



AMS Accelerate Concepts and Procedures

AMS Accelerate User Guide



Version April 4, 2024

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

AMS Accelerate User Guide: AMS Accelerate 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 AMS Accelerate?	1
Operations plans	1
Accelerate operations plan	2
Advanced operations plan	2
How it works	2
Key terms	3
Service description	9
AMS Accelerate features	9
Supported configurations	12
Supported services	14
Roles and responsibilities	15
Scope of changes performed by AMS Accelerate	24
Unsupported operating systems	25
Contact and escalation	26
Contact hours	26
Business hours	27
Escalation path	28
Resource Inventory	28
Getting started	30
Onboarding	30
Onboarding prerequisites	30
Step 1. Account discovery	33
Step 2. Onboarding management resources	41
Step 3. Onboarding features with default policies	51
Step 4. Customize features	61
Using the AMS consoles	63
AMS patterns	64
How AMS patterns work	65
AMS patterns	65
Automated instance configuration	68
How it works	69
SSM Agent automatic installation	70
Automated instance configuration changes	71
Offboarding	75

Offboarding considerations	75
Getting offboarding assistance	77
Notification settings	77
Tagging	79
Tags	80
What are tags?	81
How tagging works	81
Customer-managed tags	81
Accelerate-managed tags	85
Customer-provided tags	86
Tag management tools	87
Resource Tagger	87
CloudFormation	107
Terraform	111
Incident reports, service requests, and billing questions	113
Incident management	113
What is incident management?	114
How incident response and resolution work	115
Working with incidents	116
Service request management	119
How service request management works	120
Working with service requests	121
Incident report and service request testing	125
Billing questions	125
Operations On Demand	127
Requesting AMS Operations On Demand	134
Making changes to Operations on Demand offerings	135
Reports and options	136
On-request reporting	137
AMS host management	137
AMS Backup reporting	138
AWS Config Control Compliance report	142
AMS Config Rules Response Configuration report	143
Incidents Prevented report	145
Billing Charges Details report	146
Self-service reporting	147

Daily Patch reports	148
Monthly billing report	157
Daily backup report	160
Weekly Incident report	163
Data retention policy	166
Offboarding from SSR	166
Access management	167
Accessing the console	167
Permissions to use features	167
Why and when we access your account	182
Access Triggers	182
Access IAM roles	183
How we access your account	185
How and when to use root	186
Security management	187
Using the Log4j SSM Document to Discover Occurrences	188
Infrastructure security	189
Using service-linked roles	190
AWS managed policies	203
Data protection	211
Amazon Macie	212
GuardDuty	212
Data encryption	214
AWS Identity and Access Management	214
Authenticating with identities in AMS Accelerate	214
Managing access using policies	221
Security Incident Response	222
How it works	222
Prepare	223
Detect	226
Analyze	228
Contain	231
Eradicate	234
Recover	236
Post Incident Report	237
Security Incident Response Runbooks	238

Security event logging and monitoring	244
Configuration compliance	244
AMS Config Rule library	245
Responses to violations	268
Creating rule exceptions	270
Customized findings responses	271
Incident response	273
Incident response and onboarding	273
Resilience	274
Security control for end-of-support operating systems	274
Security best practices	275
Change request security reviews	275
Customer Security Risk Management process	275
AMS Accelerate technical standards	276
Standard controls in AMS Accelerate	276
Changes that introduce high or very high security risks in your environment	290
Monitoring and event management	292
What is monitoring?	292
How monitoring works	293
Alert notification	295
Tag-based alert notification	295
Alerts from baseline monitoring in AMS	296
Alarm Manager	309
How Alarm Manager works	309
Getting started with Alarm Manager	311
Alarm Manager tags	311
Alarm Manager configuration profiles	316
Creating additional CloudWatch alarms	333
Viewing the number of resources monitored by Alarm Manager	333
AMS automatic remediation of alerts	335
EC2 status check failure remediation automation	336
EC2 volume usage remediation automation	336
Amazon RDS low storage event remediation automation	337
AMS Event Router	338
Amazon EventBridge Managed Rules deployed by AMS	338
Creating Managed Rules for AMS	340

