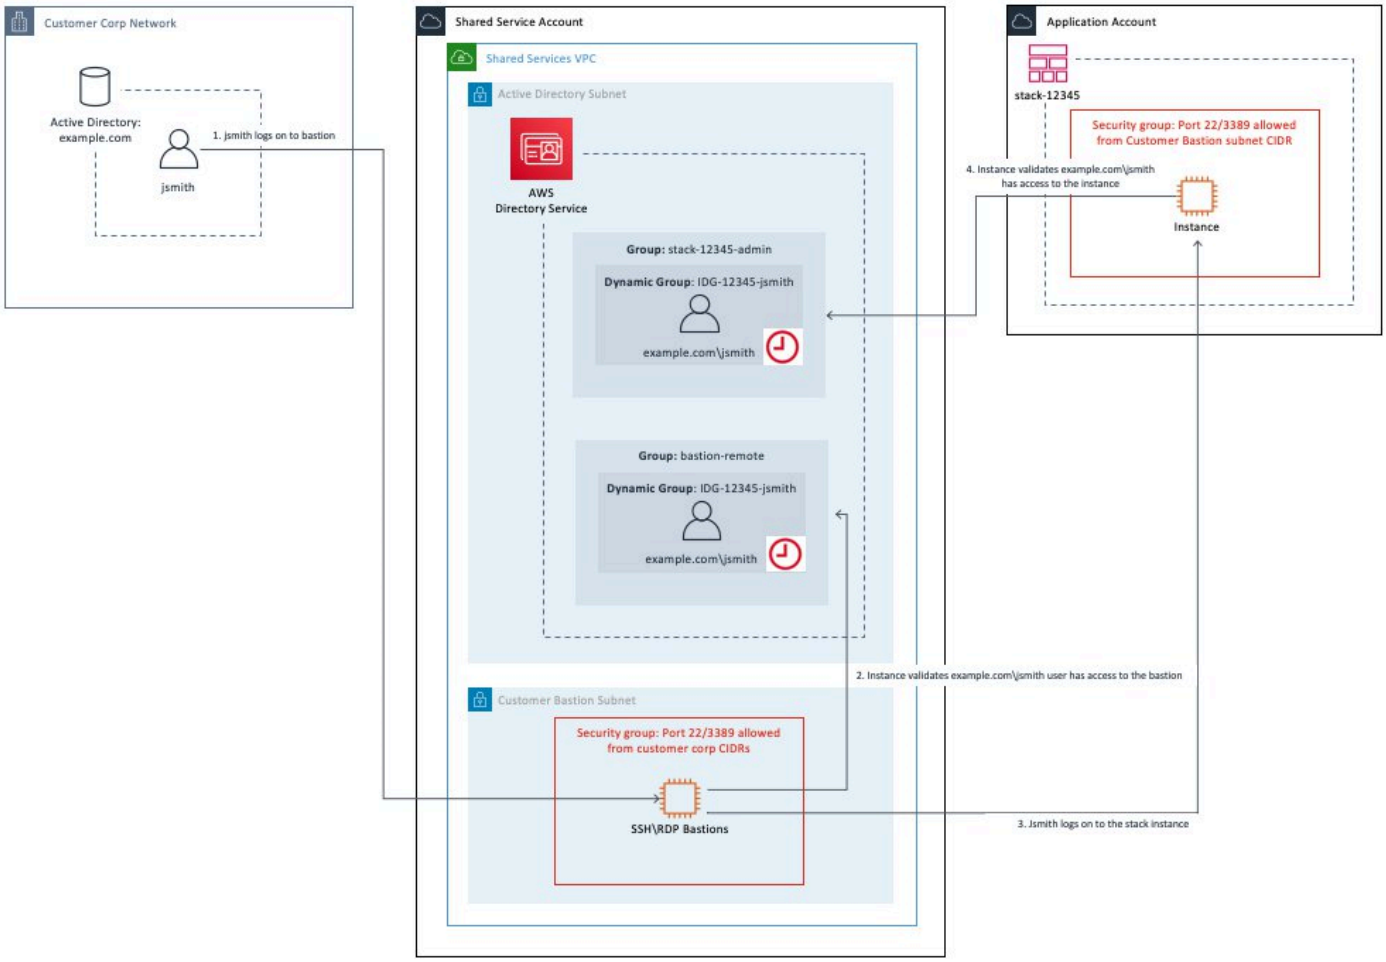

Accessing the AWS Management console and the AMS console

During onboarding, you're provided a login to the AWS Management console (with limited privileges: you can write to the AMS console, and some fields in your customer information page). You can access the AMS console by selecting the **Managed Services** link in the AWS Management console. Either federated access or shared credentials (user name/password) are prepared as

agreed with your IT administration team. For further account or group creation, submit a service request to AMS.

For information about getting access to the AWS Management console, see Working with the AWS Management console.

For some tips on using the AMS console, see Using the AMS console.

Temporary AMS console access

If you haven't yet set up an identity provider (for instance, SAML) to authenticate to AMS, you can get temporary access to the AMS console. Contact your CSDM to have a Deployment | Advanced stack components | Identity and Access Management (IAM) | Create entity or policy change request (ct-3dpd8mdd9jn1r) submitted on your behalf with these values:

- UserName: A name for the IAM user entity that you're creating
- AccessType: "Console access"
- UserPermissions: "Temporary AMS console access for *USERNAME* (the person that you want to have temporary access)"
- Email notifications: Your email address, so you can approve the request when AMS requests you to

> **Note**
>
> This RFC for temporary AMS Console access requires a security review and acceptance by both your internal security team and AMS Global Security.

After this request has been completed, and you're able to log in, you're required to approve the RFC that was created, to track the approval and allow the AMS team to close out the work. To approve the RFC, find it in the RFC's list page (there will be a Pending Approval flag next to it), select it to open the RFC details page for that RFC, and then choose **Approve**. Note that you won't be able to use AMS until the RFC is approved.

When the RFC successfully completes, AMS operations provides you with the new IAM user and a password. Then follow these steps:

1. Go to the AWS Management console and log in with provided credentials. You'll be asked to create a new password. You must also, upon login, set up multi-factor authentication (MFA); to learn more about doing that, see Using Multi-Factor Authentication (MFA) in AWS.
2. In the AWS Management console, change to the provided IAM role (`customer_CustomerCode_readonly_user_role`).
3. Open the AMS **Managed Services** Console.

> **Note**
>
> Temporary access defaults to sixty days; however, you can request a thirty-day extension by contacting your CSDM.

Accessing instances using bastions

All access to resources inside AMS-managed accounts, for both customers and AMS operators, is gated by the use of bastion hosts. We maintain both Linux and Windows RDP bastions for access for both Multi-account landing zone (MALZ) and Single-account landing zone (SALZ) AMS Advanced accounts.

Your bastions are accessible only over your private connection (VPN or AWS Direct Connect)DX. In addition to firewalling to prevent inbound traffic, bastions are regularly re-provisioned (with existing credentials) on a fixed schedule.

> **Note**
>
> For information on moving files to an EC2 instance, see File transfer: Local Windows or MAC PC to Linux Amazon EC2.

MALZ

You access your account instances by logging in to a bastion instance with your Active Directory (AD) credentials. Amazon uses bastions located in the perimeter network VPC (networking account), and you use your customer bastions, located in your Customer Bastions subnet in the shared services account.

When your AMS environment is initially onboarded, you have two SSH bastions and two RDP bastions depending on your choice.

SALZ

You access your account instances by logging in to a bastion instance with your Active Directory (AD) credentials. AMS uses bastions located in the perimeter network subnets, and you use bastions located in your private subnets.

When your account is initially onboarded, you have two RDP and two SSH bastions, by default.

> **Note**
>
> As part of the single-account landing zone, AMS provides both RDP (Windows) and SSH (Linux) bastions to access your stacks; however, you can choose whether you want only RDP bastions or only SSH bastions. To request that only RDP, or only SSH bastions are maintained, submit a service request.

In order to access an instance, you need:

- Access granted to the stack. To get access granted to a stack, see Stack Admin Access | Grant or Stack Read-Only Access | Grant.
- The stack ID that you want to access so you can be granted access to the instance. To find a stack ID, see Find stack IDs in AMS.
- The instance IP that you want to access. To find an instance IP, see Find instance IDs or IP addresses in AMS.
- The DNS friendly bastion name or the bastion IP. How to use DNS friendly bastion names and how to find a bastion IP are described next.

DNS friendly bastion names

AWS Managed Services (AMS) uses DNS friendly bastion names.

MALZ

For Multi-account landing zone (MALZ), DNS records are created for the bastions in the FQDN of the AMS-managed Active Directory. AMS replaces Linux and Windows bastions as required.

For example, if there is a new bastion AMI that must be deployed, the bastion DNS records dynamically update to point to new, valid bastions.

1. To access SSH (Linux) bastions, use DNS records like this: `sshbastion(1-4).Your_Domain.com`

   For example, where the domain is `Your_Domain`:

   - `sshbastion1.Your_Domain.com`
   - `sshbastion2.Your_Domain.com`
   - `sshbastion3.Your_Domain.com`
   - `sshbastion4.Your_Domain.com`

2. To access RDP (Windows) bastions, use DNS records like this: `rdp-Username.Your_Domain.com`.

   For example, where the user name is `alex`, `test`, `demo`, or `bob`, and the domain is `Your_Domain.com`:

   - `rdp-alex.Your_Domain.com`
   - `rdp-test.Your_Domain.com`
   - `rdp-demo.Your_Domain.com`
   - `rdp-bob.Your_Domain.com`

SALZ

Single-account landing zone (SALZ) replaces Linux and Windows bastions as required. For example, if there is a new bastion AMI that must be deployed, the bastion DNS records dynamically update to point to new, valid bastions.

1. To access SSH (Linux) bastions, use DNS records like this: `sshbastion(1-4).AAccountNumber.amazonaws.com.`

   For example, where `123456789012` is the account number:

   - `sshbastion1.A123456789012.amazonaws.com`
   - `sshbastion2.A123456789012.amazonaws.com`
   - `sshbastion3.A123456789012.amazonaws.com`

- `sshbastion4.A123456789012.amazonaws.com`

2. To access RDP (Windows) bastions, use DNS records like this: `rdpbastion`*`(1-4)`*`.A`*`ACCOUNT_NUMBER`*`.amazonaws.com`.

   For example, where `123456789012` is the account number:

   - `rdpbastion1.A123456789012.amazonaws.com`
   - `rdpbastion2.A123456789012.amazonaws.com`
   - `rdpbastion3.A123456789012.amazonaws.com`
   - `rdpbastion4.A123456789012.amazonaws.com`

Saving costs on Single-account landing zone (SALZ) bastions

AMS provides two SSH bastions and two RDP bastions in the default configuration for you to connect to your Amazon EC2 instances, and also deploys two DMZ bastions in the default configuration for service operations. The bastions use m4. large Amazon EC2 instances by default. You have an option to change the Amazon EC2 instances used for bastions to t3.small, and save cost.

If you are using on-demand instances, or spot instances, or a savings plan, you should consider this feature, and save costs. If you use Reserved Instances consider if using t3.small instances might lower your costs. To change the instance type, submit an RFC with Management | Advanced stack components | EC2 instance stack | Resize (ct-15mazjj88xc69) CT from your AMS account.

Contact your cloud service delivery manager (CSDM) for additional questions, or to check if you can benefit from this feature.

Using bastion IP addresses

AMS customers can use SSH and RDP bastions, either the DNS friendly bastion names described previously, or bastion IP addresses.

To find bastion IP addresses, SSH and RDP, for your account:

1. For multi-account landing zone only: Log in to the Shared Services account.
2. Open the EC2 Console and choose **Running Instances**.

   The **Instances** page opens.

3. In the filter box at the top, enter either **ssh-bastion** or **rdp-bastion**.

   In the filter box at the top, enter either **customer-ssh** or **customer-rdp**.

   The SSH and/or RDP bastions for your account display.

   Note that in addition to your SSH bastions, you may see AMS perimeter network bastions in the list, which are unavailable for this.

4. Select an SSH or RDP bastion. If you're using a Windows computer and want to log in to a Linux instance, you use an SSH bastion. If you want to log in to a Windows instance, you use an RDP bastion. If you're on a Linux OS and want to log in to a Windows instance, you use an SSH bastion through an RDP tunnel (this is so you can access the Windows desktop). To access a Linux instance from a Linux OS, you use an SSH bastion.

Instance access examples in AMS

These examples show how to log in to an instance in your AMS account by using a bastion after you've been granted access through an RFC. For information about getting access granted, see Requesting instance access.

> **Note**
>
> For information on moving files to an EC2 instance, see File transfer: Local Windows or MAC PC to Linux Amazon EC2.

Required data:

- **Bastion DNS friendly name or IP address**: Use a DNS friendly name as described in DNS friendly bastion names or find bastion IP addresses as described in Using bastion IP addresses.

  > **Note**
  >
  > An Amazon EC2 instance created through an Amazon EC2 Auto Scaling group will have an IP address that cycles in and out and you have to use your Amazon EC2 console to find that IP address.

- **User name** (for example *DOMAIN_FQDN**USERNAME*) and **Password**: Credentials for the account. The *USERNAME* must be your Active Directory user name.

  Note that a user name in the format username@customerdomain.com can be used but can cause trouble with your PBIS setup.

- **Stack IP address**: Find this by looking at the run output for the RFC that you submitted to launch the stack, or look up the Amazon EC2 instance IP address in the Amazon EC2 console. For a single Amazon EC2 instance, you can also use the AMS SKMS command ListStackSummaries to find the stack ID and then GetStack to find the stack IP address. For the AMS SKMS API reference, see the **Reports** tab in the AWS Artifact Console.

Access the bastion IP address, either SSH or RDP, as appropriate, and log in using one of the following procedures.

> **Note**
>
> RDP bastions only allow two simultaneous connections. So, in the best case scenario, only 4 admins are able to connect to windows stacks at the same time. If you require more connections for RDP, see AMS Bastion Options during Application Migrations/Onboarding in the *AMS onboarding guide*.

Linux computer to Linux instance

Use SSH to connect to the SSH bastion and then to the Linux instance.

MALZ

For more information about the friendly bastion names, see DNS bastions.

In order to connect to the Linux instance, you must first connect to an SSH bastion.

1. Open a shell window and enter:

```
ssh Domain_FQDN\\Username@SSH_bastion_name
    or SSH_bastion_IP
```

Which would look like this if your Domain_FQDN is "corp.domain.com", your account number is "123456789123", Your_Domain is "amazonaws.com", you choose bastion "4", and your user name is "JoeSmith":

```
ssh corp.domain.com\\JoeSmith sshbastion4.A123456789123.amazonaws.com
```

2. Log in with your corporate Active Directory credentials.
3. When presented with a Bash prompt, SSH in to the instance, and then enter:

```
ssh Domain_FQDN\\Username@Instance_IP
```

Or, you can use the Login flag (-l):

```
ssh -l Domain_FQDN\\Username@Instance_IP
```

SALZ

For more information about the friendly bastion names, see DNS bastions.

In order to connect to the Linux instance, you must first connect to an SSH bastion.

1. Open a shell window and enter:

```
ssh DOMAIN_FQDN\\USERNAME@SSH_BASTION_name
    or SSH_BASTION_IP
```

Which would look like this if your account number is 123456789123, you choose bastion 4, and your user name is JoeSmith:

```
ssh corp.domain.com\\JoeSmith sshbastion1.A123456789123.amazonaws.com
```

2. Log in with your corporate Active Directory credentials.
3. When presented with a Bash prompt, SSH in to the instance, and then enter:

```
ssh DOMAIN_FQDN\\USERNAME@INSTANCE_IP
```

Or, you can use the Login flag (-l):

```
ssh -l DOMAIN_FQDN\\USERNAME@INSTANCE_IP
```

Linux computer to Windows instance

Use an SSH tunnel and an RDP client to connect to a Windows instance from your Linux computer.

MALZ

This procedure requires a Remote Desktop Connection client for Linux; the example uses Microsoft Remote Desktop (an open source UNIX client for connecting to Windows Remote Desktop Services). Rdesktop is an alternative.

> **Note**
>
> How you log in to Windows instances might change based on the remote desktop client being used.

First you establish an SSH tunnel, and then log in.

For more information about the friendly bastion names, see DNS friendly bastion names.

Before you begin:

- Request access to the instance that you want to connect to; for information, see Access requests.
- Choose a friendly DNS SSH bastion name to connect to; for example:

  ```
  sshbastion(1-4).Your_Domain
  ```

  Which would look like this if your Domain_FQDN is "corp.domain.com", your AMS-managed Your_Domain is "amazonaws.com", you choose bastion "4", and your user name is "JoeSmith":

  ```
  ssh corp.domain.com\\JoeSmith sshbastion4.amazonaws.com
  ```

- Find the IP address of the instance that you want to connect to; for information, see Finding an instance ID or IP address.

1. Set up RDP over an SSH tunnel from a Linux desktop to a Windows instance. In order to issue the `ssh` command with the right values, there are a couple of ways to proceed:

   - In the Linux shell, set the variables, and then enter the SSH connection command:

     ```
     BASTION="sshbastion(1-4).Your_Domain""
     WINDOWS="Windows_Instance_Private_IP"
     AD="AD_Account_Number"
     USER="AD_Username"
     ssh -L 3389:$WINDOWS:3389 A$AD\\\\$USER@$BASTION
     ```

     Example, if the following values are used:

     ```
     BASTION="sshbastion4.A123456789123.amazonaws.com"

     WINDOWS="172.16.3.254"

     AD="ACORP_example"

     USER="john.doe"
     ```

   - Add the variable values directly to the `ssh` command.

   In either case, this is what the rendered request would be (assuming the same set of variable values):

   ```
   ssh -L 3389:172.16.3.254:3389 ACORP_example\\\\john.doe@myamsadomain.com
   ```

2. Either: Open your Remote Desktop Client, enter the loopback address and port, 127.0.0.1:3389, and then open the connection.

   Or, log in to the Windows instance from a new Linux desktop shell. If you use RDesktop, the command looks like this:

   ```
   rdesktop 127.0.0.1:3389
   ```

   A remote desktop window for the Windows instance appears on your Linux desktop.

Tip

If the remote desktop session fails to start, verify that network connectivity to the Windows instance from the SSH bastion is allowed on port 3389 from the shell in step 1 (replace `private_ip_address_of_windows_instance` appropriately):

```
nc private_ip_address_of_windows_instance 3389 -v -z
```

Success:

```
nc 172.16.0.83 3389 -v -z
        Connection to 172.16.0.83 3389 port [tcp/ms-wbt-server] succeeded
        netstat -anvp | grep 3389
        tcp    0     0 172.16.0.253:48079 172.16.3.254:3389 ESTABLISHED
```

SALZ

This procedure for a single-account landing zone requires a Remote Desktop Connection client for Linux; the example uses Microsoft Remote Desktop (an open source UNIX client for connecting to Windows Remote Desktop Services). Rdesktop is an alternative.

Note

How you log in to Windows instances might change based on the remote desktop client being used.

First you establish an SSH tunnel, and then log in.

For more information about the friendly bastion names, see DNS friendly bastion names.

Before you begin:

- Request access to the instance that you want to connect to; for information, see Access requests.
- Choose a friendly DNS SSH bastion name to connect to; for example:

```
sshbastion(1-4).AAMSAccountNumber.amazonaws.com
```

Which would look like this if your account number is 123456789123 and you choose bastion 4:

```
sshbastion4.A123456789123.amazonaws.com
```

- Find the IP address of the instance that you want to connect to; for information, see Finding an instance ID or IP address.

1. Set up RDP over an SSH tunnel from a Linux desktop to a Windows instance. In order to issue the `ssh` command with the right values, there are a couple of ways to proceed:

   - In the Linux shell, set the variables, and then enter the SSH connection command:

     ```
     BASTION="sshbastion(1-4).AAMSAccountNumber.amazonaws.com"
     WINDOWS="WINDOWS_INSTANCE_PRIVATE_IP"
     AD="AD_ACCOUNT_NUMBER"
     USER="AD_USERNAME"
     ssh -L 3389:$WINDOWS:3389 A$AD\\\\$USER@$BASTION
     ```

     Example, if the following values are used:

     ```
     BASTION="sshbastion4.A123456789123.amazonaws.com"

     WINDOWS="172.16.3.254"

     AD="ACORP_example"

     USER="john.doe"
     ```

   - Add the variable values directly to the `ssh` command.

   In either case, this is what the rendered request would be (assuming the same set of variable values):

   ```
   ssh -L 3389:172.16.3.254:3389 ACORP_example\\\
   \john.doe@sshbastion4.A123456789123.amazonaws.com
   ```

2. Either: Open your Remote Desktop Client, enter the loopback address and port, 127.0.0.1:3389, and then open the connection.

   Or, log in to the Windows instance from a new Linux desktop shell. If you use RDesktop, the command looks like this:

   ```
   rdesktop 127.0.0.1:3389
   ```

   A remote desktop window for the Windows instance appears on your Linux desktop.

   Tip

   If the remote desktop session fails to start, verify that network connectivity to the Windows instance from the SSH bastion is allowed on port 3389 from the shell in step 1 (replace `private_ip_address_of_windows_instance` appropriately):

   ```
   nc private_ip_address_of_windows_instance 3389 -v -z
   ```

   Success:

   ```
   nc 172.16.0.83 3389 -v -z
           Connection to 172.16.0.83 3389 port [tcp/ms-wbt-server] succeeded
           netstat -anvp | grep 3389
           tcp    0     0 172.16.0.253:48079 172.16.3.254:3389 ESTABLISHED
   ```

Windows computer to Windows instance

Use Windows Remote Desktop Connection client to connect to a Windows instance from your Windows computer.

MALZ

For more information about the friendly bastion names, see DNS friendly bastion names.

1. Open the Remote Desktop Connection program, a standard Windows program, and enter the friendly DNS name of the Windows bastion in the hostname field.

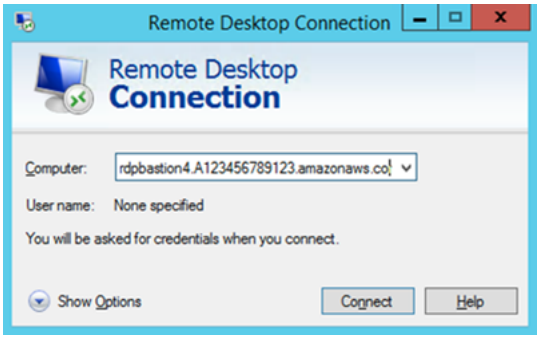

2. Choose **Connect**. The Remote Desktop Connection attempts an RDP connection to the bastion.

   If successful, a credentials dialog box opens. To gain access, use your corporate Active Directory credentials, as you would with the Windows instance.

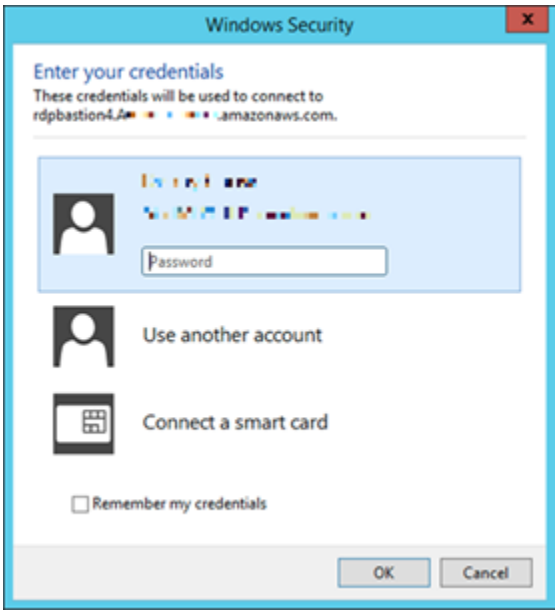

3. Open the Remote Desktop Connection program on the bastion and enter the IP address of the Windows instance you would like to connect to (for example, 10.0.0.100), and then choose **Connect**. Your corporate Active Directory credentials are again required before you connect to the Windows instance.

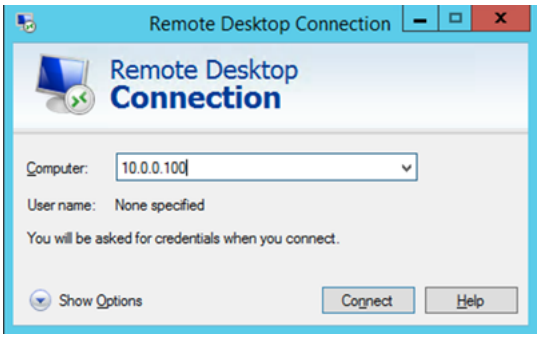

SALZ

For more information about the friendly bastion names, see DNS friendly bastion names.

1. Open the Remote Desktop Connection program, a standard Windows program, and enter the friendly DNS name of the Windows bastion in the hostname field; for example, `rdpbastion(1-4).AAMSAccountNumber.amazonaws.com`, which would look like this if your account number is 123456789123 and you choose bastion 4, `rdpbastion4.A123456789123.amazonaws.com`.

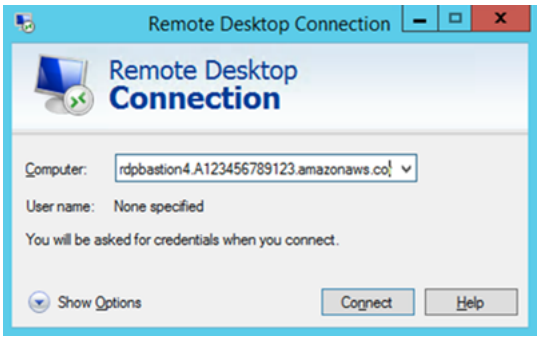

2. Choose **Connect**. The Remote Desktop Connection attempts an RDP connection to the bastion.

   If successful, a credentials dialog box opens. To gain access, use your corporate Active Directory credentials, as you would with the Windows instance.

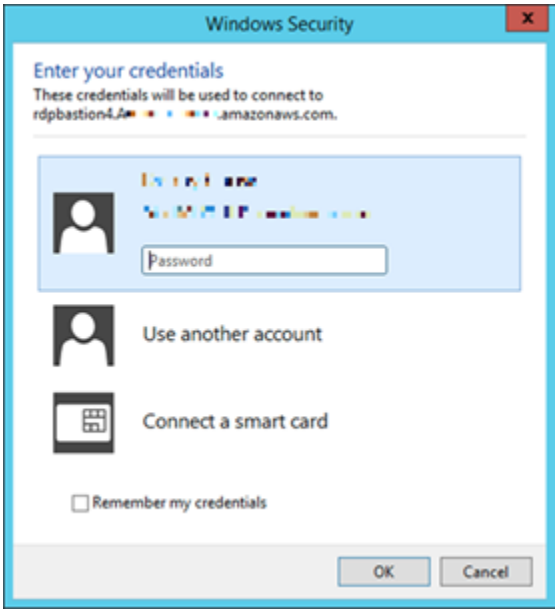

3. Open the Remote Desktop Connection program on the bastion and enter the IP address of the Windows instance you would like to connect to (for example, 10.0.0.100), and then

choose **Connect**. Your corporate Active Directory credentials are again required before you connect to the Windows instance.

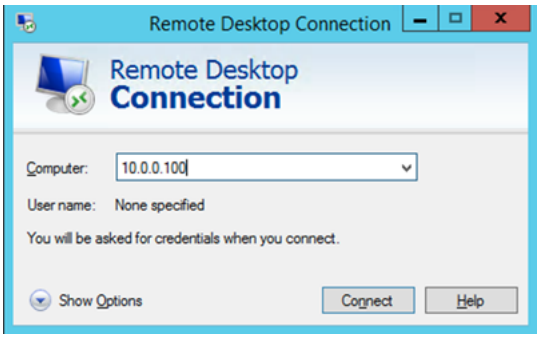

Windows computer to Linux instance

To RDP to an SSH bastion from a Windows environment, follow these steps.

MALZ

Before you begin:

- Request access to the instance that you want to connect to; for information, see Access requests.
- Choose a friendly DNS SSH bastion name to connect to; for example:

```
sshbastion(1-4).YOUR_DOMAIN
```

Which would look like this if YOUR_DOMAIN is myamsaddomain.com" and you choose bastion 4:

```
sshbastion4.myamsaddomain.com
```

- Find the IP address of the instance that you want to connect to; for information, see Finding an instance ID or IP address.

In order to connect to the Linux instance from your Windows machine, you must first connect to an SSH bastion.

Use the native Windows OpenSSH client or install PuTTY on your local machine. To learn more about OpenSSH, see OpenSSH in Windows.

1. Use the native Windows or open PuTTY and enter the SSH bastion hostname or the IP address of the SSH bastion. For example, 10.65.2.214 (22 is the port used for SSH; it will be set by default).
2. OpenSSH or PuTTY attempts an SSH connection to the bastion and open a shell window.
3. Use your corporate Active Directory credentials as you would with the RDP hosts to gain access.
4. When presented with a Bash prompt, SSH into the instance. Enter:

```
ssh DOMAIN_FQDN\USERNAME@INSTANCE_IP
```

SALZ

Before you begin:

- Request access to the instance that you want to connect to; for information, see Access requests.
- Choose a friendly DNS SSH bastion name to connect to; for example:

```
sshbastion(1-4).AAMSAccountNumber.amazonaws.com
```

Which would look like this if your account number is 123456789123 and you choose bastion 4:

```
sshbastion4.A123456789123.amazonaws.com
```

- Find the IP address of the instance that you want to connect to; for information, see Finding an instance ID or IP address.

In order to connect to the Linux instance from your Windows machine, you must first connect to an SSH bastion.

Use the native Windows OpenSSH client or install PuTTY on your local machine. To learn more about OpenSSH, see OpenSSH in Windows.

1. Use the native Windows or open PuTTY and enter the SSH bastion hostname or the IP address of the SSH bastion. For example, 10.65.2.214 (22 is the port used for SSH; it will be set by default).

2. OpenSSH or PuTTY attempts an SSH connection to the bastion and open a shell window.
3. Use your corporate Active Directory credentials as you would with the RDP hosts to gain access.
4. When presented with a Bash prompt, SSH into the instance. Enter:

```
ssh DOMAIN_FQDN\USERNAME@INSTANCE_IP
```

Team, or role, based access control in an AMS account

Scenario: Two application teams A, and B, use a single AMS account for their apps "AA", and "BB", respectively. Team A wants access only to resources for app "AA", and team B wants access only to resources for app "BB". How do I set that up?

Use their ITSM's tools to implement team-based access controls (TBAC). For example, you could use the AMS ServiceNow Connector App for integration with AMS APIs. Contact your CSDM for high level guidance of this implementation.

Automated instance configuration in AMS Advanced

The AMS Advanced automated instance configuration service runs daily and automatically scans and updates the SSM and CloudWatch agents and configuration files on your managed EC2 instances. The updates apply, as needed to:

- SSM and CloudWatch agents
- CloudWatch configuration files

These updates allow AMS to access your AMS-managed EC2 instances, and to configure your instances to emit appropriate logs and metrics.

Topics

- Prerequisites for automated instance configuration
- SSM Agent automatic installation
- Automated changes

Prerequisites for automated instance configuration

For AMS Advanced customers who deploy instances with Change Management, the following prerequisites must be met:

- The SSM Agent is installed, and in a managed state.
- The instance is tagged as a managed instance. (The `aws:cloudformation:stack-name` tag has a value starting with `stack-` or `sc-`.)

If the SSM Agent is not already installed on your instance, you can install it using the AMS SSM Agent auto installation feature. For more information, see SSM Agent automatic installation.

Or, you can install the SSM Agent manually. For more information, see the following:

- Linux: Manually install SSM Agent on EC2 instances for Linux - AWS Systems Manager
- Windows: Manually install SSM Agent on EC2 instances for Windows Server - AWS Systems Manager

For more information on SSM agent, see the AWS documentation Working with SSM Agent.

SSM Agent automatic installation

For AMS to manage your Amazon Elastic Compute Cloud (Amazon EC2) instances, install AWS Systems Manager SSM Agent on each instance. If your instances don't have SSM Agent installed, then you can use AMS SSM Agent auto-installation feature.

> **Note**
>
> This feature is only available for non-ASG (Auto Scaling group) EC2 instances running Linux operating systems supported by AMS. AMS SSM Agent auto-installation feature is disabled by default. To enable it, reach out to your CA/CSDM.

How to request auto installation of SSM Agent on your instances

Configure a patch maintenance window (MW) for the instances. A working SSM Agent is required to complete the patch process. If SSM Agent is missing on an instance, then AMS tries to automatically install it during the patch MW.

> **Note**
>
> AMS stops and starts each instance as part of the auto-installation process. When an instance is stopped, data stored in instance store volumes and data stored on the RAM is lost. For more information, see What happens when you stop an instance.

How SSM Agent auto installation works

AMS uses EC2 user data to run the installation script on your instances. To add the user data script and run it on your instances, AMS must stop and start each instance.

If your instance already has an existing user data script, then AMS completes the following steps during the auto installation process:

1. Creates a backup of the existing user data script.
2. Replaces the existing user data script with the SSM Agent installation script.

3. Restarts the instance to install SSM Agent.
4. Stops the instance and restores the original script.
5. Restarts the instance with the original script.

Automated changes

The AMS Advanced automated instance configuration service makes changes, as needed, to your EC2 instances.

What Changes:

- Automatically update code on Linux instances
- Automatically update PBIS on Linux instances
- Automatically update the minimum version of SSM and CloudWatch agents
- CloudWatch configuration files, update details
- Automatically configured logs

Automatically update code on Linux instances

AMS automatically updates on instance code on Linux instances. This helps to improve operational stability and security of the AMS components and environment altogether.

FAQ:

What's included in the On Instance Code (OIC) on Linux?

> OIC includes ams-toolkit package along with some configuration files and cron jobs. AMS require these files and packages for integration (Active Directory, CloudFormation and other dependencies). We pre-bake these files into AMS-provided AMIs or install onto your instance during workload ingestion.

When will AMS update OIC?

> AMS update OIC when we release a new version with bug fixes or other improvements. The workflow to check the OIC version and update runs daily.

Automatically update PBIS on Linux instances

AMS uses the Power Broker Identity Service (PBIS) module to join Linux instances into AMS-managed Active Directory.

AMS automatically updates PBIS on Linux instances.

FAQ:

When will AMS update PBIS?

AMS turns on PBIS update at reboot. If there is a new PBIS version available, then AMS attempts to install the new version during the next instance reboot.

Can PBIS update be turned off?

You can turn off PBIS update at the instance or account levels:

- **Account level:** Create a parameter in the SSM parameter store: Name: `/ams/skip-pbis-update`, Value: `true` (any case).

> **Note**
>
> The instance profile must have permissions to read SSM parameters. If the flag is missing, then the default behavior is to run the update.

- **Instance level:**
  - Tag-based: Add the following tag to the instance: Key: `kip_pbis_updates`, Value: `true` (any case).
  - Config file: Add the following flag to the `setc/ams.conf.d/state.ini` file: `skip_pbis_update = true.`

> **Note**
>
> Tag has a higher priority than the SSM parameter. You can turn off the PBIS update at the account level through the parameter, but turn it it for a single (or multiple) instance(s) by adding a tag `Key:skip_pbis_update`, `Value: false`.

To configure any of the described options, follow the standard change management process in your AMS environment.

Automatically update the minimum version of SSM and CloudWatch agents

The AMS Advanced *minimum version* (of the SSM or CloudWatch agents) is the version that has been tested by AMS service team and pre-approved for your operating system. We try to stay proactive and run the latest stable and compatible version, so the version number changes over time. You can find the current minimum version by raising a service request to AMS.

- **SSM Agent Management**

  The Amazon SSM Agent is responsible for running remote commands on the instance. The instance configuration automation ensures that the SSM Agent is running the minimum version.

- **Cloudwatch Agent Management**

  The Amazon CloudWatch Agent is responsible for emitting OS logs and metrics. Automated instance configuration performs the following:

  - If needed, disables the legacy CloudWatch Log agent and migrates the configuration to the new unified CloudWatch agent
  - If your instance is running the legacy CloudWatch Log Agent, automated instance configuration disables the legacy CloudWatch Log agent service and migrates its configuration to the unified CloudWatch agent.
  - Customizes your CloudWatch configuration to emit appropriate logs and metrics.

  Affected files and directories:

  - Windows
    - %ProgramData%\Amazon\AmazonCloudWatchAgent\
    - %ProgramData%\Amazon\AmazonCloudWatchAgent\Configs\
  - Linux
    - /opt/aws/amazon-cloudwatch-agent/etc/
    - /opt/aws/amazon-cloudwatch-agent/etc/amazon-cloudwatch-agent.d/
    - /opt/aws/ams/opt/aws/amazon-cloudwatch-agent/etc/amazon-cloudwatch-agent.json

CloudWatch configuration files, update details

We read your custom CloudWatch configurations (JSON only) from the following CloudWatch directories (see recommended directories), and merge them with the standard AMS CloudWatch configuration:

- CloudWatch Files
  - On the instance:
    - Windows
      - %ProgramData%\Amazon\AmazonCloudWatchAgent\Configs\
      - %ProgramFiles%\WindowsPowerShell\Modules\AWSManagedServices.Logging.Utilities\Files\Config.json
    - Linux
      - /opt/aws/ams/opt/aws/amazon-cloudwatch-agent/etc/amazon-cloudwatch-agent.json
  - On Amazon S3:
    - Windows:
      - https://ams-configuration-artifacts-*REGION_NAME*.s3.*REGION_NAME*.amazonaws.com/configurations/cloudwatch/latest/windows-cloudwatch-config.json
    - Linux:
      - https://ams-configuration-artifacts-*REGION_NAME*.s3.*REGION_NAME*.amazonaws.com/configurations/cloudwatch/latest/linux-cloudwatch-config.json

Automatically configured logs

We configure your instance to write the following logs.

- Windows:
  - AmazonSSMAgentLog
  - AmazonCloudWatchAgentLog
  - AmazonSSMErrorLog
  - AmazonCloudFormationLog
  - ApplicationEventLog
  - EC2ConfigServiceEventLog

- MicrosoftWindowsAppLockerEXEAndDLLEventLog
- MicrosoftWindowsAppLockerMSIAndScriptEventLog
- MicrosoftWindowsGroupPolicyOperationalEventLog
- SecurityEventLog
- SystemEventLog

- Linux:
  - /var/log/amazon/ssm/amazon-ssm-agent.log
  - /var/log/amazon/ssm/errors.log
  - /var/log/audit/audit.log
  - /var/log/cloud-init-output.log
  - /var/log/cloud-init.log
  - /var/log/cron
  - /var/log/dpkg.log
  - /var/log/maillog
  - /var/log/messages
  - /var/log/secure
  - /var/log/spooler
  - /var/log/syslog
  - /var/log/yum.log
  - /var/log/zypper.log

Monitoring and event management in AMS

Topics

- What is monitoring?
- What does the AMS monitoring system monitor?
- How monitoring works
- Viewing the monitoring configuration for an AMS account
- Changing the monitoring configuration for an AMS account
- Using OpsCenter in AMS
- Alert notifications from AMS
- Creating additional CloudWatch alarms in AMS
- Creating custom CloudWatch metrics and alarms in AMS
- Using CloudWatch Application Insights for .Net and SQL server in AMS
- Using Amazon EventBridge Managed Rules in AMS

The AWS Managed Services (AMS) monitoring system monitors your AMS resources for failures, performance degradation, and security issues. The AMS monitoring system relies on AWS services such as Amazon CloudWatch(CloudWatch), Amazon GuardDuty, Amazon Macie, and AWS Health. In addition to the monitoring system, AMS also deploys TrendMicro DeepSecurity for protection against malware on Amazon Elastic Compute Cloud (Amazon EC2) instances, for information about endpoint security (EPS) defaults, see Endpoint Security (EPS).

AMS monitoring provides these benefits:

- A monitoring baseline so that you have a default level of protection even if you don't configure any other monitoring for your managed accounts. For information, see Alerts from baseline monitoring in AMS.
- Investigation alerts to determine the appropriate action. For example, if GuardDuty finds activity indicating brute forcing attempts against an Amazon EC2 instance, AMS analyzes VPC flowlogs to understand the origin and context of the activity.
- Remediation of alerts, when possible, to prevent or reduce the impact for your applications. For example, if you are using a standalone Amazon EC2 instance and it fails the System health check, AMS attempts to recover the instance by stopping and restarting it. For more information, see AMS automatic remediation of alerts.

- Transparency into active, and previously resolved, alerts using OpsCenter. For example, if you have an unexpected high CPU utilization on an Amazon EC2 instance, you can request access to the AWS Systems Manager console (includes access to the OpsCenter console) and view the OpsItem directly in the OpsCenter console.

What is monitoring?

The AMS monitoring system monitors your AWS resources for failures, performance degradation, and security issues. As a managed account, AMS configures and deploys alarms for applicable AWS resources, monitors them, and performs remediation when applicable.

The AMS monitoring system generates alerts based on the monitoring configuration in your account. The monitoring configuration of an account refers to all the resource parameters in the account that create an alert; for information about the resource parameters, see Alerts from baseline monitoring in AMS. The monitoring configuration of an account includes CloudWatch Alarm definitions, and CloudWatch Event Rules that generate the alert (alarm or event).

The baseline monitoring configuration is the set of alarm definitions (Alerts from baseline monitoring in AMS) curated by AMS for monitoring resources in your managed account. The monitoring configuration of an account may differ from the baseline configuration, as a result of changes requested by you.

A notification of imminent, on-going, receding, or potential failures, performance degradation, or security issues generated by the baseline monitoring configured in an account, is called an alert. Examples of alerts are an Amazon CloudWatch Alarm, an Amazon CloudWatch Event, an Event, or a Finding from AWS service such as Amazon GuardDuty, and an event, or an alert, from Trend Micro Deep Security.

Alerts from security-related AWS services such as Amazon GuardDuty, Amazon Macie, or Trend Micro Deep Security are called security alerts to differentiate them from other types of alerts.

AMS monitoring provides these benefits:

- The ability to customize the baseline resource alarms to meet your requirements.
- Automatic remediation of alerts, when possible, to prevent or reduce the impact for your applications. For example, if you are using a standalone Amazon EC2 instance and it fails the system health check, AMS attempts to recover the instance by stopping and restarting it. For more information, see AMS automatic remediation of alerts.

- Transparency into active, and previously resolved, alerts using OpsCenter. For example, if you have an unexpected high CPU utilization on an Amazon EC2 instance, you can request access to the AWS Systems Manager console (which includes access to the OpsCenter console) and view the OpsItem directly in the OpsCenter console.
- Investigating alerts to determine the appropriate actions.
- Alerts generated based on the configuration in your account and supported AWS services. The monitoring configuration of an account refers to all the resource parameters in the account that create an alert. The monitoring configuration of an account includes CloudWatch Alarm definitions, and EventBridge (formerly known as CloudWatch Events) that generate the alert (alarm or event). For more information about resource parameters, see Alerts from baseline monitoring in AMS.
- Notification of imminent, on-going, receding, or potential failures; performance degradation; or security issues generated by the baseline monitoring configured in an account (known as an alert). Examples of alerts include a CloudWatch Alarm, an Event, or a Finding from an AWS service, such as GuardDuty or AWS Health.

What does the AMS monitoring system monitor?

In keeping with the AWS Managed Services (AMS) shared services responsibility model, the AMS monitoring system monitors your AWS infrastructure. For details on baseline monitoring in AMS, including AWS resources monitored and the type of alerts for each resource, see Alerts from baseline monitoring in AMS. For Amazon EC2 instances, AMS monitors the operating system and provides baseline monitoring based on OS metrics such as CPU utilization and root volume usage.

We recommend supplementing AMS monitoring with additional monitoring using AWS services tailored to your application. For guidance on monitoring for availability see the "Monitoring and Alarming" section in this whitepaper Reliability Pillar. You can configure your own monitoring to suit your operational needs; how to do this is discussed in Creating additional CloudWatch alarms in AMS and Creating custom CloudWatch metrics and alarms in AMS.

Single-Account Landing Zone proactive monitoring of Active Directory Trust in AMS

AMS single-account landing zone (SALZ) monitors the status of the one-way trust(s) between the Managed Active Directory (AD) in your AMS managed account and your company domain. The one-way trust with Managed AD is critical for access requests and instance logon requests. With this

new monitoring, AMS now proactively responds to trust related issues, and reduces the mean time to detect access related incidents.

This feature is automatically enabled in your AWS Managed Services (AMS) accounts.

There is a small cost impact. The feature uses four AWS CloudWatch metrics, and two AWS CloudWatch alarms for one trust.

How monitoring works

See the following graphics on monitoring architecture in AWS Managed Services (AMS).

The following diagram depicts a high-level overview of the **AMS multi-account landing zone** and **AMS single-account landing zone** monitoring workflow.

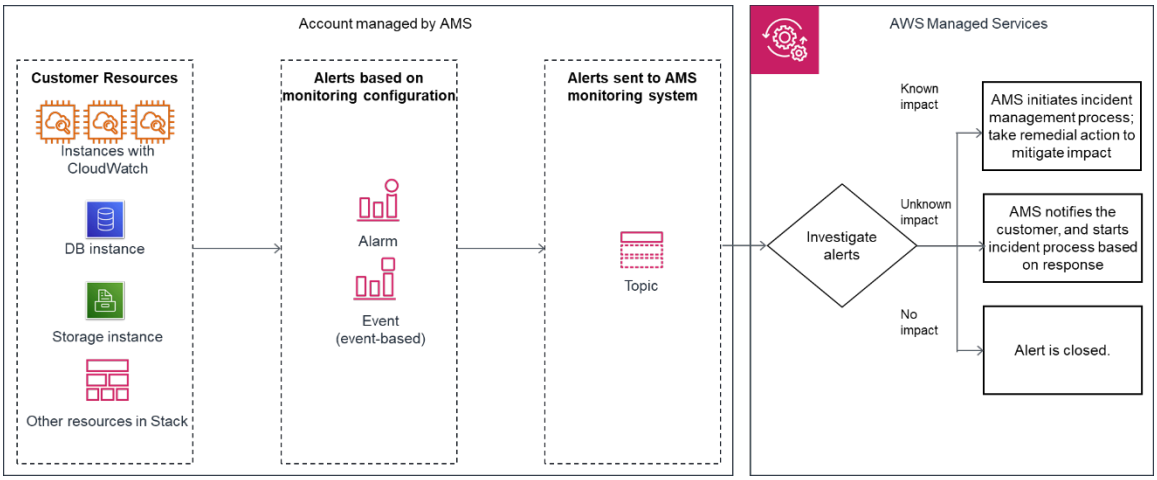

- Generation: At the time of account onboarding, AMS configures baseline monitoring (a combination of CloudWatch (CW) alarms, and CW event rules) for all your resources created in a managed account. The baseline monitoring configuration generates an alert when a CW alarm is triggered or a CW event is generated.
- Aggregation:
  - **Multi-Account Landing Zone**: Alerts are generated by your resources within Application and Core Organizational Unit accounts and sent to the AMS monitoring system by directing them through the Security account.
  - **Single-Account Landing Zone**: All alerts generated by your resources are sent to the AMS monitoring system by directing them to an SNS topic in the account.
- Processing: AMS analyzes the alerts and processes them based on their potential for impact. Alerts are processed as described next.

- Alerts with known customer impact: These lead to the creation of a new incident report and AMS follows the incident management process; for information about incident management, see AMS incident response.

  Example alert: An Amazon EC2 instance fails a system health check, AMS attempts to recover the instance by stopping and restarting it.

- Alerts with uncertain customer impact: For these types of alerts, AMS sends an incident report, in many cases asking you to verify the impact before we can take action.

  For example: An alert for >85% CPU utilization for more than 10 minutes on an Amazon EC2 instance can't be immediately categorized as an incident since this behavior may be expected based on usage. For such alerts, AMS sends an alert notification with the details and checks if the alert needs mitigating action. Alert notifications are discussed in detail in this section. We offer options for mitigating actions in the notification, and your reply that confirms that the alert is an incident triggers the creation of a new incident report and the AMS incident management process. Any service notification that receives a response of "no customer impact," or no response at all for three days, is marked as resolved and the corresponding alert is marked as resolved.

- Alerts with no customer impact: If, after evaluation, AMS determines that the alert doesn't have customer impact, the alert is closed.

  For example, AWS Health notifies of an EC2 instance requiring replacement but that instance has since been terminated.

Alert notification

As a part of the alert processing, based on the impact analysis, AWS Managed Services (AMS) creates an incident and initiates the incident management process for remediation, when impact can be determined. If impact can't be determined, AMS sends an alert notification to the email address associated with your account by way of a service notification; see the diagram on AMS monitoring architecture for alert handling process in How monitoring works.

Tag-based alert notification

We recommend tag-based alert notifications because notifications sent to a single email address can cause confusion when multiple teams use the same account. You can use tags to get alert notifications for different resources sent to different email addresses. For resources with alerts

that need to be sent to a specific email address, tag that resource with the key = OwnerTeamEmail, value = EMAIL_ADDRESS (use a group email; do not put personal information in tags). You can also use a custom tag key, but you must provide the custom tag key name to your CSDM with your explicit consent to use it in an email in order to activate automated notification for the tag-based communication. We recommend using the same tagging strategy for contact tags across all your instances and resources.

> **Note**
>
> The tag key value **OwnerTeamEmail** does not have to be in Pascal case; however, tags are case sensitive and it's best to use the recommended format. The email address must be specified in full, with the "at sign" (@) to separate the local part from the domain. Examples of invalid email addresses: Team.AppATabc.xyz or john.doe. For general guidance on your tagging strategy, see Tagging AWS resources. Do not add personally identifiable information (PII) in your tags, use distribution lists or aliases wherever possible.

Viewing the monitoring configuration for an AMS account

There are two key parts to the monitoring configuration of an account that you can view:

- CloudWatch Alarms: You can view all the CW alarms in the account by going to the CloudWatch console and selecting different services of interest.
- CloudWatch Events:
  - **Multi-Account Landing Zone**: CloudWatch Events monitored in the account can be found by filtering for all CW event rules with the string `"ams-"`.
  - **Single-Account Landing Zone**: CloudWatch Events monitored in the account can be found by filtering for all CW event rules with the string `"mc-"`.

Changing the monitoring configuration for an AMS account

You can change your baseline monitoring configuration for Amazon EC2 resources. For the alerts that can be configured, see Alerts from baseline monitoring in AMS. You can change the alarm definition, alarm destination, or opt-out of the alarm notification for the baseline monitors so that the alerts meet your application's operational requirements. You can request any or all of

the previously mentioned changes by submitting a Management | Other | Other | Update CT (ct-0xdawir96cy7k) with the following details.

- Instance ids [optional, if not mentioned, all instances in the account will be in-scope]
- CloudWatch metric name, for example, CPU utilization / swap free / IOwait
- Target - email id / phone number for SMS / SNS topic

To learn more about the type of changes you can request in the baseline monitoring configuration, see the Amazon CloudWatch Documentation.

Using OpsCenter in AMS

The AWS Managed Services (AMS) Operations team uses AWS Systems Manager OpsCenter for diagnosing and remediating many alerts related to your resources.

Using OpsCenter reduces mean time to resolution (MTTR), while providing a transparent view into the operational queues of the AMS operations teams.

With OpsCenter, AMS provides you with a transparent view of operational work items, also known as OpsItems, actively being worked upon by AMS teams, in addition to automated solutions.

To learn more about OpsCenter and OpsItems, see AWS Systems Manager OpsCenter. For information about getting access to the AWS Management Console, see Working with the AWS Management Console. From the AWS Management Console you can navigate to the AWS Systems Manager Console, and OpsCenter; to learn more, see AWS Systems Manager Session Manager. OpsCenter also provides an API that you can use; for information, see Learn More About OpsCenter.

OpsCenter is a priced feature with ~1000 OpsItems that cost under $10. For information, see AWS Systems Manager pricing.

Alert notifications from AMS

As a part of the alert processing, based on the impact analysis, AWS Managed Services (AMS) creates an incident and initiates the incident management process for remediation, when impact can be determined. In case impact cannot be determined, AMS sends an alert notification to the email address associated with your account through a service notification; see the diagram on AMS monitoring architecture for alert handling process in How monitoring works.

Receiving alerts generated by AMS

AWS Managed Services (AMS) enables you to receive alert notifications for Amazon EC2 resources directly to reduce communication delays. To receive Amazon EC2 alerts directly, subscribe your target (preferred email) to the SNS topic **Direct-Customer-Alerts** using the Management | Monitoring and notification | SNS | Subscribe change type (ct-3rcl9u1k017wu).

> **Note**
>
> Not all baseline alerts are sent to the **Direct-Customer-Alerts** topic by default. To see all alerts that are generated by the AMS monitoring system, subscribe to the SNS topic for the AMS monitoring system (in the request, ask for the "AMS Monitoring Topic"), specify a subscription channel that should receive the alerts (lambda, SQS, HTTP/S, email, or SMS), and specify the endpoints (for example, email addresses, if you choose the email protocol) that should receive the alerts.
> The AMS monitoring topic gets alerts that are used by AMS shared services, so it can be noisy.

To learn more, see the Amazon Simple Notification Service Subscribe API reference, Subscribe section.

For a list of baseline alerts AMS provides, see Alerts from baseline monitoring in AMS.

Tag-based alert notifications in AMS

Using tag-based alert notifications are recommended because notifications sent to a single email address can cause confusion when multiple developer teams use the same account. AWS Managed Services (AMS) allows you to use tags to get alert notifications for different resources sent to different email addresses. For resources with alerts that need to be sent to a specific email address, tag that resource with the `key = OwnerTeamEmail`, `value = EMAIL_ADDRESS`.

> **Important**
>
> Tag-based alert notifications only work for notifications related to Amazon EC2, Amazon EBS, Elastic Load Balancing, Network Load Balancer, Application Load Balancer, Amazon RDS, Amazon Redshift, and OpenSearch.
> Do not add personally identifiable information (PII) in your tags.

You can also use a custom tag key, but you must provide the custom tag key name to your CSDM with an email consent, to activate automated notification for the tag-based communication. We recommended using the same tagging strategy for contact tags across all your instances/resources.

The email address must be specified in full, with the "at sign" (@) to separate the local part from the domain. Examples of invalid email addresses: *`Team.AppATabc.xyz`* or *`john.doe`*. For general guidance on your tagging strategy, see Using tags. To tag an existing resource, submit an RFC with the Deployment | Advanced stack components | Tag | Create (auto) (ct-3cx7we852p3af) change type.

> **Note**
>
> The key value *`OwnerTeamEmail`* does not have to be in camel case. However, tags are case sensitive and it's best to use the recommended format.

AMS automatic remediation of alerts

Some alerts are automatically remediated by AWS Managed Services (AMS). This section describes how this remediation works and the conditions that must be met for the remediation to take place.

Alert name	Description	Remediation
Status Check Failed	This alarm indicates that the instance is running on degraded hardware or entered a fault state.	Our remediation first validates instance accessibility. If confirmed that accessibility is impacted, it stops the instance and starts it again so it can be migrated to new underlying hardware.
AMSLinuxDiskUsage	This alarm indicates that a mount point of your Linux EC2 instance is filling up.	The remediation first deletes temporary files. If this does not free up required space, it extends the volume to prevent downtime if the volume were to get full.

Alert name	Description	Remediation
AMSWindowsDiskUsage	This alarm indicates that a drive of your Windows EC2 instance is filling up.	The remediation first deletes temporary files. If this does not free up required space, it extends the volume to prevent downtime if the volume were to get full.
RDS-EVENT-0089	This alarm indicates that the DB instance has consumed more than 90% of its allocated storage.	The remediation first validates the DB is in a modifiable and available /storage-full state. It will attempt to increase the allocated storage, IOPS and storage throughput via cloudformation changeset, if stack drift is already detected it will fall back to RDS API to prevent downtime.
RDS-EVENT-0007	This alarm indicates that the allocated storage for the DB instance has been exhausted.	The remediation first validates the DB is in a modifiable and available /storage-full state. It will attempt to increase the allocated storage, IOPS and storage throughput via cloudformation changeset, if stack drift is already detected it will fall back to RDS API to prevent downtime.

EC2 status check failure remediation automation

These are some notes about how AWS Managed Services (AMS) auto-remediation works with EC2 status check failure issues.