Editing Managed Rules for AMS	341
Deleting Managed Rules for AMS	341
Continuity management	342
How continuity management works	342
Select an AMS backup plan	343
Default AMS backup plan	343
Enhanced backup plan	344
Data Sensitive backup plan	345
AMS Accelerate onboarding backup plan	345
Tag your resources for backup	346
View backups in AMS vaults	347
Monitoring and reporting for backups	349
Patch management	350
Patching recommendations	351
Patch responsibility recommendations	351
Guidance for application teams	352
Guidance for security operations teams	352
Guidance for governance and compliance teams	353
Example design for high availability Windows application	353
Patch recommendations FAQs	354
Create patch window	355
Create Patch Tuesday patch window: AMS console	355
Create patch window: AWS CloudFormation	356
Create patch window: Systems Manager console	358
Create patch window: Systems Manager CLI	359
AMS Accelerate patch baseline	361
Default patch baseline	361
Custom patch baseline	362
On-demand patching permissions	362
Patch monitoring and failure remediation	363
Patch notification	363
Patch remediation	367
Monitoring and Incident Management for EKS	369
What is Monitoring and Incident Management for EKS?	369
How Monitoring and Incident Management for EKS works	370
AMS responsibility matrix (RACI)	371

Baseline alerts	373
Alerts and actions	373
Requirements	377
Onboarding	378
Offboarding	379
Cost optimization with AMS Resource Scheduler	381
Using resources with Resource Scheduler	382
Onboarding Resource Scheduler	383
Customizing Resource Scheduler	384
Using Resource Scheduler	385
Working with periods and schedules	388
Tagging resources	396
Cost estimator	396
Alarm suppressor	397
Log management	399
Log management — AWS CloudTrail	399
Accessing and auditing CloudTrail logs	400
Protecting and retaining CloudTrail logs	401
Accessing Amazon EC2 logs	401
Retaining Amazon EC2 logs	401
Log management — Amazon EC2	401
Log management — Amazon VPC Flow Logs	402
Tracking changes	404
Viewing your change records	405
Default queries	405
Modifying the datetime filter in queries	413
Change record permissions	414
AWS Systems Manager in Accelerate	416
Available AMS Accelerate SSM documents	416
AMS Accelerate SSM document versions	417
Systems Manager pricing	417
Document history	418
AWS Glossary	445

What is AMS Accelerate?

Welcome to AMS Accelerate for Amazon Web Services (AWS). AMS Accelerate provides a range of operational services to help you achieve operational excellence on AWS. Whether you're just getting started in the cloud, looking to augment your current team, or need a long-term operational solution, Accelerate can help you meet your operational goals in the cloud. Leveraging AWS services and a library of automations, configurations, and run books, we provide an end-to-end operational solution for both new and existing AWS environments.

The Accelerate service leverages a suite of native AWS services and features to provide a comprehensive set of infrastructure management capabilities. Within these AWS services Accelerate creates and maintains curated sets of monitoring controls, detection guardrails, automations, and runbooks to operate infrastructure in a compliant and secure way.

Topics

- [AMS operations plans](#)
- [Using the AMS Accelerate operations plan](#)
- [AMS key terms](#)
- [Service description](#)
- [Capabilities for unsupported operating systems in Accelerate](#)
- [Contact and escalation](#)
- [Resource Inventory](#)

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.

Using the AMS Accelerate operations plan

AMS Accelerate is the AMS operations plan that can operate AWS infrastructure supporting workloads. Whether your workloads are already in an AWS account or you're planning to migrate new ones, 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 can continue using your existing access and change mechanisms while AMS consistently applies proven practices that help scale your team, optimize costs, increase security and efficiency, and improve resiliency.

While AMS Accelerate can simplify your operations, you remain responsible for application development, deployment, test and tuning, and management. AMS Accelerate only makes changes in your account as a result of incidents, alarms, remediation, and some service requests. AMS Accelerate doesn't provision resources in the account on your behalf. AMS Accelerate provides troubleshooting assistance for infrastructure issues that impact applications, but AMS Accelerate doesn't access or validate your application configurations without your knowledge and approval. AMS Accelerate services and changes are provided directly in the AWS console and APIs, so you continue to leverage your existing accounts with AWS and available AWS marketplace solutions. AMS Accelerate doesn't modify code in your infrastructure-as-code templates (for example, AWS CloudFormation templates), but can guide your teams on which changes are required to follow best operational and security practices.

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

AMS Accelerate is a service for managing operations of your AWS infrastructure.

AMS Accelerate features

AMS Accelerate offers the following features:

- **Incident management:**

Incident management is the process the AMS service uses to respond to your reported incidents.

AMS Accelerate proactively detects and responds to incidents and assists your team in resolving issues. You can reach out to AMS Accelerate operations engineers 24x7 using AWS Support Center, with response time SLAs depending on the level of response you selected for your account.

- **Monitoring:**

Monitoring is the process the AMS service uses to track your resources.

Accounts enrolled in AMS Accelerate are configured with a baseline deployment of CloudWatch events and alarms that have been optimized to reduce noise and to identify a possible upcoming incident. After receiving the alerts, the AMS team uses automated remediations, people, and processes, to bring the resources back to a healthy state and engage with your teams when appropriate to provide insights into learnings on the behavior and how to prevent it. If remediation fails, AMS starts the incident management process. You can change the baselines by updating the default configuration file.

- **Security:**

Security management is the process the AMS service uses to protect your resources. AWS Managed Services protects your information assets and helps keep your AWS infrastructure secure by using multiple controls, including AWS Config Rules and Amazon GuardDuty.

AMS Accelerate maintains a library of AWS Config Rules and remediation actions to ensure that all your accounts comply with industry standards for security and operational integrity. AWS Config Rules continuously tracks the configuration change among your recorded resources. If a change violates any rule conditions, AMS reports its findings, and allows you to remediate violations automatically or by request, according to the severity of the violation. AWS Config Rules facilitate compliance with standards set by: the Center for Internet Security (CIS), the National Institute of Standards and Technology (NIST) Cloud Security Framework (CSF), the Health Insurance Portability and Accountability Act (HIPAA), and the Payment Card Industry (PCI) Data Security Standard (DSS).

In addition, AMS Accelerate leverages Amazon GuardDuty to identify potentially unauthorized or malicious activity in your AWS environment. GuardDuty findings are monitored 24x7 by AMS. AMS collaborates with you to understand the impact of the findings and remediations based on best practice recommendations. AMS also supports Amazon Macie to protect your sensitive data such as personal health information (PHI), personally identifiable information (PII), and financial data.

- **Patch management:**

Patch management is the process the AMS service uses to update your resources.

For an AWS account with the patch add-on, AWS Managed Services applies and installs vendor updates to Amazon EC2 instances for supported operating systems during your chosen maintenance windows. AMS creates a snapshot of the instance prior to patching, monitors the patch installation, and notifies you of the outcome. If the patch fails, then AMS investigates

the failure and recommends a course of action for you to remediate the issue. Or, AMS restores the instance to rollback, if requested. AMS provides reports of patch compliance coverage and advises you of the recommended course of action for your business.

- **Backup management:**

AMS uses backup management to take snapshots of your resources.

AWS Managed Services creates, monitors, and stores snapshots for AWS services supported by AWS Backup. You define the backup schedules, frequency, and retention period by creating AWS Backup plans while onboarding accounts and applications. You associate the plans to resources. AMS tracks all backup jobs, and, when a backup job fails, alerts our team to run a remediation. AMS leverages your snapshots to perform restoration actions during incidents, if needed. AMS provides you with a backup coverage report and a backup status report.

- **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.

- **Designated experts:**

AMS Accelerate also designates a Cloud Service Delivery Manager (CSDM) and a Cloud Architect (CA) to partner with your organization and drive operational and security excellence. Your CSDM and CA provide you guidance during and after configuration and onboarding AMS Accelerate, deliver a monthly report of your operational metrics, and work with you to identify potential cost savings using tools such as AWS Cost Explorer, Cost and Usage Reports, and Trusted Advisor.

- **Operations tools:**

AMS Accelerate can provide ongoing operations for your workload's infrastructure in AWS. Our patch, backup, monitoring, and incident management services depend on having resources tagged, and the AWS Systems Manager (SSM) and CloudWatch agents installed and configured on your Amazon EC2 instances with an IAM instance profile that authorizes them to interact with the SSM and Amazon CloudWatch services. AMS Accelerate provides tools like Resource Tagger to help you tag your resources based on rules, and automated instance configuration to install the required agents in your Amazon EC2 instances. If you're following immutable infrastructure

practices, you can complete the prerequisites directly in the console or infrastructure-as-code templates.

- **Cost optimization:**

AMS Resource Scheduler automates the starting and stopping of Amazon Elastic Compute Cloud (Amazon EC2) instances, Amazon Relational Database Service (Amazon RDS) instances and Amazon EC2 Auto Scaling groups. AMS Resource Scheduler helps you reduce operational costs by stopping the resources that are not in use and starting them back when their capacity is needed.

- **Logging and Reporting:**

AWS Managed Services aggregates and stores logs generated as a result of operations in CloudWatch, CloudTrail, and Amazon VPC Flow Logs. Logging from AMS helps in faster incident resolution and system audits. AMS Accelerate also 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) designated 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.

All AMS Accelerate customers start with incident management, monitoring, security monitoring, log recording, prerequisite tools, backup management, and reporting capabilities. You can add the AMS Patch management add-on at an additional price.