- Your EC2 instance has become unreachable. In order to recover it, it must be stopped and started again so it's migrated to new hardware.
- The automation is not able to recover your instance if the root of the problem is within the OS. For example, missing devices in fstab, kernel corruption, and so on.

- If your instance belongs to an Auto Scaling group, the automation takes no action. The autoscaling replaces the instance.
- The remediation doesn't take action if EC2 Auto Recovery is enabled for this instance.

EC2 volume usage remediation automation

How AWS Managed Services (AMS) auto-remediation works with EC2 volume usage issues.

- Before extending a volume, the automation performs cleanup tasks (Windows: Disk Cleaner Linux: Logrotate + Simple Service Manager Agent Log removal) on the instance to try to free up space.

> **Note**
>
> The cleanup tasks are not run on EC2 "T" family instances due to their reliance on CPU credits for continued functionality.

- On Linux, the automation only supports extending file systems of type EXT2, EXT3, EXT4 and XFS.
- On Windows, the automation only supports New Technology File System (NTFS) and Resilient File System (ReFS).
- The automation doesn't extend volumes that are part of Logical Volume Manager (LVM) or a RAID array.
- The automation does not extend *instance store* volumes.
- The automation does not take action if the affected volume is already bigger than 2 TiB.
- The capacity expansion portion of the automation occurs once every 6 hours with a 3-time weekly and 5-time lifetime volume expansion limit.

When these rules prevent the automation from taking action, AMS reaches out to you through an outbound service request to determine the next actions to take.

Amazon RDS low storage event remediation automation

How AWS Managed Services (AMS) auto-remediation works with Amazon RDS low storage event issues.

- Before trying to extend the Amazon RDS instance storage, the automation performs several checks to ensure the Amazon RDS instance is in a modifiable and available, or storage-full, state.
- Where CloudFormation stack drift is detected, remediation occurs through Amazon RDS API.
- The remediation action does not run in the following scenarios:
  - The Amazon RDS instance status is not "available" or "storage-full".
  - The Amazon RDS instance storage is not currently modifiable (such as when the storage has been modified in the last 6 hours).
  - The Amazon RDS instance has auto-scaling storage enabled.
  - The Amazon RDS instance is not a resource within a CloudFormation stack.
- Remediation is limited to 1 expansion per 6 hours and no more than 3 expansions within a rolling fourteen day period.
- Where the above states are met, AMS reaches out to you with an outbound incident to determine next actions.

Creating additional CloudWatch alarms in AMS

You can create new CloudWatch alarms using the AWS Managed Services (AMS) Deployment | Monitoring and notification | CloudWatch | Create alarms change type.

> **Important**
>
> AMS does not monitor CloudWatch alarms created by you.

Using custom CloudWatch metrics and alarms for Amazon EC2 instances (works only for mutable deployments that do not rely on updated AMIs deployed to Auto Scaling groups):

1. Produce your application monitoring script and custom metric. For more information and access to example scripts, see Monitoring Memory and Disk Metrics for Amazon EC2 Linux Instances. The Amazon CloudWatch monitoring scripts for Linux Amazon EC2 instances demonstrate how to produce and consume Amazon CloudWatch custom metrics. These sample Perl scripts comprise a fully functional example that reports memory, swap, and disk space utilization metrics for a Linux instance.
2. Upload your monitoring script. To upload the monitoring script to your Auto Scaling group or Amazon EC2 instance configuration, you can use UserData when configuring your Auto Scaling

group or Amazon EC2 instance, or, if your application was deployed with CodeDeploy, you can modify the configuration with a Deployment | Applications | CodeDeploy application | Deploy CT (ct-2edc3sd1sqmrb).

3. Publish your custom metric to CloudWatch (the first time you publish a data point for a new custom metric, it is created), see Publishing Custom Metrics.
4. Create the CloudWatch alarm, see Create a CloudWatch Alarm for an Instance.

> **Important**
>
> Monitoring data must be sent to this path [`infra/INSTANCE_ID/YOUR_CUSTOM_METRIC`]

To modify or delete a CloudWatch alarm, submit an RFC with the Management | Other | Other | Create change type (ct-1e1xtak34nx76) with the parameters required to complete the action as described in the Amazon CloudWatch API reference PutMetricAlarm.

You can use the CloudWatch event stream. AMS is integrated with CloudWatch and you can request that any AWS API call trigger a CloudWatch event.

To do this, submit a Management | Other | Other | Update CT (ct-0xdawir96cy7k) with the API calls that you are interested in. An AMS operator will talk to you to gather requirements. To learn more, see the Amazon CloudWatch Documentation.

To get access to the CloudWatch event stream, submit a Management | Other | Other | Update CT (ct-0xdawir96cy7k) to add a party to the SNS notification topic. An AMS operator will talk to you to gather requirements.

Creating custom CloudWatch metrics and alarms in AMS

You can store your business and application metrics in Amazon CloudWatch. You can view graphs, and set alarms based on these metrics, just as you can for the metrics that CloudWatch already stores for your AWS Managed Services (AMS) resources. To learn more about CloudWatch, see Amazon CloudWatch Concepts.

Amazon SNS allows applications to send time-critical messages to multiple subscribers through a "push" mechanism against the AMS Managed Monitoring System or MMS, Amazon SNS (SNS) topic that the alarms are published to; in this case, MMS and your SQS queues. You can use CloudWatch

to create custom metrics and, through an SNS topic, have AMS alarm you appropriately. To do this, follow these steps.

> **Note**
>
> This process doesn't work for immutable deployments that rely on updated AMIs deployed to Auto Scaling groups, it is suitable for mutable application (not ASG) deployments. Setting up a custom metric within the limitations of AMS Advanced, is a complex task. For an example from CloudWatch, see Example: Count occurrences of a term.

1. Produce your application monitoring script and custom metric (such as the count occurrences example). For more information and access to example scripts, see Monitoring Memory and Disk Metrics for Amazon EC2 Linux Instances.
2. Upload your monitoring script. To upload the monitoring script to your Auto Scaling group or Amazon EC2 instance configuration, you can use **UserData** when configuring your Auto Scaling group or Amazon EC2 instance, or, if your application was deployed with CodeDeploy, you can modify the configuration with a Deployment | Applications | CodeDeploy application | Deploy CT (ct-2edc3sd1sqmrb).
3. Publish your custom metric to CloudWatch (the first time you publish a data point for a new custom metric, it is created), see Publish Custom Metrics.
4. To integrate your customer metric to your application monitoring system, request AMS create an SNS topic for the metric by submitting an RFC with the Management | Other | Other | Create CT (ct-1e1xtak34nx76).
5. Create the CloudWatch alarm, see Creating Amazon CloudWatch Alarms.

> **Important**
>
> Monitoring data must be sent to this path [infra/*INSTANCE_ID*/*YOUR_CUSTOM_METRIC*].

Using CloudWatch Application Insights for .Net and SQL server in AMS

You can use Amazon CloudWatch Application Insights to set up the monitors for your AWS Managed Services (AMS) application resources to continuously analyze data for signs of problems

with your applications and reduce your mean time to repair (MTTR) when troubleshooting application issues. For details about CloudWatch Application Insights, see CloudWatch Application Insights for .NET and SQL Server.

> **Important**
>
> AMS does not monitor problems from CloudWatch Application Insights because they are for application code controlled by you.

To use CloudWatch Application Insights, submit an RFC with the Management | Other | Other | Create change type (ct-1e1xtak34nx76) with a request to create an IAM role that provides you with permission to configure CloudWatch Application Insights. There are two options to receive the problems identified: through an SNS topic or with a target in CloudWatch Event rules. In the RFC, specify which you want. If you plan to use CloudWatch Event rules, also specify the rule definition in the RFC. After you're set up with CloudWatch Application Insights, you receive notice of potential problems including insights that point to a possible root cause.

To learn how you can assume the role, see the AMS Onboarding Guide Federate your Active Directory with the AMS IAM Roles.

Using Amazon EventBridge Managed Rules in AMS

AMS Advanced uses Amazon EventBridge Managed Rules. A Managed Rule is a unique type of rule that is directly linked to AMS. These rules match incoming events and send them to targets for processing. Managed Rules are predefined by AMS and include event patterns that are required by the service to manage customer accounts, and unless defined otherwise, only the owning service can utilize these Managed Rules.

AMS Managed Rules are linked to `events.managedservices.amazonaws.com` service principal. These Managed Rules are managed through the `AWSServiceRoleForManagedServices_Events` service-linked role. To delete these rules a special confirmation by the customer is required. For more information see Deleting Managed Rules for AMS.

For more information about rules, see Rules in the *Amazon EventBridge User Guide*.

Amazon EventBridge Managed Rules deployed by AMS

Amazon EventBridge Managed Rules

Rule Name	Description	Definition

AMSAdvancedCoreRule	This rule forwards Amazon CloudWatch Alarms to AMS Monitoring. The Amazon CloudWatch events monitor CloudWatch Alarms.	<pre>{
		    {
		   "source": ["aws.cloudwatch"],
		    "detail-type": ["CloudWatch Alarm State Change"],
		    }
		}</pre>

Creating Managed Rules for AMS

You don’t need to manually create Amazon EventBridge Managed Rules. When you onboard to AMS in the AWS Management Console, the AWS CLI, or the AWS API, AMS creates them for you.

Editing Managed Rules for AMS

AMS doesn't allow you to edit the Managed Rules. Name and event pattern for each Managed Rule are predefined by AMS.

Deleting Managed Rules for AMS

You don’t need to manually delete the Managed Rules. When you offboard from AMS in the AWS Management Console, the AWS CLI, or the AWS API, AMS cleans up the resources and deletes all Managed Rules owned by AMS for you.

In the event AMS fails to remove the Managed Rules during offboarding, you can also use the Amazon EventBridge console, the AWS CLI or the AWS API to manually delete the Managed Rules. To do this, you must first offboard from AMS and conduct a force delete of the Managed Rules.

Log management

Topics

- What is log management?
- How AMS logging works
- Accessing your logs
- Customizing your log configuration

AMS log management collects, aggregates, and controls retention of the logs from the managed account. AWS log management aggregates logs from Amazon EC2 instances and AWS resources deployed within your account into CloudWatch Logs. The full list of services from which logs are currently aggregated can be found in AMS aggregated service logs.

What is log management?

Log management is the process of dealing with log events generated by instances, applications, and AWS services. This feature defines how AMS processes, stores, and rotates the log events generated in your managed AWS account. Infrastructure logs are used during incident resolution and to support system audits.

How AMS logging works

AMS single-account landing zone (SALZ) log management uses a variety of pre-installed agents and tools that are implemented when instances and applications are onboarded or provisioned.

Logging is configured during the account onboarding process and when a stack is launched.

AMS multi-account landing zone (MALZ) logs produced by instances and AWS services are available in CloudWatch Logs or Amazon Simple Storage Service (Amazon S3), within each account managed by AMS. AMS multi-account landing zone provides a central Logging Account that acts as a central aggregation location for some logs produced by individual application accounts.

The tables in the Accessing your logs subsections describe which logs are available in individual accounts, and which are available in the central Logging Account.

Accessing your logs

To access your logs, ensure that you have one of the required IAM roles and are in your AMS account. Then navigate to the directory shown.

Multi-Account Landing Zone (MALZ)

Provides five default IAM roles, each of which allow access to all logs within your account (all are prefaced with `AWSManagedServices`):

- `AdminRole`
- `CaseRole`
- `ChangeManagementRole`
- `ReadOnlyRole`
- `SecurityOpsRole`

Access to these roles is configured via federation, with each role being mapped to a group within your Active Directory domain.

To learn more about these roles, see IAM user role in AMS .

Single-Account Landing Zone (SALZ)

The default `Customer_ReadOnly_Role` for AMS single-account landing zone allows your access to all logs within your account. Access to the logs is controlled using AWS Identity and Access Management (IAM) roles mapped to Active Directory groups.

AMS aggregated service logs

Each AWS service logs to either CloudWatch Logs or a specific location in an Amazon S3 bucket.

> **Note**
>
> Unless specifically stated, all log locations are local to the account that generated the logs, and are not aggregated into the central Logging account.
> To find the default AMS CloudTrail trail names in SALZ and MALZ accounts, go to the AWS Console for CloudTrail and then to the **Trails** page and search for AMS. Because AMS resources have tags, you can find the trails this way. Example AMS CloudTrail tag:

```
Environment    AMSInfrastructure
```

To access your logs, ensure that you have one of the required IAM roles and are in your AMS account. Then navigate to the directory shown.

Multi-Account Landing Zone

AMS multi-account landing zone Aggregated Service Logs

	Service name	Log details	Log location
1	Amazon Aurora	General, slow query, and error logs.	CloudWatch LogGroup: /aws/rds/cluster/{*database_name* }/{*log_name*}
2	AWS CloudFormation (CFN)	API call logging only.	AWS CloudFormation API calls are documented via CloudTrail, which sends its logs to the CloudWatch LogGroup and then syncs the logs into an S3 bucket. Logs are retained for 14 days by default in the CloudWatch LogGroup, and are retained indefinitely in the S3 bucket. CloudWatch LogGroup: /CloudTrail/Landing-Zone-Logs S3 bucket [in the central Logging Account]: aws-landing-zone-logs-ams-a{*account_ID* }-log-management-{*region*} Path: /AWSLogs/{*account_ID* }/CloudTrail/

	Service name	Log details	Log location
3	Amazon CloudFront (CloudFront)	User request logging. CloudFront logging must be explicitly enabled. For information, see Enabling logging for supported services.	S3 bucket: ams-a{*account_ID* }-log-management-{*region*} Path: AWS/RedShift/{*CloudFront distribution ID* }
4	Amazon CloudWatch (CloudWatch)	API call logging only.	CloudWatch LogGroup: /CloudTrail/Landing-Zone-Logs S3 bucket [in the central Logging Account]: aws-landing-zone-logs-{*account_ID* }-{*region*} Path: /AWSLogs/{*account_ID* }/CloudTrail/
5	Amazon Elastic Block Store (Amazon EBS)	No logs are produced by the EBS service.	Not applicable
6	Amazon Elastic Compute Cloud (Amazon EC2)	System and application logs. For information, see the Amazon Elastic Compute Cloud (Amazon EC2) - system level logs.	CloudWatch Logs: /{*instance ID*}
7	Amazon Elastic File System (Amazon EFS)	API call logging only.	CloudWatch LogGroup: /CloudTrail/Landing-Zone-Logs S3 bucket [in the central Logging Account]: aws-landing-zone-logs-{*account_ID* }-{*region*} Path: /AWSLogs/{*account_ID* }/CloudTrail/

	Service name	Log details	Log location
8	Elastic Load Balancing (ELB)	Access and error log entries. Elastic load balancers log all requests sent to them, including requests that aren't routed to back-end instances. For example, if a client sends a malformed request, or there are no healthy instances to respond, the request is still logged. For more information about Elastic Load Balancing log entries, see • Classic Load Balancers: Access log entries. • Application Load Balancers: Access log entries. • Network Load Balancers: Access log entries.	API call logs: CloudWatch LogGroup: /CloudTrail/Landing-Zone-Logs S3 bucket [in the central Logging Account]: aws-landing-zone-logs-{*account_ID* }-{*region*} Path: /AWSLogs/{*account_ID* }/CloudTrail/ Access logs: S3 bucket: mc-a{*account_ID* }-logs{*region*} Path: aws/elbaccess
9	Amazon OpenSearch Service (OpenSearch Service)	Service error logs. You must explicitly enable OpenSearch logging. For information, see Enabling logging for supported services	CloudWatch LogGroup: /CloudTrail/Landing-Zone-Logs S3 bucket [in the central Logging Account]: aws-landing-zone-logs-{*account_ID* }-{*region*} Path: /AWSLogs/{*account_ID* }/CloudTrail/

	Service name	Log details	Log location
10	Amazon ElastiCache	API call logging only.	CloudWatch LogGroup: //CloudTrail/Landing-Zone-Logs S3 bucket [in the central Logging Account]: aws-landing-zone-logs-{*account_ID* }-{*region*} Path: /AWSLogs/{*account_ID* }/CloudTrail/
11	Amazon GuardDuty		
12	Amazon Inspector		
13	Amazon Macie		
14	Amazon Redshift	Connection, user, and activity logs. Logging is enabled by default when you create your Redshift cluster by invoking the Create Redshift cluster CT (ct-1malj7snzxrkr). For information, see Database Audit Logging.	S3 bucket: ams-a{*account_ID* }-log-management-{*region*} Path: /AWS/RedShift/{*CloudFront Distribution ID*}
15	Amazon Relational Database Service (RDS)	Logs specific to database type. You must explicitly enable RDS logging. For information, see Enabling logging for supported services You can only access MSSQL logs through a stored procedure; for information, see Archiving Log Files.	CloudWatch LogGroup: /aws/rds/(*instance* or *cluster*)/{*database_name* }/{*log_name*}

	Service name	Log details	Log location
16	Amazon S3 (S3)	Bucket access logs. Each access log record provides details about a single access request such as the requester, bucket name, request time, request action, response status, and error code (if any). Access log information can be useful in security and access audits. It can also help you learn about your customer base and understand your Amazon S3 bill. For more information about S3 Access Log entries, see S3 Server Access Log Format.	S3 bucket: mc-a{*account_ID* }-log-management-{*region*} Path: /aws/s3access/{*bucket_name* } S3 bucket [in the central Logging Account]: aws-landing-zone-s3-access-logs-{*account_ID* }-{*region*} Path: /
17	Amazon Simple Email Service (SES)	SES API service calls.	CloudWatch LogGroup: /CloudTrail/Landing-Zone-Logs S3 bucket [in the central Logging Account]: aws-landing-zone-logs-{*account_ID* }-{*region*} Path: /AWSLogs/{*account_ID* }/CloudTrail/
18	Amazon Virtual Private Cloud (VPC)	VPC flow data (information about the IP traffic going to and from your VPC's network interfaces).	CloudWatch LogGroup: /aws/vpcflow/{*VPC_ID*}

	Service name	Log details	Log location
19	Auto Scaling	API call logging only.	CloudWatch LogGroup: /CloudTrail/Landing-Zone-Logs S3 bucket [in the central Logging Account]: aws-landing-zone-logs-{*account_ID* }-{*region*} Path: /AWSLogs/{*account_ID* }/CloudTrail/
20	AWS Certificate Manager		
21	AWS CodeDeploy	Instance-specific deployment logs.	On Instance
22	AWS Config	AWS Config API service calls.	CloudWatch LogGroup: /CloudTrail/Landing-Zone-Logs S3 bucket [in the central Logging Account]: aws-landing-zone-logs-{*account_ID* }-{*region*} Path: /AWSLogs/{*account_ID* }/CloudTrail/
		Resource configuration changes, as tracked by AWS Config.	S3 bucket [in the central Logging Account]: aws-landing-zone-logs-{*account_ID* }-{*region*} Path: /AWSLogs/{*account_ID* }/Config/
23	AWS Database Migration Service	Database migration logs. For information, see Introducing log management in AWS Database Migration Service.	Database migration console

	Service name	Log details	Log location
24	AWS Direct Connect (DX)	API call logging only.	CloudWatch LogGroup: /CloudTrail/Landing-Zone-Logs
25	AWS Glacier		S3 bucket [in the central Logging Account]: aws-landing-zone-logs-{*account_ID* }-{*region*}
26	AWS IAM (IAM)		
27	AWS Key Management Service		Path: /AWSLogs/{*account_ID* }/CloudTrail/
28	AWS Management Console (console or AWS Console)		
29	AWS Simple Notification Service (SNS)		
30	AWS Simple Queueing Service (SQS)		

Single-Account Landing Zone

AMS single-account landing zone Aggregated Service Logs

	Service name	Log details	Log location
1	Amazon Aurora	General, slow query, and error logs.	CloudWatch LogGroup: /aws/rds/cluster/{*database_name* }/{*log_name*}
2	Amazon CloudForm	API call logging only.	CloudFormation API calls are documented via CloudTrail, which

	Service name	Log details	Log location
	ation (CloudFormation or CFN)		sends its logs to the CloudWatch LogGroup and then syncs the logs into an S3 bucket. CloudWatch LogGroup: /aws/ams/cloudtrail S3 bucket: ams-a{*account_ID* }-log-management-{*region*}
3	Amazon CloudFront (CloudFront)	User request logging. You must explicitly enable CloudFront logging. For information, see Enabling logging for supported services	S3 bucket: ams-a{*account_ID* }-log-management-{*region*} Path: AWS/RedShift/{*CloudFront_distribution_ID* }
4	Amazon CloudWatch (CloudWatch)	API call logging only.	CloudWatch LogGroup: /aws/ams/cloudtrail
5	Amazon Elastic Block Store (EBS)	No logs are produced by the EBS service.	Not applicable
6	Amazon Elastic Compute Cloud (EC2)	System and application logs. For information, see the Amazon Elastic Compute Cloud (Amazon EC2) - system level logs.	CloudWatch Logs: /{*instance_ID* }
7	Amazon Elastic File System (Amazon EFS)	API call logging only.	CloudWatch LogGroup: /aws/ams/cloudtrail

	Service name	Log details	Log location
8	Elastic Load Balancing (ELB)	Access and error log entries. Elastic load balancers log all requests sent to them, including requests that aren't routed to back-end instances. For example, if a client sends a malformed request, or there are no healthy instances to respond, the request is still logged. For more information about elastic load balancer log entries, see • Classic Load Balancers: Access log entries. • Application Load Balancers: Access log entries. • Network Load Balancers: Access log entries.	CloudWatch LogGroup: /aws/ams/cloudtrail S3 bucket: mc-a{*account_ID* }-logs-{*region*} Path: aws/elbaccess
9	Amazon OpenSearch Service (OpenSearch Service)	Service error logs. You must explicitly enable OpenSearch logging. For information, see Enabling logging for supported services	CloudWatch LogGroup: /aws/ams/cloudtrail
10	Amazon ElastiCache	API call logging only.	CloudWatch LogGroup: /aws/ams/cloudtrail
11	Amazon GuardDuty		

	Service name	Log details	Log location
12	Amazon Inspector		
13	Amazon Macie		
14	Amazon Redshift	Connection, user, and activity logs. Logging is enabled by default when you create your Redshift cluster by invoking the Create Redshift cluster CT (ct-1malj7snzxrkr). For information, see Database Audit Logging.	S3 bucket: ams-a{*account_ID* }-log-management-{*region*} Path: /AWS/RedShift/{*CloudFront_Distribution_ID* }
15	Amazon Relational Database Service (RDS)	Logs specific to database type. RDS logging must be explicitly enabled. For information, see Enabling logging for supported services You can only access MSSQL logs through a stored procedure; for information, see Archiving Log Files.	CloudWatch LogGroup: /aws/rds/(instance|cluster)/{database name}/{log name}

	Service name	Log details	Log location
16	Amazon S3 (S3)	Bucket access logs. Each access log record provides details about a single access request, such as: requester, bucket name, request time, request action, response status, and error code (if any). Access log information can be useful in security and access audits; it can also help you learn about your customer base and understand your Amazon S3 bill. For more information on S3 Access Log entries, see S3 Server Access Log Format.	S3 bucket: mc-a{*account_ID* }-log-management-{*region*} Path: /aws/s3access/{*bucket_name* }
17	Amazon Simple Email Service (SES)	SES API service calls.	CloudWatch LogGroup: /aws/ams/cloudtrail S3 bucket: ams-a{*account_ID* }-log-management-{*region*} Path: AWS/CloudTrail/AWSLogs/{*account_ID* }/CloudTrail/{*region*}
18	Amazon Virtual Private Cloud (VPC)	VPC flow data (information about the IP traffic going to and from your VPC's network interfaces).	CloudWatch LogGroup: /aws/vpcflow/{vpc_id}
19	Auto Scaling	API call logging only.	CloudWatch LogGroup: /aws/ams/cloudtrail
20	AWS Certificate Manager		
21	AWS CodeDeploy	Instance specific deployment logs.	On instance

	Service name	Log details	Log location
22	AWS Config	AWS Config API service calls.	CloudWatch LogGroup: /aws/ams/cloudtrail S3 bucket: ams-a{*account_ID*}-log-management-{*region*} Path: AWS/CloudTrail/AWSLogs/{*account_ID*}/CloudTrail/{*region*}
23	AWS Database Migration Service	Database migration logs. For information, see Introducing log management in AWS Database Migration Service.	Database migration console
24	AWS Direct Connect (DX)	API call logging only.	CloudWatch LogGroup: /aws/ams/cloudtrail
25	AWS Glacier		
26	AWS IAM (IAM)		
27	AWS Key Management Service		
28	AWS Management Console (console or AWS Console)		
29	AWS Simple Notification Service (SNS)		

	Service name	Log details	Log location
30	AWS Simple Queueing Service (SQS)		

AMS shared services logs

The following table describes the logs, and log location, for the AMS Shared Services in your account.

To access your logs, ensure that you have one of the required IAM roles and are in your AMS account. Then navigate to the directory shown.

AMS single-account landing zone Shared Services Logging

	Shared service name	Log details	Log location
1	Bastion Hosts	Information regarding users accessing the bastion host.	**Linux Bastions**: CloudWatch Logs: /{instance id}/var/log/secure CloudWatch Logs: /{instance id}/var/log/audit/audit.log **Windows Bastions**: CloudWatch Logs: /{instance id}/SecurityEventLog
2	Management Hosts	Output of scripts, which assist in automated access management actions within the account.	CloudWatch Logs: /{instance id}/ApplicationEventLog
4	EPS Hosts (DSM)	Information regarding the enrollment of instances	CloudWatch Logs: /{instance id}/var/log/DSM.log

	Shared service name	Log details	Log location
		onto the Deep Security Management platform.	
5	Directory Services	Information regarding account login, account management, detailed tracking, object access, policy change, and privilege use within the account's directory. You must explicitly enable Directory Services logging. For information, see Enabling logging for supported services.	CloudWatch Logs: /aws/directoryservice/{directory id}-{directory dns name}
6	Lambdas	Output of various lambdas, which assist in automated operational actions within the account.	CloudWatch Logs: /aws/lambda/{lambda name}

AMS multi-account landing zone Shared Services Logging

	Shared service name	Log details	Log location
1	Bastions	Output of instance logins and authentication failures.	**Linux Bastions** CloudWatch Logs: /{*instance_ID* }/var/log/secure.log **Windows Bastions** CloudWatch Logs: /{*instance_ID* }/SecurityEventLog

	Shared service name	Log details	Log location
2	Management Hosts	Output of scriptsy, which assist in automated access management actions within the account.	CloudWatch Logs: /{*instance_ID* }/ApplicationEventLog
3	EPS Hosts (DSM)	Information regarding the enrollment of instances onto the Deep Security Management platform.	CloudWatch Logs: /{*instance_ID* }/var/log/DSM.log
4	Directory Services	Information regarding account login, account management, detailed tracking, object access, policy change, and privilege use within the account's directory. You must explicitly enable Directory Services logging. For information, see Enabling logging for supported services.	CloudWatch Logs: /aws/directoryservice/{*directory_ID* }-{*directory_DNS_name* }
5	Lambdas	Output of various lambdas, which assist in automated operational actions within the account.	CloudWatch Logs: /aws/lambda/{*Lambda_name* }

Amazon Elastic Compute Cloud (Amazon EC2) - system level logs

Instance logs are collected by a CloudWatch Logs agent running on the instance and can be accessed through a CloudWatch Log group of the same name as the instance. For example, if the

instance ID is i-0123456789abcdef0 and the log file name is /var/log/messages, the Log Group would be i-0123456789abcdef0 and the Log Stream /var/log/messages.

See also AMS aggregated service logs.

To access your logs, ensure that you have one of the required IAM roles and are in your AMS account. Then navigate to the directory shown.

> **Note**
>
> The following logs are collected by default.

Amazon Linux / Red Hat Linux / Centos Linux / Ubuntu / SUSE Linux

Log file / Log stream

```
/var/log/amazon/ssm/amazon-ssm-agent.log
/var/log/amazon/ssm/errors.log
/var/log/audit/audit.log
/var/log/cloud-init-output.log
/var/log/cfn-init.log
/var/log/cfn-init-cmd.log
/var/log/cloud-init.log (Amazon Linux 1 / Amazon Linux 2 only)
/var/log/cron
/var/log/maillog
/var/log/messages
/var/log/secure
/var/log/spooler
/var/log/yum.log
/var/log/aws/ams/bootstrap.log
/var/log/aws/ams/build.log
/var/log/syslog
/var/log/dpkg.log
/var/log/auth.log
/var/log/zypper.log
```

> **Note**
>
> For information on accessing logs for Amazon Linux 2023, see Why is the /var/log directory missing logs in my EC2 Amazon Linux 2023 instance?

Windows

Log file / Log stream

```
SecurityEventLog
SystemEventLog
AmazonSSMAgentLog
MicrosoftWindowsAppLockerMSIAndScriptEventLog
MicrosoftWindowsAppLockerEXEAndDLLEventLog
AmazonCloudWatchAgentLog
EC2ConfigServiceEventLog (Windows Server 2012 R2 Only)
ApplicationEventLog
AmazonCloudFormationLog
MicrosoftWindowsGroupPolicyOperationalEventLog
AmazonSSMErrorLog
```

Integrating with Splunk

AMS supports AWS Lambda-based push to customer log analytics services, such as Splunk.

AMS leverages the Splunk Add-on for Amazon Web services, which allows AWS data to be streamed to Splunk. See Hardware and software requirements.

Refer to this Splunk blog post How to stream AWS CloudWatch Logs to Splunk (Hint: it's easier than you think). Because CloudWatch log streaming is enabled by default for AMS customers, and AMS configures the AWS Lambda function for you, though you need to configure the Splunk HTTP Event Collector (HEC) input and submit a request to AMS for the added functionality.

Here's how the data input settings might look:

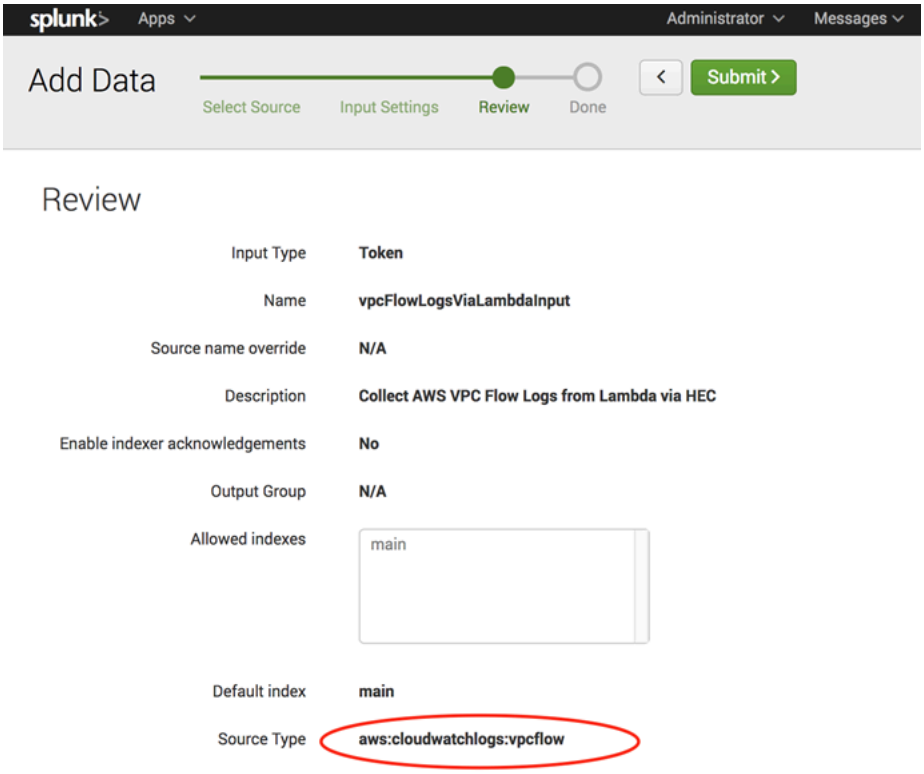

Customizing your log configuration

You can alter log data retention for CloudWatch logs, and you can enable logging for additional AWS services.

Altering CloudWatch log retention

You can change the log data retention setting for CloudWatch logs. By default, logs are kept indefinitely and never expire. You can adjust the retention policy for each log group, keeping the indefinite retention, or choosing a retention period between 10 years and one day. To view the allowed minimum retention period in AMS, see the AMS Technical Standards document available through AWS Artifact. To access AWS Artifact, contact your CSDM for instructions or go to Getting Started with AWS Artifact.

The CloudWatch Logs log retention feature deletes the log events in a stream based on retention policy. It doesn't delete log streams or log groups. For general information, see the *Amazon CloudWatch Logs User Guide* What is Amazon CloudWatch Logs?.

For information on customizing a log retention period, and to learn more, see Change Log Data Retention in CloudWatch Logs.

Enabling logging for supported services

Some services do not have logging enabled by default and require explicit enablement.

To enable logging for CloudFront, OpenSearch, Amazon RDS and Route53, submit an RFC with the Management | Other | Other | Create change type (ct-1e1xtak34nx76) with the following values, replacing *variables* as appropriate:

```
Subject: Enable logging for SERVICE_NAME
Description: Service ARN: SERVICE_ARN
```

Security management

AWS Managed Services (AMS) security management is the process by which AMS identifies an organization's assets and implements policies and procedures to protect those assets.

> **Note**
>
> AMS now has a change type (CT), Deployment | Advanced stack components | ACM certificate with additional SANs | Create (ct-3l14e139i5p50), that you can use to submit a request for an AWS Certificate Manager certificate. For information, see AWS::CertificateManager::Certificate. This CT provides for the creation of additional subject alternative name (SAN).

To better understand general AWS security, see Best Practices for Security, Identity, & Compliance.

AMS categorizes security risks as follows:

- Known risks detected by anti-malware, which the malware mitigation process handles.
- Security events including access breaches, which the security event management process handles.

Topics

- Data protection in AMS
- Identity and access management

Data protection in AMS

AMS continuously monitors your managed accounts by leveraging native AWS services such as Amazon GuardDuty, Amazon Macie (optionally), and other internal proprietary tools and processes. After an alarm is triggered, AMS assumes responsibility for the initial triage and response to the alarm. Our response processes are based on NIST standards. AMS regularly tests its response processes using Security Incident Response Simulation with you to align your workflow with existing customer security response programs.

When AMS detects any violation, or imminent threat of violation, of AWS or your security policies, we gather information, including impacted resources and any configuration-related changes. AMS provides 24/7/365 follow-the-sun support with dedicated operators actively reviewing and investigating monitoring dashboards, incident queue, and service requests across all of your managed accounts. AMS investigates the findings with our security experts to analyze the activity and notify you through the security escalation contacts listed in your account.

Based on our findings, AMS engages with you proactively. If you believe the activity is unauthorized or suspicious, AMS works with you to investigate and remediate or contain the issue. There are certain finding types generated by GuardDuty that require you to confirm the impact before AMS is able to take any action. For example, the GuardDuty finding type **UnauthorizedAccess:IAMUser/ConsoleLogin**, indicates that one of your users has logged in from an unusual location; AMS notifies you and asks that you review the finding to confirm if this behavior is legitimate.

Amazon Macie

We recommend, and AMS Accelerate supports, Macie to detect a large and comprehensive list of sensitive data, such as personal health information (PHI), personally identifiable information (PII), and financial data.

Macie can be configured to run periodically on any Amazon S3 bucket, automating the evaluation of any new or modified objects within a bucket over time. As security findings are generated, AMS will notify you and work with you to remediate as needed.

For more information, see Analyzing Amazon Macie findings.

Amazon Macie security

Macie is an artificial intelligence/AI powered security service that helps you prevent data loss by automatically discovering, classifying, and protecting sensitive data stored in AWS. Macie uses machine learning to recognize sensitive data such as personally identifiable information (PII) or intellectual property, assigns a business value, and provides visibility into where this data is stored and how it is being used in your organization. Macie continuously monitors data access activity for anomalies, and delivers alerts when it detects risk of unauthorized access or inadvertent data leaks. Macie service supports Amazon S3 and AWS CloudTrail data sources.

AMS continuously monitors for alerts from Macie and, if alerted, takes quick actions to protect your resources and account. With the addition of Macie to the list of services AMS supports, we are also now responsible for enabling and configuring Macie in all of your accounts, per your instructions. You can view Macie alerts and our actions as they unfold in the AWS console or

supported integrations. During account onboarding, you can indicate accounts that you use to store PII. For all new accounts with PII, we recommend using Macie. For existing accounts with PII, contact us and we will turn it on in your account. As a result, you can have an added layer of protection available and enjoy all the benefits of Macie in your AWS environment managed by AMS.

AMS Macie FAQs

- Why do I need Macie when all AMS accounts have Trend Micro and GuardDuty enabled?

  Macie helps you protect your data in Amazon S3 by helping you classify what data you have, the value that data has to the business, and the behavior associated with access to that data. Amazon GuardDuty provides broad protection of your AWS accounts, workloads, and data by helping to identify threats such as threat actor reconnaissance, instance issue, and problematic account activity. Both services incorporate user behavior analysis, machine learning, and anomaly detection to detect threats in their respective categories. Trend Micro does not focus on identifying PII and threats from them.
- How do I turn Macie on in my AMS account?

  If you have PII/PHI stored in your accounts or are planning to store it, contact your CSDM or raise a service request to enable Macie for your new or existing accounts managed by AMS.
- What are the cost implications of enabling Macie in my AMS account?

  Macie pricing works for AMS similar to other services such as Amazon Elastic Compute Cloud (Amazon EC2). You pay for Amazon Macie based on usage and an AMS uplift based on your SLAs. Macie fees are based on usage, see Amazon Macie Pricing, measured based on AWS CloudTrail events and Amazon S3 storage. Please note that Macie charges tend to flatten out from the second month after it's enabled because it charges based on incremental data added to Amazon S3 buckets.

To learn more about Macie, see Amazon Macie.

GuardDuty

GuardDuty is a continuous security monitoring service that uses threat intelligence feeds, such as lists of malicious IP addresses and domains, and machine learning to identify unexpected and potentially unauthorized and malicious activity within your AWS environment. This can include issues like escalations of privileges, uses of exposed credentials, or communication with malicious

IP addresses, or domains. GuardDuty also monitors Amazon Web Services account access behavior for signs of compromise, such as unauthorized infrastructure deployments, like instances deployed in a Region that has never been used, or unusual API calls, like a password policy change to reduce password strength. For more information, refer to the GuardDuty User Guide.

To view and analyze your GuardDuty findings, use the following procedure.

1. Open the GuardDuty console.
2. Choose **Findings**, and then choose a specific finding to view details. The details for each finding differ depending on the finding type, resources involved, and nature of the activity.

For more information on available finding fields, see GuardDuty finding details.

GuardDuty security

Amazon GuardDuty offers threat detection that enables you to continuously monitor and protect your AWS accounts and workloads. Amazon GuardDuty analyzes continuous streams of meta-data generated from your account and network activity found in AWS CloudTrail Events, Amazon VPC flow logs, and Domain Name System (DNS) logs. It also uses integrated threat intelligence such as known malicious IP addresses, anomaly detection, and machine learning to identify threats more accurately. GuardDuty is a monitored AMS service. To learn more about Amazon GuardDuty monitoring, see GuardDuty monitoring. To learn more about GuardDuty, see Amazon GuardDuty.

All new AMS accounts have GuardDuty enabled by default. AMS configures GuardDuty during account onboardings. You can submit change requests to modify the settings at any time. GuardDuty pricing works for AMS similarly to other services such as Amazon Elastic Compute Cloud (Amazon EC2). You pay for GuardDuty based on usage and an AMS uplift based on your SLAs. GuardDuty fees are based on usage (Amazon GuardDuty Pricing), measured based on AWS CloudTrail events and volume of your Amazon VPC Flow log.

For GuardDuty in AMS, the following primary detection categories are enabled:

- Reconnaissance -- Activity suggesting reconnaissance by a threat actor, such as unusual API activity, intra-VPC port scanning, unusual patterns of failed login requests, or unblocked port probing from a known bad IP.
- Instance issue -- Problematic instance activity, such as cryptocurrency mining, malware using domain generation algorithms (DGA), outbound denial of service activity, unusually high volume of network traffic, unusual network protocols, outbound instance communication with a known

malicious IP, temporary Amazon EC2 credentials used by an external IP address, and data exfiltration using DNS.

- Account activity -- Common patterns indicative of account activity include API calls from an unusual geolocation or anonymizing proxy, attempts to disable AWS CloudTrail logging, unusual instance or infrastructure launches, infrastructure deployments in an unusual AWS Region, and API calls from known malicious IP addresses.

AMS uses GuardDuty in your managed accounts to continuously monitor for findings and alerts from GuardDuty and, if alerted, AMS operations takes proactive actions to protect your resources and account. You can view GuardDuty findings and our actions as they unfold in the AWS console or supported integrations.

GuardDuty works with Trend Micro Deep Security Manager in your account. Trend Micro Deep Security Manager provides host-based Intrusion Detection / Intrusion Prevention services. Trend Micro Web Reputation services have some overlap with GuardDuty in the ability to detect when a host is attempting to communicate with a host or web service known to be a threat. However, GuardDuty provides additional threat detection categories and accomplishes this by monitoring network traffic, a method which is complementary to Trend Micro's host-based detection. Network-based threat detection allows for increased security by not allowing controls to fail if the host has been exhibiting problematic behavior. AMS recommends using GuardDuty in all your AMS accounts.

To learn more about Trend Micro, see Trend Micro Deep Security Help Center; note that non-Amazon links may change without notice to us.

GuardDuty monitoring

GuardDuty informs you of the status of your AWS environment by producing security findings that AMS captures and can alert on.

Amazon GuardDuty monitors the security of your AWS environment by analyzing and processing VPC flow logs, AWS CloudTrail event logs, and Domain Name System logs. You can expand this monitoring scope by configuring GuardDuty to also use your own custom, trusted IP lists, and threat lists.

- Trusted IP lists consist of IP addresses that you have allowed for secure communication with your AWS infrastructure and applications. GuardDuty does not generate findings for IP addresses on trusted IP lists. At any given time, you can have only one uploaded trusted IP list per AWS account per region.

- Threat lists consist of known malicious IP addresses. GuardDuty generates findings based on threat lists. At any given time, you can have up to six uploaded threat lists per AWS account per region.

To implement GuardDuty, use the AMS CT Deployment | Monitoring and notification | GuardDuty IP set | Create (ct-08avsj2e9mc7g) to create a set of approved IP addresses. You can also use the AMS CT Deployment | Monitoring and notification | GuardDuty threat intel set | Create (ct-25v6r7t8gvkq5) to create a set of denied IP addresses.

For a list of the services that AMS monitors, see What does the AMS monitoring system monitor?.

AWS Certificate Manager (ACM) certificate

AMS has a CT, Deployment | Advanced stack components | ACM certificate with additional SANs | Create (ct-3l14e139i5p50), that you can use to submit a request for an AWS Certificate Manager certificate, with up to five additional Subject alternative names (SAN) (such as example.com, example.net, and example.org). For details, see What Is AWS Certificate Manager? and ACM Certificate Characteristic.

> **Note**
>
> This timeout setting isn't just about the run, but also your validation of the ACM certificate through email validation. Without your validation, the RFC fails.

Data encryption in AMS

AMS uses several AWS services for data encryption, notably Amazon Simple Storage Service, AWS Key Management Service (AWS KMS), Amazon Elastic Block Store, Amazon Relational Database Service, Amazon Redshift, Amazon ElastiCache, AWS Lambda, and Amazon OpenSearch Service.

Amazon S3

Amazon S3 offers several object encryption options that protect data in transit and at rest. Server-side encryption encrypts your object before saving it on disks in its data centers and then decrypts it when you download the objects. As long as you authenticate your request and you have access permissions, there is no difference in the way you access encrypted or unencrypted objects. For more information, see Data protection in Amazon S3.

Amazon EBS

With Amazon EBS encryption, you don't need to build, maintain, and secure your own key management infrastructure. Amazon EBS encryption uses AWS KMS keys when creating encrypted volumes and snapshots. Encryption operations occur on the servers that host Amazon EC2 instances. This is done to make sure that both data-at-rest and data-in-transit between an instance and its attached Amazon EBS storage is secure. You can attach both encrypted and unencrypted volumes to an instance simultaneously. For more information, see Amazon EBS Encryption.

Amazon RDS

Amazon RDS can encrypt your Amazon RDS DB instances. Data that's encrypted at rest includes the underlying storage for DB instances, its automated backups, read replicas, and snapshots. Amazon RDS-encrypted DB instances use the industry standard AES-256 encryption algorithm to encrypt your data on the server that hosts your Amazon RDS DB instances. After your data is encrypted, Amazon RDS handles authentication of access and decryption of your data transparently with a minimal impact on performance. You don't need to modify your database client applications to use encryption. For more information, see Encrypting Amazon RDS resources.

Amazon Simple Queue Service

In addition to the default Amazon SQS managed server-side encryption (SSE) option, Amazon SQS-managed SSE (SSE-SQS) allows you to create custom managed server-side encryption that uses Amazon SQS-managed encryption keys to protect sensitive data that's sent over message queues. Server-side encryption (SSE) allows you to transmit sensitive data in encrypted queues. SSE protects the content of messages in queues using Amazon SQS-managed encryption keys (SSE-SQS) or keys that are managed in AWS KMS (SSE-KMS). For information about managing SSE using the AWS Management Console, see Encryption at rest.

Data encryption at rest

OpenSearch Service domains offer encryption of data at rest, a security feature that helps prevent unauthorized access to your data. The feature uses AWS Key Management Service (AWS KMS) to store and manage your encryption keys and the Advanced Encryption Standard algorithm with 256-bit keys(AES-256) to perform the encryption. For more information, see Encryption of Data at Rest for Amazon OpenSearch Service.

Key management

AWS KMS is a managed service that makes it easy for you to create and control customer master keys (CMKs), the encryption keys used to encrypt your data. AWS KMS CMKs are protected

by hardware security modules (HSMs) that are validated by the FIPS 140-2 Cryptographic Module Validation Program except in the China (Beijing) and China (Ningxia) Regions. For more information, see What is AWS Key Management Service?

Identity and access management

AWS Identity and Access Management (IAM) is a web service that helps you securely control access to AWS resources. You use IAM to control who is authenticated (signed in) and authorized (has permissions) to use resources. During AMS onboarding, you are responsible for creating cross-account IAM Admin roles within each of your managed accounts.

Multi-Account Landing Zone (MALZ) IAM safeguards

AMS multi-account landing zone (MALZ) requires an Active Directory (AD) trust as a primary design goal of AMS access management to allow each organization (both AMS, and customer) management of their own identities' life cycles. This avoids the need to have credentials in one another's directory. The one-way trust is configured, so that the Managed Active Directory within the AWS account trusts the customer owned or managed AD to authenticate users. Because the trust is only one way, it doesn't mean that the Managed AD is trusted by the Customer Active Directory.

In this configuration, the customer directory that manages user identities is known as the User Forest, and the Managed AD to which Amazon EC2 instances are attached is known as the Resource Forest. This is a commonly-leveraged Microsoft design pattern for Windows authentication; for more information, see Forest Design Models.

This model allows both organizations to automate their respective lifecycles and allows both AMS and you to rapidly revoke access if an employee leaves the organization. Without this model, if both organizations used a common directory (or created users/groups in one another's directories), then both organizations would have to put in additional workflows, and user syncs, to account for employees starting and leaving. This introduces risk as that process has latency and can be error-prone.

MALZ access pre-requisites

MALZ Identity Provider Integration for access to the AWS/AMS console, CLI, SDK.

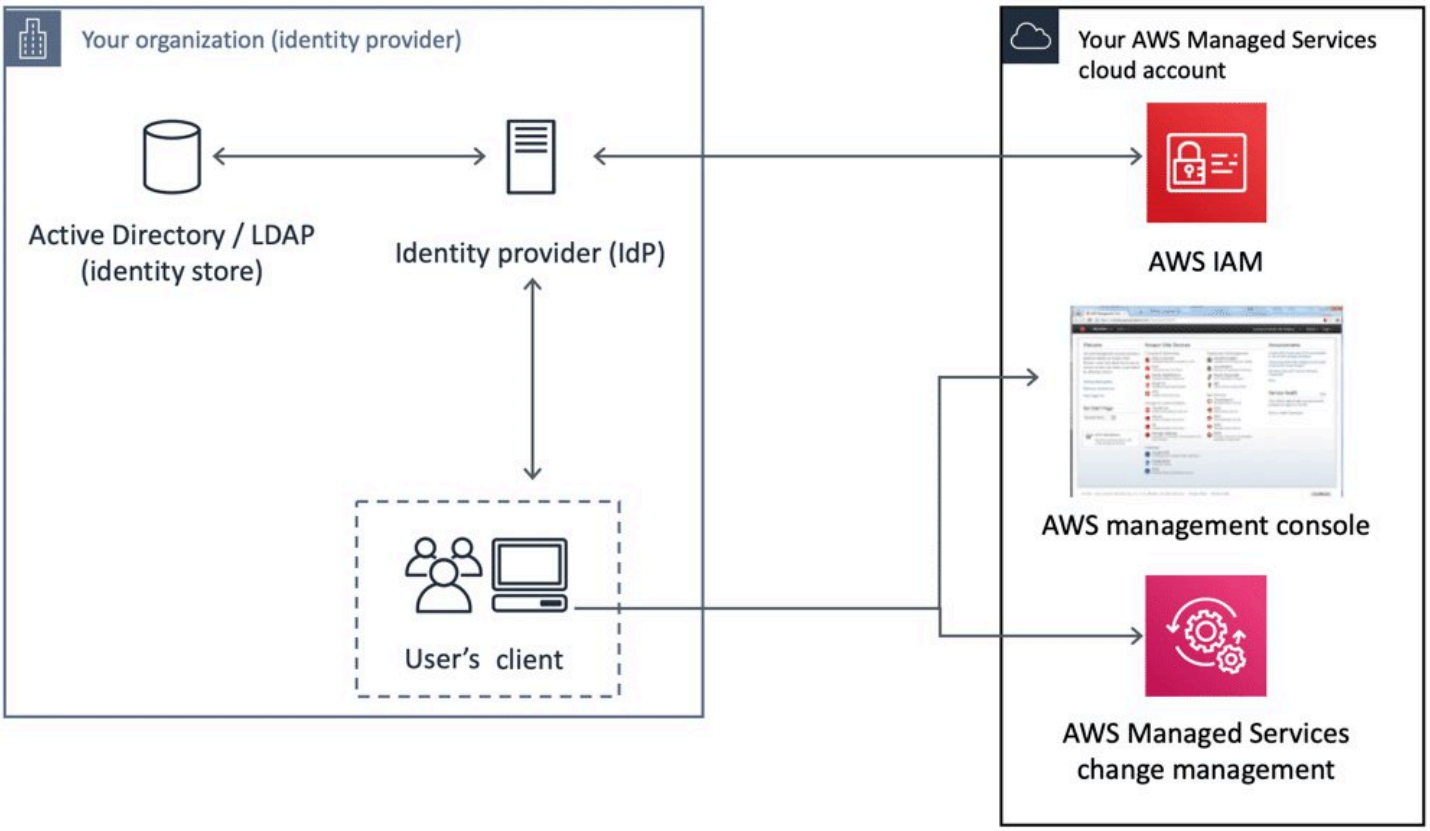

One-way trust for Amazon EC2 instances in your AMS account.

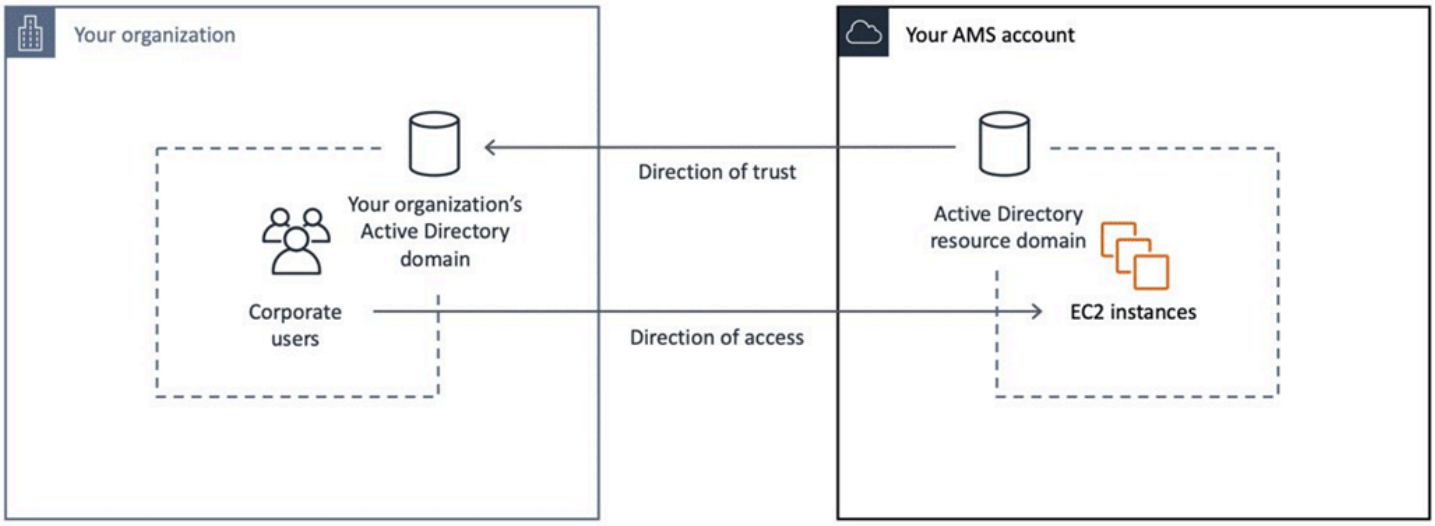

Authenticating with identities

AMS uses IAM roles, which is a type of IAM identity. An IAM role is very similar to a user, in that it is an identity with permission policies that determine what the identity can and cannot do in AWS. However, a role doesn't have credentials associated with it and, instead of being uniquely

associated with one person, a role is intended to be assumable by anyone who needs it. An IAM user can assume a role to temporarily take on different permissions for a specific task.

Access roles are controlled by internal group membership, which is administered and periodically reviewed by Operations Management.

IAM user role in AMS

An IAM role is similar to an IAM user, in that it is an AWS identity with permission policies that determine what the identity can and can't do in AWS. However, instead of being uniquely associated with one person, a role is intended to be assumable by anyone who needs it.

Currently there is one AMS default user role, `Customer_ReadOnly_Role`, for standard AMS accounts and an additional role, `customer_managed_ad_user_role` for AMS accounts with Managed Active Directory.

The role policies set permissions for CloudWatch and Amazon S3 log actions, AMS console access, read-only restrictions on most AWS services, restricted access to account S3 console, and AMS change-type access.

Additionally, the `Customer_ReadOnly_Role` has mutative, reserved-instances permissions that allow you to reserve instances. It has some cost-saving values, so, if you know that you're going to need a certain number of Amazon EC2 instances for a long period of time, you can call those APIs. To learn more, see Amazon EC2 Reserved Instances.

> **Note**
>
> The AMS service level objective (SLO) for creating custom IAM policies for IAM users is four business days, unless an existing policy is going to be reused. If you want to modify the existing IAM user role, or add a new one, submit an IAM: Update Entity or IAM: Create Entity RFC, respectively.

If you're unfamiliar with Amazon IAM roles, see IAM Roles for important information.

Multi-Account Landing Zone (MALZ): To see the AMS multi-account landing zone default, un-customized, user role policies, see MALZ: Default IAM User Roles, next.

MALZ: Default IAM User Roles

JSON policy statements for the default multi-account AMS multi-account landing zone user roles.

> **Note**
>
> The user roles are customizable and may differ on a per-account basis. Instructions on finding your role are provided.

These are examples of the default MALZ user roles. To make sure that you have the policies set that you need, run the AWS command `get-role` or sign in to the AWS Management -> IAM console and choose **Roles** in the navigation pane.

Core OU account roles

A core account is an MALZ-managed infrastructure account. AMS multi-account landing zone Core accounts include a management account and a networking account.