Supported configurations

AMS Accelerate supports the following configurations:

- Language: English.
- Regions: See the AWS Regions supported by AWS Managed Services in the [AWS Regional Services](#) webpage.

Note

AWS Regions introduced before March 20, 2019 are considered "Original" Regions and are enabled by default. Regions introduced after this date are "Opt-in" Regions and are disabled by default. If your account uses multiple Regions and you onboard AMS Accelerate to an account with an enabled "Opt-in" Region as the default Region, the AMS Reporting feature is only available in that Region. If you do not set a default Region, the last Region you visited is your default Region.

To enable a Region, see [Enabling a Region](#). To set a default Region, see [Choosing a Region](#). For a list of the Opt-in status for each Region, see [Available Regions](#) in the *Amazon Elastic Compute Cloud User Guide*.

- Operating system architecture (x86-64 or ARM64): any supported by both [Systems Manager](#) and [CloudWatch](#).
- 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.
 - Ubuntu 20.04, 22.04
 - 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 OS follow the [security controls](#) as specified by AMS in the Accelerate 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 OS, 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
- If you use AWS Control Tower to manage your multi-account environment, then make sure that you're running the latest version of AWS Control Tower for compatibility with Accelerate. Environments that use AWS Control Tower versions earlier than 2.7 (released in April 2021), aren't supported. For information on how to update AWS Control Tower, see [Update Your Landing Zone](#).

Supported services

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

- Incidents: All AWS services
- Service request: All AWS services
- Patching: Amazon EC2
- Backups and Restoration: All AWS services supported by AWS Backup. For a list of services supported by AWS Backup, see [AWS Backup supported resources](#).
- Resource Scheduler: Amazon Elastic Compute Cloud (Amazon EC2) instances, Amazon Relational Database Service (Amazon RDS) and Amazon EC2 AutoScaling groups (ASGs)

- Services monitored by CloudWatch alarms: Application Load Balancer, Aurora, Amazon EC2, Elastic Load Balancing, Amazon FSx for NetApp ONTAP, Amazon FSx for Windows File Server, NAT gateway (a Network Address Translation (NAT) service), OpenSearch, Personal Health Dashboard, RedShift, Amazon Relational Database Service (Amazon RDS), Site-to-Site VPN. To learn more about what AMS Accelerate is monitoring as part of a service, see [Alerts from baseline monitoring in AMS](#).
- Services monitored by security Config Rules: AWS Account, GuardDuty, Macie, Amazon API Gateway, AWS Certificate Manager, AWS Config, CloudTrail, CloudWatch, AWS CodeBuild, AWS Database Migration Service, Amazon DynamoDB, Amazon EC2, Amazon ElastiCache, Amazon Elastic Block Store (Amazon EBS), Amazon Elastic File System (Amazon EFS), Amazon Elastic Kubernetes Service (Amazon EKS), Elastic Load Balancing, Amazon OpenSearch Service, Amazon EMR, AWS Identity and Access Management (IAM), AWS Key Management Service, AWS Lambda, Amazon Redshift, Amazon Relational Database Service, Amazon S3, Amazon SageMaker, AWS Secrets Manager, Amazon Simple Notification Service, AWS Systems Manager, Amazon VPC (Security group, volume, Elastic IP address, VPN connection, Internet gateways), Amazon VPC Flow Logs. For more details, see [Configuration compliance in Accelerate](#) and [Data protection in AMS Accelerate](#). You can find additional AMS security information in our private Security Guide that can be accessed through AWS Artifact, on the **Reports** tab, for AWS Managed Services.

Roles and responsibilities

The AMS Accelerate responsible, accountable, consulted, and informed, or RACI, matrix assigns primary responsibility either to the customer or AMS for a variety of activities. The table describes your (the "Customer") responsibilities versus our ("AMS Accelerate") responsibilities.

The [Scope of changes performed by AMS Accelerate](#) section lists the specific circumstances when AMS is authorized to make changes to your account; and some types of changes that AMS never makes.

AMS Accelerate RACI Matrix

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

- **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 who is informed on progress, often only on completion of the task.