Core OU account: Common roles and policies

Role	Policy or policies
AWSManagedServicesReadOnlyRole	ReadOnlyAccess (Public AWS Managed Policy).
AWSManagedServicesCaseRole	ReadOnlyAccess
	AWSSupportAccess (Public AWS Managed Policy).
AWSManagedServicesChangeManagementRole (Core account version)	ReadOnlyAccess
	AWSSupportAccess
	AMSChangeManagementReadOnlyPolicy
	AMSChangeManagementInfrastructurePolicy

Core OU account: Management account roles and policies

Role	Policy or policies
AWSManagedServicesBillingRole	AMSBillingPolicy (AMSBillingPolicy).

Role	Policy or policies
AWSManagedServicesReadOnlyRole	ReadOnlyAccess (Public AWS Managed Policy).
AWSManagedServicesCaseRole	ReadOnlyAccess
	AWSSupportAccess (Public AWS Managed Policy).
AWSManagedServicesChangeManagementRole (Management account version)	ReadOnlyAccess
	AWSSupportAccess
	AMSChangeManagementReadOnlyPolicy
	AMSChangeManagementInfrastructurePolicy
	AMSMasterAccountSpecificChangeManagementInfrastructurePolicy

Core OU Account: Networking account roles and policies

Role	Policy or policies
AWSManagedServicesReadOnlyRole	ReadOnlyAccess (Public AWS Managed Policy).
AWSManagedServicesCaseRole	ReadOnlyAccess
	AWSSupportAccess (Public AWS Managed Policy).
AWSManagedServicesChangeManagementRole (Networking account version)	ReadOnlyAccess
	AWSSupportAccess
	AMSChangeManagementReadOnlyPolicy
	AMSChangeManagementInfrastructurePolicy

Role	Policy or policies
	AMSNetworkingAccountSpecificChangeManagementInfrastructurePolicy

Application Account Roles

Application account roles are applied to your application-specific accounts.

Application account: Roles and policies

Role	Policy or policies
AWSManagedServicesReadOnlyRole	ReadOnlyAccess (Public AWS Managed Policy).
AWSManagedServicesCaseRole	ReadOnlyAccess
	AWSSupportAccess (Public AWS Managed Policy). This policy provides access to all support operations and resources. For information, see Getting Started with AWS Support.
AWSManagedServicesSecurityOpsRole	ReadOnlyAccess
	AWSSupportAccess Example This policy provides access to all support operations and resources.
	AWSCertificateManagerFullAccess information, (Public AWS Managed Policy)
	AWSWAFFullAccess information, (Public AWS Managed policy). This policy grants full access to AWS WAF resources.
	AMSSecretsManagerSharedPolicy

Role	Policy or policies
AWSManagedServicesChangeManagementRole (Application account version)	ReadOnlyAccess
	AWSSupportAccess (Public AWS Managed Policy). This policy provides access to all support operations and resources. For information, see Getting Started with AWS Support.
	AMSSecretsManagerSharedPolicy
	AMSChangeManagementPolicy
	AMSReservedInstancesPolicy
	AMSS3Policy
AWSManagedServicesAdminRole	ReadOnlyAccess
	AWSSupportAccess
	AMSChangeManagementInfrastructurePolicy
	AWSMarketplaceManageSubscriptions
	AMSSecretsManagerSharedPolicy
	AMSChangeManagementPolicy
	AWSCertificateManagerFullAccess
	AWSWAFFullAccess
	AMSS3Policy
	AMSReservedInstancesPolicy

Policy Examples

Examples are provided for most policies used. To view the ReadOnlyAccess policy (which is pages long as it provides read-only access to all AWS services), you can use this link, if you have an active AWS account: ReadOnlyAccess. Also, a condensed version is included here.

AMSBillingPolicy

`AMSBillingPolicy`

The new Billing role can be used by your accounting department to view and change billing information or account settings in the Management account. To access information such as Alternate Contacts, view the account resources usage, or keep a tab of your billing or even modify your payment methods, you use this role. This new role comprises of all the permissions listed in the AWS Billing IAM actions web page.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": [
                "aws-portal:ViewBilling",
                "aws-portal:ModifyBilling"
            ],
            "Resource": "*",
            "Effect": "Allow",
            "Sid": "AllowAccessToBilling"
        },
        {
            "Action": [
                "aws-portal:ViewAccount",
                "aws-portal:ModifyAccount"
            ],
            "Resource": "*",
            "Effect": "Allow",
            "Sid": "AllowAccessToAccountSettings"
        },
        {
            "Action": [
                "budgets:ViewBudget",
                "budgets:ModifyBudget"
            ],
            "Resource": "*",
```

```
            "Effect": "Allow",
            "Sid": "AllowAccessToAccountBudget"
        },
        {
            "Action": [
                "aws-portal:ViewPaymentMethods",
                "aws-portal:ModifyPaymentMethods"
            ],
            "Resource": "*",
            "Effect": "Allow",
            "Sid": "AllowAccessToPaymentMethods"
        },
        {
            "Action": [
                "aws-portal:ViewUsage"
            ],
            "Resource": "*",
            "Effect": "Allow",
            "Sid": "AllowAccessToUsage"
        },
        {
            "Action": [
                "cur:DescribeReportDefinitions",
                "cur:PutReportDefinition",
                "cur:DeleteReportDefinition",
                "cur:ModifyReportDefinition"
            ],
            "Resource": "*",
            "Effect": "Allow",
            "Sid": "AllowAccessToCostAndUsageReport"
        },
        {
            "Action": [
                "pricing:DescribeServices",
                "pricing:GetAttributeValues",
                "pricing:GetProducts"
            ],
            "Resource": "*",
            "Effect": "Allow",
            "Sid": "AllowAccessToPricing"
        },
        {
            "Action": [
                "ce:*",
```

```
                "compute-optimizer:*"
            ],
            "Resource": "*",
            "Effect": "Allow",
            "Sid": "AllowAccessToCostExplorerComputeOptimizer"
        },
        {
            "Action": [
                "purchase-orders:ViewPurchaseOrders",
                "purchase-orders:ModifyPurchaseOrders"
            ],
            "Resource": "*",
            "Effect": "Allow",
            "Sid": "AllowAccessToPurchaseOrders"
        },
        {
            "Action": [
                "redshift:AcceptReservedNodeExchange",
                "redshift:PurchaseReservedNodeOffering"
            ],
            "Resource": "*",
            "Effect": "Allow",
            "Sid": "AllowAccessToRedshiftAction"
        },
        {
            "Action": "savingsplans:*",
            "Resource": "*",
            "Effect": "Allow",
            "Sid": "AWSSavingsPlansFullAccess"
        }
    ]
}
```

AMSChangeManagementReadOnlyPolicy

`AMSChangeManagementReadOnlyPolicy`

Permissions to see all AMS change types, and the history of requested change types.

```
{
 "Version": "2012-10-17",
 "Statement": [{
  "Sid": "AMSCoreAccountsCMAndSKMSReadOnlyAccess",
  "Effect": "Allow",
```

```
  "Action": [
   "amscm:GetChangeTypeVersion",
   "amscm:GetRfc",
   "amscm:ListChangeTypeCategories",
   "amscm:ListChangeTypeClassificationSummaries",
   "amscm:ListChangeTypeItems",
   "amscm:ListChangeTypeOperations",
   "amscm:ListChangeTypeSubcategories",
   "amscm:ListChangeTypeVersionSummaries",
   "amscm:ListRestrictedExecutionTimes",
   "amscm:ListRfcSummaries",
   "amsskms:GetStack",
   "amsskms:GetSubnet",
   "amsskms:GetVpc",
   "amsskms:ListAmis",
   "amsskms:ListStackSummaries",
   "amsskms:ListSubnetSummaries",
   "amsskms:ListVpcSummaries"
  ],
  "Resource": "*"
 }]
}
```

AMSMasterAccountSpecificChangeManagementInfrastructurePolicy

`AMSMasterAccountSpecificChangeManagementInfrastructurePolicy`

Permissions to request the Deployment | Managed landing zone | Management account | Create application account (with VPC) change type.

```
{
 "Version": "2012-10-17",
 "Statement": [{
  "Sid": "AMSMasterAccountAccess",
  "Effect": "Allow",
  "Action": [
   "amscm:ApproveRfc",
   "amscm:CancelRfc",
   "amscm:CreateRfc",
   "amscm:RejectRfc",
   "amscm:SubmitRfc",
   "amscm:UpdateRfc",
   "amscm:UpdateRfcActionState",
```

```
    "amscm:UpdateRestrictedExecutionTimes"
   ],
   "Resource": [
    "arn:aws:amscm:global:*:changetype/ct-1zdasmc2ewzrs:*"
   ]
  }]
}
```

AMSNetworkingAccountSpecificChangeManagementInfrastructurePolicy

`AMSNetworkingAccountSpecificChangeManagementInfrastructurePolicy`

Permissions to request the Deployment | Managed landing zone | Networking account | Create application route table change type.

```
{
 "Version": "2012-10-17",
 "Statement": [{
  "Sid": "AMSNetworkingAccountAccess",
  "Effect": "Allow",
  "Action": [
   "amscm:ApproveRfc",
   "amscm:CancelRfc",
   "amscm:CreateRfc",
   "amscm:RejectRfc",
   "amscm:SubmitRfc",
   "amscm:UpdateRfc",
   "amscm:UpdateRfcActionState",
   "amscm:UpdateRestrictedExecutionTimes"
  ],
  "Resource": [
   "arn:aws:amscm:global:*:changetype/ct-1urj94c3hdfu5:*"
  ]
 }]
}
```

AMSChangeManagementInfrastructurePolicy

`AMSChangeManagementInfrastructurePolicy` (for Management | Other | Other CTs)

Permissions to request the Management | Other | Other | Create, and Management | Other | Other | Update change types.

```
{
 "Version": "2012-10-17",
 "Statement": [{
  "Sid": "AMSCoreAccountsAccess",
  "Effect": "Allow",
  "Action": [
   "amscm:CancelRfc",
   "amscm:CreateRfc",
   "amscm:SubmitRfc",
   "amscm:UpdateRfc",
   "amscm:UpdateRfcActionState",
   "amscm:UpdateRestrictedExecutionTimes",
  ],
  "Resource": [
   "arn:aws:amscm:global:*:changetype/ct-1e1xtak34nx76:*",
   "arn:aws:amscm:global:*:changetype/ct-0xdawir96cy7k:*",
  ]
 }]
}
```

AMSSecretsManagerSharedPolicy

`AMSSecretsManagerSharedPolicy`

Permissions to view secret passwords/hashes shared by AMS through AWS Secrets manager (e.g. passwords to infrastructure for auditing).

Permissions to create secret password/hashes to share with AMS. (e.g. license keys for products that need to be deployed).

```
{
 "Version": "2012-10-17",
 "Statement": [{
   "Sid": "AllowAccessToSharedNameSpaces",
   "Effect": "Allow",
   "Action": "secretsmanager:*",
   "Resource": [
    "arn:aws:secretsmanager:*:*:secret:ams-shared/*",
    "arn:aws:secretsmanager:*:*:secret:customer-shared/*"
   ]
  },
  {
```

```
    "Sid": "DenyGetSecretOnCustomerNamespace",
    "Effect": "Deny",
    "Action": "secretsmanager:GetSecretValue",
    "Resource": "arn:aws:secretsmanager:*:*:secret:customer-shared/*"
   },
   {
    "Sid": "AllowReadAccessToAMSNameSpace",
    "Effect": "Deny",
    "NotAction": [
     "secretsmanager:Describe*",
     "secretsmanager:Get*",
     "secretsmanager:List*"
    ],
    "Resource": "arn:aws:secretsmanager:*:*:secret:ams-shared/*"
   }
  ]
}
```

AMSChangeManagementPolicy

`AMSChangeManagementPolicy`

Permissions to request and view all AMS change types, and the history of requested change types.

```
{
 "Version": "2012-10-17",
 "Statement": [{
  "Sid": "AMSFullAccess",
  "Effect": "Allow",
  "Action": [
   "amscm:*",
   "amsskms:*"
  ],
  "Resource": [
   "*"
  ]
 }]
}
```

AMSReservedInstancesPolicy

`AMSReservedInstancesPolicy`

Permissions to manage EC2 reserved instances; for pricing information, see Amazon EC2 Reserved Instances.

```
{
 "Version": "2012-10-17",
 "Statement": [{
  "Sid": "AllowReservedInstancesManagement",
  "Effect": "Allow",
  "Action": [
   "ec2:ModifyReservedInstances",
   "ec2:PurchaseReservedInstancesOffering"
  ],
  "Resource": [
   "*"
  ]
 }]
}
```

AMSS3Policy

`AMSS3Policy`

Permissions to create and delete files from existing S3 buckets.

> **Note**
>
> These permissions do not grant the ability to create S3 buckets; that must be done with the Deployment | Advanced stack components | S3 storage | Create change type.

```
{
 "Version": "2012-10-17",
 "Statement": [{
  "Effect": "Allow",
  "Action": [
   "s3:AbortMultipartUpload",
   "s3:DeleteObject",
   "s3:PutObject",
  ],
  "Resource": "*"
 }]
```

```
}
```

AWSSupportAccess

`AWSSupportAccess`

Full access to AWS Support. For information, see Getting Started with AWS Support. For Premium Support information, see AWS Support.

```
{
 "Version": "2012-10-17",
 "Statement": [{
  "Effect": "Allow",
  "Action": [
   "support:*"
  ],
  "Resource": "*"
 }]
}
```

AWSMarketplaceManageSubscriptions

`AWSMarketplaceManageSubscriptions` (Public AWS Managed Policy)

Permissions to subscribe, unsubscribe, and view AWS Marketplace subscriptions.

```
{
 "Version": "2012-10-17",
 "Statement": [{
  "Action": [
   "aws-marketplace:ViewSubscriptions",
   "aws-marketplace:Subscribe",
   "aws-marketplace:Unsubscribe"
  ],
  "Effect": "Allow",
  "Resource": "*"
 }]
}
```

AWSCertificateManagerFullAccess

`AWSCertificateManagerFullAccess`

Full access to AWS Certificate Manager. For Certificate Manager information, see AWS Certificate Manager.

AWSCertificateManagerFullAccess information, (Public AWS Managed Policy).

```
{
 "Version": "2012-10-17",
 "Statement": [{
  "Effect": "Allow",
  "Action": [
   "acm:*"
  ],
  "Resource": "*"
 }]
}
```

AWSWAFFullAccess

AWSWAFFullAccess

Full access to AWS Web Application Firewall (WAF). For WAF information, see AWS WAF - Web Application Firewall.

AWSWAFFullAccess information, (Public AWS Managed policy). This policy grants full access to AWS WAF resources.

```
{
 "Version": "2012-10-17",
 "Statement": [{
  "Action": [
   "waf:*",
   "waf-regional:*",
   "elasticloadbalancing:SetWebACL"
  ],
  "Effect": "Allow",
  "Resource": "*"
 }]
}
```

ReadOnlyAccess

ReadOnlyAccess (actions a-l only)

Read-only access to all AWS services and resources on the AWS console.

When AWS launches a new service, AMS updates the ReadOnlyAccess policy to add read-only permissions for the new service. The updated permissions are applied to all principal entities that the policy is attached to.

This doesn't grant the ability to log into EC2 hosts or database hosts.

```
{
 "Version": "2012-10-17",
 "Statement": [{
  "Action": [
   "aws-marketplace:ViewSubscriptions",
   "aws-marketplace:Subscribe",
   "aws-marketplace:Unsubscribe"
  ],
  "Effect": "Allow",
  "Resource": "*"
 }]
}
```

Single-Account Landing Zone (SALZ): To see the AMS single-account landing zone default, uncustomized, user role policies, see SALZ: Default IAM User Role, next.

SALZ: Default IAM User Role

JSON policy statements for the default AMS single-account landing zone user role.

> **Note**
>
> The SALZ default user role is customizable and may differ on a per-account basis. Instructions on finding your role are provided.

This is an example of the default SALZ user role, but to make sure that you have the policies set for you, run the AWS command `get-role` or sign in to the AWS Management -> IAM console at https://console.aws.amazon.com/iam/. In the IAM console, in the navigation pane, choose **Roles**.

The customer read-only role is a combination of multiple policies. A breakdown of the role (JSON) follows.

Managed Services Audit Policy:

```
{"Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "BasicConsoleAccess",
      "Effect": "Allow",
      "Action": [
        "aws-portal:View*",
        "ec2-reports:View*",
        "support:*"
      ],
      "Resource": [
        "*"
      ]
    },
    {
      "Sid": "AuditAccessToAWSServices",
      "Effect": "Allow",
      "Action": [
        "acm:Describe*",
        "acm:List*",
        "appstream:Get*",
        "autoscaling:Describe*",
        "cloudformation:Describe*",
        "cloudformation:Get*",
        "cloudformation:List*",
        "cloudformation:ValidateTemplate",
        "cloudfront:Get*",
        "cloudfront:List*",
        "cloudsearch:Describe*",
        "cloudsearch:List*",
        "cloudtrail:DescribeTrails",
        "cloudtrail:GetTrailStatus",
        "cloudtrail:LookupEvents",
        "cloudwatch:Describe*",
        "cloudwatch:Get*",
        "cloudwatch:List*",
        "codecommit:Get*",
        "codecommit:List*",
        "codedeploy:BatchGet*",
        "codedeploy:Get*",
        "codedeploy:List*",
        "codepipeline:Get*",
```

```
        "codepipeline:List*",
        "config:Describe*",
        "config:Get*",
        "datapipeline:Describe*",
        "datapipeline:EvaluateExpression",
        "datapipeline:GetPipelineDefinition",
        "datapipeline:ListPipelines",
        "datapipeline:ValidatePipelineDefinition",
        "directconnect:Describe*",
        "ds:Describe*",
        "dynamodb:Describe*",
        "dynamodb:List*",
        "ec2:Describe*",
        "ec2:Get*",
        "ecs:Describe*",
        "ecs:List*",
        "elasticache:Describe*",
        "elasticache:List*",
        "elasticbeanstalk:Check*",
        "elasticbeanstalk:Describe*",
        "elasticbeanstalk:List*",
        "elasticbeanstalk:RequestEnvironmentInfo",
        "elasticbeanstalk:RetrieveEnvironmentInfo",
        "elasticfilesystem:Describe*",
        "elasticloadbalancing:Describe*",
        "elasticmapreduce:Describe*",
        "elasticmapreduce:List*",
        "elastictranscoder:List*",
        "events:Describe*",
        "events:Get*",
        "events:List*",
        "guardduty:Get*",
        "guardduty:List*",
        "kinesis:Describe*",
        "kinesis:List*",
        "kms:List*",
        "lambda:Get*",
        "lambda:List*",
        "macie:Describe*",
        "macie:Get*",
        "macie:List*",
        "opsworks:Describe*",
        "opsworks:Get*",
        "rds:Describe*",
```

```
        "rds:Download*",
        "rds:List*",
        "redshift:Describe*",
        "redshift:View*",
        "route53:Get*",
        "route53:List*",
        "route53domains:CheckDomainAvailability",
        "route53domains:Get*",
        "route53domains:List*",
        "sdb:Get*",
        "sdb:List*",
        "ses:Get*",
        "ses:List*",
        "sns:Get*",
        "sns:List*",
        "sqs:Get*",
        "sqs:List*",
        "ssm:ListCommands",
        "ssm:ListCommandInvocations",
        "storagegateway:Describe*",
        "storagegateway:List*",
        "swf:Count*",
        "swf:Describe*",
        "swf:Get*",
        "swf:List*",
        "tag:get*",
        "trustedadvisor:Describe*",
        "waf:Get*",
        "waf:List*",
        "waf-regional:Get*",
        "waf-regional:List*"
      ],
      "Resource": [
        "*"
      ]
    },
    {
      "Sid": "AWSManagedServicesFullAccess",
      "Effect": "Allow",
      "Action": [
        "amscm:*",
        "amsskms:*"
      ],
      "Resource": [
```

```
        "*"
      ]
    }
  ]
}
```

Managed Services IAM ReadOnly Policy

```
{
  "Statement": [
    {
      "Action": [
        "iam:GenerateCredentialReport",
        "iam:GetAccountAuthorizationDetails",
        "iam:GetAccountPasswordPolicy",
        "iam:GetAccountSummary",
        "iam:GetCredentialReport",
        "iam:GetGroup",
        "iam:GetGroupPolicy",
        "iam:GetInstanceProfile",
        "iam:GetPolicy",
        "iam:GetPolicyVersion",
        "iam:GetRole",
        "iam:GetRolePolicy",
        "iam:GetUser",
        "iam:GetUserPolicy",
        "iam:ListAccountAliases",
        "iam:ListAttachedRolePolicies",
        "iam:ListEntitiesForPolicy",
        "iam:ListGroupPolicies",
        "iam:ListGroups",
        "iam:ListGroupsForUser",
        "iam:ListInstanceProfiles",
        "iam:ListInstanceProfilesForRole",
        "iam:ListMFADevices",
        "iam:ListPolicies",
        "iam:ListPolicyVersions",
        "iam:ListRolePolicies",
        "iam:ListRoles",
        "iam:ListSAMLProviders",
        "iam:ListUsers",
        "iam:ListVirtualMFADevices"
      ],
```

```
      "Effect": "Allow",
      "Resource": [
        "*"
      ],
      "Sid": "IAMReadOnlyAccess"
    },
    {
      "Action": [
        "iam:*"
      ],
      "Effect": "Deny",
      "Resource": [
        "arn:aws:iam::*:group/mc-*",
        "arn:aws:iam::*:group/mc_*",
        "arn:aws:iam::*:policy/mc-*",
        "arn:aws:iam::*:policy/mc_*",
        "arn:aws:iam::*:role/mc-*",
        "arn:aws:iam::*:role/mc_*",
        "arn:aws:iam::*:role/Sentinel-*",
        "arn:aws:iam::*:role/Sentinel_*",
        "arn:aws:iam::*:user/mc-*",
        "arn:aws:iam::*:user/mc_*"
      ],
      "Sid": "DenyAccessToIamRolesStartingWithMC"
    }
  ],
```

Managed Services User Policy

```
"Version": "2012-10-17"
}
 {
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowCustomerToListTheLogBucketLogs",
      "Effect": "Allow",
      "Action": [
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::mc-a*-logs-*"
      ],
```

```
      "Condition": {
        "StringLike": {
          "s3:prefix": [
            "aws/*",
            "app/*",
            "encrypted",
            "encrypted/",
            "encrypted/app/*"
          ]
        }
      }
    },
    {
      "Sid": "BasicAccessRequiredByS3Console",
      "Effect": "Allow",
      "Action": [
        "s3:ListAllMyBuckets",
        "s3:GetBucketLocation"
      ],
      "Resource": [
        "arn:aws:s3:::*"
      ]
    },
    {
      "Sid": "AllowCustomerToGetLogs",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject*"
      ],
      "Resource": [
        "arn:aws:s3:::mc-a*-logs-*/aws/*",
        "arn:aws:s3:::mc-a*-logs-*/encrypted/app/*"
      ]
    },
    {
      "Sid": "AllowAccessToOtherObjects",
      "Effect": "Allow",
      "Action": [
        "s3:DeleteObject*",
        "s3:Get*",
        "s3:List*",
        "s3:PutObject*"
      ],
      "Resource": [
```

```
        "*"
      ]
    },
    {
      "Sid": "AllowCustomerToListTheLogBucketRoot",
      "Effect": "Allow",
      "Action": [
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::mc-a*-logs-*"
      ],
      "Condition": {
        "StringEquals": {
          "s3:prefix": [
            "",
            "/"
          ]
        }
      }
    },
    {
      "Sid": "AllowCustomerCWLConsole",
      "Effect": "Allow",
      "Action": [
        "logs:DescribeLogStreams",
        "logs:DescribeLogGroups"
      ],
      "Resource": [
        "arn:aws:logs:*:*:log-group:*"
      ]
    },
    {
      "Sid": "AllowCustomerCWLAccessLogs",
      "Effect": "Allow",
      "Action": [
        "logs:FilterLogEvents",
        "logs:GetLogEvents"
      ],
      "Resource": [
        "arn:aws:logs:*:*:log-group:/aws/*",
        "arn:aws:logs:*:*:log-group:/infra/*",
        "arn:aws:logs:*:*:log-group:/app/*",
        "arn:aws:logs:*:*:log-group:RDSOSMetrics:*:*"
```

```
      ]
    },
    {
      "Sid": "AWSManagedServicesFullAccess",
      "Effect": "Allow",
      "Action": [
        "amscm:*",
        "amsskms:*"
      ],
      "Resource": [
        "*"
      ]
    },
    {
      "Sid": "ModifyAWSBillingPortal",
      "Effect": "Allow",
      "Action": [
        "aws-portal:Modify*"
      ],
      "Resource": [
        "*"
      ]
    },
    {
      "Sid": "DenyDeleteCWL",
      "Effect": "Deny",
      "Action": [
        "logs:DeleteLogGroup",
        "logs:DeleteLogStream"
      ],
      "Resource": [
        "arn:aws:logs:*:*:log-group:*"
      ]
    },
    {
      "Sid": "DenyMCCWL",
      "Effect": "Deny",
      "Action": [
        "logs:CreateLogGroup",
        "logs:CreateLogStream",
        "logs:DescribeLogStreams",
        "logs:FilterLogEvents",
        "logs:GetLogEvents",
        "logs:PutLogEvents"
```

```
      ],
      "Resource": [
        "arn:aws:logs:*:*:log-group:/mc/*"
      ]
    },
    {
      "Sid": "DenyS3MCNamespace",
      "Effect": "Deny",
      "Action": [
        "s3:*"
      ],
      "Resource": [
        "arn:aws:s3:::mc-a*-logs-*/encrypted/mc/*",
        "arn:aws:s3:::mc-a*-logs-*/mc/*",
        "arn:aws:s3:::mc-a*-logs-*-audit/*",
        "arn:aws:s3:::mc-a*-internal-*/*",
        "arn:aws:s3:::mc-a*-internal-*"
      ]
    },
    {
      "Sid": "ExplicitDenyS3CfnBucket",
      "Effect": "Deny",
      "Action": [
        "s3:*"
      ],
      "Resource": [
        "arn:aws:s3:::cf-templates-*"
      ]
    },
    {
      "Sid": "DenyListBucketS3LogsMC",
      "Action": [
        "s3:ListBucket"
      ],
      "Effect": "Deny",
      "Resource": [
        "arn:aws:s3:::mc-a*-logs-*"
      ],
      "Condition": {
        "StringLike": {
          "s3:prefix": [
            "auditlog/*",
            "encrypted/mc/*",
            "mc/*"
```

```
          ]
        }
      }
    },
    {
      "Sid": "DenyS3LogsDelete",
      "Effect": "Deny",
      "Action": [
        "s3:Delete*",
        "s3:Put*"
      ],
      "Resource": [
        "arn:aws:s3:::mc-a*-logs-*/*"
      ]
    },
    {
      "Sid": "DenyAccessToKmsKeysStartingWithMC",
      "Effect": "Deny",
      "Action": [
        "kms:*"
      ],
      "Resource": [
        "arn:aws:kms::*:key/mc-*",
        "arn:aws:kms::*:alias/mc-*"
      ]
    },
    {
      "Sid": "DenyListingOfStacksStartingWithMC",
      "Effect": "Deny",
      "Action": [
        "cloudformation:*"
      ],
      "Resource": [
        "arn:aws:cloudformation:*:*:stack/mc-*"
      ]
    },
    {
      "Sid": "AllowCreateCWMetricsAndManageDashboards",
      "Effect": "Allow",
      "Action": [
        "cloudwatch:PutMetricData"
      ],
      "Resource": [
        "*"
```

```
      ]
    },
    {
      "Sid": "AllowCreateandDeleteCWDashboards",
      "Effect": "Allow",
      "Action": [
        "cloudwatch:DeleteDashboards",
        "cloudwatch:PutDashboard"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

Customer Secrets Manager Shared Policy

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowSecretsManagerListSecrets",
      "Effect": "Allow",
      "Action": "secretsmanager:listSecrets",
      "Resource": "*"
    },
    {
      "Sid": "AllowCustomerAdminAccessToSharedNameSpaces",
      "Effect": "Allow",
      "Action": "secretsmanager:*",
      "Resource": [
        "arn:aws:secretsmanager:*:*:secret:ams-shared/*",
        "arn:aws:secretsmanager:*:*:secret:customer-shared/*"
      ]
    },
   {
      "Sid": "DenyCustomerGetSecretCustomerNamespace",
      "Effect": "Deny",
      "Action": "secretsmanager:GetSecretValue",
      "Resource": "arn:aws:secretsmanager:*:*:secret:customer-shared/*"
    },
    {
```

```
      "Sid": "AllowCustomerReadOnlyAccessToAMSNameSpace",
      "Effect": "Deny",
      "NotAction": [
        "secretsmanager:Describe*",
        "secretsmanager:Get*",
        "secretsmanager:List*"
      ],
      "Resource": "arn:aws:secretsmanager:*:*:secret:ams-shared/*"
    }
  ]
}
```

Customer Marketplace Subscribe Policy

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowMarketPlaceSubscriptions",
      "Effect": "Allow",
      "Action": [
        "aws-marketplace:ViewSubscriptions",
        "aws-marketplace:Subscribe"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

Security Incident Response in AMS

Security is the top priority at AWS Managed Services (AMS). AMS deploys resources and controls in your accounts to manage them. AWS has a shared responsibility model: AWS manages the security of the cloud, and you are responsible for security in the cloud. AMS protects your data and assets and helps keep your AWS infrastructure secure by using security controls and active monitoring for security issues. These capabilities help you establish a security baseline for applications running in the AWS Cloud. AMS collaborates with you through Security Incident Response to assess the effect, and then carry out containment and remediations based on best practice recommendations.

When a deviation from the baseline occurs, such as by a misconfiguration or a change in external factors, you need to respond and investigate. To successfully do so, you need to understand the basic concepts of Security Incident Response within your AMS environment. You must also understand the requirements to prepare, educate, and train cloud teams before security issues occur. It is important to know the controls and capabilities that you can use, prepare response plans for common security issues such as a user account compromise or a misuse of privileged accounts, and identify remediation methods that use automation to improve response speed and consistency. Additionally, you need to understand your compliance and regulatory requirements as they relate to building a Security Incident Response program to fulfill those requirements.

Security Incident Response can be complex, but by implementing an iterative approach you can simplify the process and allow the incident response team to keep asset stakeholders satisfied by providing early and continuous detection and response. In this guide, we provide you with the methodology that AMS uses for incident response, the AMS responsibility matrix (RACI), how you can be prepared for a security event, how to engage AMS during security incidents, and some of the incident response runbooks that AMS uses.

How AMS Security Incident Response works

AWS Managed Services aligns to the NIST 800-61 Computer Security Incident Handling Guide for Security Incident Response. By aligning to this industry standard, we provide a consistent approach to security event management and adhere to best practices in securing and responding to security incidents in your cloud.

Incident response lifecycle

When detection identifies and generates a security alert, or you request security assistance, the AWS Managed Services Operations team makes sure that there is a timely investigation, executes automations to perform data collection, triages and analyzes, informs you of the analysis, performs investigation and any containment activities, and then posts event analysis.

The data collection, triage, analysis, and containment activities performed during the incident response vary depending on the type of security event being investigated. Example Security Incident Response workflows for select scenarios are at the end of this document.

During incidents, AMS determines the correct course of action dynamically, which might result in documented steps being re-ordered or bypassed as appropriate to make sure that the right outcome occurs.

Prepare

As the threat landscape evolves, AMS continues to expand detection and response capabilities. As new detections are added, AMS incorporates the alerts from these new detections into the detection and response platform. AMS security responders are trained to investigate and partner with you throughout the Security Incident Response lifecycle.

Because of this partnership approach, it's important that your security and application teams are prepared to engage with AMS to handle security events as these events occur. This documentation explains what to expect during a security event and helps you prepare for rapid response when a security incident occurs.

This documentation uses the NIST 800-61 definition of an **event** as any observable occurrence in a system or network and an **incident** as a violation or imminent threat of violation of policies, acceptable use policies, or standard security practices.

Preparation checklist

Work through the following checklist with your AMS cloud solution delivery manager (CSDM) and AMS cloud architect (CA):

- Understand what workloads are running in which accounts.
- Understand what internal teams are responsible for the various workloads and tag them appropriately in the workloads.
- Maintain contact details internally for other teams who might be required during a security event investigation and for containment decisions.
- Confirm that security contacts are up to date and added to all managed AWS accounts. The contacts are managed on a per account basis.
- Know how to raise security incident to AMS, and be familiar with the severity and expected response times.

- Make sure that when security notifications are received, they are routed to the appropriate people and systems such as pagers or your security operations center.
- Understand what log sources are available to you, where these are stored in your accounts and who has access to them.
- Understand how to use CloudWatch Insights to Query Logs during investigations.
- Understand the containment options available to you by resource (EC2, IAM, S3, and son on) and the consequences on your workload availability when in containment.

RACI

Function	Customer	AMS
Communications		
Provide Customer Security contact details for AMS to use during Security Events notifications and Security Escalations	RA	I
Store and manage the supplied Customer Security contact details to use during Security Events and Security Escalations	CI	RA
Maintain Customer internal contact details, escalatio n matrix, responsible for managing workloads and systems to be engaged during a security incident in AMS accounts	RA	-
Provide contact information for the customers to reach out to AMS	I	RA

Function	Customer	AMS
Training		
Provide customer with documentation to support AMS during Incident Response Process	I	RA
Practice shared responsib ility during Incident response processes through security gamedays	RI	RAC
Resource management		
Maintain asset (AWS resources) inventory, and know the asset value and criticality of assets. This information is helpful during incident containment strategy	RA	CI
Employ AWS Tags to identify resources and workloads	RA	CI
Initiate and maintain periodical backups of AWS resources and AMIs	RCI	RA
Define and configure Log retention and archival	ACI	R
Secure baselining of AWS account, configura tions, policies and access management	ACI	RAC

Detect

During the management of your AWS accounts, AMS monitors for anomalies in user behavior, account activities and potential security events using data collected from detection sources and controls including but not limited to Amazon CloudWatch, Amazon GuardDuty, VPC Flow Logs, Amazon Macie, AWS Config and Amazon internal Threat Intelligence feeds.

AMS uses both native AWS services and other detection technologies to respond to security events created by:

- Config Conformance Finding Types
- GuardDuty Finding Types
- Macie Finding Types
- AMS Security events (cloud watch alarms)

Additional findings are added as services, products and threat ecosystems evolves.

Report security events to AMS

Raise an incident through the AMS Support Portal or AWS Support Center to notify AMS of a security incident or to request investigations.

RACI

Function	Customer	AMS
Logging, indicators and monitoring		
Enable and configure logging, monitoring to enable event management for the instance	CI	RA
Enable and configure logging, monitoring to enable event management for the Accounts	CI	RA

Function	Customer	AMS
Notify customer of events through outbound messaging	I	RA
Route notification and any subsequent updates to the decision makers for specific accounts and workloads to improve incident response time.	RA	CI
Deploy and maintain Endpoint Protection	CI	RA
Define, deploy, and maintain AMS standard detection services (for example, Amazon GuardDuty, Amazon Macie, and AWS Config)	I	RAC
Record AWS infrastructure change logs	I	RAC
Enable and configure logging, monitoring to enable event management for the applicati on	RA	C
Implement and maintain an allow-list, deny-list, and custom detections on AWS security services	RCI	RA

Function	Customer	AMS
Implement and maintain allow-list, deny-list, and custom detections on other security services supported by AMS	RACI	RA
Security event reporting		
Notify AMS of a suspicious activity or an active security investigation	RA	CI
Security event and incident notification to the customer	I	RA
Notification of planned event that might trigger Security Incident Response process	RA	CI

Analyze

After a security event is identified and reported, the next step is to analyze whether the reported event is a false positive or a real incident. AMS uses automation and manual investigative techniques to handle security events. The analysis includes investigation of logs from different detection sources such as network traffic logs, host logs,CloudTrail events, AWS service logs and so on. The analysis also looks for patterns that show an anomalous behavior by correlation.

Your partnership is required to understand context specific to the account environment and to establish what is normal for your account and workloads. This helps AMS identify an anomaly faster and to an accelerated incident response.

Handle communications from AMS about security events

AMS keeps you informed during the investigation by engaging your security contacts through an incident ticket. Your AMS cloud service delivery manager (CSDM) and AMS cloud architect (CA) are the point of contacts to reach out to for any communication during an active security investigation.

Communication includes automated notification when a security alert is generated, communication after event analysis, establishing call bridges and the ongoing delivery of artifacts such as log files, snapshot of infected resources, and getting investigation results to you during the security event.

Standard fields included in AMS security alert notifications are listed below. These fields provide you with information so that you can route events to the appropriate teams within your organization for remediation.

- Finding Type
- Finding Identifier (Where relevant)
- Finding Severity
- Finding Description
- Finding created Date & Time
- AWS Account Id
- Region (Where relevant)
- AWS Resources (IAM user/role/policy, EC2, S3, EKS)

Additional fields are provided depending on the Finding Type, for example EKS Findings include Pod, Container, and Cluster details.

RACI

Function	Customer	AMS
Investigation and analysis		
Perform initial response for any security alert generated by a detection source	I	RAC
False/True positive assessmen t based on available data	RAI	RAC
Generate a snap shot of affected instances to be shared with the customer if needed	I	RA

Function	Customer	AMS
Forensics tasks such as chain of custody, file system analysis, memory forensics, and binary analysis.	RA	CI
Application log collection to aid investigation	RA	I
Data and log collection to aid investigation on security alerts	RACI	RAC
Engage SMEs within AWS services on security investiga tions	CI	R
Engage third party vendors during investigation (for example, for EPS antimalwa re investigation and engaging with TrendMicro support team)	RACI	-
Share investigation logs to customers during an investiga tion	AI	R
Communication		
Alert and notifications from AMS detection sources for managed resources	I	RA
Alert and notifications on application security events	RA	I

Function	Customer	AMS
Engage customer leadership and C-Level executives during a security incident investiga tion	RA	I

Contain

AMS's approach to containment is partnership with you. You understand your business and the workload impacts that might occur from containment activities, such as network isolation, IAM user or role de-provisioning, instance re-building, and so forth.

An essential part of containment is decision-making. For example, shut down a system, isolate a resource from the network, or turn off access or end sessions. These decisions are easier to make if there are predetermined strategies and procedures to contain the incident. AMS provides the containment strategy and then implements the solution after you have considered the risk involved with implementing the containment actions.

There are different containment options depending on the resources under analysis. AMS expects multiple types of containment to be simultaneously deployed during an incident investigation. Some of these examples include:

- Apply protection rules to block unauthorized traffic (Security group, NACL, WAF Rules, SCP rules, Deny listing, setting signature action to quarantine or block)
- Resource Isolation
- Network Isolation
- Disabling IAM users, roles and policies
- Modifying/Reducing IAM user, role privilege
- Terminating / Suspending / Deleting compute resources
- Restricting public access from affected resource
- Rotating access keys, API keys, and passwords
- Scrubbing disclosed credentials and sensitive information

AMS encourages you to consider the type of containment strategies for each major incident type that is within their risk appetite, with criteria clearly documented to help with decision making in the event of an incident. Criteria to determine the appropriate strategy include:

- Potential damage to resources
- Preservation of evidence
- Service unavailability (for example, network connectivity, services provided to external parties)
- Time and resources needed to implement the strategy
- Effectiveness of the strategy (For example, partial containment, full containment)
- Permanence of the solution (For example, one-way door vs two-way door decisions)
- Duration of the solution (For example, emergency workaround to be removed in four hours, temporary workaround to be removed in two weeks, permanent solution).
- Apply security controls that you can turn on to lower the risk and allow time to define and implement a more effective containment.

The speed of containment is critical, AMS advises a staged approach to achieve efficient and effective containment by strategizing short-term and long-term approaches.

Use this guide to consider your containment strategy that involves different techniques based on the resource type.

- Containment Strategy
  - Can AMS identify the scope of the security incident?
    - If yes, identify all the resources (users, systems, resources).
    - If no, investigate in parallel with executing the next step on identified resources.
  - Can the resource be isolated?
    - If yes, then proceed to isolate the affected resources.
    - If no, then work with system owners and managers to determine further actions necessary to contain the problem.
  - Are all affected resources isolated from non-affected resources?
    - If yes, then continue to the next step.
    - If no, then continue to isolate affected resources until short-term containment is accomplished to prevent the incident from escalating further.
- System Backup

- Were backup copies of affected systems created for further analysis?
- Are the forensic copies encrypted and stored in a secure location?
  - If yes, then continue to the next step.
  - If no, encrypt the forensic images, then store them in a secure location to prevent accidental usage, damage, and tampering.

RACI

Function	Customer	AMS
Containment strategy and execution		
Decide on the execution of the agreed containment strategy and agree with the consequences that might affect the availability of services during the containment window	RA	ACI
Make a backup of affected systems for further analysis	CI	R
Contain applications and workloads (through application specific configuration or response activity)	RA	CI
Define the containment strategy based on the security incident and the affected resource	CI	RA
Encryption and secure storage of forensic copies of affected systems	CI	R

Function	Customer	AMS
Contain AWS Resources including EC2 Instance, Network and IAM containme nt.	CI	RA

Eradicate

After an incident is contained, eradication might be necessary to eliminate sources of threat altogether to secure the system before you proceed to the next recovery stage. Eradication steps might include deleting malware and removing compromised user accounts, as well as identifying and mitigating all vulnerabilities that were exploited. During eradication, it's important to identify all affected accounts, resources, and instances within the environment so that they can be remediated.

It's a best practice that eradication and recovery is done in a phased approach so that remediation steps are prioritized. For large-scale incidents, recovery might take months. The intent of the early phases must be to increase the overall security with relatively quick (days to weeks) high value changes to prevent future incidents. The later phases must focus on longer-term changes (for example, infrastructure changes) and ongoing work to keep the enterprise as secure as possible.

For some incidents, eradication is either not necessary or is performed during recovery.

Consider the following:

- Can the system be re-imaged and then hardened with patches or other countermeasures to prevent or reduce the risk of attacks?
- Are all malware and other artifacts left behind by the attackers removed and the affected systems hardened against further attacks?

RACI

Function	Customer	AMS
Eradication strategy and execution		

Function	Customer	AMS
Define eradication options based on the security incident and the affected resource on customer application workloads	RA	CI
Decide on the agreed eradication strategy, timing of eradication execution and the consequences.	RA	CI
Define eradication steps based on the security incident and the affected resource on AMS managed workloads	CI	RA
Eradicate and harden AWS resources including Amazon EC2instances, network, and IAM eradication.	CI	RA
Eradicate and harden applications and workloads (through application specific configuration or response activity)	RA	I

Recover

AMS partners with you to restore systems to normal operation, confirm that the systems are functioning normally, and (as applicable) remediate vulnerabilities to prevent similar incidents.

Consider the following:

- Are the affected system(s) patched and hardened against the recent attack and possible future attacks?

- What day and time is feasible to restore the affected systems back into production?
- What tools will you use to test, monitor, and verify that the systems that you restore to production aren't vulnerable to the initial attack techniques?

RACI

Function	Customer	AMS
Recovery preparation and execution		
Backing up AWS resources based on default AMS backup strategy or customer provided request.	CI	RA
Review backup plans to fully restore AMS managed workloads	CI	RA
Restore AMS managed resources from AWS Backup	CI	RA
Back up customer applicati on, APP configuration, and deployment settings	RA	I
Review backup plans to fully restore customer applications and workloads post-incident	RA	I
Restore applications and customer workloads (through application specific restorati on steps)	RA	I

Post Incident Report

Post event, AMS runs an investigation review process for all security incidents. And, AMS initiates a correction of error (COE) process to address security incidents caused by a system or a procedural miss that plausibly has room for improvement. AMS partners with you to continuously-improve security investigation experience. The COE process helps AMS identify the contributing factors of customer-impacting events and connects those causes to next actions items that can prevent similar events from recurring, or helps mitigate the duration or level of impact.

The investigation review process for security incidents addresses the following items to identify opportunities for improvement:

- What was the elapsed time from the beginning of the incident to incident discovery, to the initial impact assessment, and to each stage of the incident handling process (for example, containment, recovery)?
- How long did it take the incident response team to respond to the initial report of the incident?
- How long did it take to do an initial impact analysis?
- Was this preventable and how? Is there a tool or process that could have prevented this?
- Could we have detected this sooner and how?
- What could have made the investigation go faster?
- Were the documented Incident Response Procedures followed? Were they adequate?
- Was the information sharing with other stakeholders done in a timely manner How could it be improved?
- Was the collaboration with other teams (AWS Security, account teams, AWS Development team and customer security team's) effective? If not, what could be improved?
- What preparation steps were missing that might have helped, escalation matrices, RACI’s, shared responsibility models, and so on? Is there a need to update any Runbooks?
- What was the difference between the initial impact assessment and the final impact assessment? What can we do to improve accuracy of assessments earlier in the incident response?
- What are the Action Items from the Lessons Learned?

RACI

Function	Customer	AMS
Post incident reporting		
Conduct post incident analysis and creat action items as appropriate	-	RA
Prepare a correction of error (COE) document and perform internal reviews within AMS	-	R
Share appropriate lessons learned and action items with customer post incident as required	I	R

Security Incident Response Runbooks in AMS

This section contains two runbooks:

- Response to root user activity
- Response to malware events

Response to root user activity

The root user is the superuser within your AWS account. Note that AMS monitors root usage. It's a best practice to use the root user only for the few tasks that require it, such as to change your account settings, activate AWS Identity and Access Management (IAM) access to billing and cost management, change your root password, and turn on multi-factor authentication (MFA). For more information, see Tasks that require root user credentials.

For more information on how to inform AMS of planned root usage, see When and how to use the root account in AMS.

When root user activity is detected, either failed attempts to login that might indicate a brute force attack or activity in the account after a successful login, an event generates and an incident sent to your defined security contacts.

AWS Managed Services Operations investigates unplanned root user activity, perform data collection, triage and analysis, and perform containment activities at your direction, followed by post event analysis.

If you have the AMS Advanced operating model, you receive additional communications from AMS CSDM and AMS Ops engineers that confirm unplanned root user activity due AMS's responsibility to secure root user credentials. AMS investigates root user activity until you confirm a path forward.

Prepare

Advise AMS of any planned use of root user by submitting an AMS service request with data and times of planned event to prevent unnecessary incident response activities.

Periodically conduct GameDays with AMS to validate AMS's customer incident response processes, people and systems are current, and build muscle memory with responsible individuals to achieve faster incident response.

Phase A: Detect

AMS monitors for root activity in the accounts through detection sources including GuardDuty and AMS monitoring.

If you have AMS Accelerate, the operating model responds to the incident requesting investigation for unexpected root user activity. When this occurs, AMS Operations initiates the Compromised Account runbook.

If you have AMS Advanced, the operating model responds to the incident, or informs the CSDM of any planned root user activity to terminate an active Account Compromise investigation.

Phase B: Analyze

AMS performs a thorough investigation of the root user events when it's determined that the activity isn't authorized. Using both automations and AMS security response team, logs and events are analyzed for anomalies and unexpected behavior for root users. Logs are provided to you to help determine if the activity is unknown, or if it's an authorized root user event, or if it requires further investigation.

Some examples of the information provided during the investigation to support internal checks includes:

- Account information: What account was the root account used on?
- E-mail address for root user: Each root user is associated with an e-mail address from your organization
- Authentication details: Where and when did the root user access your environment from?
- Activity records: What did the user do when logged in as root? These records are in the form of CloudWatch events. Understanding how to read these logs aids in investigation.

It's a best practice that you are prepared to receive the analysis information and have a plan for how to reach authorized points of contact for accounts within your organization. Because root users aren't named as individuals, determining who has access to the root e-mail address used for the account within your organization helps to quickly route questions internally.

Phase C: Contain and Eradicate

AMS partners with your security teams to perform containment at the direction of your authorized Customer Security contacts. Containment options include:

- Rotating appropriate credentials and keys.
- Terminating active sessions to accounts and resources.
- Eradicating resources created.

During the containment activities AMS works closely with your security team to ensure any disruption to your workloads are minimized and the root credentials are appropriately secured.

After the containment plan is completed, you work with AMS Operations team for any recovery actions as required.

Post Incident Report

As required, AMS initiates the investigation review process to identify any lessons learned. As part of completing a COE, AMS communicates any relevant findings to affected customers to help them improve their incident response process.

AMS documents all final details of the investigation, collects appropriate metrics, and then reports the incident to any AMS internal teams that require information, including your assigned CSDM and CA.

Response to malware events

Amazon EC2 instances are used to host a variety of workloads including third-party software and custom-developed software deployed by application teams within organizations. AMS provides and encourages you to deploy your workloads on images that are patched and maintained on an ongoing basis by AMS.

During the operation of instances, AMS monitors for anomalies in behavior or activity through a variety of security detection controls, including Amazon GuardDuty, Endpoint Protection, Network Traffic, and Amazon internal Threat Intelligence feeds.

AMS customers with the AMS Advanced operating model automatically have the endpoint security (EPS) monitoring client installed on provisioned resources. This makes sure that the resources are monitored and supported 24x7, including the creation of a security incident when an event is detected.

AMS also monitors GuardDuty Malware Findings. These are available on both AMS Advanced and AMS Accelerate, if enabled. See Malware Protection in Amazon GuardDuty for more information.

> **Note**
>
> If you opted for Bring Your Own EPS, then the process for incident response differs from what's outlined on this page. For more information, see the referenced documentation.

When malware is detected, an incident is created and you are notified of the event. This notification is followed by any remediation activities that occurred. AMS Operations investigates, performs data collection, triage and analysis, and then performs containment activities at your direction, followed by post event analysis.

Phase A: Detect

AMS monitors for events on instances with GuardDuty and end point security solution monitoring. AMS determines the appropriate enrichment and triage activities to help you make containment or risk acceptance decisions based on the finding or alert type.

Data collection is performed based on the finding type. Data collection involves querying multiple data sources both inside and outside of the affected account to build a picture of the activity observed or the configurations of concern.

AMS performs correlation of the finding with any other alarms and alerts or telemetry from any impacted accounts or AMS threat intelligence platforms.

Phase B: Analyze

After data is collected, it's analyzed to identify any activity or indicators of concern. During this phase of the investigation, AMS partners with you to integrate business and domain knowledge of the instances and workloads to help understand what's expected and what's out of the ordinary.

Some examples of the information provided during the investigation to support internal checks includes:

- Account Information: What account was the malware activity observed on?
- Instance Details: What instance(s) are implicated with the malware events?
- Event timestamp: When did the alert trigger?
- Workload Information: What is running on the instance?
- Malware details, if relevant: Families of malware and Open Source information about the malware.
- Users or Role Details: What users or roles are affected by and involved in the activity?
- Activity Records: What activities are recorded on the instance? These are in the form of CloudWatch events, and system events from the instance. Understanding how to read these logs will aid you in investigation
- Network Activity: What endpoints are connecting to the instance, what the instance is connecting to, and what is the traffics analysis?

It's a best practice to be prepared to receive investigation information, and have a plan about how to contact the appropriate points of contact for accounts, instances and workloads within your organization. Understanding your network topology and expected connection can help accelerate impact analysis. Knowledge of planned penetration testing in the environment and recent deployments performed by application owners can also speed up the investigation.

If you determine that the activity is planned and authorized, then the incident is updated and the investigation ends. If compromise is confirmed, then you and AMS determine the appropriate containment plan.

Phace C: Contain and Eradicate

AMS partners with you to determine appropriate containment activities based on the data collected and information known. Containment options include but are not limited to:

- Preserving data through snapshots
- Modifying network rules to limit traffic in or out of instances
- Modifying SCP, IAM user and role policies to limit access
- Terminating, Suspending or Turning off Instances
- Terminating any persistent connections
- Rotating appropriate credentials/keys

If you opt to perform eradication activity against the instance, then AMS supports you in achieving this. Options include, but are not limited to:

- Removing any unwanted software
- Rebuilding the instance from a clean fully patched image and redeploying applications and configuration
- Restoring the instance from a previous backup
- Deploying applications and services on to another instance within your account that might be suitable to host the workloads.

It's important to determine how the malware was delivered and run on the instance before restoration of service to make sure that any additional controls are applied to prevent reoccurrence of the malware on the instance. AMS provides additional insights or information to your forensics partners or teams as necessary to support forensics.

At this point, you work with AMS Operations for the recovery activities. AMS works closely with you to minimize disruption to the workloads and secure the instances.

Post Incident Report

As required, AMS initiates the investigation review process to identify lessons learned. As part of completing a COE, AMS communicates relevant findings to you to help you improve your incident response process.

AMS documents the final details of the investigation, collects appropriate metrics, and reports the incident to AMS internal teams that require information, including your assigned CSDM and CA.

Security event logging and monitoring

AMS continuously monitors the managed environment for security threats. Security events might be detected by AMS or by you. AMS regularly updates its automation process—based on the Computer Security Incident Handling Guide by the National Institute of Standards and Technology (NIST)—to better detect security threats.

Endpoint Security (EPS)

Resources that you provision in your AMS Advanced environment automatically include the installation of an endpoint security (EPS) monitoring client. This process ensures that the AMS Advanced-managed resources are monitored and supported 24x7. In addition, AMS Advanced monitors all agent activity, and an incident is created if any security event is detected.

> **Note**
>
> Security incidents are handled as incidents; for more information, see Incident response.

Endpoint security provides anti-malware protection, specifically, the following actions are supported:

- EC2 instances register with EPS
- EC2 instances deregister from EPS
- EC2 instances real-time anti-malware protection
- EPS agent-initiated heartbeat
- EPS restore quarantined file
- EPS event notification
- EPS reporting

AMS Advanced uses Trend Micro for endpoint security (EPS). These are the default EPS settings. To learn more about Trend Micro, see the Trend Micro Deep Security Help Center; note that non-Amazon links may change without notice to us.

AMS Advanced Multi-Account Landing Zone (MALZ) default settings are described in the following sections; for non-default AMS multi-account landing zone EPS settings, see AMS Advanced Multi-Account Landing Zone EPS non-default settings.

> **Note**
>
> You can bring your own EPS, see AMS bring your own EPS.

General EPS settings

Endpoint security general network settings.

EPS defaults

Setting	Default
Firewall Ports (Instances' Security Group)	EPS Deep Security Manager agents (DSMs) must have port 4120 open for the Agent/Relay to Manager communication, and port 4119 for the Manager Console. EPS Relays must have port 4122 open for the Manager/Agent to Relay communication. No specific ports should be open for customer instance inbound communication because agents initiate all requests.
Communication Direction	Agent/Appliance Initiated
Heartbeat Interval	Ten minutes
Number of missed heartbeats before an alert	Two
Maximum allowed drift (difference) between server times	Unlimited

Setting	Default
Raise offline errors for inactive (registered, but not online) virtual machines	No
Default policy	Base policy (described next)
Activation of multiple computers with the same host name	Is allowed
Alerts for pending updates are raised	After seven days
Update source	Trend Micro Update Server (https://ipv6-iaus.trendmicro.com/iau_server.dll/)
Event or log data deletion	Events and logs are deleted from the DSM database after seven days.
Agent software versions are held	Up to five
Most recent rule updates are held	Up to ten
Logs storage	By default, log files are stored securely in Amazon S3, but you can also archive them to Amazon Glacier to help meet audit and compliance requirements.

Base policy

Endpoint security base policy default settings.

EPS base policy

Setting	Default
Enabled Modules	Anti-Malware
Disabled Modules	Web Reputation
	Firewall

Setting	Default
	Intrusion Protection
	Integrity Monitoring
	Log Inspection
	Application Control

Anti-malware

Endpoint security anti-malware settings.

EPS anti-malware defaults

Setting	Default	Notes
Real-Time Scan	Scan everything	Quarantine all suspected viruses. Enable IntelliTrap and spyware/grayware protection. Spyware and Grayware trigger Anti-Malware and result in a quarantine of the item.
	Every Day/All Day (24 hours)	
Manual Scan	Scan everything	Must be requested, then follows default real-time scan configuration.
Scheduled Scan	Scan everything	Set for the last Sunday of every month, 6am.

Setting	Default	Notes
Smart Protection	Disabled	N/A
Quarantined Files	Trend Micro Deep Security Manager (DSM)	Appx 1GB of disk reserved for quarantine.
Scan Limitation	Trend Micro DSM	Scan files of all sizes.
Allowed Spyware or Grayware	None	N/A
Local Event Notification	Yes	N/A

Malware mitigation process

AMS uses Trend Micro's Deep Security Platform (anti-malware system) to detect and respond to malware on your AMS-managed instances. By default, the Trend Micro detection agent runs on all Amazon EC2 instances, including those in the shared services and private subnets, for both Windows and Linux operating systems. The anti-malware system is connected to AMS monitoring so that an event is generated whenever malware is detected. If there is customer impact, the event is escalated to the incident management process (for details, see AMS incident response). While AMS assesses the impact, you are notified, and attempts are made to mitigate the impact.

Trend Micro anti-malware definitions are updated automatically when Trend Micro publishes updates.

During application onboarding, you indicate the action you want AMS to take when malware is found on an instance:

- Make sure the quarantined file is on the allow list, removing it from the quarantine and releasing it back to the file system.
- Delete the quarantined file, removing it from the instance.
- Suspend the instance and replace it. The suspended instance is then available to you to mount for forensic research.

After application onboarding:

- When the anti-malware system discovers malware on an instance, AMS automatically quarantines the malware. This triggers an event and a follow-up investigation.
- AMS notifies you of the event through a service notification and starts following the default mitigation action that you selected.
- If you haven't chosen a default action, AMS asks you which action to take. After receiving your instructions, AMS runs the selected action and notifies you. AMS notifies you again after the action is complete, including details needed for forensic analysis, if applicable.

Enable IDS and IPS in Trend Micro Deep Security

You can request that AMS enable Trend Micro Intrusion Detection System (IDS) and Intrusion Protection Systems (IPS), non-default features, for your account.

To do this, submit an update request (Management | Other | Other | Update) and include a list of email addresses to receive IDS and IPS notifications. These addresses are added to an SNS topic in your account, which AMS creates for you.

> **Note**
>
> AMS cannot add any Trend Micro service that might interfere with our ability to provide other AMS services.

Full system malware scans

The Payment Card Industry Data Security Standard (PCI DSS) requires full system malware scans, which are enabled on your AMS-managed VPC by default. Full system scans are set to occur at 2AM (on the time zone set on the server) because they use a lot of CPU. Full system scans are in addition to regular malware scans that do not use a lot of CPU.

There is a new Management change type (CT), **Disable malware scans**, that allows you to disable full system malware scans. You can find the CT in the Management | Host security | Full system scan | Disable classification, change ID ct-1pybwg08h8qsz. To re-enable scans, use the Management | Other | Other | Update CT. Disabling full system scans does not disable your regular malware scans.

Amazon Inspector security

The Amazon Inspector service monitors the security of your AMS-managed stacks. Amazon Inspector is an automated security assessment service that helps identify gaps in the security and compliance of infrastructure deployed on AWS. Amazon Inspector security assessments enable you to automatically assess stacks for exposure, vulnerabilities, and deviations from best practices by checking for unintended network accessibility and vulnerabilities in your Amazon EC2 instances. After performing an assessment, Amazon Inspector produces a detailed list of security findings prioritized by level of severity. Amazon Inspector assessments are offered as pre-defined rules packages mapped to common security best practices and definitions. These rules are regularly updated by AWS security researchers. For more information about Amazon Inspector go to Amazon Inspector.

AMS Amazon Inspector FAQs

- Is Amazon Inspector installed to my AMS accounts by default?

  No. Amazon Inspector is not part of the default AMI build or workload ingestion.

- How do I access and install Amazon Inspector?

  Submit an RFC (Management | Other | Other | Create) to request account access and installation to Inspector and the AMS operations team will modify the Customer_ReadOnly_Role to provide Amazon Inspector console access (without SSM access).

- Does the Amazon Inspector Agent have to be installed on all of the Amazon EC2 instances I want to assess?

  No, Amazon Inspector assessments with the network reachability rules package can be run without an agent for any Amazon EC2 instances. The agent is required for host assessment rules packages. For more information about agent installation, see Installing Amazon Inspector Agents.

- Is there an additional cost for this service?

  Yes. Amazon Inspector pricing can be found on the Amazon Inspector pricing site.

- What are Amazon Inspector findings?

  Findings are potential security issues discovered during the Amazon Inspector assessment of the selected assessment target. Findings are displayed in the Amazon Inspector console or the API,

and contain both a detailed description of the security issues and recommendations for resolving them.

- Are reports of the Amazon Inspector assessment available?

  Yes. An assessment report is a document that details what is tested in the assessment run, and the results of the assessment. The results of your assessment are formatted into standard reports, which can be generated to share results within your team for remediation actions, to enrich compliance audit data, or to store for future reference. An Amazon Inspector assessment report can be generated for an assessment run once it has been successfully completed.

- Can I use tags to identify the stacks I want to run Amazon Inspector reports against?

  Yes.

- Will AMS Operations teams have access to the Amazon Inspector assessment results?

  Yes. Anyone with access to the Amazon Inspector console in AWS is able to view findings and assessment reports.

- Will AMS Operations teams recommend or take action based on the findings of the Amazon Inspector reports?

  No. If you want changes made based on the findings of the Amazon Inspector report, you must request changes through an RFC (Management | Other | Other | Update).

- Will AMS be notified when I run an Amazon Inspector report?

  When you request Amazon Inspector access, the AMS Operator running the RFC notifies your CSDM of the request.

For more information, see Amazon Inspector FAQs.

AMS incident response

AMS uses traditional IT service management (ITSM) incident management best practices to restore service, when needed, as quickly as possible.

We provide 24/7/365 follow-the-sun support through multiple operations centers around the world with dedicated operators actively monitoring dashboards and incident queues.

Our operations engineers use internal incident tracking tools to identify, log, categorize, prioritize, diagnose, resolve, and close incidents and provide updates on all of these activities to you through

the AMS console or through the AWS Support API. Our operators, many of whom have spent time in AWS Premium Support in various technology profiles and roles, leverage a variety of internal AWS Support tools to help with all of those activities. These operators are deeply familiar with AMS supported infrastructures and have expert level technical skills to address all identified support issues. In the rare case where our operators need assistance, the Premium Support and AWS Service teams are available to assist as needed.

In cases where High priority incidents are impacting your critical workloads, AMS will recommend an infrastructure restore. There is often a tradeoff between troubleshooting an issue or restoring from a known good backup, and customer risks and impacts from service downtime are the deciding factors. If you have time to devote to troubleshooting issues, AMS will assist you, but if the urgency to restore is high, we can initiate a restore right away.

> **Note**
>
> Ephemeral data that is not part of the stack template or data restore is lost. AMS uses reasonable efforts to perform infrastructure restore while AWS service offerings are unavailable. Infrastructure restore is completed once AWS service offerings are available. If you don't authorize an infrastructure restore as recommended by AMS, you won't be eligible for a service credit for the AMS service commitment for incident resolution time.

Compliance validation

AMS deploys and manages a library of AWS Config rules and remediation actions, to protect against misconfigurations that could reduce the security and operational integrity of your accounts.

As an example, when an Amazon S3 bucket is created, AWS Config can evaluate the Amazon S3 bucket against a rule that requires Amazon S3 buckets to deny public read access. If the Amazon S3 bucket policy or bucket access control list (ACL), allows public read access, AWS Config flags both the bucket and the rule as noncompliant. These AWS Config Rules mark resources as either Compliant, Noncompliant, or Not Applicable, based on the result of their evaluation. For more information about AWS Config service, see the AWS Config Developer Guide.

You can use the AWS Config console, AWS CLI, or AWS Config API to view the rules deployed in your account and the compliance state of your rules and resources. For more information, see the AWS Config documentation: Viewing Configuration Compliance.

> **Note**
>
> Additional information on this topic is available by accessing AWS Artifact reports and downloading the *AWS Managed Services (AMS) Customer Security Guide*. For more information, see Downloading reports in AWS Artifact. To access AWS Artifact, you can contact your CSDM for instructions or go to Getting Started with AWS Artifact. This information is not included in this user guide because it contains sensitive security content.

Multi-Account Landing Zone viewing the compliance status of your AWS Config Rules

AMS multi-account landing zone utilizes the AWS Config aggregator service to create a centralized view of compliance across all your accounts. This means you can see the compliance status of all AWS Config Rules across your AMS multi-account landing zone environment under the AWS Config aggregator in your security account.

The following is a sample of the AWS Config aggregator showcasing central compliance status of AWS Config Rules across accounts.

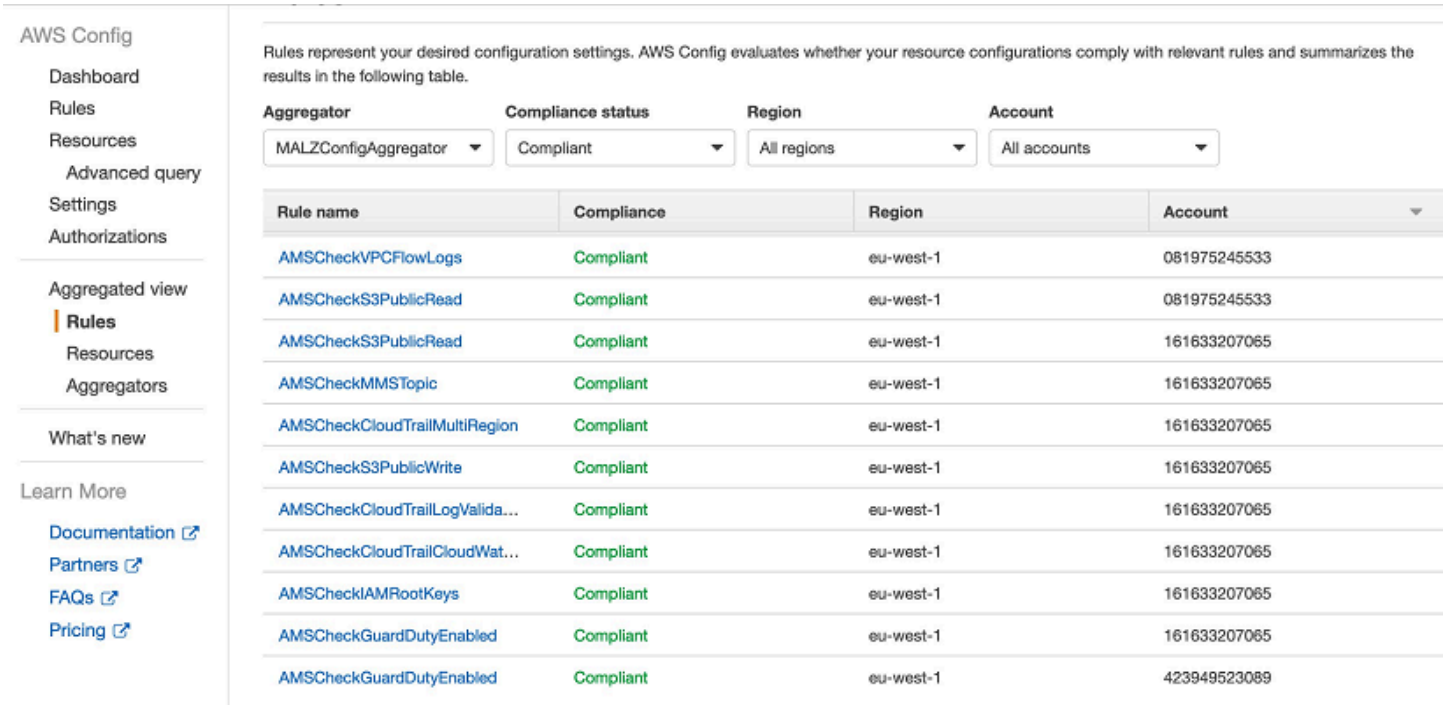

For more information, see the AWS documentation for Config Aggregator.

- How does AMS use AWS Config rules?

AMS creates AWS Config Rules to give visibility into the configuration of your AWS resources against conditions specified in the rules. If a rule is non-compliant, you can request a change and the AMS Ops team will work with you to take corrective action.

- In that case, you see the following changes appear in your AMS accounts:
  - AWS Config Rules under AWS Config > Rules
  - Custom Config rules with their Lambda functions exist in your account
  - Config Aggregator in Security account and Config Authorization in all accounts (Multi-Account Landing Zone only)

The following is a sample of AWS Config Rules and their compliance evaluation results is shown below:

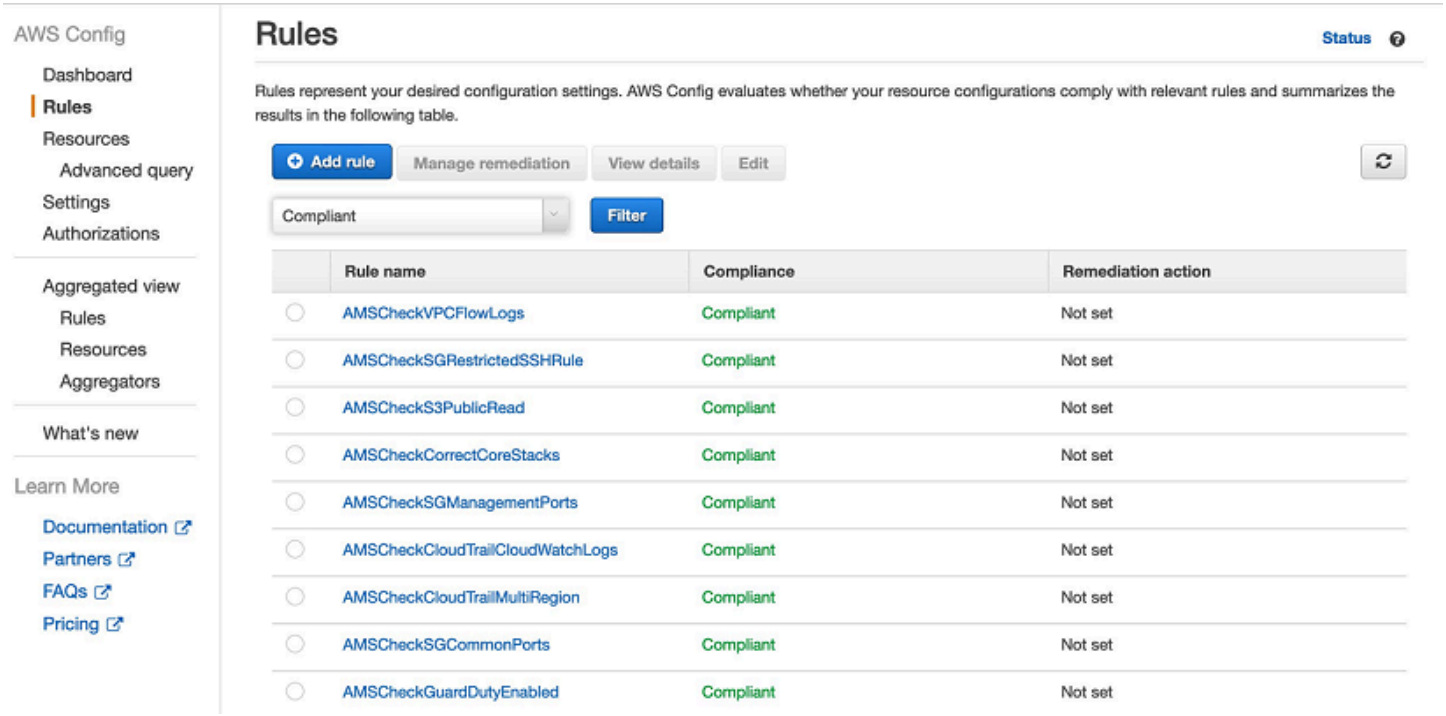

To learn more about AWS Config, see:

- AWS Config: What Is Config?
- AWS Config Rules: Evaluating Resources with Rules
- AWS Config Rules: Dynamic Compliance Checking: AWS Config Rules – Dynamic Compliance Checking for Cloud Resources
- AWS Config Aggregator: Multi-Account Multi-Region Data Aggregation

AMS multi-account landing zone service control policy restrictions

This section has been redacted because it contains sensitive AMS security-related information. This information is available through the AMS console **Documentation**. To access AWS Artifact, you can contact your CSDM for instructions or go to Getting Started with AWS Artifact.

Resilience

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. With Availability Zones, you can design and operate applications and databases that automatically fail over between zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about AWS Regions and Availability Zones, see AWS global infrastructure.

Infrastructure security

> **Note**
>
> Additional information on this topic is available by accessing AWS Artifact reports and downloading the *AWS Managed Services (AMS) Customer Security Guide*. For more information, see Downloading reports in AWS Artifact. To access AWS Artifact, you can contact your CSDM for instructions or go to Getting Started with AWS Artifact. This information is not included in this user guide because it contains sensitive security content.

Security control for end-of-support operating systems

Operating systems that are outside of the general support period of the operating system manufacturer's "end-of-support" or EOS, and do not receive security updates, have an increased security risk.

AWS offers some services to help with handling operation system end-of-support. For information about Windows end-of-support, see End-of-Support Migration Program for Windows Server.

Note

Additional information on this topic is available by accessing AWS Artifact reports and downloading the *AWS Managed Services (AMS) Customer Security Guide*. For more information, see Downloading reports in AWS Artifact. To access AWS Artifact, you can contact your CSDM for instructions or go to Getting Started with AWS Artifact. This information is not included in this user guide because it contains sensitive security content.

Using security groups

A security group acts as a virtual firewall that controls the traffic for one or more instances. AMS security groups allow you to set inbound traffic rules and outbound traffic rules on an instance-level basis. You can create a security group and specify resources in your AMS account, Amazon EC2 instances, Amazon RDS DB instances, Load Balancers, Deep Security Manager (DSM) replication instances, EFS mount targets, and ElastiCache clusters, to associate with the security group. Once associated, traffic to or from those instances is constrained by the rules set in the security group.

To better understand general AWS security, see Best Practices for Security, Identity, & Compliance and Amazon EC2 Security Groups for Linux Instances.

AMS now has a set of change types for creating and managing security groups:

- Deployment | Advanced stack components | Security group | Create (ct-1oxx2g2d7hc90)
- Management | Advanced stack components | Security group | Delete (ct-3cp96z7r065e4)
- Management | Advanced stack components | Security group | Update (ct-3memthlcmvc1b)

For examples, see Security groups.

Security groups

In AWS VPCs, AWS Security Groups act as virtual firewalls, controlling the traffic for one or more stacks (an instance or a set of instances). When a stack is launched, it's associated with one or more security groups, which determine what traffic is allowed to reach it:

- For stacks in your public subnets, the default security groups accept traffic from HTTP (80) and HTTPS (443) from all locations (the internet). The stacks also accept internal SSH and RDP traffic from your corporate network, and AWS bastions. Those stacks can then egress through any port

to the Internet. They can also egress to your private subnets and other stacks in your public subnet.

- Stacks in your private subnets can egress to any other stack in your private subnet, and instances within a stack can fully communicate over any protocol with each other.

> **Important**
>
> The default security group for stacks on private subnets allows all stacks in your private subnet to communicate with other stacks in that private subnet. If you want to restrict communications between stacks within a private subnet, you must create new security groups that describe the restriction. For example, if you want to restrict communications to a database server so that the stacks in that private subnet can only communicate from a specific application server over a specific port, request a special security group. How to do so is described in this section.

Default Security Groups

MALZ

The following table describes the default inbound security group (SG) settings for your stacks. The SG is named "SentinelDefaultSecurityGroupPrivateOnly-vpc-ID" where *ID* is a VPC ID in your AMS multi-account landing zone account. All traffic is allowed outbound to "mc-initial-garden-SentinelDefaultSecurityGroupPrivateOnly" via this security group (all local traffic within stack subnets is allowed).

All traffic is allowed outbound to 0.0.0.0/0 by a second security group "SentinelDefaultSecurityGroupPrivateOnly".

> **Tip**
>
> If you're choosing a security group for an AMS change type, such as EC2 create, or OpenSearch create domain, you would use one of the default security groups described here, or a security group that you created. You can find the list of security groups, per VPC, in either the AWS EC2 console or VPC console.

There are additional default security groups that are used for internal AMS purposes.

AMS default security groups (inbound traffic)

Type	Protocol	Port range	Source
All traffic	All	All	SentinelDefaultSecurityGroupPrivateOnly (restricts outbound traffic to members of the same security group)
All traffic	All	All	SentinelDefaultSecurityGroupPrivateOnlyEgress All (does not restrict outbound traffic)
HTTP, HTTPS, SSH, RDP	TCP	80 / 443 (Source 0.0.0.0/0) SSH and RDP access is allowed from bastions	SentinelDefaultSecurityGroupPublic (does not restrict outbound traffic)
MALZ bastions:			
SSH	TCP	22	SharedServices VPC CIDR and DMZ VPC CIDR, plus Customer-provided on-prem CIDRs
SSH	TCP	22	
RDP	TCP	3389	
RDP	TCP	3389	
SALZ bastions:			
SSH	TCP	22	mc-initial-garden-LinuxBastionSG
SSH	TCP	22	mc-initial-garden-LinuxBastionDMZSG
RDP	TCP	3389	mc-initial-garden-WindowsBastionSG
RDP	TCP	3389	mc-initial-garden-WindowsBastionDMZSG

SALZ

The following table describes the default inbound security group (SG) settings for your stacks. The SG is named "mc-initial-garden-SentinelDefaultSecurityGroupPrivateOnly-*ID*" where *ID* is a unique identifier. All traffic is allowed outbound to "mc-initial-garden-SentinelDefaultSecurityGroupPrivateOnly" via this security group (all local traffic within stack subnets is allowed).

All traffic is allowed outbound to 0.0.0.0/0 by a second security group "mc-initial-garden-SentinelDefaultSecurityGroupPrivateOnlyEgressAll-*ID*".

> **Tip**
>
> If you're choosing a security group for an AMS change type, such as EC2 create, or OpenSearch create domain, you would use one of the default security groups described here, or a security group that you created. You can find the list of security groups, per VPC, in either the AWS EC2 console or VPC console.

There are additional default security groups that are used for internal AMS purposes.

AMS default security groups (inbound traffic)

Type	Protocol	Port range	Source
All traffic	All	All	SentinelDefaultSecurityGroupPrivateOnly (restricts outbound traffic to members of the same security group)
All traffic	All	All	SentinelDefaultSecurityGroupPrivateOnlyEgressAll (does not restrict outbound traffic)
HTTP, HTTPS, SSH, RDP	TCP	80 / 443 (Source 0.0.0.0/0) SSH and RDP access is allowed from bastions	SentinelDefaultSecurityGroupPublic (does not restrict outbound traffic)
MALZ bastions:			

Type	Protocol	Port range	Source
SSH	TCP	22	SharedServices VPC CIDR and DMZ VPC CIDR, plus Customer-provided on-prem CIDRs
SSH	TCP	22	
RDP	TCP	3389	
RDP	TCP	3389	
SALZ bastions:			
SSH	TCP	22	mc-initial-garden-LinuxBastionSG
SSH	TCP	22	mc-initial-garden-LinuxBastionDMZSG
RDP	TCP	3389	mc-initial-garden-WindowsBastionSG
RDP	TCP	3389	mc-initial-garden-WindowsBastionDMZSG

Create, Change, or Delete Security Groups

You can request custom security groups. In cases where the default security groups do not meet the needs of your applications or your organization, you can modify or create new security groups. Such a request would be considered approval-required and would be reviewed by the AMS operations team.

To create a security group outside of stacks and VPCs, submit an RFC using the `Management | Other | Other | Create` CT (ct-1e1xtak34nx76).

To add or remove a user from an Active Directory (AD) security group, submit a request for change (RFC) using the `Management | Other | Other | Update` CT (ct-0xdawir96cy7k).

> **Note**
>
> When using "review required" CTs, AMS recommends that you use the ASAP **Scheduling** option (choose **ASAP** in the console, leave start and end time blank in the API/CLI) as these CTs require an AMS operator to examine the RFC, and possibly communicate with you before it can be approved and run. If you schedule these RFCs, be sure to allow at least 24

hours. If approval does not happen before the scheduled start time, the RFC is rejected automatically.

Find Security Groups

To find the security groups attached to a stack or instance, use the EC2 console. After finding the stack or instance, you can see all security groups attached to it.

For ways to find security groups at the command line and filter the output, see `describe-security-groups`.

AMS preventative and detective controls library

AWS Managed Services (AMS) provides you with a curated library/catalog of proven service control policies (SCPs) and ConfigRules that can be leveraged to improve your security posture and mitigate compliance gaps in your AMS accounts.

Topics

- Curated SCPs and Config Rules
- Custom notification for Config rules

Curated SCPs and Config Rules

Curated SCPs and Config Rules for AMS Advanced.

- **Service control policies (SCPs)**: The provided SCPs are in addition to default AMS ones.

  You can use these library controls in tandem with the default ones to meet specific security requirements.
- **Config Rules**: As a baseline measure, AMS recommends applying Conformance Packs (see Conformance Packs in the AWS Config guide) in addition to the default AMS config rules (see AMS Artifacts for default rules). The Conformance Packs cover a majority of compliance requirements and AWS regularly updates them.

  The rules listed here can be used to cover use-case specific gaps that aren't covered by Conformance Packs

> **Note**
>
> As AMS default rules and conformance packs get updated over time, you might see duplicates of these rules.
> AMS recommends doing periodic clean-up of duplicate Config Rules in general.
> For AMS Advanced, Config Rules should not use auto-remediations (see Remediating Noncompliant AWS Resources by AWS Config Rules) in order to avoid out-of-band changes.

SCP-AMS-001: Restrict EBS creation

Prevent the creation of EBS volumes if you don't have encryption enabled.