Activity	Customer	Accelerate
AMS patterns		
Create new patterns	I	R
Deploy and customize patterns	R	C, I
Test and remove patterns	R	I
Application lifecycle		
Application development	R	I
Application infrastructure requirements, analysis, and design	R	I
Application deployment	R	I
AWS resource deployment	R	I
Application monitoring	R	I
Application testing/optimization	R	I
Troubleshoot and resolve application issues	R	I
Troubleshoot and resolve problems	R	I
Monitoring supported for AWS infrastructure	C	R
Incident response for AWS network issues	C	R
Incident response for AWS resource issues	C	R
Managed Account onboarding		

Activity	Customer	Accelerate
Grant access to the AWS Managed Account for the AMS team and tools	R	C
Implement changes in the account or environment to allow the deployment of tools in the account. For example, changes in Service Control Policies (SCPs)	R	C
Install SSM agents in EC2 instances	R	C
Install and configure tooling required to provide AMS services. For example, CloudWatch agents, scripts for patching, alarms, logs, and others	I	R
Manage access and identity lifecycle for AMS engineers	I	R
Collect all required inputs to configure AMS services. For example, patch maintenance windows duration, schedule and targets	R	I
Request the configuration of AMS services and provide all required inputs	R	I
Configure AMS services as requested by the customer. For example, patch maintenance windows, resource tagger, and alarm manager	C	R
Manage the lifecycle of users and their permissions, for local directory services, used to access AWS accounts and instances	R	I
Recommend reserved instances optimization	I	R
Patch management		
Collect all required inputs to configure patch maintenance windows, patch baselines, and target	R	I
Request the configuration of patch maintenance windows and baselines, and provide all required inputs	R	I

Activity	Customer	Accelerate
Configure patch maintenance windows, patch baselines, and targets as requested by the customer	C	R
Monitor for applicable updates to supported OS and software preinstalled with supported OS for EC2 instances	I	R
Report for missing updates to supported OS and maintenance window coverage	I	R
Take snapshots of instances before applying updates	I	R
Apply updates to EC2 instances per customer configuration	I	R
Investigate failed updates to EC2 instances	C	R
Update AMIs and stacks for Auto-Scaling groups (ASGs)	R	C
Patch the Windows operating system, and Microsoft packages installed on the operating system which are governed by Windows Update	I	R
Patch installed applications, software, or application dependencies not managed by Windows Update	R	I
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
Patch installed applications, software, or application dependencies not managed by the Linux operating system's native package manager	R	I
Backup		
Collect all required inputs to configure backup plans and target resources	R	I
Request the configuration of Backup plans and provide all required inputs	R	I

Activity	Customer	Accelerate
Configure backup plans and targets as requested by the customer	C	R
Specify backup schedules and target resources	R	I
Perform backups per plan	I	R
Investigate failed backup jobs	I	R
Report for backup jobs status and backup coverage	I	R
Validate backups	R	I
Request backup restoration for resources of supported AWS services resources as part of incident management	R	I
Perform backup restoration activities for resources of supported AWS services	I	R
Restore affected custom or third-party applications	R	I
Networking		
Provisioning and configuration of Managed Account VPCs, IGWs, Direct connect, and other AWS networking Services	R	I
Configure and operate AWS Security Groups/NAT/NACL inside the Managed account	R	I
Networking configuration and implementation within customer network (for example DirectConnect)	R	I
Networking configuration and implementation within AWS network	R	I
Monitor defined by AMS for network security, including security groups	I	R
Network-level logging configuration and management (VPC flow logs and others)	I	R

Activity	Customer	Accelerate
Logging		
Record all application change logs	R	I
Record AWS infrastructure change logs	I	R
Enable and aggregate AWS audit trail	I	R
Aggregate logs from AWS resources	I	R
Monitoring and Remediation		
Collect all required inputs to configure alarm manager, resource tagger, and alarm thresholds	R	I
Request the configuration of alarm manager and provide all required inputs	R	I
Configure alarm manager, resource tagger, and alarm thresholds as requested by the customer.	C	R
Deploy AMS CloudWatch baseline metrics and alarms per customer configuration	I	R
Monitor supported AWS resources using baseline CloudWatch metrics and alarms	I	R
Investigate alerts from AWS resources	C	R
Remediate alerts based on defined configuration, or create an incident	I	R
Define, monitor, and investigate customer-specific monitors	R	I
Investigate alerts from application monitoring	R	C
Security Architecture		

Activity	Customer	Accelerate
Review AMS resources and code for security issues and potential threats	I	R
Implement security controls in AMS resources and code to mitigate security risks	I	R
Enable supported AWS services for security management of the account and its AWS resources	I	R
Manage privileged credentials for account and OS access for AMS engineers	I	R
Security Risk Management		
Monitor supported AWS services for security management, like GuardDuty and Macie	I	R
Define and create AMS-defined Config Rules to detect if AWS resources comply with Center for Internet Security (CIS) and NIST security best practices.	I	R
Monitor AMS-defined Config Rules	I	R
Report conformance status of Config Rules	I	R
Define a list of required Config Rules and remediate them	I	R
Evaluate the impact of remediating AMS-defined Config Rules	R	I
Request remediation of AMS-defined Config Rules in the AWS account	R	I
Track resources exempted from AMS-defined Config Rules	R	I
Remediate supported AMS-defined Config Rules in the AWS account	C	R
Remediate non-supported AMS-defined Config Rules in the AWS account	R	I