```
{
      "Condition": {
        "Bool": {
          "ec2:Encrypted": "false"
        }
      },
      "Action": "ec2:CreateVolume",
      "Resource": "*",
      "Effect": "Deny"
    }
```

SCP-AMS-002: Restrict EC2 launch

Prevent the launch of an EC2 instance if the EBS volume is unencrypted. This includes denying an EC2 launch from unencrypted AMIs because this SCP also applies to root volumes.

```
{
      "Condition": {
        "Bool": {
          "ec2:Encrypted": "false"
        }
      },
      "Action": "ec2:RunInstances",
      "Resource": "arn:aws:ec2:*:*:volume/*",
      "Effect": "Deny"
    }
```

SCP-ADV-001: Restrict RFC submissions

Restrict default AMS roles from submitting specific automated RFCs like **Create VPC** or **Delete VPC**. This is helpful if you want to apply more granular permissions to your federated roles.

For example, you might want the default `AWSManagedServicesChangeManagement Role` to be able to submit most of the available RFCs except the ones that allow for the creation and deletion of a VPC, creation of additional subnets, offboarding of an application account, updating or deleting SAML identity providers:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Condition": {
                "ArnLike": {
                    "aws:PrincipalARN": [
                        "arn:aws:iam::*:role/AWSManagedServicesChangeManagement Role",
                        "arn:aws:iam::*:role/AWSManagedServicesAdminRole"
                    ]
                }
            },
            "Action": [
                "amscm:ApproveRfc",
                "amscm:CancelRfc",
                "amscm:CreateRfc",
                "amscm:CreateRfcAttachment",
                "amscm:CreateRfcCorrespondence",
                "amscm:GetRfcAttachment",
                "amscm:ListRfcAttachmentSummaries",
                "amscm:ListRfcCorrespondences",
                "amscm:RejectRfc",
                "amscm:SubmitRfc",
                "amscm:UpdateRestrictedExecutionTimes",
                "amscm:UpdateRfc"
            ],
            "Resource": [
                "arn:aws:amscm:global:*:changetype/ct-1j3503fres5a5:*",
                "arn:aws:amscm:global:*:changetype/ct-2paw0y79kvr3l:*"
                "arn:aws:amscm:global:*:changetype/ct-2ha68tpd7nr3y:*"
                "arn:aws:amscm:global:*:changetype/ct-2wlfo2jxj2rkj:*"
                "arn:aws:amscm:global:*:changetype/ct-3hox8uwjgze1f:*"
                "arn:aws:amscm:global:*:changetype/ct-01zl37gmuk4q2:*"
```

```
                "arn:aws:amscm:global:*:changetype/ct-379uwo67vbvng:*"
            ],
            "Effect": "Deny",
            "Sid": "RestrictAppAccountRFCs"
        }
    ]
}
```

SCP-AMS-003: Restrict EC2 or RDS creation in AMS

Prevent creation of Amazon EC2 and RDS instances that don't have specific tags, while allowing the AMS default `AMS Backup IAM` role to do so. This is needed for disaster recover or DR.

```
{
    "Sid": "DenyRunInstanceWithNoOrganizationTag",
    "Effect": "Deny",
    "Action": [
        "ec2:RunInstances",
        "rds:CreateDBInstance"
    ],
    "Resource": [
        "arn:aws:ec2:*:*:instance/*",
        "arn:aws:ec2:*:*:volume/*",
        "arn:aws:rds:*:*:db:*"
    ],
    "Condition": {
        "Null": {
            "aws:RequestTag/organization": "true"
        },
        "StringNotLike": {
            "aws:PrincipalArn": [
                "arn:aws:iam::<Account_Number>:role/ams-backup-iam-role"
            ]
        }
    }
}
```

SCP-AMS-004: Restrict S3 uploads

Prevent uploads of unencrypted S3 objects.

```
{
            "Sid": "DenyUnencryptedS3Uploads",
```

```
            "Effect": "Deny",
            "Action": "s3:PutObject",
            "Resource": "*",
            "Condition": {
                "StringNotLike": {
                    "s3:x-amz-server-side-encryption": ["aws:kms", "AES256"]
                },
                "Null": {
                    "s3:x-amz-server-side-encryption": "false"
                }
            }
        }
    ]
}
```

SCP-AMS-005: Restrict API and console access

Prevent AWS Console and API access for requests coming from known bad IP addresses as determined customer InfoSec.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Condition": {
                "IpAddressIfExists": {
                    "aws:SourceIp": [
                        "<ip address 1>",
                        "<ip address 2>",
                        "<ip address n>"
                    ]
                }
            },
            "Action": [
                "*"
            ],
            "Resource": [
                "*"
            ],
            "Effect": "Deny",
            "Sid": "IPDeny"
        }
    ]
```

```
}
```

SCP-AMS-006: Prevent IAM entity from removing member account from the organization

Prevent an AWS Identity and Access Management entity from removing member accounts from the organization.

```
{
  "Effect": "Deny",
  "Action": ["organizations:LeaveOrganization"],
  "Resource": ["*"]
}
```

SCP-AMS-007: Prevent sharing resources to accounts outside your organization

Prevent sharing resources with external accounts outside your AWS organization

```
  {
    "Effect": "Deny",
    "Action": [
      "ram:*"
    ],
    "Resource": [
      "*"
    ],
    "Condition": {
      "Bool": {
        "ram:AllowsExternalPrincipals": "true"
      }
    }
  },
  {
    "Effect": "Deny",
    "Action": [
      "ram:CreateResourceShare",
      "ram:UpdateResourceShare"
    ],
    "Resource": "*",
    "Condition": {
      "Bool": {
        "ram:RequestedAllowsExternalPrincipals": "true"
      }
```

```
    }
  }
```

SCP-AMS-008: Prevent sharing with organizations or organizational units (OUs)

Prevent sharing resources with an account and/or OU that's in an organization.

```
{
  "Effect": "Deny",
  "Action": [
    "ram:CreateResourceShare",
    "ram:AssociateResourceShare"
  ],
  "Resource": "*",
  "Condition": {
    "ForAnyValue:StringLike": {
      "ram:Principal": [
        "arn:aws:organizations::*:account/o-${OrganizationId}/${AccountId}",
        "arn:aws:organizations::*:ou/o-${OrganizationId}/ou-${OrganizationalUnitId}"
      ]
    }
  }
}
```

SCP-AMS-009: Prevent users from accepting resource share invitations

Prevent member accounts from accepting invitations from AWS RAM to join resource shares. This API doesn't support any conditions and prevents shares only from external accounts.

```
{
  "Effect": "Deny",
  "Action": ["ram:AcceptResourceShareInvitation"],
  "Resource": ["*"]
}
```

SCP-AMS-010: Prevent account Region enable and disable actions

Prevent enabling or disabling any new AWS Regions for your AWS accounts.

```
{
  "Effect": "Deny",
  "Action": [
```

```
    "account:EnableRegion",
    "account:DisableRegion"
  ],
  "Resource": "*"
}
```

SCP-AMS-011: Prevent billing modification actions

Prevent modifications to billing and payment configuration.

```
{
  "Effect": "Deny",
  "Action": [
    "aws-portal:ModifyBilling",
    "aws-portal:ModifyAccount",
    "aws-portal:ModifyPaymentMethods"
  ],
  "Resource": "*"
}
```

SCP-AMS-012: Prevent deletion or modification to specific CloudTrails

Prevent modifications to specific AWS CloudTrail trails.

```
{
  "Effect": "Deny",
  "Action": [
    "cloudtrail:DeleteEventDataStore",
    "cloudtrail:DeleteTrail",
    "cloudtrail:PutEventSelectors",
    "cloudtrail:PutInsightSelectors",
    "cloudtrail:UpdateEventDataStore",
    "cloudtrail:UpdateTrail",
    "cloudtrail:StopLogging"
  ],
  "Resource": [
    "arn:${Partition}:cloudtrail:${Region}:${Account}:trail/${TrailName}"
  ]
}
```

SCP-AMS-013: Prevent disabling default EBS encryption

Prevent disabling of default Amazon EBS encryption.

```
{
  "Effect": "Deny",
  "Action": [
    "ec2:DisableEbsEncryptionByDefault"
  ],
  "Resource": "*"
}
```

SCP-AMS-014: Prevent creating default VPC and subnet

Prevent the creation of a default Amazon VPC and subnets.

```
{
  "Effect": "Deny",
  "Action": [
    "ec2:CreateDefaultSubnet",
    "ec2:CreateDefaultVpc"
  ],
  "Resource": "*"
}
```

SCP-AMS-015: Prevent disabling and modifying GuardDuty

Prevent Amazon GuardDuty from being modified or disabled.

```
{
  "Effect": "Deny",
  "Action": [
    "guardduty:AcceptInvitation",
    "guardduty:ArchiveFindings",
    "guardduty:CreateDetector",
    "guardduty:CreateFilter",
    "guardduty:CreateIPSet",
    "guardduty:CreateMembers",
    "guardduty:CreatePublishingDestination",
    "guardduty:CreateSampleFindings",
    "guardduty:CreateThreatIntelSet",
    "guardduty:DeclineInvitations",
    "guardduty:DeleteDetector",
    "guardduty:DeleteFilter",
    "guardduty:DeleteInvitations",
    "guardduty:DeleteIPSet",
```

```
    "guardduty:DeleteMembers",
    "guardduty:DeletePublishingDestination",
    "guardduty:DeleteThreatIntelSet",
    "guardduty:DisableOrganizationAdminAccount",
    "guardduty:DisassociateFromMasterAccount",
    "guardduty:DisassociateMembers",
    "guardduty:InviteMembers",
    "guardduty:StartMonitoringMembers",
    "guardduty:StopMonitoringMembers",
    "guardduty:TagResource",
    "guardduty:UnarchiveFindings",
    "guardduty:UntagResource",
    "guardduty:UpdateDetector",
    "guardduty:UpdateFilter",
    "guardduty:UpdateFindingsFeedback",
    "guardduty:UpdateIPSet",
    "guardduty:UpdateMalwareScanSettings",
    "guardduty:UpdateMemberDetectors",
    "guardduty:UpdateOrganizationConfiguration",
    "guardduty:UpdatePublishingDestination",
    "guardduty:UpdateThreatIntelSet"
  ],
  "Resource": "*"
}
```

SCP-AMS-016: Prevent root user activity

Prevent the root user from performing any action.

```
{
  "Action": "*",
  "Resource": "*",
  "Effect": "Deny",
  "Condition": {
    "StringLike": {
      "aws:PrincipalArn": [
        "arn:aws:iam::*:root"
      ]
    }
  }
}
```

SCP-AMS-017: Prevent creating access keys for the root user

Prevent the creation of access keys for the root user.

```
{
  "Effect": "Deny",
  "Action": "iam:CreateAccessKey",
  "Resource": "arn:aws:iam::*:root"
}
```

SCP-AMS-018: Prevent disabling S3 account public access block

Prevent disabling an Amazon S3 account public access block. This prevents any bucket in the account from becoming public.

```
{
  "Effect": "Deny",
  "Action": "s3:PutAccountPublicAccessBlock",
  "Resource": "*"
}
```

SCP-AMS-019: Prevent disabling AWS Config or modifying Config rules

Prevent disabling or modifying AWS Config rules.

```
{
  "Effect": "Deny",
  "Action": [
    "config:DeleteConfigRule",
    "config:DeleteConfigurationRecorder",
    "config:DeleteDeliveryChannel",
    "config:DeleteEvaluationResults",
    "config:StopConfigurationRecorder"
  ],
  "Resource": "*"
}
```

SCP-AMS-020: Prevent all IAM actions

Prevent all IAM actions.

```
{
```

```
  "Effect": "Deny",
  "Action": [
    "iam:*"
  ],
  "Resource": "*"
}
```

SCP-AMS-021: Prevent deleting CloudWatch Logs groups and streams

Prevent deleting Amazon CloudWatch Logs groups and streams.

```
{
  "Effect": "Deny",
  "Action": [
    "logs:DeleteLogGroup",
    "logs:DeleteLogStream"
  ],
  "Resource": "*"
}
```

SCP-AMS-022: Prevent Glacier deletion

Prevent Amazon S3 Glacier deletion.

```
{
  "Effect": "Deny",
  "Action": [
    "glacier:DeleteArchive",
    "glacier:DeleteVault"
  ],
  "Resource": "*"
}
```

SCP-AMS-023: Prevent deletion of IAM Access Analyzer

Prevent the deletion of IAM Access Analyzer.

```
{
  "Action": [
    "access-analyzer:DeleteAnalyzer"
  ],
  "Resource": "*",
```

```
  "Effect": "Deny"
}
```

SCP-AMS-024: Prevent modifications to Security Hub

Prevent the deletion of AWS Security Hub.

```
{
  "Action": [
    "securityhub:DeleteInvitations",
    "securityhub:DisableSecurityHub",
    "securityhub:DisassociateFromMasterAccount",
    "securityhub:DeleteMembers",
    "securityhub:DisassociateMembers"
  ],
  "Resource": "*",
  "Effect": "Deny"
}
```

SCP-AMS-025: Prevent deletion under Directory Service

Prevent the deletion of resources under AWS Directory Service.

```
{
  "Action": [
    "ds:DeleteDirectory",
    "ds:DeleteLogSubscription",
    "ds:DeleteSnapshot",
    "ds:DeleteTrust",
    "ds:DeregisterCertificate",
    "ds:DeregisterEventTopic",
    "ds:DisableLDAPS",
    "ds:DisableRadius",
    "ds:DisableSso",
    "ds:UnshareDirectory"
  ],
  "Resource": "*",
  "Effect": "Deny"
}
```

SCP-AMS-026: Prevent use of denylisted service

Prevent the use of denylisted services.

> **Note**
>
> Replace *service1* and *service2* with your service names. Example *access-analyzer* or *IAM*.

```
{
  "Effect": "Deny",
  "Resource": "*",
  "Action": ["service1:*", "service2:*"]
}
```

SCP-AMS-027: Prevent use of denylisted service in specific Regions

Prevent the use of denylisted services in specific AWS Regions.

> **Note**
>
> Replace *service1* and *service2* with your service names. Example *access-analyzer* or *IAM*.
> Replace *region1* and *region2* with your service names. Example *us-west-2* or *use-east-1*.

```
{
  "Effect": "Deny",
  "Resource": "*",
  "Action": ["service1:*", "service2:*"],
  "Condition": {
    "StringEquals": {
      "aws:RequestedRegion": [
        "region1",
        "region2"
      ]
    }
  }
}
```

SCP-AMS-028: Prevent tags from being modified except by authorized principals

Prevent tag modifications by any user except the authorized principals. Use authorization tags to authorize principals. Authorization tags must be associated with resources and with principals. A user/role is only considered authorized if the tag on both the resource and the principal match. For more information, see the following resources:

- Securing resource tags used for authorization using a service control policy in AWS Organizations
- Prevent tags from being modified except by authorized principals

```
{
  "Effect": "Deny",
  "Action": [
    "ec2:CreateTags",
    "ec2:DeleteTags"
  ],
  "Resource": [
    "*"
  ],
  "Condition": {
    "StringNotEquals": {
      "ec2:ResourceTag/access-project": "${aws:PrincipalTag/access-project}",
      "aws:PrincipalArn": "arn:aws:iam::{ACCOUNT_ID}:{RESOURCE_TYPE}/{RESOURCE_NAME}"
    },
    "Null": {
      "ec2:ResourceTag/access-project": false
    }
  }
},
{
  "Effect": "Deny",
  "Action": [
    "ec2:CreateTags",
    "ec2:DeleteTags"
  ],
  "Resource": [
    "*"
  ],
  "Condition": {
    "StringNotEquals": {
      "aws:RequestTag/access-project": "${aws:PrincipalTag/access-project}",
      "aws:PrincipalArn": "arn:aws:iam::{ACCOUNT_ID}:{RESOURCE_TYPE}/{RESOURCE_NAME}"
```

```
    },
    "ForAnyValue:StringEquals": {
      "aws:TagKeys": [
        "access-project"
      ]
    }
  }
},
{
  "Effect": "Deny",
  "Action": [
    "ec2:CreateTags",
    "ec2:DeleteTags"
  ],
  "Resource": [
    "*"
  ],
  "Condition": {
    "StringNotEquals": {
      "aws:PrincipalArn": "arn:aws:iam::{ACCOUNT_ID}:{RESOURCE_TYPE}/{RESOURCE_NAME}"
    },
    "Null": {
      "aws:PrincipalTag/access-project": true
    }
  }
}
```

SCP-AMS-029: Prevent users from deleting Amazon VPC Flow Logs

Prevent the deletion of Amazon VPC Flow Logs.

```
{
  "Action": [
    "ec2:DeleteFlowLogs",
    "logs:DeleteLogGroup",
    "logs:DeleteLogStream",
    "s3:DeleteBucket",
    "s3:DeleteObject",
    "s3:DeleteObjectVersion",
    "s3:PutLifecycleConfiguration",
    "firehose:DeleteDeliveryStream"
  ],
  "Resource": "*",
  "Effect": "Deny"
```

```
}
```

SCP-AMS-030: Prevent sharing VPC subnet with account other than network account

Prevent sharing Amazon VPC subnets with accounts other than the network account.

> **Note**
>
> Replace *NETWORK_ACCOUNT_ID* with your network account ID.

```
{
  "Effect": "Deny",
  "Action": [
    "ram:AssociateResourceShare",
    "ram:CreateResourceShare"
  ],
  "Resource": "*",
  "Condition": {
    "StringNotEquals": {
      "ram:Principal": "NETWORK_ACCOUNT_ID"
    },
    "StringEquals": {
      "ram:RequestedResourceType": "ec2:Subnet"
    }
  }
}
```

SCP-AMS-031: Prevent launching instances with prohibited instance types

Prevent launcing prohibited Amazon EC2 instance types.

> **Note**
>
> Replace *instance_type1* and *instance_type2* with the instance types that you want to restrict, such as *t2.micro* or a wildcard string such as **.nano*.

```
{
  "Effect": "Deny",
  "Action": "ec2:RunInstances",
```

```
  "Resource": [
    "arn:aws:ec2:*:*:instance/*"
  ],
  "Condition": {
    "ForAnyValue:StringLike": {
      "ec2:InstanceType": [
        "instance_type1",
        "instance_type2"
      ]
    }
  }
}
```

SCP-AMS-032: Prevent launching instances without IMDSv2

Prevent Amazon EC2 instances without IMDSv2.

```
[
  {
    "Effect": "Deny",
    "Action": "ec2:RunInstances",
    "Resource": "arn:aws:ec2:*:*:instance/*",
    "Condition": {
      "StringNotEquals": {
        "ec2:MetadataHttpTokens": "required"
      }
    }
  },
  {
    "Effect": "Deny",
    "Action": "ec2:RunInstances",
    "Resource": "arn:aws:ec2:*:*:instance/*",
    "Condition": {
      "NumericGreaterThan": {
        "ec2:MetadataHttpPutResponseHopLimit": "3"
      }
    }
  },
  {
    "Effect": "Deny",
    "Action": "*",
    "Resource": "*",
    "Condition": {
      "NumericLessThan": {
```

```
        "ec2:RoleDelivery": "2.0"
      }
    }
  },
  {
    "Effect": "Deny",
    "Action": "ec2:ModifyInstanceMetadataOptions",
    "Resource": "*"
  }
]
```

SCP-AMS-033: Prevent modifications to specific IAM role

Prevent modifications to specified IAM roles.

```
{
  "Action": [
    "iam:AttachRolePolicy",
    "iam:DeleteRole",
    "iam:DeleteRolePermissionsBoundary",
    "iam:DeleteRolePolicy",
    "iam:DetachRolePolicy",
    "iam:PutRolePermissionsBoundary",
    "iam:PutRolePolicy",
    "iam:TagRole",
    "iam:UntagRole",
    "iam:UpdateAssumeRolePolicy",
    "iam:UpdateRole",
    "iam:UpdateRoleDescription"
  ],
  "Resource": [
     "arn:aws:iam::{ACCOUNT_ID}:role/{RESOURCE_NAME}"
  ],
  "Effect": "Deny"
}
```

SCP-AMS-034: Prevent AssumeRolePolicy modification on specific IAM roles

Prevent modifications to the AssumeRolePolicy for specified IAM roles.

```
{
  "Action": [
    "iam:UpdateAssumeRolePolicy"
```

```
  ],
  "Resource": [
     "arn:aws:iam::{ACCOUNT_ID}:role/{RESOURCE_NAME}"
  ],
  "Effect": "Deny"
}
```

ConfigRule: Required tags

Check whether EC2 instances have custom tags that you have required. In addition to InfoSec, this is also useful for your Cost Management

```
ConfigRuleName: required-tags
      Description: >-
        A Config rule that checks whether EC2 instances have the required tags.
      Scope:
        ComplianceResourceTypes:
          - 'AWS::EC2::Instance'
      InputParameters:
        tag1Key: COST_CENTER
        tag2Key: APP_ID
      Source:
        Owner: AWS
        SourceIdentifier: REQUIRED_TAGS
```

ConfigRule: Access key rotated

Check that access keys are being rotated within the specified time period. This is usually set to be 90 days per typical compliance requirements.

```
ConfigRuleName: access-keys-rotated
      Description: >-
        A config rule that checks whether the active access keys are rotated
        within the number of days specified in maxAccessKeyAge. The rule is
        NON_COMPLIANT if the access keys have not been rotated for more than
        maxAccessKeyAge number of days.
      InputParameters:
        maxAccessKeyAge: '90'
      Source:
        Owner: AWS
        SourceIdentifier: ACCESS_KEYS_ROTATED
      MaximumExecutionFrequency: TwentyFour_Hours
```

ConfigRule: IAM root access key in AMS

Check that a root access key is not present on an account. For AMS Advanced accounts, this is expected to be compliant out-of-the-box.

```
ConfigRuleName: iam-root-access-key-check
      Description: >-
        A config rule that checks whether the root user access key is available. The
 rule is COMPLIANT if the user access key does not exist.
      Source:
        Owner: AWS
        SourceIdentifier: IAM_ROOT_ACCESS_KEY_CHECK
      MaximumExecutionFrequency: TwentyFour_Hours
```

ConfigRule: SSM managed EC2

Check that your EC2s are being managed by SSM Systems Manager.

```
ConfigRuleName: ec2-instance-managed-by-systems-manager
      Description: >-
        A Config rule that checks whether the EC2 instances in the
        account are managed by AWS Systems Manager.
      Scope:
        ComplianceResourceTypes:
          - 'AWS::EC2::Instance'
          - 'AWS::SSM::ManagedInstanceInventory'
      Source:
        Owner: AWS
        SourceIdentifier: EC2_INSTANCE_MANAGED_BY_SSM
```

ConfigRule: Unused IAM user in AMS

Check for IAM user credentials that have not been used for a specified duration. Like the key-rotation check, this usually defaults to 90 days per typical compliance requirements.

```
ConfigRuleName: iam-user-unused-credentials-check
      Description: >-
        A config rule that checks whether IAM users have passwords
        or active access keys that have not been used within the
        specified number of days provided.
      InputParameters:
        maxCredentialUsageAge: '90'
      Source:
```

```
        Owner: AWS
        SourceIdentifier: IAM_USER_UNUSED_CREDENTIALS_CHECK
      MaximumExecutionFrequency: TwentyFour_Hours
```

ConfigRule: S3 bucket logging

Check that logging has been enabled for S3 buckets in the account.

```
ConfigRuleName: s3-bucket-logging-enabled
      Description: >-
        A Config rule that checks whether logging is enabled for S3 buckets.
      Scope:
        ComplianceResourceTypes:
          - 'AWS::S3::Bucket'
      Source:
        Owner: AWS
        SourceIdentifier: S3_BUCKET_LOGGING_ENABLED
```

ConfigRule: S3 bucket versioning

Check that versioning and MFA-delete (optional) is enabled on all S3 buckets

```
ConfigRuleName: s3-bucket-versioning-enabled
      Description: >-
        A Config rule that checks whether versioning is enabled for S3
        buckets. Optionally, the rule checks if MFA delete is enabled for S3 buckets.
      Scope:
        ComplianceResourceTypes:
          - 'AWS::S3::Bucket'
      Source:
        Owner: AWS
        SourceIdentifier: S3_BUCKET_VERSIONING_ENABLED
```

ConfigRule: S3 public access

Check that public access settings (Public ACL, Public Policy, Public Buckets) are restricted across the account

```
ConfigRuleName: s3-account-level-public-access-blocks
      Description: >-
        A Config rule that checks whether the required public access block
        settings are configured from account level. The rule is only
        NON_COMPLIANT when the fields set below do not match the corresponding
```

```
        fields in the configuration item.
      Scope:
        ComplianceResourceTypes:
          - 'AWS::S3::AccountPublicAccessBlock'
      InputParameters:
        IgnorePublicAcls: 'True'
        BlockPublicPolicy: 'True'
        BlockPublicAcls: 'True'
        RestrictPublicBuckets: 'True'
      Source:
        Owner: AWS
        SourceIdentifier: S3_ACCOUNT_LEVEL_PUBLIC_ACCESS_BLOCKS
```

ConfigRule: Non-archived GuardDuty findings

Check for any non-archived GuardDuty findings that are older than the specified duration. The default duration is 30 days for low-sev, 7 days for medium-sev and 1 day for high-sev findings.