Activity	Customer	Accelerate
Define, monitor, and investigate customer-specific Config Rules	R	I
Security monitoring and response		
Configure supported security management AWS services for alerting, alerts correlation, noise reduction, and additional rules	I	R
Monitor supported AWS services for security alerts	I	R
Install, update, and maintain endpoint security tools	R	I
Monitor for malware on instances using endpoint security	R	I
Incident Management		
Notify about incidents detected by AMS in AWS resources	I	R
Notify about incidents in AWS resources	R	I
Notify about incidents for AWS resources based on monitoring	I	R
Handle application performance issues and outages	R	I
Categorize incident priority	I	R
Provide incident response	I	R
Provide incident resolution or infrastructure restore for resources with available backups	C	R
Problem Management		
Correlate incidents to identify problems	I	R
Perform root cause analysis (RCA) for problems	I	R
Remediate problems	I	R
Identify and remediate application problems	R	I

Activity	Customer	Accelerate
Service Management		
Request information using service requests	R	I
Reply to service requests	I	R
Provide cost-optimization recommendations	I	R
Prepare and deliver monthly service report	I	R
Change Management		
Change management processes and tooling for provisioning and updating resources in the managed environment	R	I
Maintenance of application change calendar	R	I
Notice of upcoming maintenance Window	R	I
Record changes made by AMS Operations	I	R
Cost Optimization		
Collect all required inputs to configure Resource Scheduler	R	I
Request the onboarding, configuration of Resource Scheduler and provide all required inputs	R	I
Deploy Resource Scheduler per customer configuration	C, I	R
Disable and enable the Resource Scheduler on customer account	R	C
Create, delete, describe, and update schedules	C	R
Create, delete, describe, and update periods	C	R
Investigate and troubleshoot issues with Resource Scheduler	I	R
Request for offboarding the Resource Scheduler	R	I

Activity	Customer	Accelerate
Offboard the Resource Scheduler from account	C, I	R

Scope of changes performed by AMS Accelerate

AMS Accelerate only makes changes for the specific purposes and situations described next. AMS makes changes only at the infrastructure level, using the console or APIs. AMS never changes your application, control, or domain layers. You can see any changes made by AMS (or other users) using our set of pre-built queries; to do this, see [Tracking changes in your AMS Accelerate accounts](#).

AWS resources

AMS Accelerate deploys or updates AWS resources only in the following situations:

- To deploy and update tools and resources required by AMS.
- As part of AMS monitoring, in response to events and alarms.
- To remediate security issues as part of [Responses to violations in Accelerate](#) (making noncompliant resources conform to security best practices).
- During remediation and restoration as part of an incident response.
- When responding to customer requests to configure AMS features, such as the following:
 - Alarm manager
 - Resource tagger
 - Patch baselines and maintenance windows
 - Resource scheduler
 - Backup plans

AMS Accelerate does not deploy or update resources outside of these situations. If you need help from AMS to make changes in other situations, consider using [Operations on Demand](#).

Operating system software

AMS Accelerate can make changes to your operating system software during unavailability situations via incident resolution as defined in our [Service Level Agreement](#). AMS can also make changes to your operating systems as part of [Automated instance configuration in AMS Accelerate](#).

Application code and configuration

AMS Accelerate never modifies your code (for example, AWS CloudFormation templates, other infrastructure-as-code templates, or Lambda functions), but can guide your teams on which changes are required to follow best operational and security practices. AMS Accelerate provides troubleshooting assistance for infrastructure issues that impact applications, but AMS Accelerate doesn't access or validate your application configurations.

Capabilities for unsupported operating systems in Accelerate

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.
Monitoring	<p>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.</p> <p>Instance status checks include: incorrect networking or startup configuration, exhausted memory, corrupted file system, and incompatible kernel.</p>
Security management	AMS monitors and responds to Amazon EC2 GuardDuty findings and AWS Config rules .

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

Contact and escalation

You have a designated cloud service delivery manager (CSDM) who provides advisory assistance across AMS Accelerate, 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 practice 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 Accelerate.
- Respond to and resolve billing and service reporting requests.
- Provide insights for financial and capacity optimization recommendations.

Contact hours

You can contact AMS Accelerate for different reasons at different times.

Feature	AMS Accelerate
	Premium Tier
Service request	24/7
Incident management (P2-P3)	24/7
Backup and recovery	24/7

Feature	AMS Accelerate
	Premium Tier
Patch management	24/7
Monitoring and alerting	24/7
Cloud service delivery manager (CSDM)	Monday to Friday: 08:00–17:00, local business hours

Business hours

AMS contact hours

Feature	AMS Accelerate		AMS Advanced	
	Plus Tier	Premium Tier	Plus Tier	Premium Tier
Service request	Monday to Friday from 8am- 6pm, local business hours	24 x 7	Monday to Friday from 8am- 6pm, local business hours	24 x 7
Incident management (P1)	24 x 7			
Incident management (P2-P3)	Monday to Friday from 8am- 6pm, local business hours	24 x 7	Monday to Friday from 8am- 6pm, local business hours	24 x 7
Backup and recovery	24 x 7			
Patch management	24 x 7			

Feature	AMS Accelerate		AMS Advanced	
	Plus Tier	Premium Tier	Plus Tier	Premium Tier
Monitoring and alerting	24 x 7			
Cloud service delivery manager (CSDM)	Monday to Friday, 9am to 5pm, local business hours			

Escalation path

AMS supports customers with Incident Management and Service Request 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 case. For details, see [Submitting an incident](#). For general information about AMS incident management, see [Incident management](#).

To ask for information or advice, or to request additional services 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](#).

Resource Inventory

All the resources that AMS Accelerate deploys to your AWS account(s) are listed in the [resource_inventory.xlsx](#) spreadsheet.

Note: In the *Resource Name* column, the prefix *CFN:* indicates a CloudFormation logical ID instead of a resource name. These are shown for unnamed resources, for example, for S3 bucket policies.

AMS deploys a set of services as described in the [Service description](#). The cost of deploying them is low when deployed to an empty account, but the cost increases as utilization grows. For example, logs are created and config rules are invoked as resources change. When multiple changes are made to the config rules, multiple config compliance invocation can be triggered, leading to higher costs. The same possibility applies for Amazon CloudWatch used for monitoring instances—the

more granular your monitoring, the higher the cost of the service. AWS Backup is another example, if you have multiple backups stored, or if you have higher retention periods, you are using more storage and the cost is higher. These numbers are hard to predict. During your monthly business review with your cloud service delivery manager (CSDM), keep track of the changes and work to identify areas of opportunity for cost reduction.

Getting Started with AMS Accelerate

If you do not have AWS Managed Services (AMS) operating an account already, start by contacting an Amazon Web Services (AWS) sales representative using our [AWS Managed Services - Contact Sales](#) page.

After you sign up for an AMS, the AMS Accelerate team guides you through the following onboarding process for each one of your AWS accounts.

Review the feature set here: [AWS Managed Services Features](#)

Note

AMS Accelerate supports GovCloud Regions. If your service will reside in an AWS GovCloud(US) Region, see also [Getting Started with AWS GovCloud \(US\)](#).

Account onboarding process

Onboarding an account into AMS Accelerate has four stages. Each stage requires inputs from you.

1. [Step 1. Account discovery in Accelerate](#) uses tools to discover what your needs are. You are assisted by AMS Accelerate personnel.
2. [Step 2. Onboarding management resources in Accelerate](#) asks you to accept the terms and conditions; and create an onboarding role for AMS Accelerate cloud architects (CAs), who will assist you with setting a security baseline, and resolving issues as needed.
3. [Step 3. Onboarding AMS features with default policies](#) for Accelerate features, such as monitoring, patching, and backup.
4. [Step 4. Customize features in Accelerate](#) ensures that resources, including EC2 instances, are correctly configured for your application.

Accelerate onboarding prerequisites

Before you start the onboarding process, it is important to understand the technical dependencies that Accelerate components rely on.

Note

To use AMS Accelerate, you must be on one of the two supported AWS Support plans: Enterprise On-Ramp or Enterprise. The Developer and Business plans are not eligible for qualifying for AMS Accelerate. To learn more about the different plans, see [Compare AWS Support Plans](#).

AMS Accelerate VPC endpoints

A VPC endpoint enables private connections between your VPC and supported AWS services and VPC endpoint services powered by AWS. If you need to filter outbound internet connectivity, configure the following VPC service endpoints to ensure that AMS Accelerate has connectivity with its service dependencies.