```
ConfigRuleName: guardduty-non-archived-findings
      Description: >-
        A Config rule that checks whether the Amazon GuardDuty has findings that
        are non archived. The rule is NON_COMPLIANT if GuardDuty has non
        archived low/medium/high severity findings older than the specified number.
      InputParameters:
        daysLowSev: '30'
        daysMediumSev: '7'
        daysHighSev: '1'
      Source:
        Owner: AWS
        SourceIdentifier: GUARDDUTY_NON_ARCHIVED_FINDINGS
      MaximumExecutionFrequency: TwentyFour_Hours
```

ConfigRule: CMK deletion

Check for any AWS Key Management Service custom master keys (CMKs) that are scheduled (aka pending) for deletion. This is crucial as unawareness around CMK deletion can lead to data being unrecoverable

```
ConfigRuleName: kms-cmk-not-scheduled-for-deletion
      Description: >-
        A config rule that checks whether customer master keys (CMKs) are not
        scheduled for deletion in AWS Key Management Service (AWS KMS). The rule is
        NON_COMPLIANT if CMKs are scheduled for deletion.
```

```
      Source:
        Owner: AWS
        SourceIdentifier: KMS_CMK_NOT_SCHEDULED_FOR_DELETION
      MaximumExecutionFrequency: TwentyFour_Hours
```

ConfigRule: CMK rotation

Check that auto-rotation is enabled for every CMK in the account

```
ConfigRuleName: cmk-backing-key-rotation-enabled
      Description: >-
        A config rule that checks that key rotation is enabled for each customer
        master key (CMK). The rule is COMPLIANT, if the key rotation is enabled
        for specific key object. The rule is not applicable to CMKs that have
        imported key material.
      Source:
        Owner: AWS
        SourceIdentifier: CMK_BACKING_KEY_ROTATION_ENABLED
      MaximumExecutionFrequency: TwentyFour_Hours
```

Custom notification for Config rules

There can be occurrences of critical non-compliant Config Rules that require raising escalated awareness directly with the your InfoSec and Leadership teams. For such scenarios, AMS recommends that you configure a non-compliance event-driven custom notification.

For example:

```
ConfigRuleName: required-tags
      Description: >-
        A Config rule that checks whether EC2 instances have the mandated tags.
      Scope:
        ComplianceResourceTypes:
          - 'AWS::EC2::Instance'
      InputParameters:
        tag1Key: COST_CENTER
        tag2Key: APP_ID
      Source:
        Owner: AWS
        SourceIdentifier: REQUIRED_TAGS
  NotificationEventRule:
    Type: 'AWS::Events::Rule'
    Properties:
```

```
      Name: CWEventForrequired-tags
      Description: >-
        SNS Notification for Non-Compliant Events of Config Rule:
        required-tags
      State: ENABLED
      EventPattern:
        detail-type:
          - Config Rules Compliance Change
        source:
          - aws.config
        detail:
          newEvaluationResult:
            complianceType:
              - NON_COMPLIANT
          configRuleARN:
            - 'Fn::GetAtt':
                - RequiredEC2Tags
                - Arn
      Targets:
        - Id: RemediationNotification
          Arn:
            Ref: SnsTopic
          InputTransformer:
            InputTemplate: >-
              "EC2 Instance <Instance_ID> is non-compliant. Please add required tags:
COST_CENTER, APP_ID, Name, and Backup."
            InputPathsMap:
              instance_id: $.detail.resourceId
 SnsTopic:
   Type: 'AWS::SNS::Topic'
   Properties:
     Subscription:
       - Endpoint: Cloud_Ops_Leaders@customer.com
         Protocol: email
     TopicName: noncompliant-instance-notification
 SnsTopicPolicy:
   Type: 'AWS::SNS::TopicPolicy'
   Properties:
     PolicyDocument:
       Statement:
         - Sid: __default_statement_ID
           Effect: Allow
           Principal:
             AWS: '*'
```

```
            Action:
              - 'SNS:GetTopicAttributes'
              - 'SNS:SetTopicAttributes'
              - 'SNS:AddPermission'
              - 'SNS:RemovePermission'
              - 'SNS:DeleteTopic'
              - 'SNS:Subscribe'
              - 'SNS:ListSubscriptionsByTopic'
              - 'SNS:Publish'
              - 'SNS:Receive'
            Resource:
              Ref: SnsTopic
            Condition:
              StringEquals:
                'AWS:SourceOwner':
                  Ref: 'AWS::AccountId'
          - Sid: TrustCWEToPublishEventsToMyTopic
            Effect: Allow
            Principal:
              Service: events.amazonaws.com
            Action: 'sns:Publish'
            Resource:
              Ref: SnsTopic
      Topics:
        - Ref: SnsTopic
```

Amazon EventBridge rule service-linked role for AMS Advanced

AMS Advanced uses the service-linked role (SLR) named **AWSServiceRoleForManagedServices_Events** – This role trusts one of the AWS Managed Services service principals (events.managedservices.amazonaws.com) to assume the role for you. The service uses the role to create EventBridge managed rule. This rule is the infrastructure required in your AWS account to deliver alarm state change information from your account to AWS Managed Services.

Permissions for EventBridge SLR for AMS Advanced

The **AWSServiceRoleForManagedServices_Events** service-linked role trusts the following services to assume the role:

- events.managedservices.amazonaws.com

Attached to this role is the **AWSManagedServices_EventsServiceRolePolicy** AWS managed policy (see AWS managed policy: AWSManagedServices_EventsServiceRolePolicy). The service uses the role to deliver alarm state change information from your account to AWS Managed Services. 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 *AWS Identity and Access Management* User Guide.

You can download the JSON **AWSManagedServices_EventsServiceRolePolicy** in this ZIP: EventsServiceRolePolicy.zip.

Creating an EventBridge SLR for AMS Advanced

You don't need to manually create a service-linked role. When you Onboard to AMS in the AWS Management Console, the AWS CLI, or the AWS API, then AMS Advanced creates the service-linked role for you.

> **Important**
>
> This service-linked role can appear in your account if you were using the AMS Advanced service before February 7, 2023, when it began supporting service-linked roles then AMS Accelerate created the AWSServiceRoleForManagedServices_Events role in your account. 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 Onboard to AMS, AMS Advanced creates the service-linked role for you again.

Editing an EventBridge SLR for AMS Advanced

AMS Advanced does not allow you to edit the AWSServiceRoleForManagedServices_Events 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 Editing a service-linked role in the *IAM User Guide*.

Deleting an EventBridge SLR for AMS Advanced

You don't need to manually delete the AWSServiceRoleForManagedServices_Events role. When you Offboard from AMS in the AWS Management Console, the AWS CLI or the AWS API, AMS Advanced cleans up the resources and deletes the service-linked role for you.

You can also use the IAM console, the AWS CLI or the AWS API to manually delete the service-linked role. To do this, you must first manually clean up the resources for your service-linked role and then you can manually delete it.

> **Note**
>
> If the AMS Advanced 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 and try the operation again.

To delete AMS Advanced resources used by the AWSServiceRoleForManagedServices_Events service-linked role

To manually delete the service-linked role using IAM

Use the IAM console, the AWS CLI, or the AWS API to delete the AWSServiceRoleForManagedServices_Events service-linked role.

For more information, see Deleting a service-linked role in the *IAM User Guide*.

Security best practices

This section has been redacted because it contains sensitive AMS security-related information. This information is available through the AMS console **Documentation**. To access AWS Artifact, you can contact your CSDM for instructions or go to Getting Started with AWS Artifact.

AMS multi-account landing zone EPS non-default settings

This section has been redacted because it contains sensitive AMS security-related information. This information is available through the AMS console **Documentation**. To access AWS Artifact, you can contact your CSDM for instructions or go to Getting Started with AWS Artifact.

AMS Guardrails

A guardrail is a high-level rule that provides ongoing governance for your overall AMS environment.

This section has been redacted because it contains sensitive AMS security-related information. This information is available through the AMS console **Documentation**. To access AWS Artifact, you can contact your CSDM for instructions or go to Getting Started with AWS Artifact.

MALZ Service control policies

This section has been redacted because it contains sensitive AMS security-related information. This information is available through the AMS console **Documentation**. To access AWS Artifact, you can contact your CSDM for instructions or go to Getting Started with AWS Artifact.

Continuity management in AMS Advanced

As part of continuity management, AWS Managed Services (AMS) provides automated access to AWS Backup, a native service with AWS. This facilitates access to a service that supports Amazon EBS, Amazon EC2, Amazon RDS, Amazon EFS, and more.

To learn more, see AWS Backup: How It Works.

Topics

- What is continuity management?
- How continuity management works
- Disaster recovery response
- Disaster recovery planning

What is continuity management?

Continuity management is the process AMS uses to provide backups and snapshots for your account.

AMS provides access to AWS Backup through change types that you use to create and manage backup jobs and plans.

How continuity management works

AMS uses AWS Backup for continuity management.

When starting to work with AWS Backup in AMS:

1. Run an on-demand backup
2. Create a backup plan (optional, AMS provides default backup plans)
3. Use the default AMS a backup vaults (optional)
4. Manage (run, refine, delete, and so forth) your backup plans and recovery points

AMS backup plans

A backup plan is a policy expression that defines when and how you want to back up supported AWS resources, such as RDS databases, EBS volumes, DynamoDB tables, and EFS file systems. Scheduling and retention policies are managed via custom backup plans, which you can create using a change type (CT) with AMS Advanced or using AWS Backup with AMS Accelerate. Assign resources to your backup plans using tags and AWS Backup automatically backs up and retains backups for assigned resources according to the defined backup plan. You can create multiple backup plans if you have workloads with different backup requirements.

A backup plan can have up to six backup rules that define a schedule and a retention period, among other details. The backup schedule determines when AWS Backup initiates a backup job and how often a backup is created. You can choose a frequency of hourly, daily, weekly, or monthly. The deletion days setting determines how many days the snapshot is stored before being automatically deleted.

> **Note**
>
> AMS Advanced: If you are migrated from the legacy AMS backup system, AMS creates a default backup plan for backwards compatibility. The **key:value** pair in this scenario is **Backup:True**. To support backwards compatibility, the value here is case insensitive, so **Backup:True** or **Backup:TRUE** are all valid tags. All other key:value pairs are case sensitive. AWS Backup can operate at the EBS volume level or at the Amazon EC2 instance level, but do not do both at the same time, as this can lead to a race condition where the backups may clash.

Default backup plans, multi-account landing zone

During the new **Account creation** RFC, AMS ensures that there is an overarching default backup plan at the account level to safeguard your workloads. The values for mandatory fields are set by default, as shown in the following section:

Default AMS backup plan

default-backup-plan

TAG key: Backup

TAG value: `True`

```
RuleForDailyBackups schedule expression: cron(30 23 ? * *) (a daily backup for 23:30
 UTC time)
RuleForDailyBackups delete after days: 31 days
RuleForWeeklyBackups schedule expression: cron(30 23 ? * 7 *) (a weekly backup for
 23:30 UTC time only on Saturday)
RuleForWeeklyBackups delete after weeks: 6 weeks
RuleForMonthlyBackups schedule expression: cron(30 23 * ? *) (a monthly backup for
 23:30 UTC time on day 1 of the month)
RuleForMonthlyBackups delete after weeks: 26 weeks
RuleForYearlyBackups schedule expression: cron(30 23 1 1 ? *) (a yearly backup for
 23:30 UTC time on day 1 of the month, only in January)
RuleForYearlyBackups delete after years: 2 years
```

Default AMS backup plan	Start Time	Retention
hourly backup	N/A	N/A
daily backup	daily 11:30PM UTC	7 days
weekly backup	weekly 11:30PM UTC, only on Saturday	4 weeks
monthly backup	monthly 11:30 PM UTC, on day 1 of the month	26 weeks
yearly backup	11:30 PM UTC, on day 1 of the month	2 years

Enhanced default AMS backup plan

This plan is a blueprint for AWS Backup best practices to protect against ransomware attacks. It implements a daily, weekly, monthly, and yearly backup strategy. AWS Backup continuous backup is enabled with maximum retention (31 days) on supported resources.

ams-enhanced-default-backup-plan

TAG key: `backup-orchestrator-enhanced`

TAG value: `true`

```
RuleForDailyBackups schedule expression: cron(0 0 4 ? * * ) (a daily backup for 04:00
 UTC time)
RuleForDailyBackups delete after days: 31 days
RuleForDailyBackups continuous backup: true
RuleForWeeklyBackups schedule expression: cron(0 0 2 ? * 7) (a weekly backup for 02:00
 UTC time only on Saturday)
RuleForWeeklyBackups delete after weeks: 6 weeks
RuleForMonthlyBackups schedule expression: cron(0 2 1 * ? *) (a monthly backup for
 02:00 UTC time on day 1 of the month)
RuleForMonthlyBackups delete after weeks: 26 weeks
RuleForYearlyBackups schedule expression: cron(0 2 1 1 ? *) (a yearly backup for 02:00
 UTC time on day 1 of the month, only in January)
RuleForYearlyBackups delete after years: 2 years
```

Enhanced AMS backup plan	Start Time	Retention
hourly backup	N/A	N/A
daily backup	daily 4:00 UTC	31 days
weekly backup	Saturday, 2:00 UTC	6 weeks
monthly backup	1st of the month, 2:00 UTC	26 weeks
yearly backup	Jan 1st, 2:00 UTC	2 years

Data sensitive AMS backup plan

This plan is a blueprint for AWS Backup best practices to protect against ransomware attacks for data-sensitive applications. It implements an hourly, daily, weekly, monthly, and yearly backup strategy. AWS Backup continuous backup is enabled with maximum retention (31 days) on supported resources.

ams-data-sensitive-backup-plan

TAG key: `backup-orchestrator-data-sensitive`

TAG value: `true`

```
RuleForHourlyBackups schedule expression: cron(0 * ? * * *) (an hourly backup at the
 hour mark)
```

```
RuleForHourlyBackups delete after days: 7 days
RuleForDailyBackups schedule expression: cron(0 0 4 ? * * ) (a daily backup for 04:00
 UTC time)
RuleForDailyBackups delete after days: 31 days
RuleForWeeklyBackups schedule expression: cron(0 0 2 ? * 7) (a weekly backup for 02:00
 UTC time only on Saturday)
RuleForWeeklyBackups delete after weeks: 6 weeks
RuleForMonthlyBackups schedule expression: cron(0 2 1 * ? *) (a monthly backup for
 02:00 UTC time on day 1 of the month)
RuleForMonthlyBackups delete after weeks: 26 weeks
RuleForYearlyBackups schedule expression: cron(0 2 1 1 ? *) (a yearly backup for 02:00
 UTC time on day 1 of the month, only in January)
RuleForYearlyBackups delete after years: 2 years
```

Data Sensitive AMS backup plan	Start Time	Retention
hourly backup	at the hour mark	7 days
daily backup	daily 4:00 UTC	31 days
weekly backup	Saturday, 2:00 UTC	6 weeks
monthly backup	1st of the month, 2:00 UTC	26 weeks
yearly backup	Jan 1st, 2:00 UTC	2 years

AMS backup vaults

AWS Backup organizes snapshots into logical storage units called vaults.

To view a list of your AMS backups, open the AWS Backup console. In the navigation pane, choose **Backup vaults** and select the one of the AMS backup vaults from the following tables. In the **Backups** section, view the list of all the backups in the backup vault. Select a backup to edit, delete, or restore.

Vaults for AMS backup plans

AMS Vault Name	Description
ams-automated-backups	This vault receives all recovery points taken by the AMS Advanced default AWS Backup plan **default-backup-plan**.
ams-automated-enhanced-backups	This vault receives all recovery points taken by AMS Advanced enhanced default AWS Backup plan **ams-enhanced-default-backup-plan**.
ams-automated-data-sensitive-backups	This vault receives all recovery points taken by AMS Advanced AWS Backup plan **ams-data-sensitive-backup-plan**.
ams-manual-backups	This is the default location for all backups from Start Backup Job RFC (ct-2hhud 2lx01tq7) backup plans, if no vault name is defined.
ams-custom-backups	This is the default location for the snapshots AMS takes prior to patching an instance using Patch Orchestrator or the monthly patch activities. These are automatically removed according to the AMS patch lifecycle default policy of 60 days.

AMS backup change types

AMS provides several CTs for you to create and use backup plans.

> **Important**
>
> Do not edit your AMS default backup plans as your changes may be lost. Instead, create new plans for your custom configurations.

- Backup plan: Create

- Backup Job: Start
- Backup Job: Stop
- Recovery Point: Delete
- DynamoDB | Create from Backup
- EBS Volume: Create From Backup
- Amazon Elastic File System (EFS): Create From Backup

AMS backup monitoring and reporting

> **Important**
>
> AMS backup monitoring and reporting are only available in AMS-supported regions. Those are US East (Virginia), US West (N. California), US West (Oregon), US East (Ohio), Canada (Central), South America (São Paulo), EU (Ireland), EU (Frankfurt), EU (London), EU (Paris), Asia Pacific (Mumbai), Asia Pacific (Seoul), Asia Pacific (Singapore), Asia Pacific (Sydney), Asia Pacific (Tokyo).

AMS generates daily self-service reports as well as monthly reports on resource coverage and backup job status. The monthly reports are shared in Monthly Business Reviews (MBRs). To learn more about daily backup reports, see Daily backup report .

AMS experts monitor all your backup tasks that are configured using AWS Backup. In case of backup failures, AMS investigates the failure and notifies you with the root cause and remediation options, if available. To avoid alert noise, during events that cause a high number of backup failures in your accounts, AMS makes a collective recommendation, through your CSDM, instead of notifying you for each individual failure.

Note that AMS does not monitor any backups configured using an AWS service's standalone backup feature.

Disaster recovery response

In addition to the options described in the following sections, it is good for you to know what steps to take to initiate a disaster recovery (DR) with AMS.

If you experience a disaster and need to initiate a recovery, follow these general guidelines:

1. Open a **High** priority incident with the **Availability** category. AMS will open a conference bridge and invite your team to join.
2. Know the list of resources you need to recover.
3. Know the target landing zone (LZ) you need to recover to (for example, the same account, different AZ or different account and different region).
4. Submit recover requests for each resource in the target landing zone. Follow your existing DR plan or see the options in the following section (for example, Disaster protection for EC2 with EBS snapshots on AMS, or Disaster protection for EC2 with CloudEndure on AMS).
5. Restore the application functionality and use AMS assistance to troubleshoot infrastructure-related issues.

AMS can help you with preparing for this event and with creating a DR plan for your organization to cover these questions. For more details, contact your cloud service delivery manager (CSDM) or cloud architect (CA).

Disaster recovery planning

Disaster recovery (DR) is a critical service for enterprise business continuity and compliance. AMS partners with you to help you plan, implement and maintain your DR strategy on AMS.

AMS landing zone (LZ), multi-account and single-account, provides native, multi-AZ, high-availability for AMS infrastructure components that meet most disaster protection scenarios. However, depending on your business's geographical coverage, you might need regional protection. For cross-region availability and DR, another AMS account is required in a different region (this is so for both multi-account landing zone and single-account landing zone).

AMS aligns with AWS DR guidance as described in this blog, Rapidly recover mission-critical systems in a disaster, and supports the following four options:

- Multi Site (or Highly Available)
- Warm Standby
- Pilot Light
- Backup and Restore

These options and AMS support for them are described in the following sections.

Multi-site or highly available (HA)

The HA solution is usually provided by the application's built-in functionality, such as clustering or synchronous replication. Users are directed to both Prod and HA/DR nodes. DNS points either to the nodes directly or through an elastic load balancer (ELB).

Your AMS cloud architect (CA) will work with you as part of your Well-Architected-Review and DR planning.

HA DR utilizes application and AWS-native services and features, as illustrated in the following graphic:

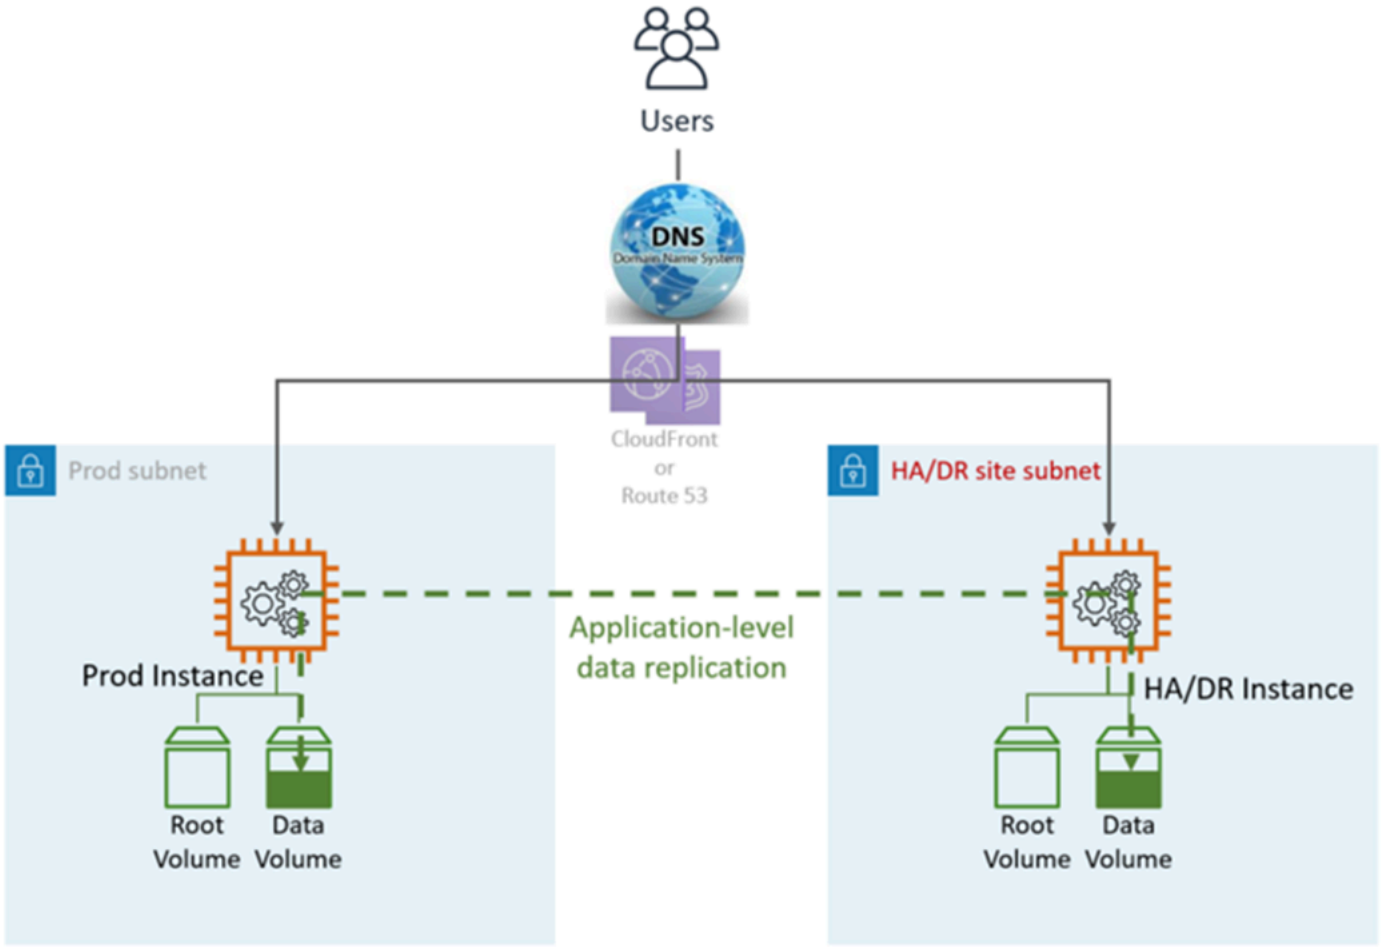

The DR site can be in the same or different AWS Region.

> **Note**
>
> Different region (Cross-Region) will have a different Active Directory environment.

DR (failover) steps: Automatic failover, no manual steps are required. In case of a failure in the primary LZ, the users will be automatically re-routed to the DR/HA node. This is achieved by both DNS and application configuration.

HA DR metrics:

- Recovery Point Objective (RPO): <5 min
- Recovery Time Objective and (RTO): <5 min
- Maintenance: High (Synchronous changes are required in both environments, like Application configuration, patching, SG or ALB, certificates, and so on).
- Cost: High

Warm standby

The term "warm standby" is used to describe a disaster recovery (DR) scenario in which a scaled-down version of the environment is running in the cloud.

Data replication is handled by the application layer, usually asynchronously, to an online instance, while the rest of the instances (for example, Application and Web tier) might be turned off to save the cost. Users are directed only to the Production site. Other AWS resources like elastic load balancer (ELB) may be pre-provisioned in the DR site as well.

Your AMS Cloud Architect (CA) will work with you as part of your Well-Architected-Review and DR planning.

Warm Standby DR utilizes application and AWS-native services and features, as illustrated in the following graphic:

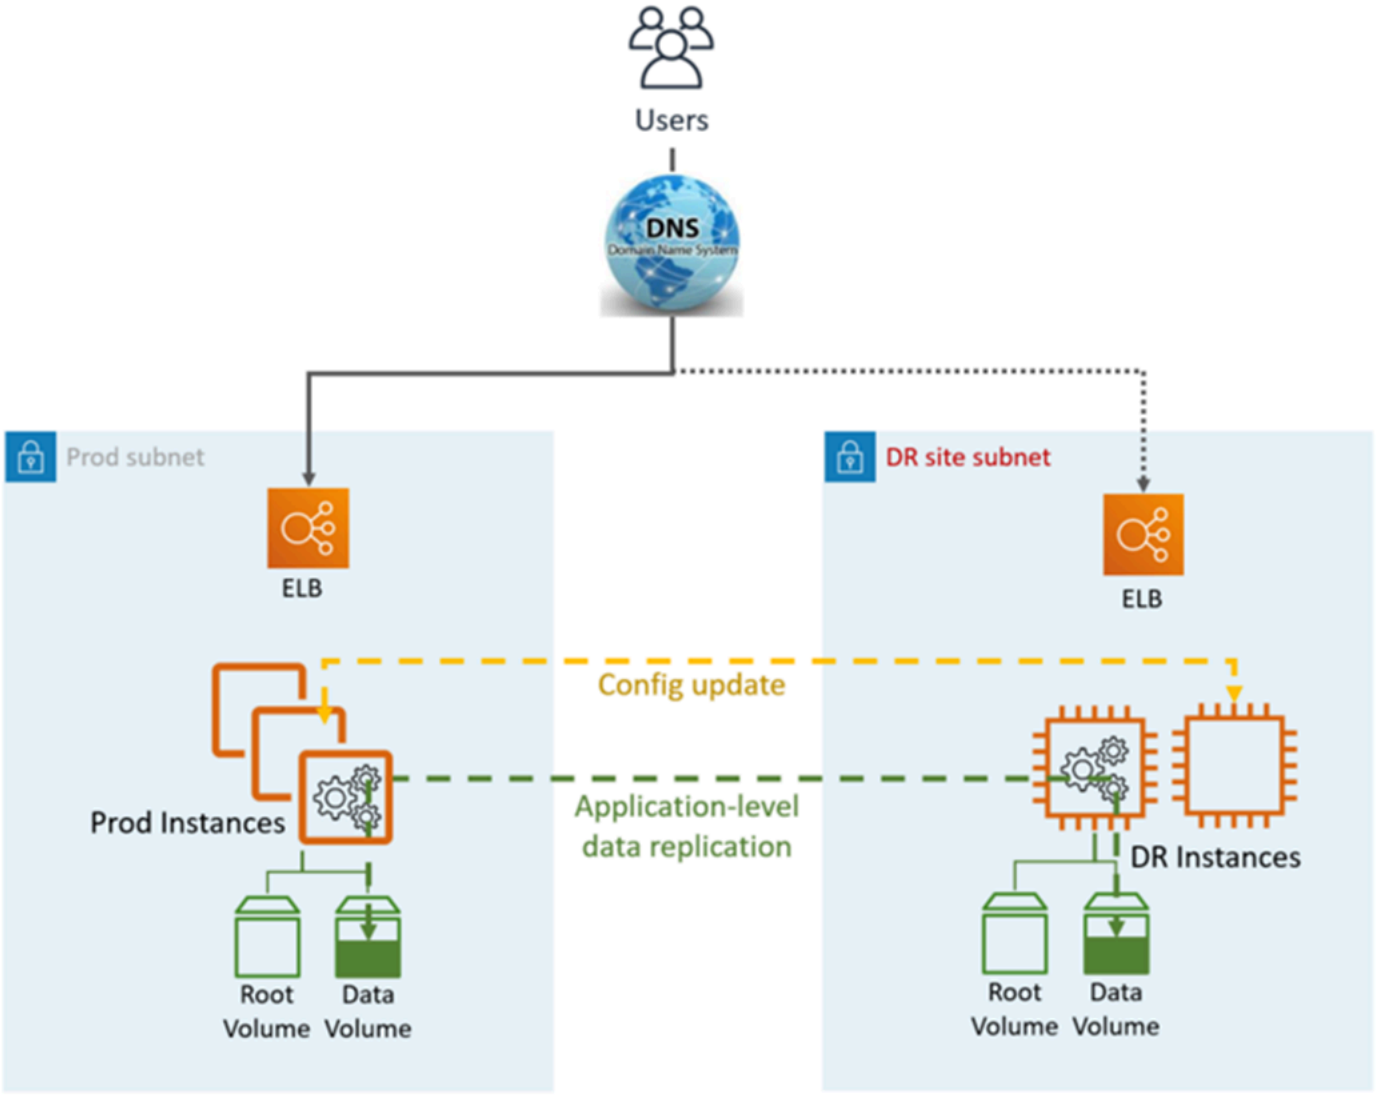

DR site can be in the same or different AWS Region.

> **Note**
>
> Different region (Cross-Region) will have a different Active Directory environment.

DR (failover) steps:

1. Brake the data replication and make the data instance in the DR site the master
2. Update application configuration as required (new IP, server name, and so on)
3. Redirect DNS to the DR site (ELB)

4. AD Dependencies if required (Service accounts, SPNs, GPOs, and so on)

HA DR metrics:

- Recovery Point Objective (RPO): <1hr
- Recovery Time Objective and (RTO): <1 hr (depends on the number of instances and orchestration)
- Maintenance: High (Synchronous changes are required in both environments, like Application configuration, patching, security groups (SG) or application load balancer (ALB), certificates, and so on).
- Cost: Medium

Pilot light

In this disaster recovery (DR) approach, you replicate part of your Prod environment for a limited set of core services. A small part of your infrastructure is always running, simultaneously syncing mutable data (such as databases or documents), while other parts of your infrastructure are switched off and used only during testing. Unlike a backup and recovery approach, you must ensure that your most critical core elements are already configured and running in the DR landing zone (the pilot light).

Your AMS Cloud Architect will work with you as part of your Well-Architected-Review and DR planning.

Pilot Light DR utilizes application and AWS-native services and features, as illustrated in the following graphic:

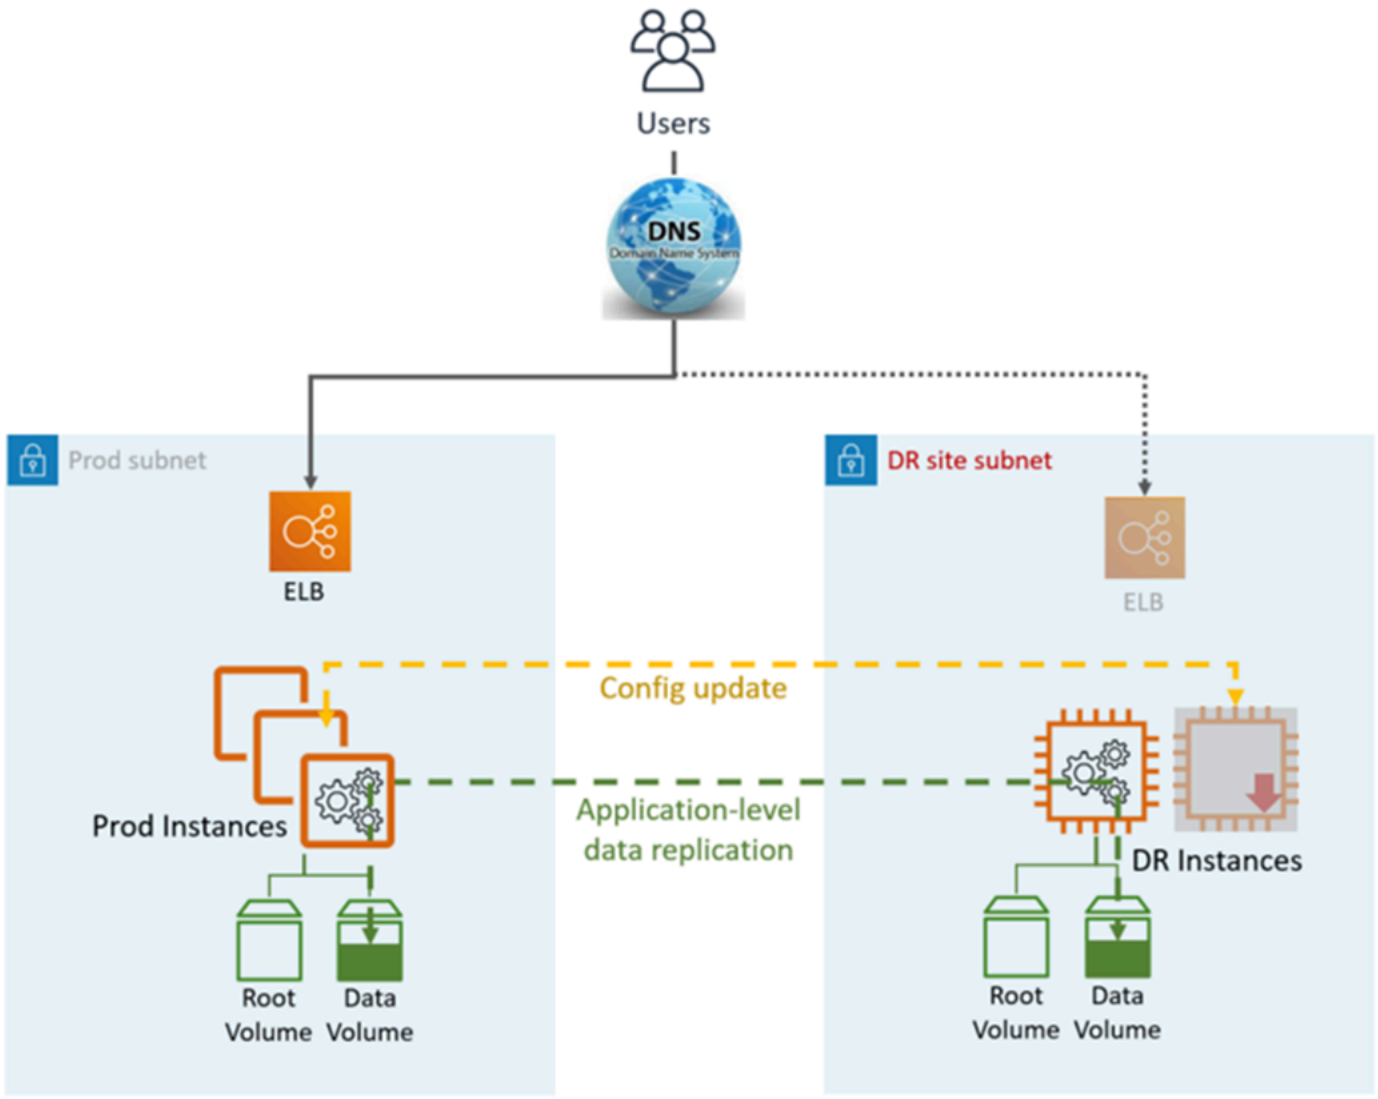

DR site can be in the same or different AWS Region.

> **Note**
>
> Different region (Cross-Region) will have a different Active Directory environment.

DR (failover) steps:

1. Brake the data replication and make the data instance in the DR site the master
2. Start the turned off instances and infrastructure
3. Update application configuration as required (new IP, server name, and so on)

4. Add the instances to the ELB as required
5. Redirect DNS to the DR site (ELB)
6. AD Dependencies, if required (Service accounts, SPNs, GPOs, and so on)

Pilot Light DR metrics:

- Recovery Point Objective (RPO): <1hr
- Recovery Time Objective and (RTO): ~1 hr (depends on the number of instances and orchestration)
- Maintenance: Medium
- Cost: Medium

Backup and restore

This simple and low cost disaster recovery (DR) approach backs up your data and applications from anywhere to the DR landing zone for use during recovery from a disaster.

Your AMS Cloud Architect will work with you as part of your Backup and DR planning.

Backup and Restore DR utilizes AMS automated tooling and processes, as illustrated in the following graphic:

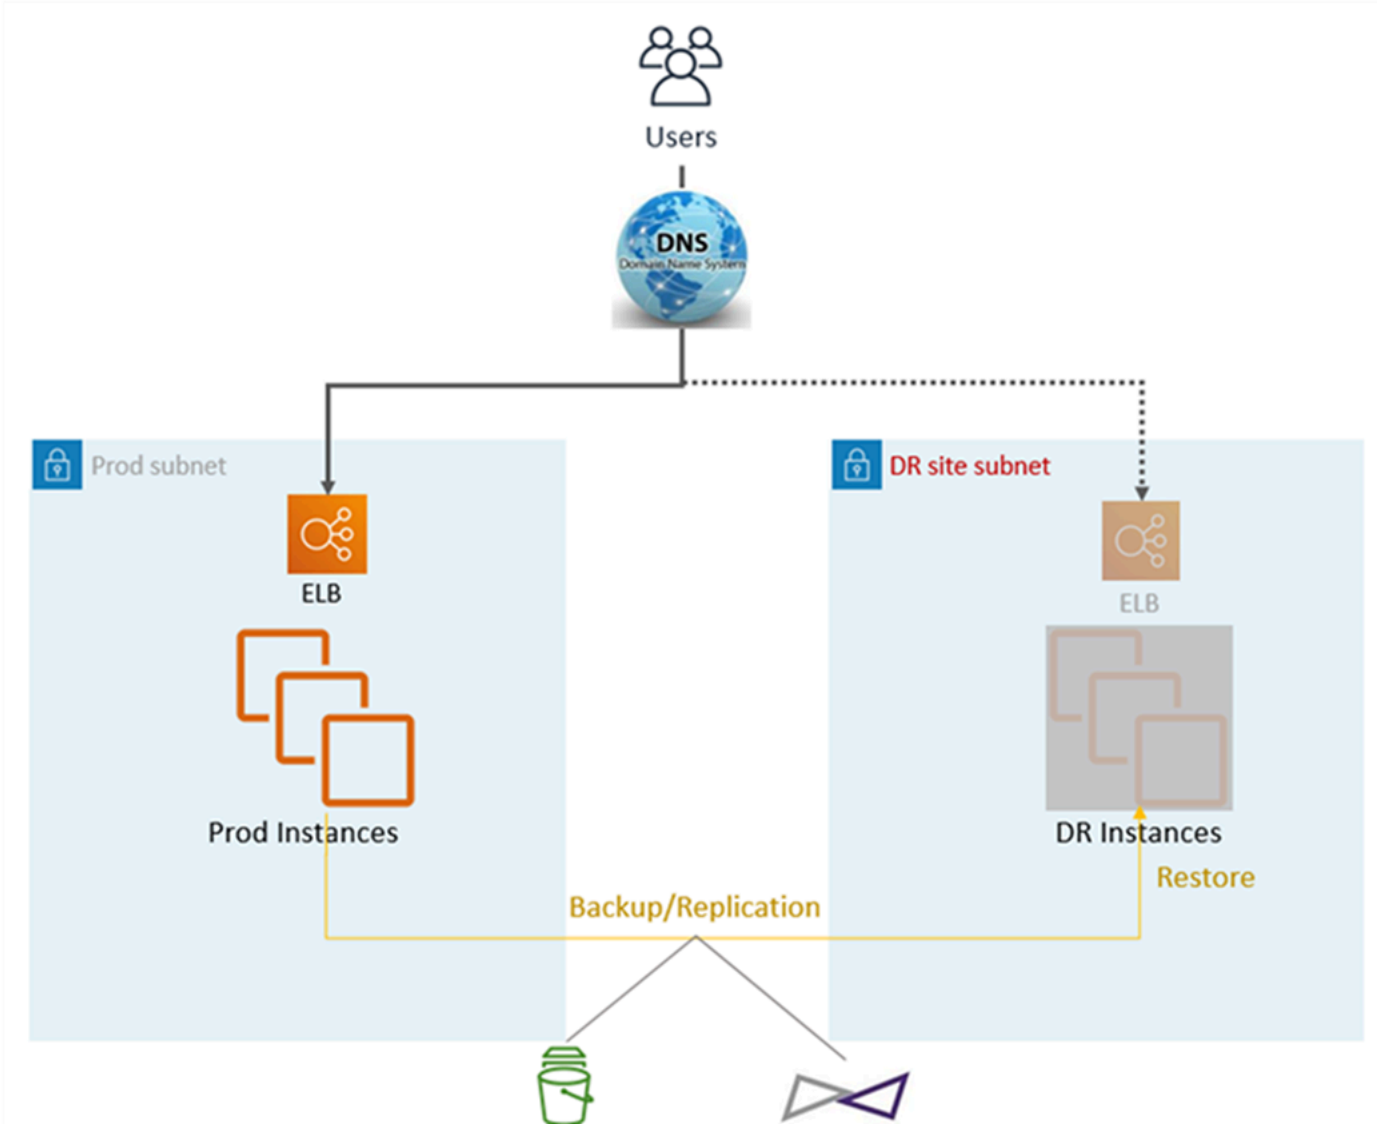

Two backup and replication methods can be used:

- EBS snapshot (Recovery Point Objective (RPO) > 1hr), known as "EBS"
- CloudEndure (Recovery Point Objective (RPO) ~ 0.25hrs), known as "CE"

The DR site can be in the same or different AWS Region.

> **Note**
>
> Different region (Cross-Region) will have a different Active Directory environment.

DR (failover) steps:

1. Restore the instances from snapshots (two-step process with placeholder instance first)
2. Update application configuration (new IP, server namea,nd so on)
3. Set up other infrastructure as required (SG, ELB, and so on)
4. Redirect DNS to the DR site (ELB)
5. Update or restore AD dependencies if required (service accounts, service principal names (SPNs), group policy objects (GPOs), and so on)

Backup and Restore DR metrics:

- Recovery Point Objective (RPO): >1hr or ~0.25hrs (depends on the solution selected - EBS or CE)
- Recovery Time Objective and (RTO): ~1 hr (depends on the number of instances and orchestration)
- Maintenance: High (Synchronous changes are required in both environments, like application configuration, patching, security groups or application load balancers, certificates, and so on.
- Cost: Medium

Disaster protection for EC2 with EBS snapshots on AMS

Prerequisites:

- AMS Prod Landing Zone (source)
- AMS DR Landing Zone (DR target)
- EBS snapshots are enabled for EC2 instances (AWS Backup)

Snapshot replication solution:

- **Cross AZ**: Not applicable - EBS snapshot are highly available within the Region by design
- **Cross-Region**: AWS Backup

The following diagram represents the EC2 restore process from EBS snapshots on AMS:

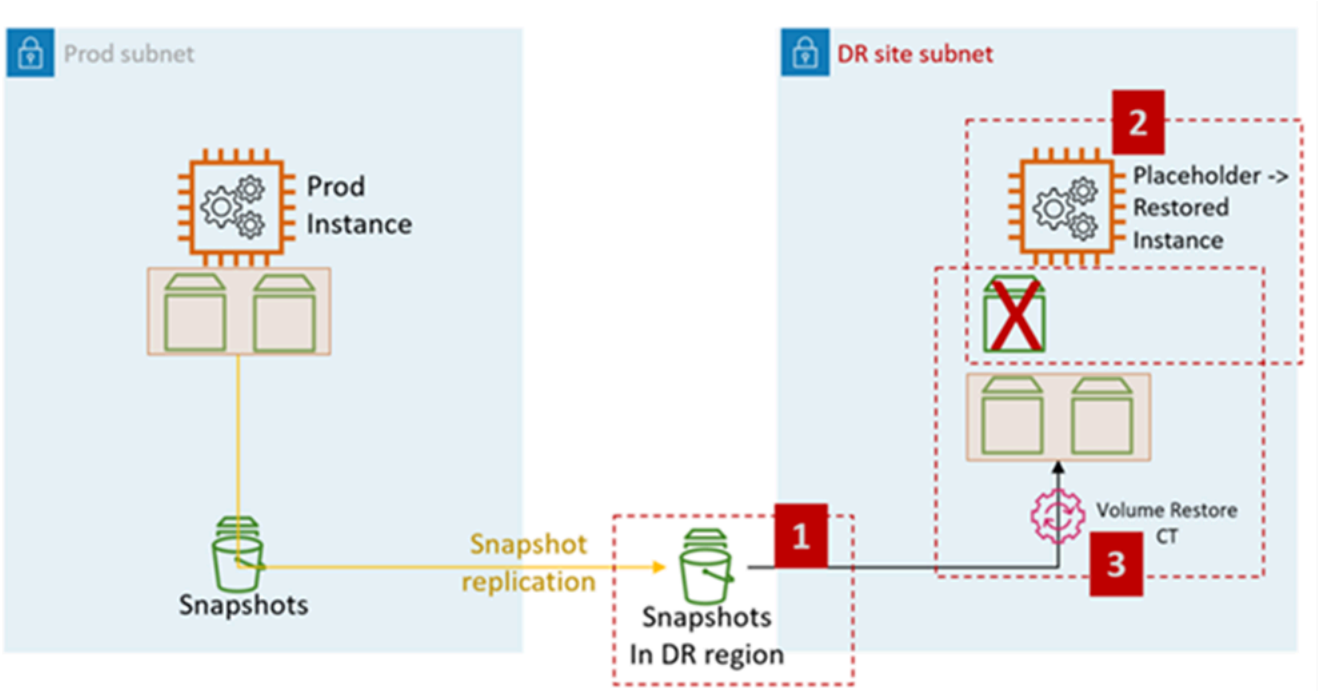

EC2 DR steps on AMS:

1. Raise an RFC to share the EBS snapshots with the target account (required for Cross-Region DR).

   : Management, Advanced Stack Components, EBS Snapshot, Share
2. Create a placeholder EC2 AMS stack in the destination subnet (DR site subnet). The recommendation is to use CFN ingestion to create the stack as the customer can combine the steps of assigning security groups and other (like adding the instance to an ELB) in the same stack.

   Change type: Deployment, Ingestion, Stack from CloudFormation Template, Create
3. Raise an RFC to perform EC2 stack volume restore.

   Change type: Management, Advanced Stack Components, EC2 instance stack, Restore volumes.

   The CT restores the volumes from the snapshots shared in step 1 and attaches to the placeholder instance created in step 2.

Volume Restore CT functionality:

- Shut the placeholder instance down
- Restore volumes from the snapshots

- Swap out the volumes
- Start the instance
- Leave the old domain
- Change the hostname
- Reboot. AMS bootstrap scripts join the instance to the target (DR) domain upon start up

Volume restore CT input:

- InstanceId (placeholder instance ID)
- RootDeviceSnapshotId, the EBS snapshot for the restored root volume
- KMSKeyId, the KMS key identifier, or ARN, to encrypt all restored volumes on the EC2 instance
- DeviceNames, up to 25 (optional)
- SnapshotIds, up to 25 (optional). List of snapshots of the volumes to be restored

Disaster protection for EC2 with CloudEndure on AMS

Prerequisites:

- AMS Prod Landing Zone (source)
- AMS DR Landing Zone (DR target)

Create an IAM user in your DR landing zone (LZ). See Migration CloudEndure documentation for the setup details (Note: The Migration subnet is not needed for the DR setup), the main difference will be that the setup is permanent as opposed to temporary like in case of migration. Other differences:

- No Migration subnet, skip steps 2, 3, 4, 5
- IAM user should have access to application subnets in the DR zone
- Step 8 is not performed, skip
- Important: Source instance should have Post Launch script configured. See this CE guide. The script will be provided by your AMS Cloud Architect.
- Destination (Placeholder) instance and its EBS volume should have a tag key: "CloudEndure creation time", value: <Anything>. Otherwise, CloudEndure won't be able to restore on top of the Placeholder instance.

The diagram below represents the CloudEndure setup and restore process for EC2 on AMS.

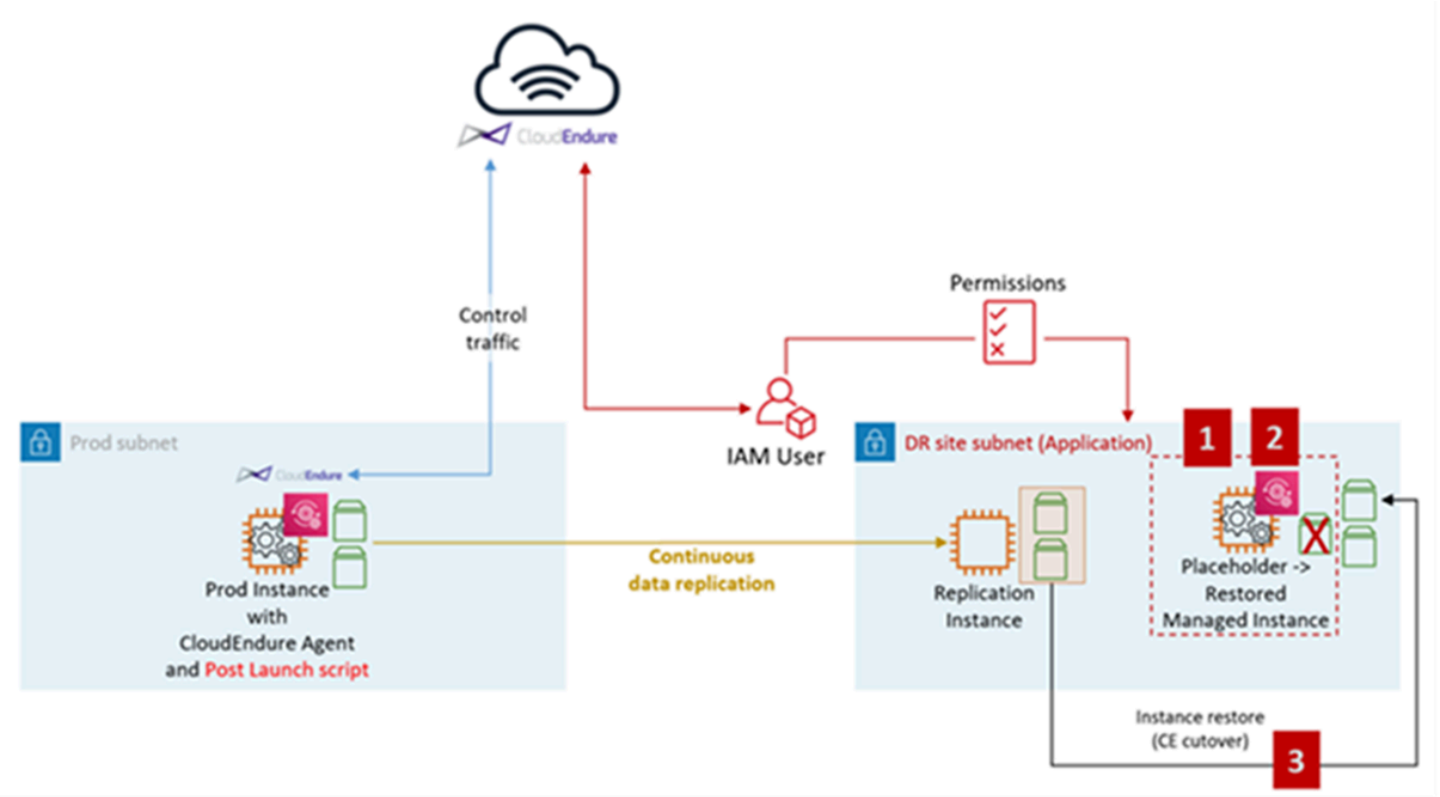

EC2 DR steps with CloudEndure on AMS:

1. Create a placeholder EC2 AMS stack in the destination subnet (DR site subnet) with proper tags, for more information, see the previous section. We recommend using CFN ingestion to create the stack as you can combine the steps of assigning security groups and tagging the instance, EBS volume, and other (like adding the instance to an ELB) in the same stack.

   Change type: Deployment, Ingestion, Stack from CloudFormation Template, Create

2. Stop the placeholder instance.

   Change type: Management, Advanced stack components, EC2 instance, Stop

3. If not done in step 1, tag the placeholder instance and its EBS volume with key: "CloudEndure creation time", value: <Anything>.

   Change type: Management, Advanced stack components, Tag, Update.

4. Initiate instance restore (cutover) from the CloudEndure console. Use the placeholder instance from step 1 as the target.

> **Note**
>
> The placeholder instance volumes are retained in the account. To delete these volumes, submit a Management | Other | Other change type at the end of the disaster recovery operation.

CloudEndure restore (cutover) workflow:

- The target (placeholder) instance needs to be in the stopped state
- Swap out the volumes and delete the source (placeholder) root volume
- Start the instance
- Run the Post Launch script, which will:
  - Leave the old domain
  - Change the hostname
  - Reboot. AMS bootstrap scripts join the instance to the target (DR) domain during the startup.

Patch management in AMS

Topics

- AMS Patch Orchestrator: a tag-based patching model
- Using Patch Orchestrator
- On-demand patching
- AMS standard patching
- Patching service commitments

In AMS, patch management is a service that helps you maintain OS vendor updates on your Amazon Elastic Compute Cloud (Amazon EC2) instances. You have the freedom to customize the frequency and process of patching your Amazon EC2 instances.

You configure patch management during onboarding, and you can update it by using the RFC process. Stacks created using the change management system and a patch-compatible template (for Amazon EC2, Auto Scaling group, HA one-tier or two-tier stack) are subscribed to patch management automatically.

AMS provides a feature, Patch Orchestrator – tag-based patching, for configuring patching.

For definitions of patching terms, see AMS key terms.

Important

- It's not possible for stacks or a stack's constituent instances to opt out of patch management, if the AMS template from which the stack is created is compatible with patch management. Currently, patching is compatible with the following stack templates:
  - Amazon EC2 stack | Create, and Amazon EC2 stack | Create (with additional volumes)
  - Amazon EC2 instance launched with AWS CloudFormation ingest
  - Auto Scaling group | Create (the Amazon EC2 instances in the group are patched)
  - High Availability One-Tier stack | Create, and High Availability Two-Tier stack | Create
- If there is an ongoing incident that affects a stack, AMS operators can reschedule or cancel scheduled patching.

- By default, all instances within a particular patch-compatible stack are patched in-place. To patch Auto Scaling groups with an Amazon Machine Image (AMI) replacement using the latest/patched AMS AMI, submit a service request. Updated AMIs are shared to accounts every month.

Tip

AMS recommends that you enable backups for instances that have valuable applications or services. For information about enabling backups, see Continuity management in AMS Advanced.

AMS Patch Orchestrator: a tag-based patching model

If you have been onboarded to the new AMS Patch Orchestrator tag-based patching model, you can use tags to apply your patch configuration to a precise set of resources, called a *patch group*, ranging from one instance to all of your instances. For information about AMS tags, see Using tags. Instructions on setting up Patch Orchestrator tags are provided in the following section.

Patches are installed during the patch windows you define with the SSM Patch Window | Create. Each patch window is an AWS Systems Manager maintenance window that runs on a schedule of your choice, has a configured duration, and applies to one patch group. Instances that are not part of an explicit patch window are patched during the default maintenance window that you define when you onboard to Patch Orchestrator.

Important

If multiple patch maintenance windows are scheduled to run at the same time, they must have fewer than 1001 instances being processed at any given time. This is an AWS Systems Manager limitation. AMS recommends at least one hour per every fifty instances.

By default, all operating system (OS) vendor-provided patches are installed during a maintenance window or an on-demand patch. This is called the *default patch baseline*. If you would like to restrict which patches are installed, you can define a custom patch baseline with one of the patch baseline create CTs (per OSes), see Patching subcategory. For example, you can use a custom patch

baseline so that only critical and important security updates are installed for one or more patch groups.

After patches are installed on an instance, the instance is rebooted. Patch notifications are sent before and after patching, and an additional reminder is sent within 96 hours before the scheduled start. In addition, AMS applies updates to infrastructure management tools (such as the AWS SSM agent) during the selected maintenance window.

> **Important**
>
> AMS is deprecating the monthly patch compliance reporting of instances with missing patches, and will not be sending monthly reports. This change has been made in view of the recently released self-serve operational reports that refresh every 24 hours and are available to you on demand and provide the most recent and granular data. To learn more about the reports, see Self-service reporting. To learn more about the reports, see Self-service reporting.

For more information on the notifications, see Patch notifications.

Using Patch Orchestrator

Enable AMS Patch Orchestrator for your account by submitting a service request that includes the following details:

- **Category**: Other
- **Subject**: Onboard to Patch Orchestrator
- **CC Emails**: CC email addresses receive notifications when the status of this onboarding RFC changes
- **Details**: Paste the following information into the email and provide your values. Note that the ThirdTagKey is optional. For recommendations and examples, see the following table.

```
Default maintenance window Schedule:
Default Maintenance Window Schedule TimeZone:
Default Maintenance Window Duration:
Default Maintenance Window Cutoff:
Default Patch Backup Retention In Days:
Default Maintenance Window Notification Emails:
First Tag Key:
```

```
Second Tag Key:
Third Tag Key:
```

The following table describes the format and recommendations for your provided values.

Patch orchestrator tag-based patching configurations

Name of parameter	Information	Recommendation or example
Default Maintenance Window Schedule	The schedule of the default maintenance window in the form of a cron or rate expression. For example: • `cron(0 3 ? * 6L *)`: 03:00 am on the last Friday of every month • `rate(7 days)`: Every seven days For more information about creating cron expressions, and links to cron and rate expression resources, see Cron and rate expressions for maintenance windows.	We recommend having the window run at least once per month on a consistent weekday.
Default Maintenance Window Schedule Time Zone	The time zone that the default maintenance window runs are based on, in Internet Assigned Numbers Authority (IANA) format.	For example: • America/Los_Angeles • etc/UTC
Default Maintenance Window Duration	The duration of the default maintenance window in hours.	At least 1 hour per every 50 instances, plus 2 hours for cutoff.

Name of parameter	Information	Recommendation or example
Default Maintenance Window Cutoff	The number of hours before the end of the Default Maintenance Window in which no new patching commands are started. This interval exists to allow enough time for patching to complete before the window ends.	At least 2 hours.
Default Patch Backup Retention In Days (optional)	The default time in days to keep the EBS restore points created before patching instances.	We recommend keeping the default, which is 60.
Default Maintenance Window Notification Emails	One to five email addresses or distribution lists to receive notifications about default maintenance window patching status.	We recommend using group distribution lists instead of individual emails.
First Tag Key	The first tag-key to use for creating your Patch Group tag values.	For example, AppId. Specify **null** if you already have defined your own patch groups with a Patch Group tag.
Second Tag Key	The second tag-key to use for creating your Patch Group tag values.	For example, Environme nt. Specify **null** if you have already defined your own patch groups with a Patch Group tag.

Name of parameter	Information	Recommendation or example
Third Tag Key (optional)	The optional third tag-key to use for creating your Patch Group tag values.	For example, Group.

After you're onboarded to the new Patch Orchestrator patching service model, all appropriately tagged instances in your account belong to a patch group with a Patch Group tag. Patch Orchestrator uses either your existing Patch Group tag, or an AMS-created tag consisting of the two or three concatenated tag values that you specified during Patch Orchestrator onboarding. For example, {*Tag Value 1*}-{*Tag Value 2*}-{*Tag Value 3*}. AMS updates these AMS-applied Patch Group tags every 12 hours. If needed, you can update your Patch Group tag values with the Tag | Update (Review Required) or Tag | Update (Review Required) change types.

For example, if your Amazon EC2 instance has the following tag key:value pairs:

- `AppId:MyApplication`
- `Environment:Production`
- `Group:1`

During onboarding you specified the following tag keys:

- `First Tag Key = AppId`
- `Second Tag Key = Environment`
- `Third Tag Key = Group`

AMS creates the following Patch Group tag and applies it to your instances: `Patch Group:MyApplication-Production-1`.

> **Note**
>
> Patch failure alerts aren't created for instances that have unsupported operating systems, or that are stopped during the maintenance window.

Patch Orchestrator prerequisites

Patch Orchestrator workflow targets Amazon EC2 instances that are patched by latest version of System Manager Automation Document: AWSManagedServices-PatchInstanceFromMaintenanceWindow.

As part of the document workflow, the run command document "AWS-RunPatchBaseline" is run against each of the Amazon EC2 instances out of patch group members. To learn more, see About the SSM document AWS-RunPatchBaseline.

Requirements:

- Amazon EC2 instance deployed from AMS-provided Amazon Machine Image (AMI), or on an AMI through the "Stack from migration partner migrated instance" CT (ct-257p9zjk14ija).
- Egress internet connection enabled. For firewall/proxy solutions the requirement is to allow Windows update endpoint and/or Linux repository mirror endpoints, AWS system manager proxy settings, and metadata proxy configuration. For more information, see Configure SSM Agent to use a proxy and Using an HTTP proxy
- IAM role matching minimum permissive access for the SSM service of `customer-mc-ec2-instance-profile` IAM role.
- We recommend 10 GB available root partition space. For Linux OS, at least 2 GB available in the `/var` partition.
- Working and valid Certificate Authority for update downloads.
- Windows Server Update Services (WSUS) - Registry including but not limited to: DisableWindowsUpdateAccess, NoWindowsUpdate; Automatic Updates must not impair operation of Windows Update process.

Validation:

- For Linux OS instances using yum package manager you can validate availability of updates by running `#yum check-update`
- For Linux OS RedHat 5.7 and newer, 6.1 and newer, and 7.0 and newer; Amazon EC2 instances migrated to your AMS account via the "Stack from migration partner migrated instance" CT (ct-257p9zjk14ija), you need to validate subscription manager status for update performance.
- On Windows OS, enable Windows Server Update Services (WSUS). No local policy should block WSUS ability to scan or install updates. Once logged as administrator you can validate it by

performing a scan for available updates from Windows Update Service console. Windows Server OS releases including 2012R2, 2016 and 2019 have default Windows Update settings to download and install. You can configure desired settings prior to scan. On later releases of OS, this operation can trigger installation; configure desired behavior beforehand.

- Request validation from the AMS Operations team by submitting a service request: "AWSManagedServices-CheckPatchingPrerequisites Automation document to run against Amazon EC2 instance for assessment of patch readiness."

> **Note**
>
> Patch failure alerts aren't created for instances that have unsupported operating systems, or that are stopped during the maintenance window.

Patch windows

Instances in a specific patch group are patched during one or more patch windows. Patch windows run on a schedule defined as a cron or rate expression, and have a configurable duration intended to keep patching-related disruption within a chosen time interval. AMS recommends creating multiple patch windows that collectively cover all of your instances, to match your organization's specific patching routines, and to use the default maintenance window as a fallback. Patch windows are created with the RFC change type Deployment | Patching | SSM patch window | Create (ct-0el2j07llrxs7). All instances that are not part of a patch window are patched during the default maintenance window created during onboarding.

Normally, a patch window does not need to be updated to include new instances. Typically, this is done by modifying instance tags. For example, consider the following sequence of events:

1. Two instances are tagged with `AppId:MyApplication`, `Environment:Production`, `Group:1`.

   This produces a tag on these instances, assuming First Tag Key = AppId, Second Tag Key = Environment, Third Tag Key = Group and a patch window for MyApplication-Production-1 patch group is created.
2. Three more instances are created and tagged with `AppId:MyApplication`, `Environment:Production`, `Group:1`.

This produces a tag for Patch Group:MyApplication-Production-1.

No change to the patch window is needed because it picks up all five instances at the time of the next scheduled run.

For a more detailed discussion and a walkthrough on using this change type, see SSM Patch Window | Create.

Patch notifications

The subscribed email addresses (up to five) receive an email similar to the following just before the patch maintenance window start:

```
Dear Customer,
The AMS Patch Maintenance Window THE_MAINTENANCE_WINDOW_NAME was started at:
 2020-02-21T12:02:18.196Z.
Details:
    Maintenance Window AccountId:   YOUR_ACCOUNT_ID
    Maintenance Window Region:      YOUR_ACCOUNT_REGION
    Maintenance Window Id:          THE_MAINTENANCE_WINDOW_ID
    Maintenance Window Name:        THE_MAINTENANCE_WINDOW_NAME
    Maintenance Window Description: MaintenanceWindow for patching patch
 Group PATCH_GROUP_NAME
    Maintenance Window Patch Group: PATCH_GROUP_NAME
    Maintenance Window ExecutionId: THE_EXECUTION_ID
Targets:
    InstanceId           InstanceName        StackId
    -------------------  --------------      ----------------------
    THE_INSTANCE_ID      THE_INSTANCE_NAME   THE_STACK_NAME

A follow-up message with a detailed report is sent as soon as the maintenance window is
 over.

Please raise a service request if you have any inquires about AMS Patch Orchestrator by
 following this URL:
https://console.aws.amazon.com/managedservices/servicerequest/new

Kind Regards,

Amazon Web Services
Amazon Managed Services
```

```
Patch Team
```

At the end of the patch activity, the subscribed email addresses receive an email similar to the following:

```
Dear Customer,
The AMS Patch Maintenance Window THE_MAINTENANCE_WINDOW_NAME ended at:
 2020-02-21T12:03:20.058Z, with status: SUCCESS.
Details:
    Maintenance Window AccountId:   YOUR_ACCOUNT_ID
    Maintenance Window Region:      YOUR_ACCOUNT_REGION
    Maintenance Window Id:          THE_MAINTENANCE_WINDOW_ID
    Maintenance Window Name:        THE_MAINTENANCE_WINDOW_NAME
    Maintenance Window Description: MaintenanceWindow for patching patch
 Group PATCH_GROUP_NAME
    Maintenance Window Patch Group: PATCH_GROUP_NAME
    Maintenance Window ExecutionId: THE_EXECUTION_ID
Targets:
    RfcId                        InstanceId        InstanceName          StackId
       Status
    ----------------------  ------------------  --------------
 ----------------------  --------
    THE_RFC_ID        THE_INSTANCE_ID     THE_INSTANCE_NAME THE_STACK_NAME
  STATUS

You can view the current Patch Compliance of your Amazon EC2 Instances by following
 this URL:
https://console.aws.amazon.com/systems-manager/compliance?region=YOUR_ACCOUNT_REGION

Please raise an Incident if an issue is impacting one of your production applications
 by following this URL:
https://console.aws.amazon.com/managedservices/incident/new

Kind Regards,

Amazon Web Services
Amazon Managed Services
Patch Team
```

Every 96 hours AMS patching system identifies all upcoming patch managed maintenance windows within that 96 hours and sends a reminder notification to all subscribed email addresses that fall

within that 96 hour window. This could be as little as one hours before the window, or the full 96 hours. For example:

```
Dear Customer,
The AMS Patch Maintenance Window THE_MAINTENANCE_WINDOW_NAME will start at:
 2020-05-06T16:35:36.523Z.
Details:
    Maintenance Window AccountId:   YOUR_ACCOUNT_ID
    Maintenance Window Region:      YOUR_ACCOUNT_REGION
    Maintenance Window Id:          THE_MAINTENANCE_WINDOW_ID
    Maintenance Window Name:        THE_MAINTENANCE_WINDOW_NAME
    Maintenance Window Description: MaintenanceWindow for patching patch
 Group PATCH_GROUP_NAME
    Maintenance Window Patch Group: PATCH_GROUP_NAME
    Maintenance Window Next Start Time: 2020-05-06T16:35:36.523Z
    Maintenance Window Schedule:        rate(24 hours)
    Maintenance Window Timezone:        THE_TIMEZONE
At this time, these are the instances in the "PATCH_GROUP_NAME" Patch Group:
    InstanceId           InstanceName    StackId                 InstanceState

-------------------------------------------------------------------------------------------

    THE_INSTANCE_ID      THE_INSTANCE_NAME   THE_STACK_NAME  running/stopped
    THE_INSTANCE_ID      THE_INSTANCE_NAME   THE_STACK_NAME  running/stopped
    THE_INSTANCE_ID      THE_INSTANCE_NAME   THE_STACK_NAME  running/stopped

A notification message is sent as soon as the maintenance window starts.

You can view the current Patch Compliance of your Amazon EC2 Instances by following
 this URL:
https://console.aws.amazon.com/systems-manager/compliance?region=YOUR_ACCOUNT_REGION

If you would like to disable this maintenance window or you have inquires about the AMS
 Patch Orchestrator click on the following URL:
https://console.aws.amazon.com/managedservices/servicerequest/new

If you would like to delete this maintenance window, you can run the CT with id
 "ct-0q0bic0ywqk6c" against the stack id "stack-rctyznutkyj4tkkzq".

Kind Regards,

Amazon Web Services
Amazon Managed Services
```

```
Patch Team
```

Patch baselines

By default, all operating system (OS) vendor-provided patches are installed using the AMS-default patch baseline. If you want to restrict which patches are installed, you can optionally create a patch baseline using the RFC change type Deployment | Patching | SSM patch baseline | Create *OS* (CT ID varies per operating system).

For information about using these change types, see Patching subcategory.

Patch Orchestrator reserved tags

Patch Orchestrator also generates the following tags that can't be modified:

- **AMSPatchGroup** – This tag is used for Patch Group tag value generation. You shouldn't modify the AMSPatchGroup. You can modify the "Patch Group" tag if you want to use a custom "Patch Group" value. Patch Orchestrator continues generating a value for AMSPatchGroup based on the tag-keys provided during onboarding, but won't modify the "Patch Group" tag value if it has been set to a custom value by you. To stop using a custom "Patch Group" value, you can set the value of "Patch Group" to match the AMSPatchGroup tag value.
- **AMSDefaultPatchGroup** – This tag indicates whether an instance is part of the default maintenance window, with a value of either True or False. If an instance's Patch Group is not assigned to a maintenance window this value is set to True.

On-demand patching

AMS has a change type that works with your patch baseline, to enable you to run a patch on instances on demand. This can be either the default baseline you set at on boarding, or the Patch Orchestrator Systems Manager patch baseline that you set with the Patch Baseline change type (CT ID varies per operating system).

You can use the on-demand patching change type with or without Patch Orchestrator.

For information about using this change type, see On Demand Patching | Run.

> **Note**
>
> You can't use instances that are part of an Auto Scaling group in an on-demand patching change type.

AMS standard patching

AMS supports existing customers using the AMS standard patching model, but this model is not available for new customers and is being retired in favor of AMS Patch Orchestrator.

Typical patch contents for AMS standard patching include vendor updates for supported operating systems and software preinstalled with supported operating systems (for example, IIS and Apache Server).

During AMS onboarding, you specify patching requirements, policy, frequency, and preferred patch windows. These configurations mean you can avoid taking applications offline all at once for infrastructure patching, so you can control what infrastructure gets patched when.

> **Note**
>
> The patching process described in this topic applies only to your stacks. AMS infrastructure is patched during a separate process. The AWS Managed Services Maintenance Window (or Maintenance Window) performs maintenance activities for AWS Managed Services (AMS) and recurs the second Thursday of every month from 3 PM to 4 PM Pacific Time. AMS may change the maintenance window with 48 hours notice. You configure the AMS patch window at onboarding, or you approve or reject the monthly patch service notification.

AMS regularly scans managed Amazon EC2 instances for updates available through the operating system update functionality. We also provide regular updates to the AMS base Amazon Machine Images (AMIs) supported in our environment.

After they are validated, AMS AMI releases are shared with all AMS accounts. You can view the available AWS AMI releases by using the DescribeImages Amazon EC2 API call or using the Amazon EC2 console. To find available AMS AMIs, see Find AMI IDs, AMS.

AMS performs ad hoc patching schedules only when requested by you.Previously AMS would send a notification; currently, a notification is not sent.

Note

By default, AMS uses Systems Manager to apply patches by having the package manager (Linux) or System Update service (Windows) query its default repository to see which new packages are available. If, during the course of your day-to-day operations, you have installed a package on a Linux host using the default package manager, that package manager also picks up new packages for that software when they're available. In such a case, you may want to take a patching action (described in this section) to opt-out for that instance.

Supported operating systems

- Amazon Linux 2023
- Amazon Linux 2 (**expected AMS support end date June 30, 2025**)
- CentOS 7.x
- Oracle Linux 7: minor versions 7.5 and above. Oracle Linux 8: minor versions up to 8.3.
- Red Hat Enterprise Linux (RHEL) 8.x, 7.x
- SUSE Linux Enterprise Server 15 SP5 and SAP specific versions, SUSE Linux Enterprise Server 12 SP5 and SAP specific versions.
- Microsoft Windows Server 2022, 2019, 2016

Supported patches

AWS Managed Services supports patching primarily at the operating system level. The patches that are installed may differ by operating system.

Important

All updates are downloaded from the Systems Manager patch baseline service remote repositories configured on the instance, and described later in this topic. The instance must be able to connect to the repositories so the patching can be performed.
To opt-out of the patch baseline service for repositories that deliver packages that you want to maintain yourself, run the following command to disable the repository:

```
yum-config-manager DASHDASHdisable REPOSITORY_NAME
```

Retrieve the list of currently configured repositories with the following command:

```
yum repolist
```

- **Amazon Linux** preconfigured repositories (usually four):

Repository ID	Repository name
amzn-main/latest	amzn-main-Base
amzn-updates/latest	amzn-updates-Base
epel/x86_64	Extra Packages for Enterprise Linux 6 - x86_64
pbis	PBIS Packages Updates

- **Red Hat Enterprise Linux** preconfigured repositories (five for Red Hat Enterprise Linux 7 and five for Red Hat Enterprise Linux 6):

Repository ID	Repository name
rhui-REGION-client-config-server-7/x86_64	Red Hat Update Infrastructure 2.0 Client Configuration Ser
rhui-REGION-rhel-server-releases/7Server/x86_64	Red Hat Enterprise Linux Server 7
rhui-REGION-rhel-server-releases/7Server/x86_64	Red Hat Enterprise Linux Server 7 RH Common(RPMs)
epel/x86_64	Extra Packages for Enterprise Linux 7 - x86_64

Repository ID	Repository name
pbis	PBIS Packages Updates

Repository ID	Repository name
rhui-REGION-client-config-server-6	Red Hat Update Infrastructure 2.0
rhui-REGION-rhel-server-releases	Red Hat Enterprise Linux Server 6 (RPMs)
rhui-REGION-rhel-server-rh-common	Red Hat Enterprise Linux Server 6 RH Common (RPMs)
epel	Extra Packages for Enterprise Linux 6 - x86_64
pbis	PBIS Packages Updates

- **CentOS 7** preconfigured repositories (usually five):

Repository ID	Repository Name
base/7/x86_64	CentOS-7 - Base
updates/7/x86_64	CentOS-7 - Updates
extras/7/x86_64	CentOS-7 - Extras
epel/x86_64	Extra Packages for Enterprise Linux 7 - x86_64
pbis	PBIS Packages Updates

- For **Microsoft Windows Server**, all updates are detected and installed using the Windows Update Agent, which is configured to use the Windows Update catalog (this doesn't include updates from Microsoft Update).

  On Microsoft Windows operating systems, Patch Manager uses Microsoft's cab file wsusscn2.cab as the source of available operating system security updates. This file contains information

about the security-related updates that Microsoft publishes. Patch Manager downloads this file regularly from Microsoft and uses it to update the set of patches available for Windows instances. The file contains only updates that Microsoft identifies as being related to security. As the information in the file is processed, Patch Manager also removes updates that have been replaced by later updates. Therefore, only the most recent update is displayed and made available for installation. For example, if KB4012214 replaces KB3135456, only KB4012214 is made available as an update in Patch Manager.

To read more about the wsusscn2.cab file, see the Microsoft article Using WUA to Scan for Updates Offline.

Patching and infrastructure design

AMS employs different patching methods depending on your infrastructure design: mutable or immutable (for detailed definitions, see AMS key terms).

With mutable infrastructures, patching is done using a traditional in-place methodology of installing updates directly to the Amazon EC2 instances, individually, by AMS operations engineers. This patching method is used for stacks that are not Auto Scaling groups, and contain a single Amazon EC2 instance or a few instances. In this scenario, replacing the AMI that the instance or stack was based on would destroy all of the changes made to that system since it was first deployed, so that is not done. Updates are applied to the running system, and you may experience system downtime (depending on the stack configuration) due to application or system restarts. This can be mitigated with a Blue/Green update strategy. For more information, see AWS CodeDeploy Introduces Blue/Green Deployments.

With immutable infrastructures, the patching method is AMI replacement. Immutable instances are updated uniformly using an updated AMI that replaces the AMI specified in the Auto Scaling group configuration. AMS releases updated (that is, patched) AMIs every month, usually the week of Patch Tuesday. The following section describes how this works.

How AMS standard patching works

AMS uses the Systems Manager Run Command service for regularly scheduled monthly and as-needed critical patching, with two principal patching methods, in-place and AMI replacement, depending on your infrastructure deployment strategy (mutable vs. immutable). This section describes the AMS patching service, types, methods, and processes.

AMS defines two patch types, which are scheduled differently:

- *Critical patching*: Updates are applied as quickly as possible, after acceptance of the notice.
- *Standard patching*: Regular OS vendor updates and applied monthly.

Patches are applied through either in-place patching or AMI replacement (upon request).

Update scanning

AMS uses the Amazon EC2 Run Command Service to contact your Amazon EC2 stacks and deploy the required scanning and patching scripts. AMS uses the native package management component already installed on the supported operating system to perform all the required scanning and patching behavior on the Amazon EC2 stack. For Red Hat and Amazon Linux, the service uses yum. For Windows, the service uses the Windows Update Agent.

Scans are performed daily using SSM Maintenance Windows and the AMS default AWS-RunPatchBaseline document. Every reachable Amazon EC2 stack is scanned, using the update repositories for Linux and Windows. The AMS patching process detects all reachable Amazon EC2 stacks and then performs the scans in a batch process so that the stack always remains in a healthy state, even if a failure occurs while running the scan. The scan results are then saved for each Amazon EC2 stack.

To view the scan results for a stack or instance, submit a service request with the stack ID or instance ID.

The default AMS patching process is to install all available patches regardless of patch classification or severity (for example, critical versus standard). The exception to this are patches that you have explicitly excluded for the stack (patches defined as mandatory by AMS should not be excluded).

You're sent a patching service notification 14 days before the proposed maintenance window. This gives you time to test the proposed patches and accept or reject them. If you don't reply to the patching service notification, your instances aren't patched. When the time comes to install the patches, AMS creates a Request for Change (RFC) for each stack, and that RFC appears in your account's RFC list.

AMS configured maintenance window and notice

With AMS configured patching, each account has a monthly maintenance window, which you define when you onboard your account. The AWS Managed Services Maintenance Window (or Maintenance Window) performs maintenance activities for AWS Managed Services (AMS) and

recurs the second Thursday of every month from 3 PM to 4 PM Pacific Time. AMS may change the maintenance window with 48 hours notice.

The patching window is different. The patching outbound service request (also known as a *service notification*) includes a suggested patch window.

> **Note**
>
> For information about replying to the patching service notification, see Actions you can take in AMS standard patching.

The patching service notification is sent by email to the contact email address on file for your account. The notification includes a link to the AWS Support console where you can respond to it. You can also respond to the notification using the AMS Service Request page. The service notification includes:

- A list of update IDs (CSUs, IUs, and OUs) that apply to the stack, and those updates that you have requested be excluded from patching (if any).
- IDs of instances that will be affected.
- A proposed patching window when the updates will be applied. You can request a different patching window.
- A request that you accept the proposed patching, or ask for additional information. AMS gives you time to test the impact of the updates and approve or reject the patching, or ask that specific updates be excluded. If you need more time to test, and want the updates to be applied after your testing, respond to the service notification and describe what you want, or submit a service request for a new patch RFC based on the details of the previous RFC. If you don't reply to the service notification at all, no patching action is taken and the RFC is cancelled.

If you approve the service notification, AMS runs the patch RFC and applies the updates within the agreed-to patch window, as per the service commitment.

When patching is finished, AMS sends you a correspondence in the Service Request, with a summary of the outcome of the patching activity (that is, success or failed).

In-place patching

In-place patching refers to a method where AMS logs into each stack instance and applies patches.

In-place patching occurs on mutable infrastructures using Amazon EC2 instances running a supported operating system. Patching applies all non-excluded updates available up to that point. When critical patches are released, there is an additional critical patching process.

Standard patching: in-place

Standard patching occurs on the agreed-to patch schedule suggested in the patch service notification, and includes regular patch updates that are not deemed critical.

Prior to the proposed patching window, and with your affirmative response to the notification, a patch RFC is created and appears in your RFC dashboard.

Critical patching: in-place

When an OS vendor releases a critical security update, AMS notifies you of the patch RFC by sending you a service notification (to the contact email for your account) for each stack, according to the AMS service commitment. The service notification includes the following for each update:

- Update release date
- Update criticality
- Update details (KB reference, etc.)
- IDs of stacks affected

You can test the updates listed in the notification, and approve or reject the patches by replying to the service notification. If you approve the notification, you need to provide a specific patch window per stack for installing the updates.

> **Note**
>
> Patch windows that are within 24 hours of reply to the service notification may be rescheduled based on available capacity.

If you don't reply within 10 days or if you reject the proposed patching, the patching is canceled.

If you want to apply the updates after the allowed period (provided in the notification), submit a service request for a new patch schedule based on the details of the previous notification.

If you approve the service notification, AMS applies the updates within your specified patch window, according to the service commitment.

In the case of multiple updates, you can exclude specific updates from the patching by specifying the updates to be excluded in your response to the service notification.

AMS sends you a service notification for each stack, of the outcome of each update (that is, success or fail).

AMI updates patching (using patched AMIs for Auto Scaling groups)

AMI-replacement patching is done on immutable infrastructures by updating the AMI ID that is configured to deploy new Amazon EC2 instances in an Auto Scaling group.

Amazon Machine Images (AMIs) are released on a regular basis for the supported operating systems. Operating system vendors release new patches on a periodic basis. AMS takes the Amazon-provided AMI, updates it with the latest patches, and then adds the appropriate components to enable it to operate in the AMS environment. Then, it makes the new AMS AMI available to all AMS customers by sharing the AMI to the accounts. Your Auto Scaling group stacks can be refreshed on a monthly basis with these newly released AMS AMIs. The following graphic illustrates how AMIs are used in AMS your environments.

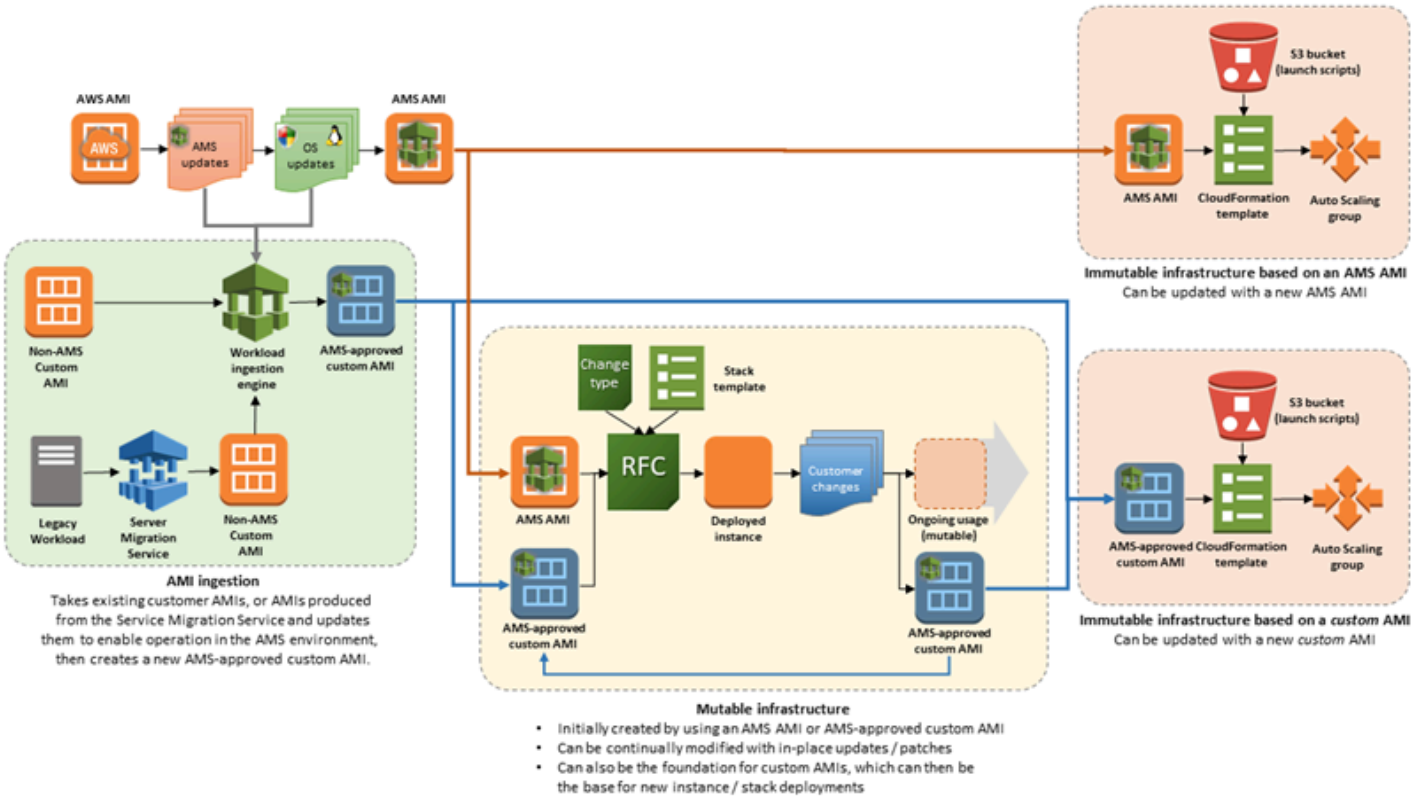

Auto Scaling groups create their instances based on the configured AMI for the Auto Scaling group. When AMS shares updated AMIs, you have the following options depending on how you are managing AMI updates:

- If you are using an application deployment tool (for example, UserData, CodeDeploy, and so forth) that customizes your instances automatically after they are created, you can do the following:
  - Reply to the patching service notification, or submit a service request, for the latest AMS AMI to replace your current Auto Scaling group's configuration AMI. After the AMI ID in your Auto Scaling groups' configuration is replaced, AMS kicks off rolling updates of your instances and your Auto Scaling group instance configurations (for example, installing applications, boot scripts, etc.) are applied to the new instances created with the new AMS AMI automatically.
- If you are using a custom/golden AMI in your Auto Scaling groups' configuration, you can:
  - Create an instance with the new AMS AMI, customize the instance and create a new golden AMI. Share the new golden AMI with AMS using the Amazon EC2 console, and submit a service request to AMS to update your Auto Scaling groups' configuration to use your new custom AMI.
  - Share your existing golden AMI with AMS by using the Amazon EC2 console, and submit a service request for AMS to update your golden AMI. To do this, AMS creates an instance from your golden AMI, applies the patches to that instance, creates a new golden AMI for you, and then updates your Auto Scaling groups' configuration to use the new AMI. The drawback here is that AMS cannot test that your new custom AMI works the way you want it to. Instead, you should test the instance created with the new AMI and verify that everything works correctly before creating a new golden AMI, sharing it, and requesting that AMS update your Auto Scaling groups. AMS does not recommend this option.

Standard patching: AMI updates

Every month AMS releases new Amazon Machine Images (AMIs) with service improvements and new patches that apply to the AMIs.

> **Note**
>
> New AMS AMIs are generated after Patch Tuesday from updated AWS AMIs. Then, AMS tests them before making them available. After the new AMIs pass testing, AMS shares updated AMIs to managed accounts.

Critical patching: AMI updates

When needed, AMS provides AMIs updated with critical security patches released since the last monthly AMI release.

The process for critical security updates to immutable infrastructures is identical to the monthly AMI process for immutable infrastructures, except that a new AMS AMI is created outside the normal schedule (Patch Tuesday), based on the release of new critical updates. AMS makes available a new AMI with the critical security patches according to the service level agreements (SLAs) defined for your account. AMS updates of Auto Scaling groups by request only. Use a service request to submit AMI replacement requests.

AMS standard patching failures

In case of failed updates, AMS performs an analysis to understand the cause of failure and communicates the outcome of the analysis to you. If the failure is attributable to AMS, we retry the updates if it's within the maintenance window. Otherwise, AMS creates service notifications for the failed instance update and waits for your instructions.

For failures attributable to your system, you can submit a service request with a new patch RFC to update the instances.

Actions you can take in AMS standard patching

In addition to testing new AMIs, there are several actions you can take to manage the patching of your infrastructure:

- If it took longer to test the updates than the patch window allowed, you can request that AMS apply the updates that were canceled when you're ready by submitting a service request (use the details in the original service notification as the basis).
- You can request that an important update (IU) or other update (OU) be applied before the next automated update window by submitting a service request providing a list of the updates, the applicable instances, and other details as appropriate. Since this CT is not automated, it takes longer to schedule and run. Check the service level objectives (SLOs) for the appropriate time. For more information, see AMS service level objectives (SLOs).

Additionally, you can use existing, patched, AMS AMIs to create custom AMIs. For information, see AMI | Create.

> **Note**
>
> You can't request a new AMS AMI based on an important update or other update before the next maintenance window because the AMS AMI release process follows a uniform cadence for the benefit of all AMS customers.

Changing what gets patched/opting out

With AMS configured patching, in your response to the patching service notification or in a Service Request, you can change what resources get patched. You can do the following:

- Define a list of patches that should be excluded from remediation, per stack and per operating system.
- Define a list of resources that should be excluded from certain patches or all patching.
- Define a list of resources that should be always be excluded from all patching.
- Define a list of resources that should be patched on a certain day and certain time (good if you haven't defined a maintenance window).

To exclude one or more patches, submit a service request, or respond to the patching service notification using the template provided next. Do not submit an RFC. Include in the request the patch name or names that you want excluded and why. Include this information in a Service Request as follows:

- Name: The name of the patch. For Windows patches, this is the KB name, such as KB3145384. For Linux patches, this is the package name, such as `openssh-6.6.1p1-25.61.amzn1.x86_64`.
- Reason: A comment indicating why the patch is being excluded.
- Expiration Time: The date/time when the exclusion expires.

If an excluded patch is already installed, it is removed.

The request is reviewed by an operator who will discuss it with you if excluding those patches poses a significant security risk. The expiry date for excluded patches is also negotiated. After the agreed upon expiry date, the exclusion expires, and the patch is installed on any subsequent patching.

Patches on the exclusion list are still returned in scan results, if applicable.

> **Note**
>
> Unlike Windows, Linux patches are version-specific. This distinction is important because new versions of an excluded patch are not automatically excluded. It is your responsibility to notify AMS to exclude new versions of a Linux patch if that's what you want to do.

Patch service notification reply templates

You must reply to patching service notifications, using the specified format, in order for patching to be performed on your instances. You should do this if you haven't already set a maintenance window with AMS.

When you reply to a service notification, use the format given.

If no maintenance window is set, let us know when to patch what as shown following:

```
UTC                 StartTime       StackId                     InstanceId (Optional)
2019-04-01          15:00           stack-123456789012          i-1234566789
2019-04-01          15:00           stack-123456789013          i-1234566784
2019-04-01          15:00           stack-123456789014          i-1234566783
2019-04-01          15:00           stack-123456789015          i-1234566782
```

If you have a set maintenance window and want certain resources to be excluded from certain patches, use the following format:

```
StackId                     InstanceId (Optional)       Exclude Patches
stack-123456789012          i-1234566789                PATCH
stack-123456789013          i-1234566784                PATCH
stack-123456789014          i-1234566783                PATCH
stack-123456789015          i-1234566782                PATCH
```

If you have a set maintenance window and want certain resources to always be excluded from all patching, use the following format:

```
StackId                     InstanceId (Optional)       Exclude Patches
stack-123456789012          i-1234566789                ALL
stack-123456789015          i-1234566782                ALL
```

Preparing for patching

To prepare your environment for automated patching, we recommend the following:

- Be sure you have a complete inventory of all instances to be patched.
- Ensure that your resources are backed up regularly as part of your Continuity of Business strategy. Additional backups are created as part of the patch sequence, and these are automatically deleted according to your configured Patch Orchestrator retention policy (default is 60 days).
- Ensure that all relevant licenses are up to date.
- Modify your stack maintenance windows to stagger patching so that testing stacks are patched before production stacks. That way, any errors with patching are found in the testing stacks and can be identified before production stacks are patched.

Viewing patch settings

To find out what your current patching configuration is you can do the following:

- Submit a service request to AMS with the query.
- Wait for a patch service notification. The patching notice advises you of all patches to be applied and instances to be patched, and also suggests a patch window.

You can submit a service request to modify the following:

- Scan Interval: The amount of time, in minutes, between compliance scans performed on instances of this stack.

  Default is `240` (4 hours).
- NotificationWindow: How far in advance (in minutes) of a scheduled change (patch) the notification should be sent to you.
  Default is `10080` (7 days).

AMS standard patching FAQs

This section provides answers to some frequently asked questions.

- Q: How do I opt out of patching globally?

A: To globally opt out of patching, file a service request. Note that you can't opt out of AMS mandatory patches. All stacks will continue to be scanned so that we can report on vulnerabilities.

- Q: How do I exclude specific stacks from patching?

  A: To permanently exclude specific stacks from patching, submit a service request. To exclude certain stacks from a particular patch cycle, respond to the upcoming patching notice with the list of stacks to exclude. For information, see Changing what gets patched/opting out. Note that you can't opt out of mandatory patches.

- Q: What happens if I don't approve a patching service notification?

  A: You have 14 days to approve a standard patching service request and 10 days to approve a critical patching notice. If you don't approve the service request within the time period, the service commitment is nullified and no patching occurs. In the case of mandatory patching, patches are applied regardless of response to the service request.

- Q: How do I exclude specific patches and packages from being installed?

  A: To permanently exclude specific patches or packages, submit a service request. To exclude certain patches or packages from a particular patch cycle, respond to the upcoming patching notice with the list of patches or packages to exclude. For details, see Changing what gets patched/opting out. Note that you can't opt out of mandatory patches.

- Q: What happens if a system fails as a result of patching?

  A: AMS monitors each system. AMS sends a service notification to you of the outcome of each update (that is, success or fail) per stack and instance. If a failure is detected, AMS investigates, works to restore the instance, and then an AMS operations engineer attempts to manually patch. For information, see AMS standard patching failures.

- Q: What updates are managed by AMS?

A: AMS manages operating system level updates that AMS is notified of by the vendor. For more information, see Supported patches.

- Q: What updates are not managed by AMS?

  A: Application-level updates are not managed by AMS.

- Q: How are Auto Scaling groups updated?

  A: Auto Scaling groups are updated with an AMI replacement in the Auto Scaling group configuration and preform a rolling update. A rolling update observes the HealthyHostThreshold setting of your patching configuration, which determines how many Amazon EC2 instances in a stack must be maintained active during patching. For more information, see AMI updates patching (using patched AMIs for Auto Scaling groups).

- Q: How do I get updates installed outside the normal cycle?

  A: For OS-level updates that you want installed outside of the normal patching schedule, submit a service request by using the patching notification that you received. This might happen if your testing of a proposed patch took longer than 21 days (for a standard patch) or 14 days (for a critical patch). Out-of-band patching can be done in-place for standalone Amazon EC2 instances.

- Q: How are newly deployed stacks or instances patched?

  A: When creating a new Amazon EC2 stack instance or Auto Scaling group, you should always specify the latest AMS AMI, which will have the latest patches on it already. For mutable infrastructures, inline patching should be performed as soon as the stack is deployed.

Patching service commitments

Based on your type of infrastructure deployment, and criticality of the update, we provide service commitments for critical security updates for mutable and immutable infrastructures, and important updates for mutable and immutable infrastructures.

Standard patching

These are AMS service commitments for standard patching.

Standard patching, mutable infrastructure (in-place patching)

Event/Action	Service commitment measurement
Important Updates are released in a month.	Clock starts
Fourteen days from when the standard patch notification is created, AMS notifies you of upcoming planned patching through a service notification and by email for each stack. The service notification includes: • A list of update IDs (CSUs and IUs) that are applicable (needed and not applied) for the stack, and those updates excluded from patching • IDs of stacks affected • The maintenance window when the updates will be applied	Clock stops after service notification is sent.
You test the impact of the updates and approve or reject the RFC. If you do not reply within 14 days or you reject the patching in your response to the service notification, no action is taken. If you take longer than 14 days to test, and want the updates to be applied after the 14-day period, submit a service request for a new	If you don't approve or reply within 14 days, the pending change is canceled and the service commitment for the updates is not applicable.

Event/Action	Service commitment measurement
patch RFC based on the details of the previous RFC.	
If you approve the service notification within 14 days, AMS applies the updates. You can choose to exclude specific updates from an RFC by specifying the updates to be excluded in your response to the service notification.	The clock starts if you approve service notification within 14 days of the receipt. The clock stops after the update installation has been attempted.
AMS sends a service notification to you of the outcome of each update that was attempted. The service notification includes the following details: • Amazon EC2 instance ID • Update 1 Success/Failed: ARN a1, ARN a2... • Update 2 Success/Failed: ARN b1, ARN b2... • Update N Success/Failed: ARN c1, ARN c2...	Not applicable.
In case of failed updates, AMS performs an analysis to understand the cause of failure and communicates the outcome of the analysis to you. If the failure is attributable to AMS, AMS retries the updates if within the maintenan ce window, otherwise AMS creates service notifications for the failed instance-update combination and waits for your instructions on a maintenance window.	Not applicable.
For failures attributable to you, submit a service request for a new patch RFC to update the instances.	Not applicable.

Critical patching

These are AMS service commitments for critical security updates.

Critical security updates, mutable infrastructure

Event/Action	Service commitment measurement
CSU is released.	Clock starts
AMS notifies you of the patch RFC through a service notification (which also sends an email) for each stack. The service notification includes: • Updates release date • Update criticality • Update details: KB reference, and so on • IDs of stacks affected	The clock stops after the service notification is sent.
You test the updates listed in the RFC, and approve or reject the RFC within 10 days by replying to the service notification. You provide a specific maintenance window (per stack) for installing the updates. Maintenance windows specified that are within 24 hours of reply to the service notification may be rescheduled based on available capacity. If you don't reply within 10 days, or if you reject the patch RFC, the pending action is canceled. If you want to apply the updates after the 14-day period, submit a service request for a new patch RFC based on the details of the previous RFC.	If you don't approve or reply within 14 days, the pending change is canceled and the service commitment for the update is not applicable.

Event/Action	Service commitment measurement
If you approve the service notification, AMS applies the updates. For multiple updates, you can choose to exclude specific updates from the change by specifying the updates to be excluded in your response to the service notification.	If the desired maintenance window is not within the service commitment time frame, the service commitment for the update is missed only if the RFC is not run within the desired maintenance window.
AMS sends a service notification to you of the outcome of each update that was applied. The service notification includes the following details: • Amazon EC2 instance ID • Update Success: ARN a1, ARN a2... • Update Failed: ARN c1, ARN c2...	Not applicable.
In case of failed updates, AMS performs an analysis to understand the cause of failure and communicates the outcome of the analysis to you. If the failure is attributable to AMS, AMS retries the updates if within the maintenan ce window, otherwise AMS creates service notifications for the failed instance-update combination and waits for your instructions on a new maintenance window.	Not applicable.
For failures attributable to you, submit a service request for a new patch RFC to update the instances.	Not applicable.

Critical security updates, immutable infrastructure

Event/Action	Service commitment measurement
CSU is released.	Clock starts
AMS notifies you of the following via a service notification: • Update release date • Update criticality • Update details (KB reference,and so on) • AMS Amazon Machine Images (AMI) impacted • Anticipated release date and time for new updated AMIs	Clock continues to run.
AMS releases updated AMIs in managed account.	Clock stops.
If you approve the service notification, AMS applies the updates. AMS notifies you of the AMIs shared in your account, through a service notification and by email.	Not applicable.
If testing the new AMIs takes longer than the allotted time (one week), you can submit a service request to AMS to update your Auto Scaling groups with the new AMS AMI (as is). If you want to modify the new AMS AMI with your configurations, use an RFC with the Management | Other | Other | Update CT (ct-0xdawir96cy7k) to request that we update your Auto Scaling groups.	Not applicable.

Reports and options

AWS Managed Services (AMS) collates data from various native AWS services to provide value-added reports on major AMS offerings.

AMS offers two types of detailed reporting:

- On request reporting: Certain reports can be requested ad hoc through your cloud service delivery manager (CSDM)
- Self-service reporting: You can generate some reports yourself

> **Note**
>
> Self-service reporting isn't available for patch and backup in the Jakarta AWS Region.

Topics

- On-request reporting
- Self-service reporting

On-request reporting

Topics

- AMS Patch reporting
- AMS Backup reporting
- Billing Charges Details report

AMS collates data from various native AWS services to provide value added reports on major AMS offerings. For a copy of these reports, make a request to your Cloud Service Delivery Manager (CSDM).

AMS Patch reporting

Topics

- Patch Instance Details Summary report

- Patch Details report
- Instances That Missed Patches report
- Patching SSM Coverage report

Patch Instance Details Summary report

The Patch Instance Details Summary report provides instance details gathered for instances that are onboarded to reporting. This is an informational report that helps identify all the instances onboarded, account status, instance details, maintenance window coverage, maintenance window execution time, stack details, and platform type. This report provides the following:

1. Insights into production and non-production instances of an account. Note: Production and non-production stage is derived from the Account Name and not from the Instance Tags.
2. Insights into distribution of instances by platform type. Note: 'N/A' platform type is when AWS Systems Manager can't retrieve the platform information.
3. Insights into the distribution of state of instances, and the number of instances running, stopped, or terminating.

Field Name	Definition
Report Datetime	The date and time the report was generated.
Account Id	AWS Account ID to which the instance ID belongs
Account Name	AWS account name
Production Account	Identifier of AMS prod, non-prod accounts, depending on whether account name include value 'PROD', 'NONPROD'. Example: PROD, NONPROD, Not Available
Account Status	AMS account status. For example: ACTIVE, INACTIVE
AMS account service commitment	PREMIUM, PLUS

Field Name	Definition
Landing Zone	Flag for account landing zone type. For example: MALZ, NON-MALZ
Access Restrictions	Regions to which access is restricted. For example: US SOIL
Instance Id	ID of EC2 instance
Instance Name	Name of EC2 instance
Instance Platform Type	Operating System (OS) type. For example: Windows, Linux, and so forth
Instance Platform Name	Operating System (OS) name. For example: MicrosoftWindowsServer2012R2Standard, RedHatEnterpriseLinuxServer
Stack Name	Name of stack that contains instance
Stack Type	AMS stack (AMS infrastructure within customer account) or Customer stack (AMS managed infrastructure that supports customer applications). Examples: AMS, CUSTOMER
Auto Scaling Group Name	Name of Auto Scaling Group (ASG) that contains the instance
Instance Patch Group	Patch group name used to group instances together and apply the same maintenance window. If the patch group is unassigned the value will be "Unassigned"

Field Name	Definition
Instance Patch Group Type	Patch group type. DEFAULT: default patch group with the default maintenance window, determined by the `AMSDefaultPatchGroup:True` tag on the instance. CUSTOMER: customer created patch group. NOT_ASSIGNED: no patch group assigned
Instance State	State within the EC2 instance lifecycle . Examples: TERMINATED, RUNNING, STOPPING, STOPPED, SHUTTING-DOWN, PENDING. For more information, see Instance lifecycle.
Maintenance Window Coverage	If there is a future Maintenance Window on this instance. Examples: COVERED or NOT_COVERED
Maintenance Window Execution Datetime	Next time the maintenance window is expected to execute. If NULL, single window execution, i.e. not recurring

Patch Details report

AWS Managed Services (AMS) Patch Details report provides patch details and maintenance window coverage of various instances, including:

1. Insights on Patch groups and its types.
2. Insights on Maintenance Windows, duration, cutoff, future dates of maintenance window executions (schedule) and instances impacted in each window.
3. Insights on all the operating systems under the account and number of instances that operating system is installed.

Field Name	Definition
Report Datetime	The date and time the report was generated.
Account ID	AWS Account ID to which the instance ID belongs
Account Name	AWS account name
Instance Id	ID of EC2 instance
Production Account	Identifier of AMS prod, non-prod accounts, depending on whether account name include value 'PROD', 'NONPROD'. If data is not available value will be "Not Available"
Account Status	AMS account status. For example: ACTIVE, INACTIVE
Instance Platform Type	Operating System (OS) type. For example: Windows, Linux
Instance Platform Name	Operating System (OS) name. For example: MicrosoftWindowsServer2012R2Standard, RedHatEnterpriseLinuxServer
Stack Type	AMS stack (AMS infrastructure within a customer account) or Customer stack (AMS managed infrastructure that supports customer applications). For example: AMS, CUSTOMER
Instance Patch Group	Patch group name used to group instances together and apply the same maintenance window. If the patch group is unassigned the value will be "Unassigned"
Instance Patch Group Type	Patch group type.

Field Name	Definition
	DEFAULT: default patch group w/ default maintenance window, determined by `AMSDefaultPatchGroup:True` tag on the instance CUSTOMER: customer created patch group UNASSIGNED: no patch group assigned
Instance State	State within the EC2 instance lifecycle. For example: TERMINATED, RUNNING, STOPPING, STOPPED, SHUTTING-DOWN, PENDING For more information, see Instance lifecycle.
Maintenance Window Id	Maintenance window identifier
Maintenance Window State	Possible values are ENABLED or DISABLED.
Maintenance Window Type	Maintenance window type
Maintenance Window Next Execution Datetime	Next time the maintenance window is expected to execute. If NULL, single window execution, i.e. not recurring
Last Execution Maintenance Window	The latest time the maintenance window was executed
Maintenance Window Duration (hrs)	The duration of the maintenance window in hours
Maintenance Window Coverage	The maintenance window coverage
Patch Baseline Id	Patch baseline currently attached to instance

Field Name	Definition
Patch Status	Overall patch compliance status. For example: COMPLIANT, NON_COMPLIANT. If there is at least one missing patch, instance is considered noncompliant, otherwise compliant.
Compliant - Total	Count of compliant patches (all severities)
Noncompliant - Total	Count of noncompliant patches (all severities)
Compliant - Critical	Count of compliant patches with "critical" severity
Compliant - High	Count of compliant patches with "high" severity
Compliant - Medium	Count of compliant patches with "medium" severity
Compliant - Low	Count of compliant patches with "low" severity
Compliant - Informational	Count of compliant patches with "informat ional" severity
Compliant - Unspecified	Count of compliant patches with "unspecified" severity
Noncompliant - Critical	Count of noncompliant patches with "critical" severity
Noncompliant - High	Count of noncompliant patches with "high" severity
Noncompliant - Medium	Count of noncompliant patches with "medium" severity
Noncompliant - Low	Count of noncompliant patches with "low" severity

Field Name	Definition
Noncompliant - Informational	Count of noncompliant patches with "informational" severity
Noncompliant - Unspecified	Count of noncompliant patches with "unspecified" severity

Instances That Missed Patches report

AWS Managed Services (AMS) Instances That Missed Patches report provides details on instances that missed patches during the last maintenance window execution, including:

1. Insights on missing patches at the patch ID level.
2. Insights on all the instances which have at least one patch missing along with attributes such as patch severity, unpatched days, range, and release date of the patch.

Field Name	Definition
Report Datetime	The date and time the report was generated.
Account ID	AWS Account ID to which the instance ID belongs
Account Name	AWS account name
Production Account	Identifier of AMS prod, non-prod accounts, depending on whether the account name includes the value 'PROD','NONPROD'.
Account Status	AMS account status. For example: ACTIVE or INACTIVE
AMS account service tier	PREMIUM or PLUS
Instance ID	ID of EC2 instance

Field Name	Definition
Instance Platform Type	Operating System (OS) type. For example: Windows
Instance State	State of the EC2 instance lifecycle. For example: TERMINATED, RUNNING, STOPPING, STOPPED, SHUTTING-DOWN, PENDING For more information, see Instance lifecycle.
Patch ID	ID of released patch. For example: KB3172729
Patch Severity	Severity of patch per publisher. For example: CRITICAL, IMPORTANT, MODERATE, LOW, UNSPECIFIED
Patch Classification	Classification of patch per publisher. For example: CRITICALUPDATES, SECURITYUPDATES, UPDATEROLLUPS, UPDATES, FEATUREPACKS
Patch Release Datetime (UTC)	Release date of patch per publisher
Patch Install State	Install state of patch on instance per SSM. For example: INSTALLED, MISSING, NOT APPLICABLE
Days Unpatched	Number of days instance unpatched since last SSM scanning
Days Unpatched Range	Bucketing of days unpatched. For example: <30 DAYS, 30-60 DAYS, 60-90 DAYS, 90+ DAYS

Patching SSM Coverage report

The AMS Patching SSM Coverage report informs you whether or not the EC2 instances in the account have the SSM Agent installed.

Field Name	Definition
Customer Name	Customer name for situations where there are multiple sub-customers
Resource Region	AWS Region where the resource is located
Account name	The name of the account
AWS Account ID	The ID of the AWS account
Resource Id	ID of EC2 instance
Resource Name	Name of EC2 instance
Compliant flag	Indicates if the resource has the SSM Agent installed ("Compliant") or not ("NON_COMPLIANT")

AMS Backup reporting

Topics

- Backup Snapshot Success / Failure report
- Backup Summary report
- Backup Snapshot Aged report
- Backup Summary / Coverage report

Backup Snapshot Success / Failure report

AWS Managed Services (AMS) Backup Snapshot Success / Failure report provides:

1. Insights on number of distinct snapshots taken.
2. The backup success rate.

Field Name	Definition
Report Datetime	The date and time the report was generated
AWS Account ID	AWS Account ID to which the resource belongs
Account Name	AWS account name
Backup Type	The type of backup if there is a plan
Backup Plan Name	User defined backup plan name
Backup Vault Name	The name of the backup vault
Resource Type	The type of resource that is being backed up
# of Resources	The number of resources that were backed up
Resource Region	The AWS Region of the backed up resource
Backup State	The state of the backup
Recovery Point ID	The unique identifier of the recovery point

Backup Summary report

AWS Managed Services (AMS) Backup Summary report provides insights on important backup metrics.

Field Name	Definition
Customer Name	Customer name for situations where multiple sub-customers are
Backup Month	Month of the backup
Backup Year	Year of the backup
Resource Type	The type of resource that is being backed up

Field Name	Definition
# of Resources	The number of resources that were backed up
Distinct Snapshots	Number of distinct snapshots
Backup Success Rate	The rate of successful backups
Max Snapshot Age	The maximum snapshot age
Backups Greater Than 30 Days Old	The count of backups that are over 30 days old

Backup Snapshot Aged report

AWS Managed Services (AMS) Backup Snapshot Aged report describes:

1. Aging of backup snapshots.
2. Classify backup snapshots into different aging buckets.
3. Understand which resources are out of backup compliance.

Field Name	Definition
Report Datetime	The date and time the report was generated
AWS Account ID	AWS Account ID to which the resource belongs
Account Name	AWS account name
Backup Type	The type of backup if there is a plan
Backup Plan Name	User defined backup plan name
Backup Vault Name	The name of the backup vault
Resource Type	The type of resource that is being backed up
# of Resources	The number of resources that were backed up

Field Name	Definition
Resource Region	The AWS Region of the backed up resource
Backup State	The state of the backup
Recovery Point ID	The unique identifier of the recovery point
Distinct Snapshots	The number of distinct snapshots
Snapshot Age (days)	The age in days of the snapshot
Backups Greater Than 30 Days Old	The number of backups that are over 30 days old
Backups 15-30 Days Old	The number of backups that are between 15 and 30 days
Backups Less Than 15 Days Old	The number of backups that are less than 15 days old

Backup Summary / Coverage report

AWS Managed Services (AMS) Backup Summary / Coverage report lists how many resources are not currently protected by any AWS Backup plan. CSDMs and customers can discuss an appropriate plan to increase coverage where possible, and reduce the risk of the data loss.

Field Name	Definition
Customer Name	Customer name for situations where multiple sub-customers are
Region	AWS region where the resource is located
Account name	The name of the account
AWS Account ID	The ID of the AWS account
Resource Type	Type of the resource. Respurces are the ones supported by AWS Backup (Aurora,

Field Name	Definition
	DocumentDB, DynamyDB, EBS, EC2, EFS, FSx, RDS, and S3)
Resource ARN	ARN of the resource
Resource ID	ID of the resource
Coverage	whether the resource is covered or not ("COVERED" or "NOT_COVERED")
# of resources	Number of supported resources in the account
perc_coverage	Percentage of supported resources with a backup executed in the last 30 days.

Billing Charges Details report

AWS Managed Services (AMS) Billing Charges Details report provides details about AMS billing charges with linked accounts and respective AWS services, including:

- Insights on AMS service-level charges, uplift percentages, account-level AMS service tiers and AMS fees.
- Insights on linked accounts and AWS usage charges

Field Name	Definition
Billing Month	The month and year of the service billed
Payer Account ID	The 12 digit ID identifying the account that will be responsible for paying the AMS charges
Linked Account ID	The 12 digit ID identifying the AMS account that consumes services that generates expenses
AWS Service Name	The AWS service that was used

Field Name	Definition
AWS Charges	The AWS charges for the AWS service name in AWS Service Name
Pricing Plan	The name of the pricing plan associated with the linked account
Uplift Proportion	The uplift percentage (as a decimal V.WXYZ) based on pricing_plan, SLA, and AWS service
Adjusted AWS Charges	AWS usage adjusted for AMS
Uplifted AWS Charges	The percentage of AWS charges to be charged for AMS; adjusted_aws_charges * uplift_percent
Instances EC2 RDS Spend	Spend on EC2 and RDS instances
AMS Charges	Total AMS charges for the product; uplifted_aws_charges + instance_ec2_rds_spend + uplifted_ris + uplifted_sp
Prorated Minimum Fee	The amount we charge to meet the contractual minimum
Minimum Fee	AMS Minimum Fees (if applicable)
Linked Account Total AMS Charges	Sum of all charges for the linked_account
Payer Account Total AMS Charges	Sum of all charges for payer account

Self-service reporting

AWS Managed Services (AMS) Self-Service Reporting (SSR) collects data from various native AWS services and provides access to reports on major AMS offerings. It also provides the information needed to support operations, configuration management, asset management, security management and compliance.

Use SSR to access the reports from the AMS console and report datasets through S3 buckets (one bucket per account) so you can plug it into your favorite Business Intelligence (BI) tool for customizing the reports based on your unique needs. AMS creates this S3 bucket (s3 bucket name: (ams-reporting-data-a<Account_ID>) in your primary AWS Region, and the data is shared from the AMS control plane hosted in us-east-1.

Important

You need one of the following roles to access this feature:

- Multi-Account Landing Zone: **AWSManagedServicesReadOnlyRole**
- Single-Account Landing Zone: **Customer_ReadOnly_Role**

Important

Using custom keys with AWS Glue

If want to encrypt your AWS Glue metadata with a customer-managed KMS key, you need to perform a few additional steps to allow AMS to aggregate data from the account:

1. Open the KMS console: https://console.aws.amazon.com/kms **Customer Managed Keys** page and click on the key ID that you plan to use to encrypt the AWS Glue metadata.
2. Click on the **Aliases** tab, and then **Create alias** on the right side of the screen.
3. In the text box, enter **AmsReportingFlywheelCustomKey** and click **Create alias**.

Topics

- Daily Patch reports
- Monthly billing report
- Daily backup report
- Weekly Incident report
- Data retention policy
- Offboarding from SSR

Daily Patch reports

These reports provide patching details.

Topics

- Instance details summary (Patch Orchestrator)
- Patch details
- Instances that missed patches

Instance details summary (Patch Orchestrator)

This is an informational report that helps identify all the instances onboarded to Patch Orchestrator (PO), account status, instance details, maintenance window coverage, maintenance window execution time, stack details, and platform type.

This dataset provides:

- Insights into Production and Non-Production instances of an account. Production and Non-Production stage is derived from the account name and not from the instance tags.
- Insights into distribution of instances by platform type. The 'N/A' platform type is when AWS Systems Manager (SSM) is not able to get the platform information.
- Insights into distribution of state of instances, number of instances running, stopped, or terminating.

Console Field Name	Dataset Field Name	Definition
Report Datetime	dataset_datetime	The date and time the report was generated.
Account Id	aws_account_id	AWS Account ID to which the instance ID belongs
Account Name	account_name	AWS account name
Production Account	prod_account	Identifier of AMS prod, non-prod accounts, depending on whether account name

Console Field Name	Dataset Field Name	Definition
		include value 'PROD', 'NONPROD'.
Account Status	account_status	AMS account status
	account_sla	AMS account service commitment
Landing Zone	malz_flag	Flag for MALZ-related account
Account Type	malz_role	MALZ role
Access Restrictions	access_restrictions	Regions to which access is restricted
Instance Id	instance_id	ID of EC2 instance
Instance Name	instance_name	Name of EC2 instance
Instance Platform Type	instance_platform_type	Operating System (OS) type
Instance Platform Name	instance_platform_name	Operating System (OS) name
Stack Name	instance_stack_name	Name of stack that contains instance
Stack Type	instance_stack_type	AMS stack (AMS infrastru cture within customer account) or Customer stack (AMS managed infrastru cture that supports customer applications)
Auto Scaling Group Name	instance_asg_name	Name of Auto Scaling Group (ASG) that contains the instance

Console Field Name	Dataset Field Name	Definition
Instance Patch Group	instance_patch_group	Patch group name used to group instances together and apply the same maintenance window
Instance Patch Group Type	instance_patch_group_type	Patch group type
Instance State	instance_state	State within the EC2 instance lifecycle
Maintenance Window Coverage	mw_covered_flag	If an instance has at least one enabled maintenance window with a future execution date, then it's considered covered, otherwise not covered
Maintenance Window Execution Datetime	earliest_window_execution_time	Next time the maintenance window is expected to execute

Patch details

This report provides patch details and maintenance window coverage of various instances.

This report provides:

- Insights on Patch groups and its types.
- Insights on Maintenance Windows, duration, cutoff, future dates of maintenance window executions (schedule) and Instances impacted in each window.
- Insights on all the operating systems under the account and number of instances that operating system is installed.

Field Name	Dataset Field Name	Definition
Report Datetime	dataset_datetime	The date and time the report was generated.
Account Id	aws_account_id	AWS Account ID to which the instance ID belongs
Account Name	account_name	AWS account name
Instance Id	instance_id	ID of EC2 instance
Instance Name	instance_name	Name of EC2 instance
Production Account	prod_account	Identifier of AMS prod, non-prod accounts, depending on whether account name include value 'PROD', 'NONPROD'.
Account Status	account_status	AMS account status
	account_sla	AMS account service tier
Instance Platform Type	instance_platform_type	Operating System (OS) type
Instance Platform Name	instance_platform_name	Operating System (OS) name
Stack Type	instance_stack_type	AMS stack (AMS infrastru cture within customer account) or Customer stack (AMS managed infrastru cture that supports customer applications)
Instance Patch Group Type	instance_patch_group_type	DEFAULT: default patch group w/ default maintenan ce window, determined by AMSDefaultPatchGroup:True tag on the instance

Field Name	Dataset Field Name	Definition
		CUSTOMER: customer created patch group NOT_ASSIGNED: no patch group assigned
Instance Patch Group	instance_patch_group	Patch group name used to group instances together and apply the same maintenance window
Instance State	instance_state	State within the EC2 instance lifecycle
Maintenance Window Id	window_id	Maintenance window ID
Maintenance Window State	window_state	Maintenance window state
Maintenance Window Type	window_type	Maintenance window type
Maintenance Window Next Execution Datetime	window_next _execution_time	Next time the maintenan ce window is expected to execute
Last Execution Maintenance Window	last_execution_window	The latest time the maintenance window was executed
	window_next_exec_yyyy	Year part of window_ne xt_execution_time
	window_next_exec_mm	Month part of window_ne xt_execution_time
	window_next_exec_D	Day part of window_ne xt_execution_time

Field Name	Dataset Field Name	Definition
	window_next _exec_HHMI	Hour:Minute part of window_next_execution_time
Maintenance Window Duration (hrs)	window_duration	The duration of the maintenance window in hours
Maintenance Window Coverage	mw_covered_flag	If an instance has at least one enabled maintenance window with a future execution date, then it's considered covered, otherwise not covered
Patch Baseline Id	patch_baseline_id	Patch baseline currently attached to instance
Patch Status	patch_status	Overall patch compliance status. If there is at least one missing patch, instance is considered noncompliant, otherwise compliant.
Compliant - Critical	compliant_critical	Count of compliant patches with "critical" severity
Compliant - High	compliant_high	Count of compliant patches with "high" severity
Compliant - Medium	compliant_medium	Count of compliant patches with "medium" severity
Compliant - Low	compliant_low	Count of compliant patches with "low" severity
Compliant - Informational	compliant_informational	Count of compliant patches with "informational" severity

Field Name	Dataset Field Name	Definition
Compliant - Unspecified	compliant_unspecified	Count of compliant patches with "unspecified" severity
Compliant - Total	compliant_total	Count of compliant patches (all severities)
Noncompliant - Critical	noncompliant_critical	Count of noncompliant patches with "critical" severity
Noncompliant - High	noncompliant_high	Count of noncompliant patches with "high" severity
Noncompliant - Medium	noncompliant_medium	Count of noncompliant patches with "medium" severity
Noncompliant - Low	noncompliant_low	Count of noncompliant patches with "low" severity
Noncompliant - Informational	noncompliant _informational	Count of noncompliant patches with "informational" severity
Noncompliant - Unspecified	noncompliant _unspecified	Count of noncompliant patches with "unspecified" severity
Noncompliant - Total	noncompliant_total	Count of noncompliant patches (all severities)

Instances that missed patches

This report provides details on instances that missed patches during the last maintenance window execution.

This report provides:

- Insights on missing patches at the patch id level.
- Insights on all the instances which have at-least one patch missing along with attributes such as patch severity, unpatched days, range, and release date of the patch.

Field Name	Dataset Field Name	Definition
Report Datetime	dataset_datetime	The date and time the report was generated.
Account Id	aws_account_id	AWS Account ID to which the instance ID belongs
Account Name	account_name	AWS account name
Customer Name Parent	customer_name_parent	
Customer Name	customer_name	
Production Account	prod_account	Identifier of AMS prod, non-prod accounts, depending on whether account name include value 'PROD', 'NONPROD'.
Account Status	account_status	AMS account status
Account Type	account_type	
	account_sla	AMS account service tier
Instance Id	instance_id	ID of EC2 instance
Instance Name	instance_name	Name of EC2 instance
Instance Platform Type	instance_platform_type	Operating System (OS) type
Instance State	instance_state	State within the EC2 instance lifecycle

Field Name	Dataset Field Name	Definition
Patch Id	patch_id	ID of released patch
Patch Severity	patch_sev	Severity of patch per publisher
Patch Classification	patch_class	Classification of patch per publisher
Patch Release Datetime (UTC)	release_dt_utc	Release date of patch per publisher
Patch Install State	install_state	Install state of patch on instance per SSM
Days Unpatched	days_unpatched	Number of days instance unpatched since last SSM scanning
Days Unpatched Range	days_unpatched_bucket	Bucketing of days unpatched

Monthly billing report

Billing charges details

This report provides details about AMS billing charges with linked accounts and respective AWS services.

This report provides:

- Insights on AMS service-level charges, uplift percentages, account-level AMS service tiers and AMS fees.
- Insights on linked accounts and AWS usage charges.

Important

The Monthly billing report is only available in your Management Payer Account (MPA) or your defined Charge Account. These are the accounts where your AMS monthly bill is sent. If you're unable to locate these accounts, then contact your Cloud Service Delivery Manager (CSDM) for assistance.

Field Name	Dataset Field Name	Definition
Billing Date	date	The month and year of the service billed
Payer Account Id	payer_account_id	The 12 digit id identifying the account that will be responsib le for paying the ams charges
Linked Account Id	linked_account_id	The 12 digit id identifying the AMS account that consumes services that generates expanses
AWS Service Name	product_name	The AWS service that was used
AWS Charges	aws_charges	The AWS charges for the AWS service name in AWS Service Name
Pricing Plan	pricing_plan	The pricing plan associated with the linked account
AMS Service Group	tier_uplifting_groups	AMS service group code that determines uplift percentage
Uplift Proportion	uplift_percent	The uplift percentage (as a decimal V.WXYZ) based on

Field Name	Dataset Field Name	Definition
		pricing_plan, SLA, and AWS service
Adjusted AWS Charges	adjusted_aws_usage	AWS usage adjusted for AMS
Uplifted AWS Charges	uplifted_aws_charges	The percentage of AWS charges to be charged for AMS; adjusted_aws_charges * uplift_percent
Instances EC2 RDS Spend	instances_ec2_rds_spend	Spend on EC2 and RDS instances
Reserved Instance Charges	ris_charges	Reserved instance charges
Uplifted Reserved Instance Charges	uplifted_ris	The percentage of reserved instance charges to be charged for AMS; ris_charges * uplift_percent
Savings Plan Charges	sp_charges	SavingsPlan usage charges
Uplifted Savings Plan Charges	uplifted_sp	The percentage of savings plans charges to be charged for AMS; sp_charges * uplift_percent
AMS Charges	ams_charges	Total ams charges for the product; uplifted_aws_charges + instance_ec2_rds_spend + uplifted_ris + uplifted_sp
Prorated Minimum Fee	prorated_minimum	The amount we charge to meet the contractual minimum

Field Name	Dataset Field Name	Definition
Linked Account Total AMS Charges	linked_account_total_ams_charges	Sum of all charges for the linked_account
Payer Account Total AMS Charges	payer_account_total_ams_charges	Sum of all charges for payer account
Minimum Fee	minimum_fees	AMS Minimum Fees (if applicable)
Reserved Instance and Savings Plan discount	adj_ri_sp_charges	RI/SP discount to be applied against RI/SP charges (applicable under certain circumstances)

Daily backup report

The backup report covers primary and secondary (when applicable) regions. It covers the status of backups (success/failure), and insights into snapshots taken.

This report provides:

- Backup status
- Number of snapshots taken
- Recovery point
- Backup plan and vault information

Field Name	Dataset Field Name	Definition
Report Datetime	dataset_datetime	The date and time the report was generated.
Account Id	aws_account_id	AWS Account ID to which the instance ID belongs

Field Name	Dataset Field Name	Definition
Account Name	account_name	AWS account name
Account SLA	account_sla	AMS account service commitment
	malz_flag	Flag for MALZ-related account
	malz_role	MALZ role
	access_restrictions	Regions to which access is restricted
Resource ARN	resource_arn	The Amazon resource name
Resource Id	resource_id	The unique resource identifier
Resource Region	resource_region	The resource's primary (and secondary, when applicable) regions.
Resource Type	resource_type	The type of resource
Recovery Point ARN	recovery_point_arn	The ARN of the recovery point
Recovery Point Id	recovery_point_id	The unique identifier of the recovery point
Backup snapshot scheduled start datetime	start_by_dt_utc	Timestamp when snapshot is scheduled to begin
Backup snapshot actual start datetime	creation_dt_utc	Timestamp when snapshot actually begins
Backup snapshot completion datetime	completion_dt_utc	Timestamp when snapshot is completed

Field Name	Dataset Field Name	Definition
Backup snapshot expiration datetime	expiration_dt_utc	Timestamp when snapshot expires
Backup Job status	backup_job_status	State of the snapshot
Backup Type	backup_type	Type of backup
Backup Job Id	backup_job_id	The unique identifier of the backup job
Backup Size In Bytes	backup_size_in_bytes	The backup size in bytes
Backup Plan ARN	backup_plan_arn	The backup plan ARN
Backup Plan Id	backup_plan_id	Backup plan unique identifier
Backup Plan Name	backup_plan_name	The Backup Plan name
Backup Plan Version	backup_plan_version	The backup plan version
Backup Rule Id	backup_rule_id	The backup rule id
Backup Vault ARN	backup_vault_arn	Backup vault ARN
Backup Vault Name	backup_vault_name	The backup vault name
IAM Role ARN	iam_role_arn	The IAM role ARN
Recovery Point Status	recovery_point_status	Recovery point status
Recovery Point Delete After Days	recovery_point_delete_after_days	Recovery point delete after days
Recovery point move to cold storage after days	recovery_point_move_to_cold_storage_after_days	Number of days after completion date when backup snapshot is moved to cold storage

Field Name	Dataset Field Name	Definition
Recovery Point Encryption Status	recovery_point_is_encrypted	Recovery point encryption status
Recovery Point Encryption Key ARN	recovery_point_encryption_key_arn	Recovery point encryption key ARN
Volume State	volume_state	Volume State
Instance Id	instance_id	Unique instance Id
Instance State	instance_state	Instance state
Stack Id	stack_id	Cloudformation stack unique identifier
Stack Name	stack_name	Stack Name
Tag: AMS Default Patch Group	tag_ams_default_patch_group	Tag Value: AMS Default Patch Group
Tag: App Id	tag_app_id	Tag Value: App ID
Tag: App Name	tag_app_name	Tag Value: App Name
Tag: Backup	tag_backup	Tag Value: Backup
Tag: Compliance Framework	tag_compliance_framework	Tag Value: Compliance Framework
Tag: Cost Center	tag_cost_center	Tag Value: Cost Center
Tag: Customer	tag_customer	Tag Value: Customer
Tag: Data Classification	tag_data_classification	Tag Value: Data Classification
Tag: Environment Type	tag_environment_type	Tag Value: Environment Type
Tag: Hours of Operation	tag_hours_of_operation	Tag Value: Hours of Operation

Field Name	Dataset Field Name	Definition
Tag: Owner Team	tag_owner_team	Tag Value: Owner Team
Tag: Owner Team Email	tag_owner_team_email	Tag Value: Owner Team Email
Tag: Patch Group	tag_patch_group	Tag Value: Patch Group
Tag: Support Priority	tag_support_priority	Tag Value: Support Priority

Weekly Incident report

This report provides the aggregated list of incidents along with its priority, severity and latest status, including:

- Insights on support cases categorized as incidents on the managed account
- Incident information required to visualize the incident metrics for the managed account
- Insights on incident categories and remediation status of every incident

Both visualization and data are available for the Weekly incident report.

- Visualization can be accessed through the AMS console in the account through the **Reports** page.
- Dataset with the following schema, can be accessed through S3 bucket in the managed account.
- Use the provided date fields to filter incidents based on the month, quarter, week, and/or day that the incident was created or resolved.

Field Name	Dataset Field Name	Definition
Report Datetime	dataset_datetime	The date and time the report was generated.
Account Id	aws_account_id	AWS Account ID to which the incident belongs.
Account Name	account_name	AWS account name.

Field Name	Dataset Field Name	Definition
Case Id	case_id	The ID of the incident.
Created Month	created_month	The month when the incident was created.
Priority	priority	The priority of the incident.
Severity	severity	The severity of the incident.
Status	status	The status of the incident.
Category	yuma_category	The category of the incident.
Created Day	created_day	The day when the incident was created in YYYY-MM-DD format.
Created Week	created_wk	The week when the incident was created in YYYY-WW format. Sunday to Saturday is counted as the beginning and end of a week. Week is from 01 to 52. Week 01 is always the week that contains the first day of the year. For example, 2023-12-31 and 2024-01-01 are in week 2024-01.
Created Quarter	created_qtr	The quarter when the incident was created in YYYY-Q format. 01/01 to 03/31 is defined as Q1, and so on.
Resolved Day	resolved_day	The day when the incident was resolved in YYYY-MM-DD format.

Field Name	Dataset Field Name	Definition
Resolved Week	resolved_wk	The week when the incident was resolved in YYYY-WW format. Sunday to Saturday is counted as the beginning and end of a week. Week is from 01 to 52. Week 01 is always the week that contains the first day of the year. For exmaple, 2023-12-31 and 2024-01-01 are in week 2024-01.
Resolved Month	resolved_month	The month when the incident was resolved in YYYY-MM format.
Resolved Quarter	resolved_qtr	The quarter when the incident was resolved in YYYY-Q format. 01/01 to 03/31 is defined as Q1, and so on.

Data retention policy

AMS SSR has a data retention policy per report after the period reported, the data is cleared out and no longer available.

Report name	Data Retention SSR Console	Data Retention SSR S3 Bucket
Instance Details Summary (Patch Orchestrator)	2 Months	2 Years
Patch Details	2 Months	2 Years

Report name	Data Retention SSR Console	Data Retention SSR S3 Bucket
Instances that missed patches during maintenance window execution	2 Months	2 Years
AMS Billing Charges Details	2 Years	2 Years
Daily Backup Report	1 Month	2 Years
Weekly Incident Report	2 Months	2 Years

Offboarding from SSR

To offboard from the SSR service please create a service request (SR) through the AMS console an AMS operations engineers will help you offboard from SSR. In the ticket please provide the reason for offboarding.

If you are offboarding and account and want to do a cleanup please create an SR through the AMS console and AMS operations engineers will help delete SSR S3 bucket.

If you are leaving AMS you will automatically be offboarded from the AMS SSR console. AMS will automatically stop sending data to your account. AMS deletes your SSR S3 bucket as part of offboarding process.

Incident reports, service requests, and billing questions in AMS

Topics

- Incident management
- Service request management
- Billing questions

With AWS Managed Services (AMS), you can request help with operational issues and requests at any time through the AMS console. AMS operations engineers are available to respond to your incidents and service requests 24x7, with response time Service Level Agreements (SLAs) and Service Level Objectives (SLOs), dependent on your selected account Service Tier (Plus, Premium). AMS operations engineers proactively notify you of important alerts and questions using the same mechanisms.

Incident management

Topics

- What is incident management?
- Incident management service commitments
- Incident management examples

Incidents are AWS service performance issues that impact your managed environment, as determined by AWS Managed Services (AMS) or you. Incidents identified by the AMS team are first received as "events": a change in system state captured by monitoring. If a configured threshold is breached, the event triggers an alarm, also called an alert. The AMS operations team determines if the event is non-impacting, an incident (a service interruption or degradation), or a problem (the underlying root cause of one or more resolved incidents).

The AMS team also receives incidents identified by you through the AWS Support center or programmatically using the AWS Support API with the service code `sentinel-report-incident`.

After your incident is received by the AMS operations team, it's reviewed to ensure that the incident is not better classified as a service request. If it should be classified as a service request, it's immediately reclassified and the AMS service request team takes over and you are notified. If the incident can be resolved by the receiving operator, steps are taken to immediately to resolve the incident. AMS operators consult internal documentation for a resolution and, if needed, escalate the incident to other support resources until the incident is resolved. To be kept informed at each step of the incident resolution process, be sure to fill in the **CC Emails** option, and, if you'll connect by federation, log in before following the link in the email that AMS sends. After it is resolved, the AMS operations team documents the incident and resolution for future use.

If an incident resolution requires infrastructure changes, a security review might be needed. Infrastructure changes that might require a security review include those related to IAM, or resource-based policy, or risk approvals. Those types of incidents require an AMS Operations engineer to create an RFC before making the change, and your approval to that RFC is required. For example, should the incident resolution require the update of an IAM policy, there would be an AMS security review and then an AMS Operations engineer would create an RFC with the Management | Advanced stack components | Identity and Access Management (IAM) | Update entity or policy change type (ct-27tuth19k52b4) and wait for you to approve the RFC before proceeding.

> **Note**
>
> AMS now allows incident resolution that requires infrastructure changes to be made without the additional step of RFC approval. If the changes needed to resolve the incident do NOT require a security review (the change is not related to IAM, or resource-based policy, or risk approvals), AMS can make the changes based on your approval received in the incident, without needing separate approval in an RFC.

For definitions of incident management terms, see AMS Key Terms.

To understand the escalation path of incidents, see Getting help.

For a description of AMS response to incidents, see AMS incident response.

What is incident management?

Incident management is the process AMS uses to record, act on, communicate progress of, and provide notification of, active incidents.

The goal of the incident management process is to ensure that normal operation of your managed service is restored as quickly as possible, the business impact is minimized, and all concerned parties are kept informed.

Examples of incidents include (but are not restricted to) loss of or degradation of network connectivity, a non-responsive process or API, or a scheduled task not being performed (for example, a failed backup).

The following graphic depicts the workflow of an incident reported by you to AMS.

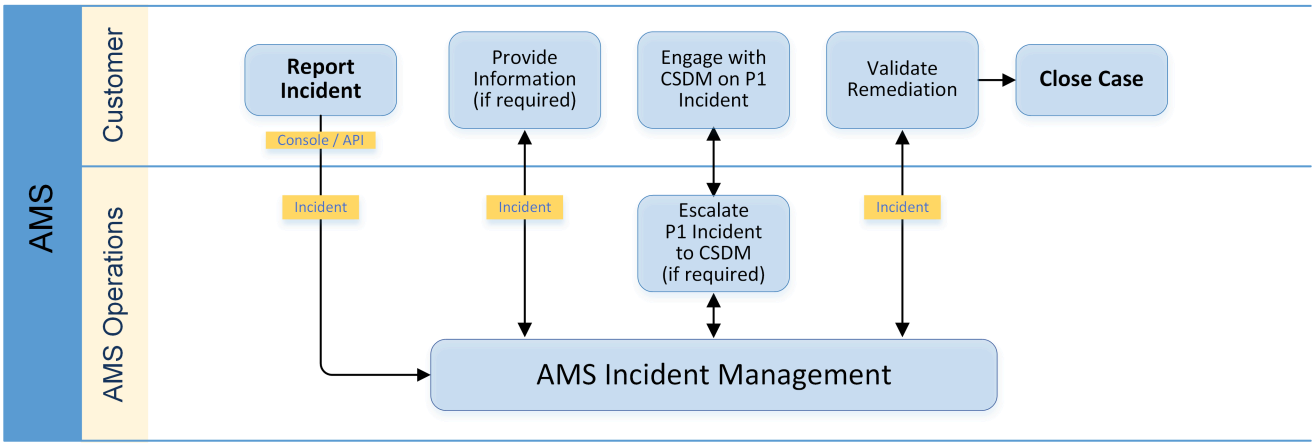

This graphic depicts the workflow of an incident reported by AMS to you.

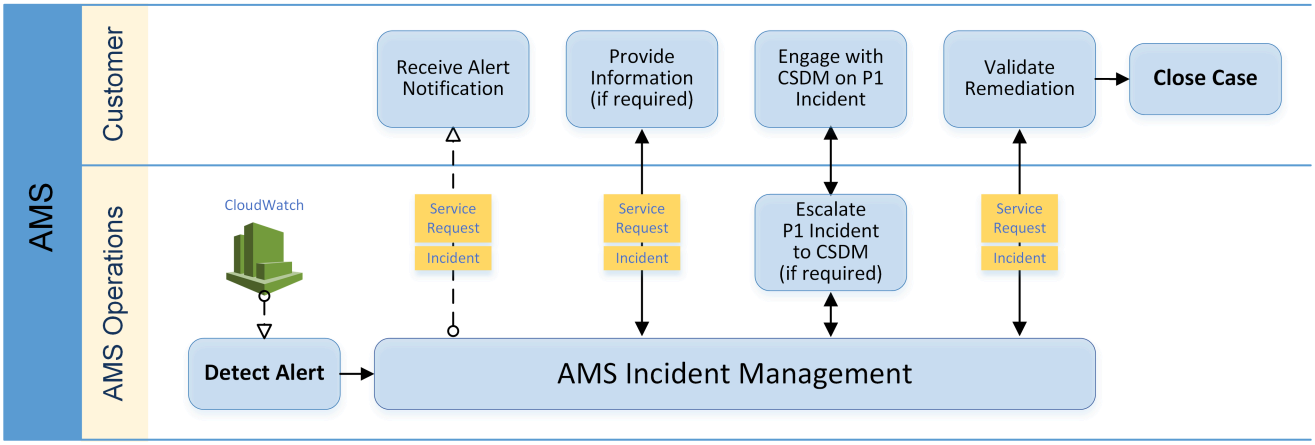

Incident priority

Incidents created in AWS Support center, console or Support API (SAPI), have different classifications than incidents created in the AMS console.

- Low: Non-critical functions of your business service, or application, related to AWS or AMS resources are impacted.

- Medium: A business service or application related to AWS and/or AMS resources is moderately impacted and is functioning in a degraded state.
- High: Your business is significantly impacted. Critical functions of your application related to AWS and/or AMS resources are unavailable. Reserved for the most critical outages affecting production systems.

> **Note**
>
> The AWS Support Console offers five levels of incident priority that we translate to the three AMS levels.

Problem vs incident

When AMS believes that an incident reveals a larger defect or misconfiguration and could recur, it is considered a problem rather than just an incident. In such cases, AMS undertakes analyses of the problem and offers suggestions to resolve the problem.

Incident management service commitments

Incident management service commitments

Event or action	Service commitment measurement
Case 1: An event with known impact is generated. AMS opens an incident and informs you. Case 2: AMS contacts you to confirm the impact of the event. You confirm the event is an incident. Case 3: You notice an issue and submit an incident report.	Clock for incident response and incident resolution starts when: Case 1: AMS creates an incident. Case 2: You confirm the alert is an incident. Case 3: You submit an incident. Service commitments depend on the priority of the incident created.
If you submit the incident, AMS sends a response to acknowledge it.	Clock for incident resolution continues ticking. Clock for incident response time stops when AMS sends the incident acknowledgement.

Event or action	Service commitment measurement
If AMS creates the incident on your behalf, a separate incident response is not sent.	**Note** Time spent waiting for inputs from you is excluded from incident resolution time calculations. For incidents that AMS creates, the initial response time is the time of the creation of the initial incident notification to you.
For the resources / services in question, AMS checks the health to verify if: • AMS detected event or customer submitted incident qualifies as an incident, and • the incident is correctly prioritized, and If an incident you submit is not correctly prioritized, AMS re-prioritizes it. If AMS changes an incident priority, a notification is sent to you along with reasoning behind the priority change. In certain cases, an issue you submitted may not qualify as an incident, depending on the cause. In those cases, AMS closes the incident and sends you a notification explaining the reason why. Irrespective of the event categorization, AMS works with you to assist as needed. To understand the rules for incident categorization, see Incident priority.	In case incident priority changes, the service commitment for the new priority is applicable; clock continues ticking. In cases when an incident is closed because it does not meet the definition of an incident, service commitments are not applicable; clock stops.

Event or action	Service commitment measurement
AMS works on the incident to resolve it within service commitment. In certain cases, if AMS determines that unavailable stack(s) or resource(s) cannot be resolved in a timely manner, AMS will offer Infrastructure Restore as an option for resolution. Infrastructure Restore involves re-deploying existing stack(s), based on the templates of the impacted stack(s), and initiating a data restore based on the last known restore point (EBS/RDS snapshot), unless otherwise specified by you. Ephemeral data on individual EC2 instances will be lost. If you do not authorize an Infrastructure Restore as recommended by AWS, you will not be eligible for a service credit for the associated Incident Resolution Time Service Commitment.	Clock stops when: • AMS has restored all Unavailable services or resources pertaining to that Incident to an available state, or • an infrastructure restore is started.
Occasionally, AMS needs clarification from, or activity by, you to keep incident resolution efforts moving forward, unless you have a pre-defined, approved action. As a result, there is communication between AMS and you in order to resolve incidents	Clock stops when: AMS is waiting for a response or action from you. Clock restarts when: AMS receives the response from you or the action AMS requires of you is completed.

Note

For a complete list of service commitments, download the AMS Service Level Agreement.

Incident management examples

Incident management examples.

Topics

- Incident testing
- Reporting incidents
- Monitoring and updating incidents
- Managing incidents with the AWS Support API
- Responding to AMS-generated incidents

The following examples describe using the AMS console to submit an incident. Once submitted, the AMS team works with you to resolve the incident per your Service Level Agreement (SLA).

Incident testing

When testing AMS incident submissions, we ask that you include in the subject text this flag: **AMSTestNoOpsActionRequired**. This flag lets AMS know that the incident submission is only for testing. When AMS operations engineers see that flag, they will not respond in any way to the incident submission.

Reporting incidents

Use the AMS console to report an incident. It's important to create a new incident for each new issue or question. When opening cases related to old inquiries, it's helpful to include the related case number so we can refer to previous correspondence.

> **Note**
>
> If case correspondence strays from the original issue, an AMS operator might ask you to report a new incident.

To report an incident using the AMS console:

1. From the left navigation, choose **Incidents**

   The **Incidents** list opens:

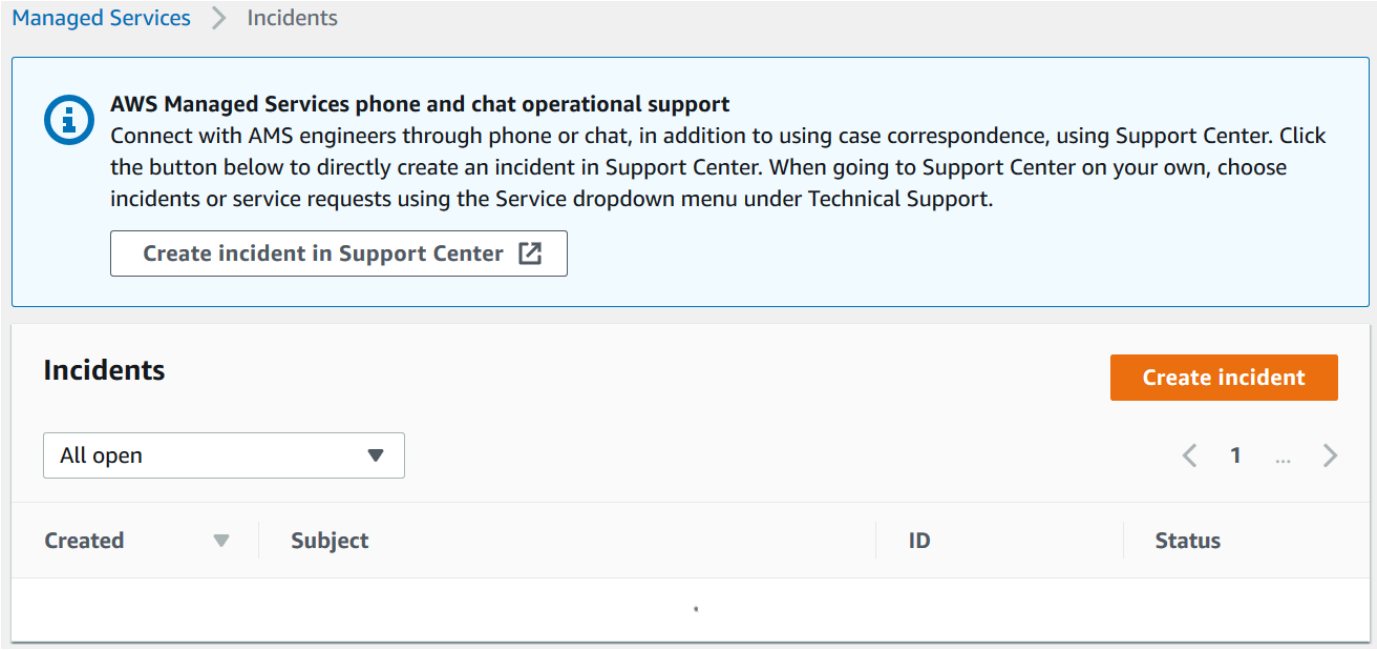

If your incident list is empty, the **Clear filter** option resets the filter to **Any status**.

If you know you want to use phone or chat, click **Create incident in Support Center** to open the incident **Create** page in the AWS Support Center, auto-populated with the AMS service type.

> **Note**
>
> Phone calls initiated with AWS Support center are recorded, to better improve response. If the call drops, you must call back through the Support Center case, AWS has no mechanism for calling you back.

> **Important**
>
> Phone and chat support is designed to help with support cases, incidents and service requests. For RFC issues, use the correspondence option on the relevant RFC details page, to reach an AMS engineer.

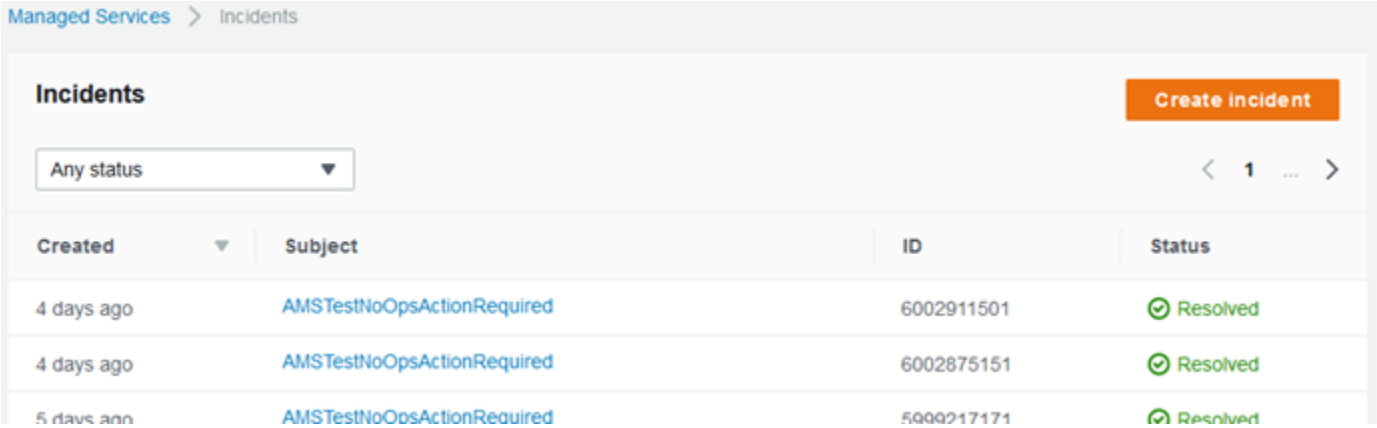

2. If you want to find an existing incident, select an incident status filter in the drop-down list.

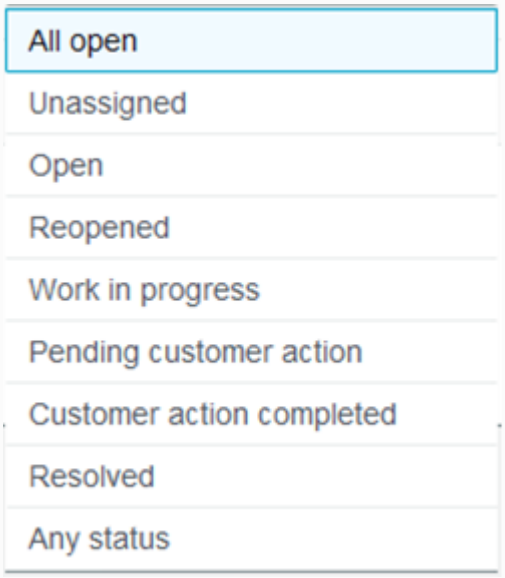

- All incidents that are not yet resolved.
- A new incident that is not yet assigned.
- An incident that has been assigned.
- An incident that you reopened.
- An assigned, complicated incident.
- Incidents that require your feedback before the next step.
- Incidents to which you have recently submitted information.
- An incident that has concluded.
- All incidents in the account.

3. Choose **Create**.

   The **Create an incident** page opens:

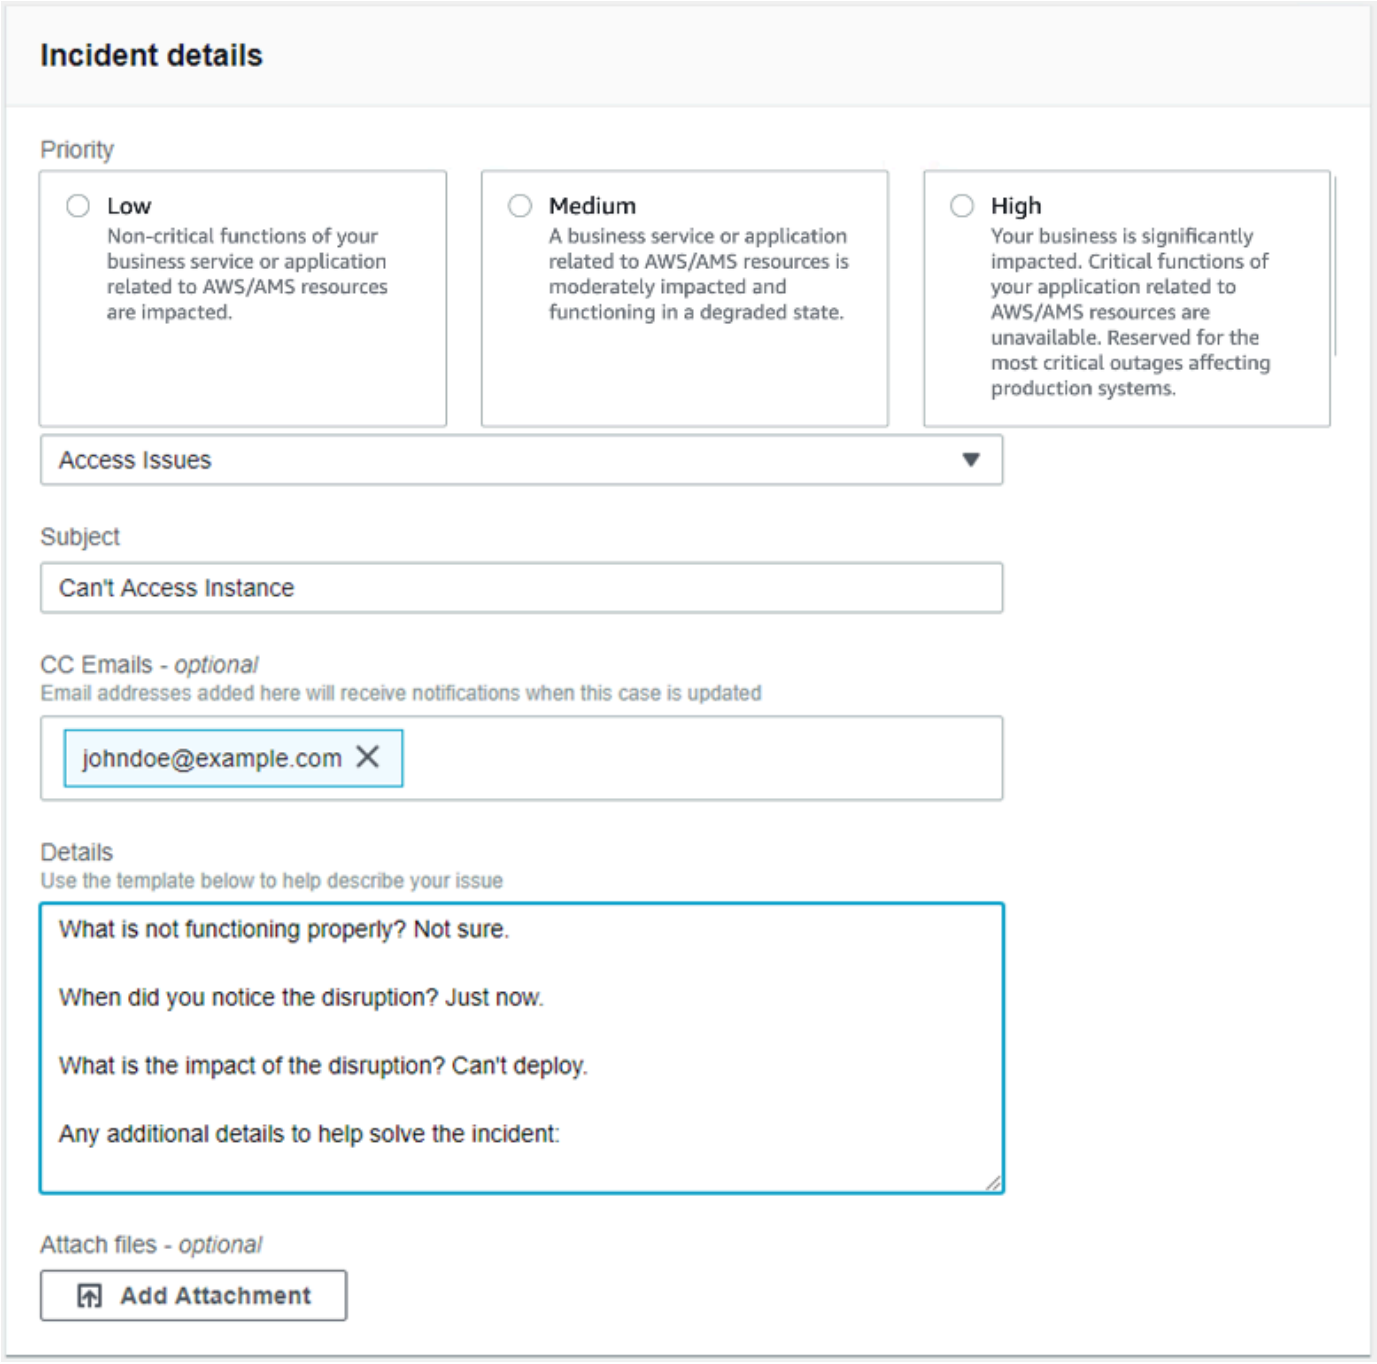

4. Select a **Priority**:

   - **Low**: Non-critical functions of your business service or application related to AWS/AMS resources are impacted.
   - **Medium**: A business service or application related to AWS/AMS resources is moderately impacted and functioning in a degraded state.

- **High**: Your business is significantly impacted. Critical functions of your application related to AWS/AMS resources are unavailable. Reserved for the most critical outages affecting production systems.

5. Select a **Category**:

   - **Access Issues**: You have a question about accessing your AMS-managed resources.
   - **Availability**: A resource appears to be unavailable.
   - **Performance Issue**: A resource seems to be under-performing.
   - **Security Related**: You have a security concern about your AMS-managed resources.
   - **Other**: None of the other categories apply.

   > **Note**
   >
   > If you are going to test incident functionality, AMS asks that you add the no-action flag (AMSTestNoOpsActionRequired) to your incident title.

6. Enter information for:

   - **Subject**: A descriptive title for the incident report.
   - **CC emails**: A list of email addresses for people you want informed about the incident report and resolution.
   - **Details**: A comprehensive description of the incident, the systems impacted, and the expected outcome of the resolution. Answer the pre-set questions, or delete them and enter any relevant information.

   To add an attachment, choose **Add Attachment**, browse to the attachment you want, and click **Open**. To delete the attachment, click the Delete icon:

7. Choose **Submit**.

   A details page opens with information on the incident—such as **Type**, **Subject**, **Created**, **ID**, and **Status**—and a **Correspondence** area that includes the description of the request you created.

   Click **Reply** to open a correspondence area and provide additional details or updates in status.

Click **Close Case** when the incident has been resolved.

Click **Load More** if there is more correspondence than will fit on one page.

Don't forget to rate the communication!

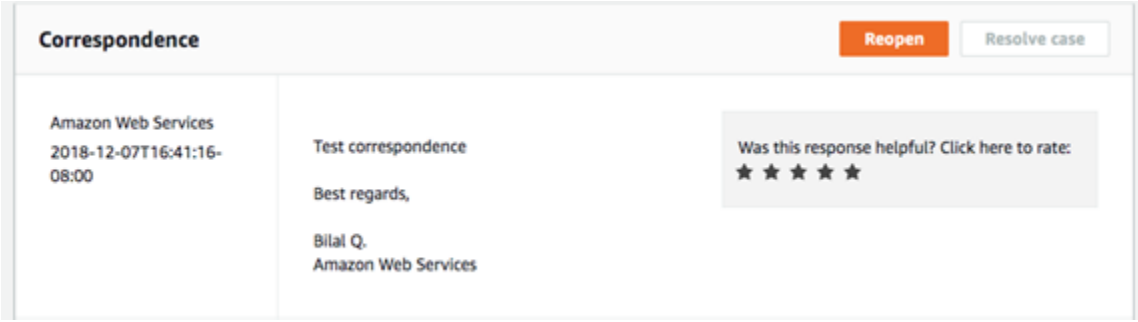

Your incident displays on the **Incidents** list page.

YouTube Video: How do I raise an incident from the AWS Managed Services console?

Monitoring and updating incidents

You can update, monitor, and review incident reports and service requests, both called cases, by using the AMS console, or programmatically using the AWS Support API. For information on using the AWS Support API, see `DescribeCases` operation.

To monitor a case, incident or service request, using the AMS console, follow these steps.

1. In the AMS console **Incident reports** or **Service requests** dashboard, browse to a case and choose the **Subject** to open a details page with current status and correspondences.

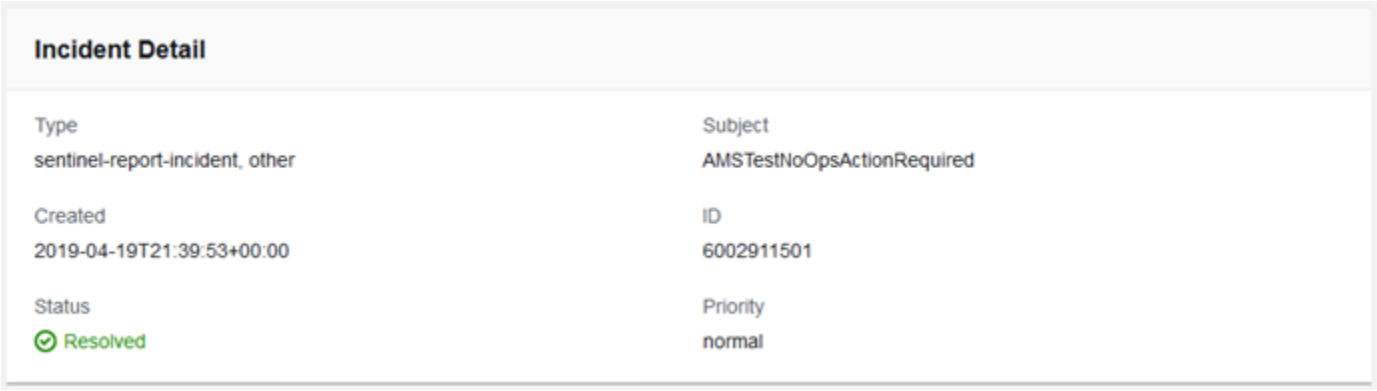

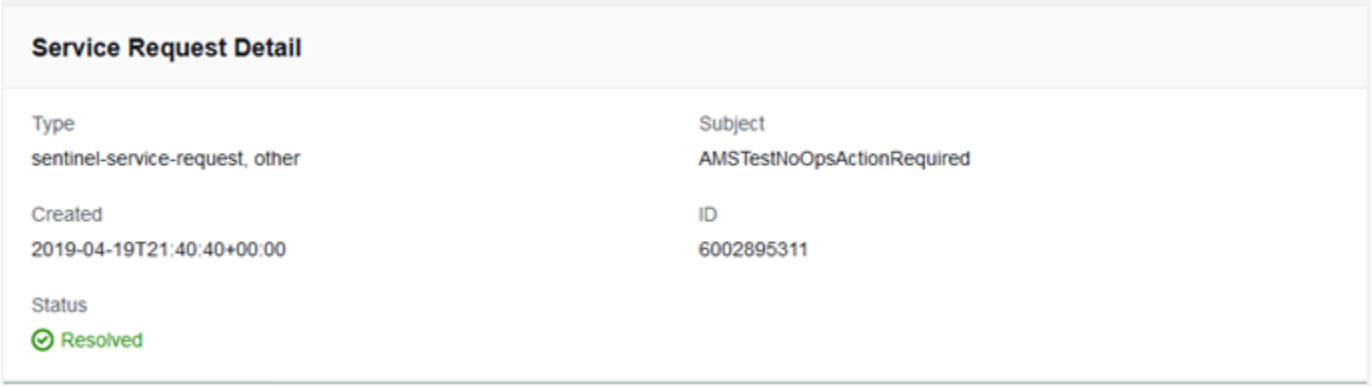

When a reported incident or service request case is updated by the AMS operations team, you receive an email and a link to the incident in the AMS console so you can respond. You can't respond to incident correspondence by replying to the email.

> **Important**
>
> You must have entered an email address to receive notifications of state change for a service request or incident case. Notifications only go to the email address added to the case when it's created.
> The link in the notification email will not work unless you are using an email server on your AMS federated network. However, you can respond to the correspondence by going to your AMS console and using the case details page.

2. If there are many cases in the list, you can use the **Filter** option:

   - **All open** (default): Use this filter to see all cases that have not been resolved.
   - **Unassigned**: Use if you've just submitted the case and have not received any notice that the case state has changed. Note, incidents and service request cases are addressed with different promptness depending on the submitted priority (incidents) or your service level agreement (service requests).
   - **Open**: Use if you have received notice that the case is "Pending Amazon" action; this means that the case has been assigned but work has not yet begun.
   - **Reopened**: Use if you have received notice that the case was reopened after having been resolved.
   - **Work in progress**: Use if you have received notice that an operator has begun to work on the case.

- **Pending customer action**: Use if you have received an operator request for action on your part.
- **Customer action completed**: Use if you have received notice that your action on the case has been processed.
- **Resolved**: Use to view cases that you know have been resolved. Resolved cases are maintained in history for twelve months.
- **Any status**: Use this filter to see all cases, regardless of status.

3. To check the latest status, refresh the page.
4. If there are so many correspondences that they do not all appear on the page, choose **Load More**.
5. To provide an update to the case status, choose **Reply**, enter the new correspondence, and then choose **Submit**.
6. To close out the case after it has been resolved to your satisfaction, choose **Close case**.

   Be sure to rate the service through the 1-5 star rating to let AMS know how we're doing!

Managing incidents with the AWS Support API

The AWS Support API enables you to create incidents and add correspondence to them throughout investigations of your issues and interactions with AWS Support staff. The AWS Support API models much of the behavior of the AWS Support Center. For more details about how you can use this AWS support service, see Programming an AWS Support Case.

> **Note**
>
> When using the AWS Support API, or SAPI, for AMS Advanced incidents, use this service code: `sentinel-report-incident`.

Responding to AMS-generated incidents

AMS proactively monitors your resources; for more information, see Monitoring and event management. Sometimes AMS identifies and creates an incident case, most often to notify you of an incident. In the event that action is required on your part to resolve an incident, AMS sends a notification to the contact information you have provided for the account. You respond to this

incident in the same way as you would any other incident. You would usually respond to incidents via the AMS console; in some cases, contact by email or phone is required.

> **Note**
>
> AMS sends communications to your primary email address on your AWS account; we recommend adding an alternate Operations contact email alias to facilitate the incident management process. This is covered during the AMS onboarding process and related onboarding documentation. If you have provided AMS with non-resource based contacts (that you informed your CSDM of) during onboarding, those contact are used. For example, you could provide a list of contacts named "SecurityContacts" to your CSDMs/CAs to use for security-related incidents or notifications. Contact tags on your instances/resources are used for AMS-generated incidents, if you have provided your consent to CSDM for using tag information.
> To learn more about this notification service, see Notifications.

Service request management

Topics

- How service request management works
- Service request management examples for AMS

Service requests are communications to AMS created by you to ask for information or advice. A good example of a standard service request is for guidance or help in configuring an AMS service, like Alarm Manager, Patch Orchestrator, and so forth. You can also receive service requests from AMS; these are called outbound service requests or service notifications. To see a list of your service requests, and outbound service requests,(service notifications), sent to you by AMS, look on the **Service requests** page of the AMS console.

To learn more about outbound service requests, see Responding to an AMS-generated service requests.

You create an AWS Managed Services (AMS) service request by using the AMS console or, programmatically, by using the AWS Support API. For details on using the API, see AWS Support API. For AMS choose the `sentinel-service-request` service code.

After your service request is received by the AMS operations team, it is prioritized according to your service level agreement. To be kept informed at each step of the service request resolution process, be sure to fill in the **CC Emails** option, and, if you will connect by federation, log in before following the link in the email AMS sends.

How service request management works

Service requests are handled by the on-call AMS operations team.

After your service request is received by the AMS operations team, it's reviewed to ensure that the request is not more properly classified as an incident. If it should be classified as an incident, it's immediately reclassified, the AMS incident management team takes over, and you're notified.

If the service request can be resolved with the submission of an RFC, the reviewing operator sends you an email requesting that you submit the appropriate RFC (details are provided).

If the AMS operator can resolve the service request, steps to do so are taken immediately. For example, if the service request is for architecture advice or other information, then the operator refers you to the appropriate resources or answers the question directly.

If the analysis of your service request identifies a bug or a feature request, then AMS sends you a notification through the service request. Since there is no ETA for feature requests or bug fixes, the original service request is closed. Contact your CSDM for follow up questions related to the original service request.

If the service request is out of scope for AMS operations, the operator either sends the request to your cloud service delivery manager so they can communicate with you, or to the appropriate AWS operations team, along with an email to you, as to what steps are being taken.

The service request is not resolved until you have indicated that you're satisfied with the outcome.

> **Note**
>
> We recommend providing a contact email, name, and phone number in all cases to facilitate communications.

Service request management examples for AMS

Using the AMS console **Create Service Request** page, you can perform the following tasks:

- Create and update a service request
- Get a list of, and detailed information about, all of your current service requests
- Narrow your search for service requests by dates and incident identifiers, including requests that have been resolved
- Add communications and file attachments to your requests, and add email recipients for case correspondence
- Resolve service requests
- Rate service request communications

The following examples describe using the AMS console to create a service request.

Testing a service request in AMS

When testing AMS service requests, we ask that you include in the subject text this flag: **AMSTestNoOpsActionRequired** to let AMS know that the service request is only for testing. When AMS operations engineers see that flag, they do not respond to the service request.

Creating a service request in AMS

To create a service request using the AWS Managed Services (AMS) console:

1. From the left navigation, choose **Service requests**.

   The **Service requests** list opens.

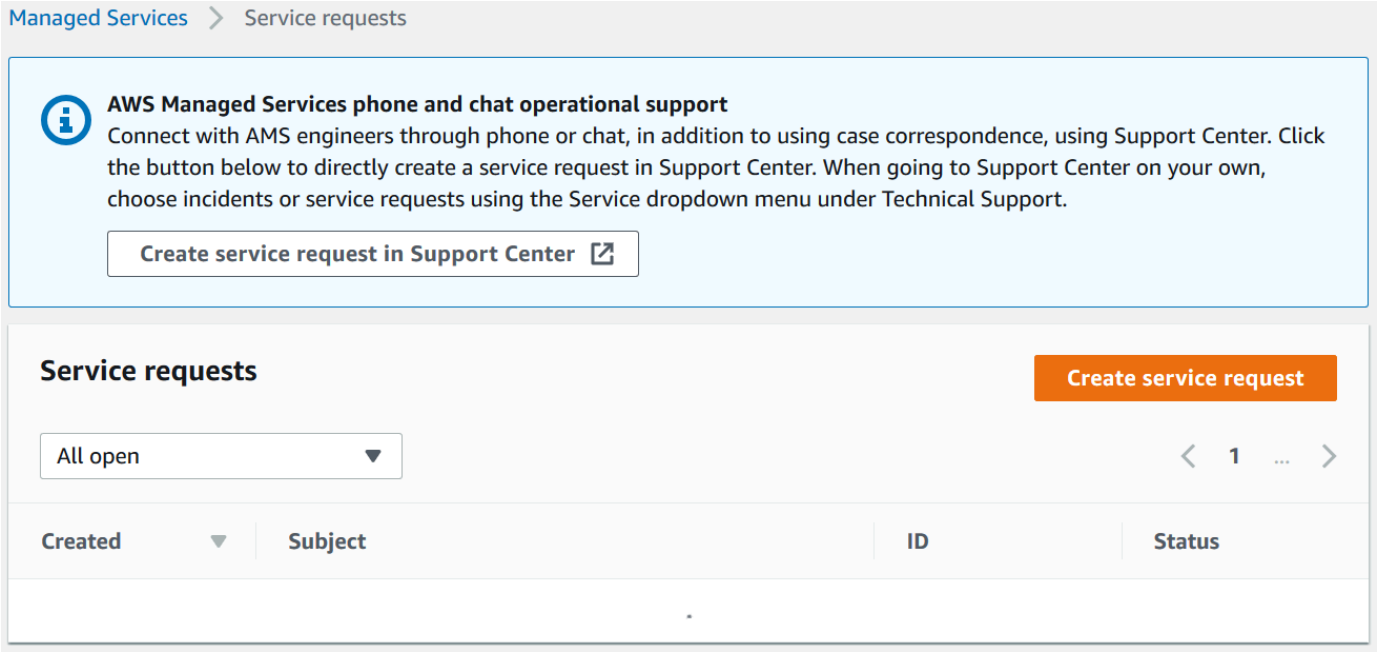

If your service request list is empty, the **Clear filter** option resets the filter to **Any status**.

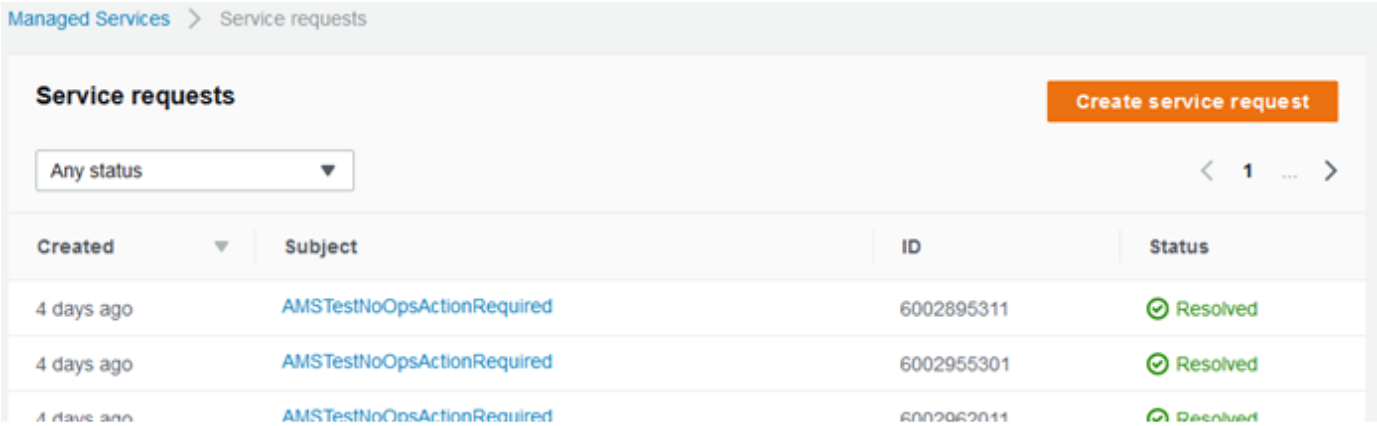

If you know you want to use phone or chat, click **Create service request in Support Center** to open the service request **Create** page in the AWS Support Center, auto-populated with the AMS service type.

> **Note**
>
> Phone calls initiated with AWS Support center are recorded, to better improve response. If the call drops, you must call back through the Support Center case, AWS has no mechanism for calling you back.

> **Important**
>
> Phone and chat support is designed to help with support cases, incidents and service requests. For RFC issues, use the correspondence option on the relevant RFC details page, to reach an AMS engineer.

2. If you want to find an existing service request, select a service request status filter in the drop-down list.

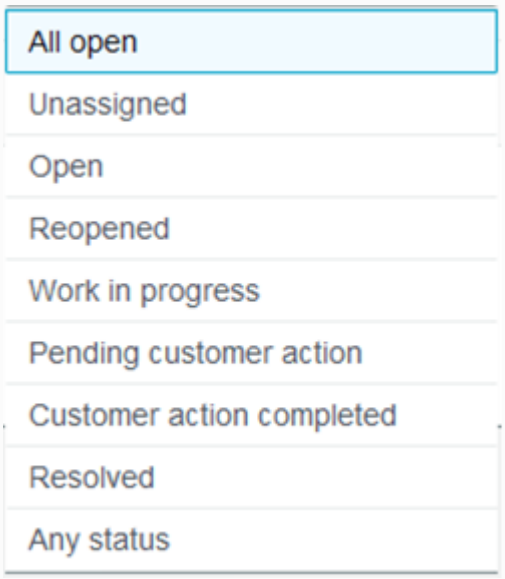

- All service requests that are not yet resolved.
- A new service request that is not yet assigned.
- A service request that has been assigned.
- A service request that you reopened.
- An assigned, complicated, service request.
- Service requests that require your feedback before the next step.
- Service requests to which you have recently submitted information.
- A service request that has concluded.
- All service requests in the account.

3. Choose **Create**.

   The **Create a service request** page opens.

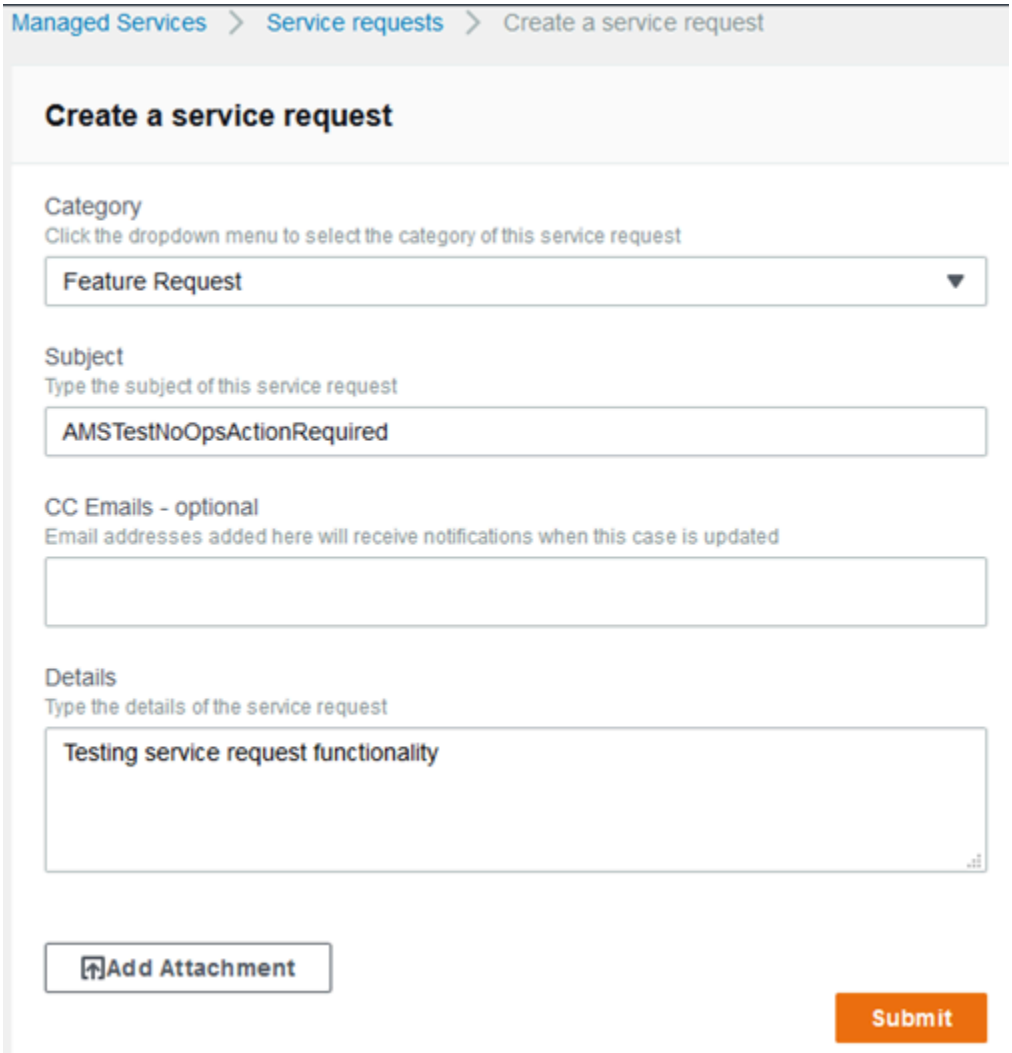

4. Select a **Category**:

   - **Access**: Use this when you have a question about accessing your AMS-managed resources. To request access to an AMS-managed resource, submit an RFC with the Access Management category.
   - **Alert notification**: Use this when you have an alert and have not heard from AMS.
   - **Feature Request**: Use this to request that AMS add a feature.
   - **General Guidance**: Use this for non-resource specific questions.
   - **Operations On Demand**: Use this for requests about the OOD feature.
   - **Security Related**: Use this when you have a security concern about your AMS-managed resources. Note that while we use encryption, you should exercise caution with the information you submit here.
   - **Service Reporting Query**: Use this to request a specific report.

- **Other**: Use this when none of the other categories apply.

> **Note**
>
> If you are going to test service request functionality, AMS asks that you add the no-action flag, `AMSTestNoOpsActionRequired`. to your service request title.

5. Enter information for:

   - **Subject**: This creates a link to the service request details on the list page.
   - **CC emails**: These emails receive correspondence in addition to your default email contacts.
   - **Details**: Provide as much information here as possible.

   To add an attachment, choose **Add Attachment**, browse to the attachment you want, and click **Open**. To delete the attachment, click the Delete icon:

6. Choose **Submit**.

   A details page opens with information on the service request--such as **Type**, **Subject**, **Created**, **ID**, and **Status**--and a **Correspondence** area that includes the description of the request you created.

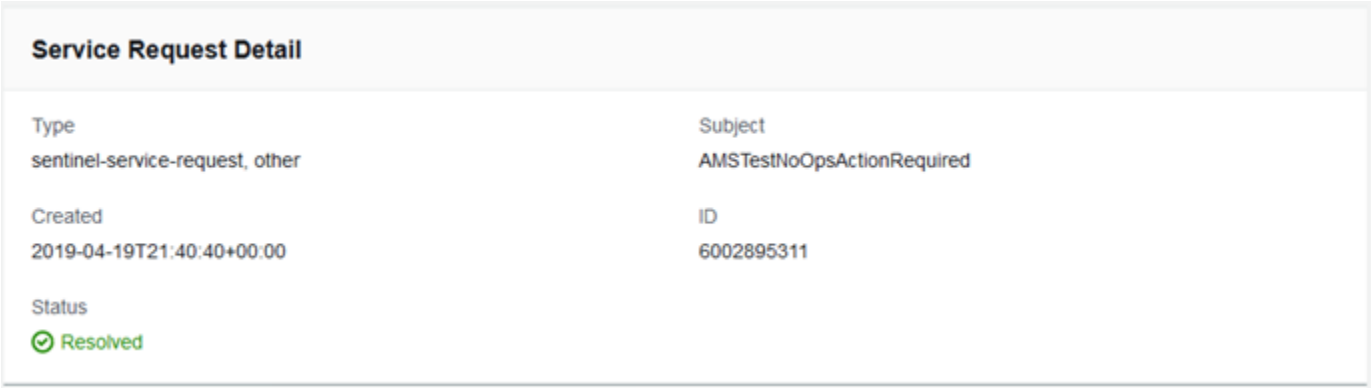

   Additionally, your service request displays on the **Service Request** list page. Use this when you have an alert but have not yet heard from AMS.

   Click **Reply** to open a correspondence area and provide additional details or status updates.

   Click **Resolve Case** when the service request has been resolved.

Click **Load More** to view additional correspondences that do not fit on the inital page.

Don't forget to rate the communication!

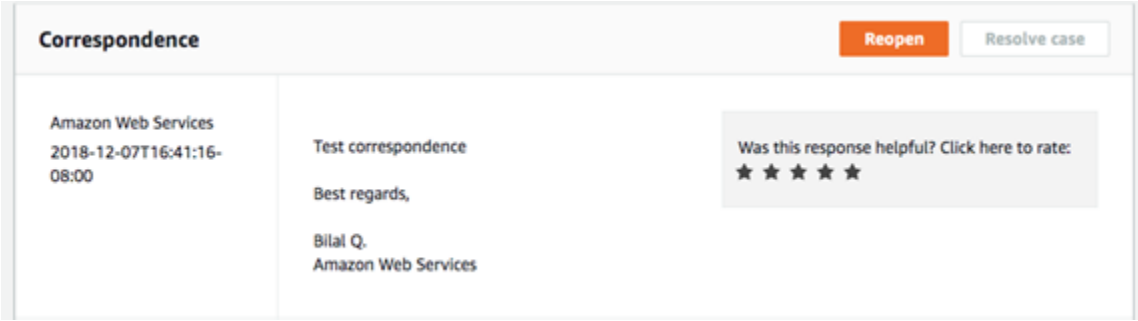

For billing-related queries, use the **Other** Category in the AMS console; the `ct-1e1xtak34nx76` change type in the AMS CM API, or the `IssueType=AMS` in the AWS Support API.

YouTube Video: How and when to raise service requests from AWS Console and what are it's Service Level Objectives?

Monitoring and updating service requests in AMS

You can update, monitor, and review incident reports and service requests, both called cases, by using the AMS console, or programmatically using the AWS Support API. For information on using the AWS Support API, see `DescribeCases` operation.

To monitor a case, incident or service request, using the AMS console, follow these steps.

1. In the AMS console **Incident reports** or **Service requests** dashboard, browse to a case and choose the **Subject** to open a details page with current status and correspondences.

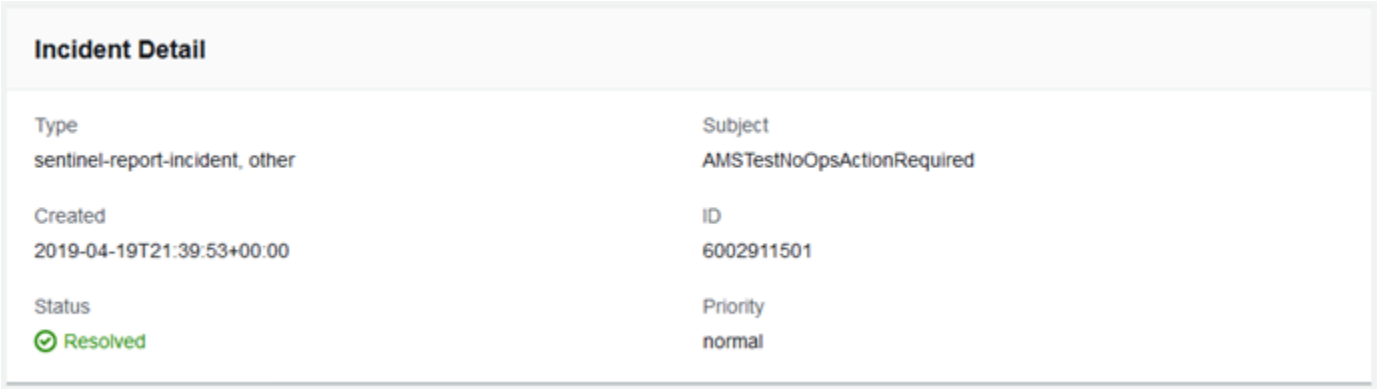

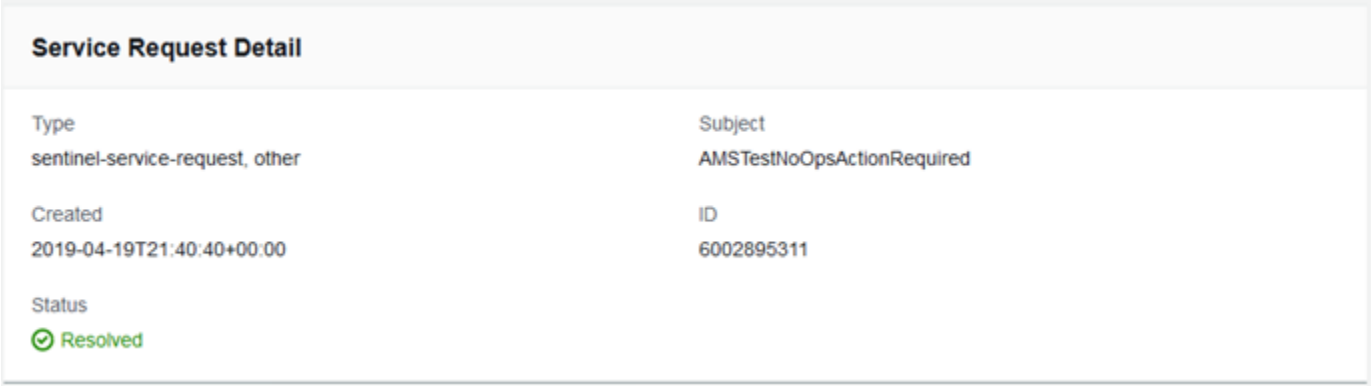

When a reported incident or service request case is updated by the AMS operations team, you receive an email and a link to the incident in the AMS console so you can respond. You can't respond to incident correspondence by replying to the email.

> **Important**
>
> You must have entered an email address to receive notifications of state change for a service request or incident case. Notifications only go to the email address added to the case when it's created.
> The link in the notification email will not work unless you are using an email server on your AMS federated network. However, you can respond to the correspondence by going to your AMS console and using the case details page.

2. If there are many cases in the list, you can use the **Filter** option:

   - **All open** (default): Use this filter to see all cases that have not been resolved.
   - **Unassigned**: Use if you've just submitted the case and have not received any notice that the case state has changed. Note, incidents and service request cases are addressed with different promptness depending on the submitted priority (incidents) or your service level agreement (service requests).
   - **Open**: Use if you have received notice that the case is "Pending Amazon" action; this means that the case has been assigned but work has not yet begun.
   - **Reopened**: Use if you have received notice that the case was reopened after having been resolved.
   - **Work in progress**: Use if you have received notice that an operator has begun to work on the case.

- **Pending customer action**: Use if you have received an operator request for action on your part.
- **Customer action completed**: Use if you have received notice that your action on the case has been processed.
- **Resolved**: Use to view cases that you know have been resolved. Resolved cases are maintained in history for twelve months.
- **Any status**: Use this filter to see all cases, regardless of status.

3. To check the latest status, refresh the page.
4. If there are so many correspondences that they do not all appear on the page, choose **Load More**.
5. To provide an update to the case status, choose **Reply**, enter the new correspondence, and then choose **Submit**.
6. To close out the case after it has been resolved to your satisfaction, choose **Close case**.

   Be sure to rate the service through the 1-5 star rating to let AMS know how we're doing!

Responding to an AMS-generated service requests

AMS patch management sends service requests (aka service notification) to you prior to the time of your set maintenance window; for more information, see AMS maintenance window. AMS also sends service notifications to you when there is a chance that your infrastructure will be impacted by an AWS service or when an EC2 instance in your account may need to be rebooted; for more information, see Service notifications.

> **Note**
>
> AMS sends communications to the primary email address on your AWS account that you have given; we recommend adding an alternate Operations contact email alias to facilitate the service request or service notification management process. Adding these emails is covered during the AMS onboarding process and related onboarding documentation.

Billing questions

To submit a billing-related question, complete the following steps:

1. Open the AWS Support Center at https://console.aws.amazon.com/support/home#/.
2. Choose **Account & billing**.

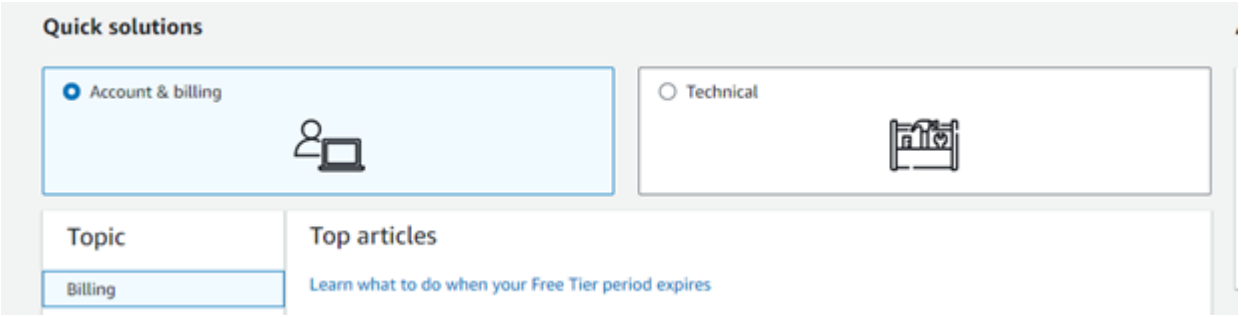

3. Choose **Create case**.

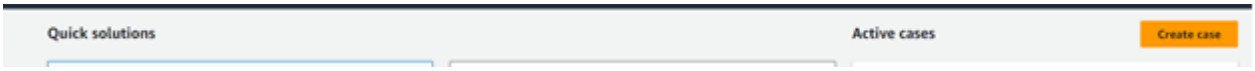

4. Choose **Account and billing**, and then follow the prompts to submit your case.

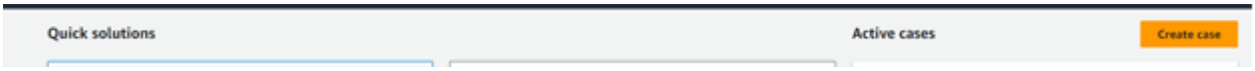

Operations On Demand

Operations on Demand (OOD) is an AWS Managed Services (AMS) feature that extends the standard scope of your AMS operations plan by providing operational services that are not currently offered natively by the AMS operations plans or AWS. Once selected, the catalog offering is delivered by a combination of automation and highly skilled AMS resources. There are no long term commitments or additional contracts, allowing you to extend your existing AMS and AWS operations and capabilities as needed. You agree to purchase blocks of hours (OOD blocks), 20 hours per block, on a monthly basis.

You can select from the catalog of standardized offerings and initiate a new OOD engagement through a service request. Examples of OOD offerings include assisting with the maintenance of Amazon EKS, operations of AWS Control Tower, and management of SAP clusters. New catalog offerings are added regularly based on demand and the operational use cases we see most often.

OOD is available for both AMS Advanced and AMS Accelerate operations plans and is available in all AWS Regions where AMS is available.

AMS performs Customer Security Risk Management (CSRM) while implementing your requested changes. To learn more about the CSRM process, see Change request security reviews.

Operations on Demand catalog of offerings

Operations on Demand (OOD) offers you the services described in the following table.

> **Note**
>
> For definitions of key terms refer to the AWS Managed Services documentation Key Terms.

Operations Plan	Title	Description	Expected Outcomes
AMS Accelerate	**Amazon EKS cluster maintenance**	AMS frees your container developers by handling the ongoing maintenance of your Amazon Elastic Kubernetes Service (Amazon EKS) deployments.	Customer teams assisted with the underlying operations work of

		AMS performs the end-to-end procedures necessary to update a cluster addressing the component s of control plane, add-ons, and nodes. AMS performs the updating to managed node types as well as a curated set of Amazon EKS and Kubernetes add-ons.	updating Amazon EKS clusters.
AMS Accelerate	**AMI Building and Vending**	AMS provides ongoing management of AMI building and vending for customers. Our engineers perform a monthly release of subscribed AMIs, release on-demand AMIs for emergent patching activities, manage changes using runbooks, and monitor AMI builds using CloudWatch Monitoring. We also provide troubleshooting assistanc e and detailed reporting for all AMIs used in designated accounts. This offering requires AMI build Pipelines to be deployed via EC2 Image builder. AMS does not support any other automation or service that interacts with EC2 Image builder.	Customer security posture improved and customer time spent on building and vending AMIs reduced.

AMS Accelerate	**Curated change execution**	Work with our skilled operations engineers to translate your business requirements into validated change requests that can be executed safely within your AWS environment. Take advantage of our unique approach to automation and knowledge of operational best practices (for example, impact assessment, roll backs, two-person rule), whether it is a simple change at scale or a complex action with downstream impacts.	Customers assisted with defining, creating, and executing custom change requests. Changes can be manual or automated (CloudFormation, SSM). Includes consultation with AWS Support for configuration guidance when necessary. Not intended for changes to application code, application installation/deployment, data migration, or OS configuration changes.

AMS Accelerate	**AWS Network Firewall Operations**	AMS collaborates with you to onboard your firewall and implement and manage the policies and rules for ongoing firewall operations. Our engineers do this by leveraging our operational best practices and automation to configure standardized policies and rules, and by enabling monitoring to detect changes made outside of the automation process. AMS quickly notifies you of unwanted changes and provides options to include them, if requested, or restore the account to a previous configuration to ensure the overall stability of your systems.	Customer teams assisted with reducing management overhead by quickly detecting unintentional network firewall changes, resulting in improved incident resolution and reduced root cause analysis time for both expected and unexpected issues.
AMS Accelerate	**AWS Control Tower operations**	Ongoing operations and management of your AWS Control Tower landing zone, including AWS Transit Gateway and AWS Organizations - providing a comprehensive landing zone solution. We handle account vending, SCP and OU management, drift remediation, SSO user management, and AWS Control Tower upgrades with our library of custom controls and guardrails.	Customer teams assisted with some of the underlying operations work of managing AWS Control Tower, AWS Transit Gateway, and AWS Organizations.

AMS Accelerate	**AWS landing zone Accelerate operations**	AMS provides ongoing operations of AWS landing zones deployed through AWS Landing Zone Accelerator (LZA). Our engineers handle configuration file changes, AWS Control Tower (CT) environment management (account vending, OU creation, CT guardrails), service contol policy (SCP) management, CT drift detection and remediation, network configuration management, and updates to CT and the LZA framework. AWS LZA provides a means to set up and govern a secure, multi-account AWS environment using operational best practices and services such as AWS Control Tower.	Customer teams assisted with ongoing operations and management of the AWS Landing Zone Accelerator solution.
AMS Accelerate	**SAP Cluster Assist**	Dedicated alarming, monitoring, cluster patching, backup, and incident remediation for your SAP clusters. This catalog item allows you to offload some of the ongoing operational work from your SAP operations team so that they can focus on capacity management and performance tuning.	Customer or partner SAP teams assisted with some of the underlying operations work. Still requires the customer to provide other SAP capabilities such as capacity management, performance tuning, DBA, and SAP basis administration.

AMS Accelerate	**SQL Server on EC2 Operations**	AMS collaborates with you to onboard, implement, and manage the ongoing operations of your SQL Server databases deployed on EC2 instances. Our engineers leverage our operational best practices and automation to free up your database teams by performin g tasks such as backup and patching, extending AMS operational support to SQL Server patching to include cluster-a ware rolling updates, backup and restore services aligned with our ransomware defense strategy, and monitoring adherence to customer-provided backup and patching controls.	SQL Server customers assisted with offloadin g patching and backup database operations to improve resilience, and security posture of their workloads , in addition to optimizing license costs by bringing their own licenses (BYOL) to EC2.
AMS Advanced	**Amazon EKS Cluster Maintenance**	AMS frees your container developers by handling the ongoing maintenance and health of your Amazon Elastic Kubernetes Service (Amazon EKS) deploymen ts. AMS performs the end-to-end procedures necessary to update a cluster addressing the component s of control plane, add-ons, and nodes. AMS performs the updating to managed node types as well as a curated set of Amazon EKS and Kubernetes add-ons.	Customer teams assisted with the underlying operations work of updating Amazon EKS clusters.

AMS Advanced	**Priority RFC Execution**	Designated AMS operations engineer capacity to prioritize the execution of your requests for change (RFC). All submissions receive a higher level of response and priority order can be adjusted by interacting directly with engineers through an Amazon Chime meeting room.	Customers receive a response SLO of 8 hours for RFCs.
AMS Advanced and AMS Accelerate	**Legacy OS Upgrade**	Avoid an instance migration by upgrading instances to a supported operating system version. We can perform an in-place upgrade on your selected instances leveraging automatio n and the upgrade capabilities of the software vendors (for example, Microsoft Windows 2008 R2 to Microsoft Windows 2012 R2). This approach is ideal for legacy applications that cannot be easily re-installed on a new instance and provides additional protectio n from known and unmitigat ed security threats on older OS versions.	Solution provided for applications that can no longer be re-installed on a new instance (for example, lost source code, ISV out of business, and so forth). Failed upgrades can be rolled back to their original state. From an operational perspective, this is preferred as it puts the instance in a more supportable state with the latest security patches.

Topics

- Requesting AMS Operations On Demand
- Making changes to Operations on Demand offerings

Requesting AMS Operations On Demand

AWS Managed Services (AMS) Operations on Demand (OOD) is available for all AWS accounts that have been onboarded to AMS. To take advantage of Operations on Demand, request additional information from your cloud service delivery manager (CSDM), Solutions Architect (SA), account manager, or Cloud Architect (CA). Available OOD offerings are listed in the preceding Operations on Demand catalog of offerings table. After the engagement scoping is completed, submit a service request to AMS Operations to initiate an engagement for OOD.

Each OOD service request must contain the following detailed information pertaining to the engagement:

- The specific OOD offerings requested, and for each specific OOD offering:
  - The number of blocks (one block is equal to 20 hours of operational resource time in a given calendar month, to be charged at AWS's then-current standard rate for the applicable Operations on Demand offering) to allocate to the specific OOD offering.
  - The account ID for each AWS Managed Services account for which the specific OOD offering is being requested.

OOD service requests must be submitted by you through either:

- The AWS Managed Services account that receives the applicable Operations on Demand offerings, or
- An AWS Managed Services account that is an AWS Organizations Management account in **all features** mode, on behalf of any of its member accounts that are AWS Managed Services accounts.

After the OOD service request is received, AMS Operations reviews and updates the accounts with their approval, partial approval, or denial.

Once the OOD offerings service request is approved, AMS and you coordinate to begin the engagement. No OOD offerings are initiated until the service request is approved and an engagement start date is agreed on.

AMS uses a monthly subscription allocation of OOD blocks. We allocate the approved number of blocks monthly, starting from the engagement start date, until you request to opt out through a

new service request. OOD blocks are valid for a calendar month. Unused blocks, or block portions, are not rolled over or carried forward to future months.

You are billed a minimum of one OOD block each month, regardless of the number of hours actually used. Any additional, allocated, OOD block in which no hours were used, is not billed.

Making changes to Operations on Demand offerings

To request changes to ongoing engagements for Operations on Demand (OOD) offerings, submit a service request containing the following information:

- The modification(s) being requested, and
- The requested date for the modifications to become effective.

After receiving the OOD service request, AMS Operations reviews the request and either updates with their approval or requests that the assigned CSDM work with you to determine the scope and implications of the modification. If the modification is determined to require a scoping effort with the CSDM, you are required to submit a second OOD service request to initiate the modified engagement following the completion of the scoping exercise.

Once approved, the most recently modified block allocation becomes and continues to stay active, superseding any prior block allocations, unless agreed otherwise by AWS and you.

Document history

The following table describes the important changes to the documentation since the last release of AMS.

- **API version: 2019-05-21**
- **Latest documentation update:** April 12, 2024

Change	Description	Link
New or Updated CTs and Walkthroughs: AMS Advanced Change Type Reference Document History.		
New AMS AMIs: AMS Amazon Machine Images (AMIs).		
April 2024		
Updated **Offboard from multi-account landing zone (MALZ) landing zone accounts** section.	Added detailed information, including how to offboard Application and Core accounts from multi-account landing zone.	See Offboard from AMS multi-account landing zone accounts
March 2024		
Updated **How service request management works** section.	Added clarification on how AMS handles service requests that contain a feature request or a bug.	See How service request management works
Updated **Incident management service commitments** section.	Added a link to the AMS Service Level Agreement.	See Incident management service commitments
Updated **Get support** section.	Updated **Get support** section to include a new **Billing questions** section.	See Billing questions

Change	Description	Link
Updated **Service request management** description.	Updated **Service request management** description in **Service description** topic.	See Service description
Updated AMS Automated IAM Provisioning	Updated AMS Automated IAM Provisioning with custom deny list information	See AMS Automated IAM Provisioning permission boundary check
February 2024		
Updated Supported Operating Systems	Updated Supported Operating Systems to include SUSE Linux Enterprise Server 15 SP5.	See Supported configurations
Added note to **Alerts from baseline monitoring in AMS**.	Added note indicating that the alarm for EC2 Non-root Volume Usage is disabled by default.	See Alerts from baseline monitoring in AMS
Added a new section **AMS Event Router** to **Monitoring and event management**.	Added a new section discussing the AMS Advanced Event Router.	See Using Amazon EventBridge Managed Rules in AMS
Updated: AMS AMI Notes	Zip file includes notes on the latest AMS Amazon machine images (AMIs) and a CSV file of the latest AMIs.	See AMIs.csv-and-notes.02.2024
February 2024		
Added a new section for Amazon EventBridge rule service-linked role for AMS Advanced	Added a new section for Amazon EventBridge rule service-linked role for AMS Advanced in the **Infrastructure Security** section.	See Amazon EventBridge rule service-linked role for AMS Advanced

Change	Description	Link
Updated Self Servicing Provision Mode section	Added a new section for the new Amazon Inspector in **Self Servicing Provision Mode**.	See Amazon Inspector Classic (AMS SSPS)
January 2024		
Updated Planned event management (PEM) section	Added additional details and an FAW to **Planned event management (PEM)**.	See Planned event management (PEM)
Added a new section for SSM Agent auto installation	Added a new section for **SSM Agent auto installation** in **Automated EC2 instance configuration**.	See SSM Agent automatic installation
Added AWS Resilience Hub (AMS SSPS)	Added a new SSPS service.	See AWS Resilience Hub (AMS SSPS)
Updated: AMS AMI Notes	Zip file includes notes on the latest AMS Amazon machine images (AMIs) and a CSV file of the latest AMIs.	See AMIs.csv-and-notes .01.2024
December 2023		
Updated Direct Change mode in AMS	Added a new subsection, **Direct Change Mode** use cases, to **Direct Change mode in AMS**.	See Direct Change mode in AMS
Updated AWS Amplify (AMS SSPS)	Updated FAQ to clarify that a Risk Acceptance is required to request Amplify.	See AWS Amplify (AMS SSPS)
New AWS Elastic Disaster Recovery (AMS SSPS)	Added a new SSPS service	See AWS Elastic Disaster Recovery (AMS SSPS)

Change	Description	Link
New Amazon Managed Service for Prometheus (AMS SSPS)	Added a new SSPS service	See Amazon Managed Service for Prometheus (AMS SSPS)
Updated How continuity management works section.	Added a new subsection, **AMS backup monitoring and reporting**.	See How continuity management works
New Amazon DevOps Guru (AMS SSPS)	Added a new SSPS service	See Amazon DocumentDB (with MongoDB compatibility) (AMS SSPS)
Updated: AMS AMI Notes	Zip file includes notes on the latest AMS Amazon machine images (AMIs) and a CSV file of the latest AMIs.	See AMIs.csv-and-notes .12.2023
November, 2023		
Updated Amazon CloudWatch Synthetics (AMS SSPS)	Updated FAQs to use the correct role names.	See Amazon CloudWatch Synthetics (AMS SSPS)
Updated Amazon API Gateway Self-service Provisioning mode	Added an additional role, `customer_apigatewa y_cloudwatch_role` , to the API Gateway section.	See Amazon API Gateway (AMS SSPS)
Added a new service to Self-service Provisioning mode	Added AWS Service Catalog AppRegistry to the **Self-Service Provisioning mode** section	See AWS Service Catalog AppRegistry
Updated: AMS AMI Notes	Zip file includes notes on the latest AMS Amazon machine images (AMIs) and a CSV file of the latest AMIs.	See AMIs.csv-and-notes .11.2023

Change	Description	Link
September, 2023		
Added a note to Using Patch Orchestrator	Added the following note to **Using Patch Orchestrator** section: "Patch failure alerts aren't created for instances that have unsupported operating systems, or that are stopped during the maintenance window"	See Patch management in AMS
Updated data encryption with additional services	Added services to Data encryption in AMS.	See Data protection in AMS
Added new paragraph to RFC error messages.	Added a new paragraph to add Create a service request link.	See Troubleshooting RFC errors in AMS
Corrected IAM role names.	Corrected the IAM rolename customer_emr_cluster_autoscaling_role.	See Self-Service Provisioning mode
Updated baselone monitoring information	Removed reference to two deprecated alarms RDSReadLatencyAlarm and RDSWriteLatencyAlarm.	See Alerts from baseline monitoring in AMS
August, 2023		
Added: AMS Security Incident Response	Added documentation for using AMS Security Incident Response.	See Security Incident Response in AMS
July, 2023		

Change	Description	Link
Added: Automated IAM Provisioning	Added documentation for using Automated IAM Provisioning.	See Automated IAM Provisioning AMS
Updated: Access roles table	Added missing roles for AMS Access.	See AMS customer account access IAM roles
June, 2023		
Updated: List of monitored RDS alerts.	Updated the list of RDS alerts for AMS baseline monitoring. 9 new RDS alert types were added and 3 existing RDS alert types were removed.	See Alerts from baseline monitoring in AMS.
Updated: Access roles table	New roles for AMS Security.	See AMS customer account access IAM roles
May, 2023		
Updated: Service Billing Start Date policy.	Updated definitions of Billing Start Date.	See AMS key terms.
April, 2023		
Updated: Monthly Billing Self-Service Report.	Added note: The Monthly Billing reports are only available in a Management Payer account (AMS Advanced multi-account landing zone), but are available for all linked AMS Accelerate-managed accounts.	See Monthly billing report.
Updated: Removed "Standard Patching" content	AMS uses Patch Orchestrator.	Patch management in AMS

Change	Description	Link
Updated: What is AMS?	Moved some topics previously under What is AMS? to be part of the AMS Service Description.	Service description
Updated: Offboarding multi-account landing zone	Made various clarifications.	Offboard from AMS multi-account landing zone accounts
Updated: AWS Transfer Family (AMS SSPS)	Added link to transfer setup tutorial.	AWS Transfer Family (AMS SSPS)
Updated Content: Self-service provisioning	Replaced "CodeSuite" with "Code services" per AWS legal.	AMS Code services (AMS SSPS)
Updated Content: CloudWatch metrics and alarms	Added link to Example: Count occurrences of a term.	Creating custom CloudWatch metrics and alarms in AMS
Updated: AMS AMI Notes	Zip file includes notes on the latest AMS Amazon machine images (AMIs) and a CSV file of the latest AMIs.	AMIs.csv-and-notes.04.2023
March, 2023		
Updated Content: Offboarding from AMS	Clarified what resources are deleted when offboard multi-account landing zone accounts	Offboard from AMS multi-account landing zone accounts
Updated: AMS AMIs	Added link to AMI ZIP file for each month in the Doc History section.	Document history
Updated: Auto remediation	Removed LVM support for EC2 volume automation.	AMS automatic remediation of alerts

Change	Description	Link
Updated: Patch RACI	Several updates and clarfications to the RACI for patching.	AMS responsibility matrix (RACI)
Updated Content: Self-service provisioning	Added an FAQ bullet. To launch a new AWS Datasync agent, WIGS ingestion is not required.	Self-Service Provisioning mode
Updated Content: Self-service provisioning	Added an FAQ bullet. To launch a new AWS Datasync agent, WIGS ingestion is not required.	Self-Service Provisioning mode
Updated: AMS AMI Notes	Zip file includes notes on the latest AMS Amazon machine images (AMIs) and a CSV file of the latest AMIs.	AMIs.csv-and-notes.03.2023
February, 2023		
Updated Content: Offboarding from AMS	Clarified how to offboard multi-account landing zone environments, VPCs, and Application accounts	Offboard AMS accounts
Updated Content: Finding ARNs	Added DynamoDB `describe-table` CLI for finding a DynamoDB table ARN	Find Amazon Resource Names (ARNs) in AMS

Change	Description	Link
Updated Content: Self-Service Provisioning	Removed the AMS "CodeSuite" option as it is not an actual SSPS. You can still use the Management | AWS service | Self-provisioned service | Add (review required) (ct-3qe6io8t6jtny) change type and request the three services: CodeBuild, CodeDeploy and CodePipeline. AMS will then provision the following IAM roles to your account: `customer_codebuild_service_role` , `customer_codedeploy_service_role` , and `aws_code_pipeline_service_role` . After provisioned in your account, you must onboard the role in your federation solution.	Self-Service Provisioning mode
Updated Content: Secrets Manager update	Corrected roles needed for multi-account landing zone (MALZ) vs single-account landing zone (SALZ).	Sharing Keys using Secrets Manager FAQs
Updated Content: AMS automatic remediation of alerts	Added support for Logical Volume Manager (LVM) volumes.	EC2 volume usage remediation automation
Updated Content: AMS Amazon Machine Images (AMIs)	Added the section Offboarding AMS AMIs with sample code to remove AMIs from your account.	AMS Amazon Machine Images (AMIs)

Change	Description	Link
Updated Content: IAM User Role	Updated the IAM policy: AMSBillingPolicy.	IAM user role in AMS
New Content: Unsupported OSes	Added information on what services AMS provides for unsupported operating systems (OSes).	Capabilities for unsupported operating systems in AMS
Updated Content: On-demand reports	Certain on-demand reports not available in AMS Advanced and were mistaken shown as available.	On-request reporting
Updated Content: Offboarding AMS Accounts	Clarified instructions for offboarding MALZ application accounts.	Offboard AMS Application accounts
Updated Content: Secrets Manager	Corrected the names of IAM roles required to use Secrets Manager.	Secrets Manager in AWS Managed Services FAQs
Updated: AMS AMI Notes	Zip file includes notes on the latest AMS Amazon machine images (AMIs) and a CSV file of the latest AMIs.	AMIs.csv-and-notes.02.2023
January, 2023		
New Content: AWS Device Farm (AMS SSPS)	Added a new SSPS service: AWS Device Farm.	AWS Device Farm (AMS SSPS)
Updated: supported Windows versions	Added support for Windows Server 2022.	AMS Amazon Machine Images (AMIs), Supported configurations, and AMS AMI notifications with SNS

Change	Description	Link
Updated: Continuity management	Updated the rules in the Default AMS backup plan.	Default backup plans, multi-account landing zone
Updated: AMS AMI Notes	Zip file includes notes on the latest AMS Amazon machine images (AMIs) and a CSV file of the latest AMIs.	AMIs.csv-and-notes.01.2023
December, 2022		
Updated: Using bastions	Fixed bad link.	Accessing instances using bastions
Updated: Resource Scheduler	Made several improveme nts and added links to AWS Instance Scheduler for more context.	AWS Managed Services Resource Scheduler
Updated: Windows AMIs and Supported Configurations (for new Windows AMIs)	Updated AMS AMI content added from EC2Launch (Windows Server 2016 and later) to EC2Launch (Windows Server 2016 and Windows Server 2019) and added EC2LaunchV2 (Windows Server 2022 and later). Updated Windows-based AMIs from Microsoft Windows Server (2012, 2012 R2, 2016, and 2019) to Microsoft Windows Server (2012, 2012 R2, 2016, 2019 and 2022).	AMS Amazon Machine Images (AMIs) and Service description
Updated: Resourced Scheduler section	Improved methods for deploying and customizing AMS Resource Scheduler.	AWS Managed Services Resource Scheduler

Change	Description	Link
Updated: Setting upu AMS: private and public DNS	Updated the DNS architecture diagram.	Setting up private and public DNS
Updated: MALZ network architecture	Updated the diagram and added guidance for Accelerate application accounts.	MALZ network architecture
Updated: Setting up: Using tags	New note: custom tagging is only supported for MALZ application accounts, not core accounts.	AMS infrastructure automatic tagging
Updated: Access managemen t: using bastions	Updated introduction to inclue RDP bastions.	Saving costs on Single-ac count landing zone (SALZ) bastions
Updated: AMS default settings: alerts	Added EC2 instance: *Non-Root Volume Usage* to the table of alerts.	Alerts from baseline monitoring in AMS
Updated: Continuity Management	Added guidance about continuous backups.	How continuity management works
Updated: Automated EC2 instance configuration	Added support for *PowerBrok er Identity Service* (PBIS) and *On Instance Code* (OIC).	Automatically update PBIS on Linux instances and Automatically update code on Linux instances
Updated: Self-Service Provisioning for Secrets Manager	Updated the CT for adding Secrets Manager to your account (under FAQs).	AWS Secrets Manager (AMS SSPS)
Updated: Log management	Updated the list of EC2 system-level logs.	Amazon Elastic Compute Cloud (Amazon EC2) - system level logs

Change	Description	Link
Updated: AMS AMI Notes	Zip file includes notes on the latest AMS Amazon machine images (AMIs) and a CSV file of the latest AMIs.	AMIs.csv-and-notes.12.2022
November, 2022		
Updated: AMS Amazon Machine Images (AMIs)	Updated supported SUSE Linux versions.	AMS Amazon Machine Images (AMIs)
Updated: MALZ accounts	Added guidance for deleting a Customer Managed applicati on account.	Customer Managed applicati on accounts
Updated: Setting up AMS	Added `customer-ams-amazon2-security-enhanced .`	AMS AMI notifications with SNS
Updated: How monitoring works	Updated explanation of service notifications and incident reports.	How monitoring works
Updated: MALZ Application account types	Improved the explanation of account types.	Application account types
Updated: Developer mode	Added a warning about Developer mode.	Before you begin with AMS Developer mode
Updated: Planned event management	Added the section: Types of PEM	Planned event management (PEM)
Updated: Amazon Machine Images (AMIs)	Updated supported SUSE Linux versions	AMS Amazon Machine Images (AMIs)

Change	Description	Link
Updated: AMS AMI Notes	Zip file includes notes on the latest AMS Amazon machine images (AMIs) and a CSV file of the latest AMIs.	AMIs.csv-and-notes.11.2022
October, 2022		
New: Automated Instance Configuration	New section describes the Automated Instance Configuration process.	Automated instance configuration in AMS Advanced
New: Only manual CT is acceptable for some SSPS	Updated over 50 self-service provisioning service FAQs to use the manual CT and not the automated CT for adding SSPS.	Self-Service Provisioning mode
Update: Setting up AMS	Added two policies to the Amazon EC2 IAM instance profiles for MALZ.	EC2 IAM instance profile
New: Library of custom detective and preventive rules.	Added a set of example service control policies (SCPs) and preventive Config rule controls based off our learnings from multiple customers.	Curated SCPs and Config Rules
Update: AWS Backup warning	Added a warning: "Do not edit AMS backup plans as your changes may be lost. Instead, create new backup plans using ct-2hyozbpa0sx0m for your custom configurations."	How continuity management works

Change	Description	Link
Update: AWS Backup caution	Added a note about adding new IAM roles to your federation.	Deploying IAM resources in AMS Advanced
Update: Monitoring management	Alerts generate incident reports, not service requests.	How monitoring works
Update: Bring your own EPS	Applies to SALZ as well as MALZ.	AMS bring your own EPS
Update: Accelerate Applicati on account	Clarified that your Accelerat e account is an Application account.	Application account types
Updated: AMS AMI Notes	Zip file includes notes on the latest AMS Amazon machine images (AMIs) and a CSV file of the latest AMIs.	AMIs.csv-and-notes.10.2022
September, 2022		
Updated: CLI command examples for finding resources	Added new example and that the `--region` option may be needed.	Finding the data you need (SKMS), AMS
Updated: Provisioning IAM roles	IAM roles can now be created and managed with the `AWSManagedServices CloudFormationAdmi nRole .`	Creating stacks using Direct Change mode
Updated: AMS Technical Standards	AMS-STD-007 Logging: (#20) Clarified forwarding requirements.	Security and compliance

Change	Description	Link
Updated: How continuity management works	Revised Start Backup Job wording to "on-demand" rather than "existing".	How continuity management works
Updated: Security and compliance	Updated description and guidance for standard AMS-STD-007 number 20: forwarding logs between accounts.	Security and compliance
Updated: Change management use cases	Removed a broken link to the legacy Change Management User Guide.	Change management use cases
Updated: AMS AMI Notes	Zip file includes notes on the latest AMS Amazon machine images (AMIs) and a CSV file of the latest AMIs.	AMIs.csv-and-notes.09.2022
August 11, 2022		

Change	Description	Link
Updated: Chapter headings for consistency and readability, moved some topic sub-sections into more appropriate sections	"MALZ network architecture" and "SALZ network architecture" are now, both, subsections of the top-level "Network architecture" section, formerly the "AMS network architecture" section "Modes for change management" is the new heading for "Change management" "Default settings" is now a subsection of "Setting up AMS" "AD FS claim rule and SAML settings" (formerly "ActiveDirectory Federation Services (ADFS) claim rule and SAML settings) is now a subsection of "Setting up AMS" "Access management" is the new heading for "Access in AMS" and is moved up in the TOC "Finding the data you need" is the new heading for "Service knowledge management" "Reports and options" is the new heading for "AMS	What is AWS Managed Services?

Change	Description	Link
	Reporting" and is lower down in the TOC "Operations on Demand" is now the last topic in the TOC	
Updated: Finding and ARN, New: Finding a resource with an ARN	Both procedures completely rewritten for usefulness.	Find Amazon Resource Names (ARNs) in AMS and Find resources by ARN in AMS.
Updated: Connecting your CMA with Transit Gateway	The automation does not support adding routes to core route domains, and the procedure needed updating.	Connecting your CMA with Transit Gateway
Updated: MALZ basic components pricing	All prices are in US Dollars, formatted with dollar signs.	AMS environment basic components
Updated: AMS AMI Notes	Zip file includes notes on the latest AMS Amazon machine images (AMIs) and a CSV file of the latest AMIs.	AMIs.csv-and-notes.08.2022
July 14, 2022		
Updated: Self-Service Reporting	Added instructions for encrypting AWS Glue metadata with KMS keys.	Self-service reporting
Updated: AMS baseline monitoring	Added DeleteRecoveryPoint backup alert.	Alerts from baseline monitoring in AMS
Updated: Supported operating systems	Added End of Support date for Amazon Linux 2.	Supported configurations
Updated: Self-Service Provisioning	Added a prerequisite for the AWS Transfer SSPS.	AWS Transfer Family (AMS SSPS)

Change	Description	Link
Updated: AMS Reporting	Added note about Opt-in Regions.	Reports and options
Updated: RFC correspondence and attachment	Clarified allowed text file types; in particular, YAML files must end in .yaml (not .yml).	RFC correspondence and attachment (console)
June 21, 2022		
Updated content	The AMS mode previousl y known as "Change Management mode" or "Standard CM mode" is now known as "RFC mode." The modes section has been expanded.	Modes overview.
New alarm	Added a AWS Backup alarm.	Alerts from baseline monitoring in AMS
June 16, 2022		
New content	Incident management. Incidents that are not a security risk can now be resolved by AMS with your approval in the incident report and do not need a separate RFC and approval.	Incident management
Updated content	MALZ: Updated network architecture diagram. Updates: The VPC Peering for the master account VPC to shared services vpc should be removed as it doesn't exist.	Networking account architect ure

Change	Description	Link
	Sagemaker self-service provisioned service (SSPS). Updated with new IAM role added at onboarding for Sagemaker's use.	Amazon SageMaker (AMS SSPS)
	To list of AMIs supported for SNS notifications: Added customer-ams-sles12, customer-ams-sles15, customer-ams-amazon1-security-enhanced, customer-ams-rhel8, customer-ams-rhel8-security-enhanced, customer-ams-ubuntu18, customer-ams-windows2012, customer-ams-windows2019, and customer-ams-windows2019-security-enhanced. Removed customer-ams-rhel6 and customer-ams-rhel6-security-enhanced AMIs.	AMS AMI notifications with SNS
	Removed escalation emails.	Getting help in AWS Managed Services
	Moved topic list to below opening paragraphs.	What is AWS Managed Services?
	Updated service logs with better links for load balaning logs, also re-formatted.	AMS aggregated service logs

Change	Description	Link
EKS self-service provision ing service (SSPS). Added information on enabling envelope secrets encryption in your cluster.	Amazon EKS on AWS Fargate (AMS SSPS)	
June 09, 2022		
Updated content, Getting help	Removed escalation path emails. AMS provides communication methods through incident reports, service requests, and RFCs.	Getting help in AWS Managed Services
May 12, 2022		
New content, Operations on Demand (OOD) subscription model	AMS has changed Operation s on Demand onboarding from the current signup and renew model, to a subscript ion allocation and default opt-in model. When you onboard an AMS account, you are automatically enrolled in Operations on Demand now.	Operations On Demand
April 14, 2022		
New content, Cost Optimizat ion	AMS provides recommend ations for cost optimization.	Cost optimization in AWS Managed Services
Updated content, Accelerate account in MALZ	An incorrect role name (CustomerDefaultAdminRole) was updated to the correct role name (AccelerateDefault AdminRole).	AMS Accelerate accounts "Accessing your Accelerate account" section.

Change	Description	Link
Updated content, AMS access IAM roles	Added other AMS IAM roles used to access your accounts.	Why and when AMS accesses your account "AMS customer account access IAM roles" section.
Updated content, AMS backup plans	Added AMS-managed backup plans.	AMS backup plans and AMS backup vaults
Updated content, AWS Secrets Manager	Updated the FAQ.	AWS Secrets Manager (AMS SSPS)
Updated content, Direct Change Mode (DCM) onboarding	AMS does not support onboarding Service Catalog customers to DCM.	Getting Started with Direct Change mode
Updated content, Service Description	Clarified the Supported Services section: • Amazon EKS on AWS Fargate -> Amazon Elastic Kubernetes Service on Fargate • Amazon ECS for Fargate -> Amazon Elastic Container Service on AWS Fargate • Amazon Kinesis -> Amazon Kinesis Data Streams	Supported AWS services
Updated content, Offboarding MALZ accounts	Updated to reference new change types for offboarding application accounts.	Offboard AMS Application accounts

Change	Description	Link
Updated content, Developer mode incident management	Updated incident SLA description to: AMS SLA does not apply for resources created or updated outside of AMS Change Managemen t (Developer Mode included) therefore, resources updated or created in Developer mode are automatically degraded to a P3 and support is best effort.	Incident management in AMS Developer mode
Updated content, DCM onboarding	The RFC template for new DCM now includes a field for your SAML Provider ARN.	Getting Started with Direct Change mode
Updated content, DCM for AWS CloudFormation	Instructions for creating and updating AWS CloudForm ation stacks now include YAML examples.	AMS Transform
Updated content, MALZ Tools account	There is a new IAM role for migrations: `AWSManage dServicesMigration Role` .	AWS Application Migration Service (AWS MGN) and Enable access to the new AMS Tools account
New content, multi-account landing zone accounts	You can create an Accelerat e account in your multi-account landing zone AMS Management account.	AMS Accelerate accounts
Updated content, API/CLI SDK installation	The installation instructions listed the wrong file name for Mac/Linux installs, and an incorrect command. This has been fixed.	Using the AMS API and CLI

Change	Description	Link
Updated content, Accelerate account in MALZ	There was an incorrect rule name (CustomerDefaultAdminRole), it's been updated to the correct one (AccelerateDefaultAdminRole).	AMS Accelerate accounts, "Accessing your Accelerate account" section
Updated content, monitoring	Root usage monitoring was revised from 85% to 95%.	Alerts from baseline monitoring in AMS
Updated content, AMI notifications	You can create many types of SNS notifications for new AMS AMIs, we've added information on creating various types.	AMS AMI notifications with SNS
Updated content, AMS default settings	Removed references to Macie *Classic*, replaced by Macie.	Alerts from baseline monitoring in AMS
Updated content, AMS reserved prefixes	Alphabetized the list of reserved prefixes.	AMS reserved prefixes
Updated content, Service Description	The features sections on change management and self-service provisioning were updated with more information on AMS modes.	AMS features
Updated content, AWS Secrets Manager	Sharing Keys using Secrets Manager.	February 10, 2022
New content, Self-service provisioning, Amazon Connect	Added an FAQ for how to request to add a list of countries for outbound or inbound calls.	February 10, 2022

Change	Description	Link
New content, Self-service provisioning, Amazon EKS on Fargate	Added an FAQ restriction that deploying EKS clusters through the AWS cloud development kit (CDK) or CloudFormation Ingest is not supported in AMS.	February 10, 2022
Changed content, Developer mode	Correction, you do not use an RFC or service request to assign users to your federatio n solution, you do that yourself depending on your solution.	February 10, 2022
Changed content, Direct Change mode (DCM)	Note that IAM is not supported in DCM.	February 10, 2022
	DCM, note validations that we do.	February 10, 2022
	DCM, clarify restrictions of different roles.	February 10, 2022
Changed content, Monitoring baseline alerts	Redshift cluster resource alerts changed.	February 10, 2022
Changed content, Self-service reporting	Added the exact s3 bucket name, (ams-reporting-data-a<Account_ID>) for customers to use to fetch the reports.	February 10, 2022

Change	Description	Link
Changed content, updated content to reference automated change types instead of manual Management | Other | Other (MOO)	Updated multi-account landing zone (MALZ) applicati on account content to reference automated change types (three, "Associating the TGW attachment to a route table", "Create routes in the TGW route tables to connect to this VPC", and "Configur ing your VPC Route tables to point at the AMS Multi-Acc ount Landing Zone transit gateway"). Receiving alerts generated by AMS Tag-based alert notifications	February 10, 2022
Changed content: AMS AMIs.	Added new information about security-enhanced AMIs. see Supported configurations, AMS Amazon Machine Images (AMIs), and Security enhanced AMIs.	January 27, 2022
New content: Self-service provisioning.	Added Amazon Fsx for OpenZFS. See Amazon FSx OpenZFS (AMS SSPS).	January 27, 2022
Changed content: Code-Depl oy self-service provisioning service (SSPS).	Additional role name, and additional restriction note. see AWS CodeDeploy (AMS SSPS).	January 27, 2022

Change	Description	Link
Changed content: Updated links.	Fixed broken links: AMS-AMIs, Finding your settings, Finding a Stack ID, Finding a VPC ID, ListVpcSummaries, ListStack Summaries, and GetStack APIs. For example, see Find stack IDs in AMS.	January 13, 2022
Changed content: EKS Support for Fargate	Added limitation to FAQs: Creating or managing EC2 nodegroups with EKS is not supported. See Amazon EKS on AWS Fargate (AMS SSPS).	January 13, 2022
Changed content: CloudFormation, Direct Change Mode (DCM)	Added instructions for creating or updating CF stacks using AmsStackTransform. See Creating stacks using Direct Change mode.	January 13, 2022
Changed content: Uniformity in AWS Service Names	AMS references to AWS services exactly match the official AWS titles or metadata. Previously, there were minor variations that complicated pattern matching. For example, see Alexa for Business (AMS SSPS).	January 13, 2022
Changed content: Self service provisioning of Elastic Container Registry (ECR)	Added an FAQ item for using ECR to manage user permissions. See Amazon Elastic Container Registry (AMS SSPS).	January 13, 2022

AWS Glossary

For the latest AWS terminology, see the AWS glossary in the *AWS Glossary Reference*.