Note

In the following list, *region* represents the identifier for an AWS Region, for example us-east-2 for the US East (Ohio) Region.

```
com.amazonaws.region.logs
com.amazonaws.region.monitoring
com.amazonaws.region.ec2
com.amazonaws.region.ec2messages
com.amazonaws.region.ssm
com.amazonaws.region.ssmmessages
com.amazonaws.region.s3
com.amazonaws.region.events
```

For information about how to configure AWS VPC endpoints, see [VPC endpoints](#).

Note

If you are creating VPC endpoints in your account for all of the above mentioned services, then see this [sample AWS CloudFormation template](#). You can update this template and remove or add VPC endpoints definition as per your use-case.

Outbound internet connectivity in Accelerate

1. Download [egressMgmt.zip](#).
2. Open the **ams-egress.json** file.
3. Find the URLs under the JSON properties:
 - WindowsPatching
 - RedHatPatching
 - AmazonLinuxPatching
 - EPELRepository
4. Allow access to these URLs.

Testing outbound connectivity in Accelerate

Test outbound connectivity using one of the following methods.

Note

Before running the script/command, replace the red *region* with your region identifier, for example, `us-east-1`.

Windows PowerShell script

```
$region = 'region'
@('logs', 'monitoring', 'ec2', 'ec2messages', 'ssm', 'ssmmessages', 's3', 'events') | `
ForEach-Object { `
Test-NetConnection ("$_" + '.' + "$region" + '.amazonaws.com') -Port 443 } | `
Format-Table ComputerName,RemotePort,RemoteAddress,PingSucceeded,TcpTestSucceeded -
AutoSize
```

Linux command

```
for endpoint in logs monitoring ec2 ec2messages ssm ssmmessages s3 events; do nc -zv
$endpoint.region.amazonaws.com 443; done
```



Amazon EC2 Systems Manager in Accelerate

You must install the AWS Systems Manager Agent (SSM Agent) on all of the EC2 instances you want AMS to manage. You also need to add the [bucket permissions](#) that SSM Agent requires. For an overview that includes EC2, see [Step 3. Onboarding AMS features with default policies](#).

Step 1. Account discovery in Accelerate

Account discovery is the stage when AMS Accelerate works with you to assess the current state of your account, evaluate that our service is a good fit for your account, and identify any major technical blockers for supporting your environment. During the Account Discovery stage, AMS Accelerate does not provide any operational services.

To assist with the analysis and discovery of an account, we ask you to run a pre-built script in the **AwsAccountDiscoveryCli** command line interface that generates a comprehensive picture of your account and resources, focusing on the services you are using. You can control which part of the report to share with AMS Accelerate and then AMS Accelerate starts an iterative process to remove technical blockers, if any, before moving to the Account-Level onboarding stage.



Account discovery with the **AwsAccountDiscoveryCli** in Accelerate

AwsAccountDiscoveryCli is a command line interface to discover AWS resources in a given account. You can perform discovery from a variety of platforms, including AWS CloudShell, Linux, and Windows.

⚠ Important

The **AwsAccountDiscoveryCli** performs read-only calls and does not transmit data to AMS Accelerate during collection. Data is stored locally on the machine that runs the commands. AMS recommends that you review the collected data with your security team to determine whether or not you can share it with AMS for further analysis. Then, work with your AMS account team to determine the process for sharing your approved data with AMS. For the latest changes, and to know if you need to update, see the [Account Discovery Changelog zip file](#).

Account discovery with AWS CloudShell (Recommended method)

AWS CloudShell is a browser-based shell that makes it easy to securely manage, explore, and interact with your AWS resources. AWS CloudShell is pre-authenticated with your console credentials when you log in. Common development and operations tools are pre-installed, so no local installation or configuration is required. With AWS CloudShell, you can quickly run scripts with the AWS Command Line Interface (AWS CLI), experiment with AWS service APIs using the AWS SDKs, or use a range of other tools to be productive. You can use AWS CloudShell right from your browser at no additional cost.

📘 Note

You can use the CloudShell AWS console from any other, or the closest, AWS Region where it is available, to perform resource discovery for all other AWS Regions. For example, to perform discovery in the Singapore region, open a CloudShell in the "US West(Oregon) us-west-2" AWS Region in the AWS Console and follow the instructions given next.

To use **AwsAccountDiscoveryCli** with AWS CloudShell:

1. (Prerequisites) Before you can use **AwsAccountDiscoveryCli**, you need the following AMS policies:

(Recommended) Access to the AWS CloudShell with read-only permissions (for more information, see [Managing AWS CloudShell access and usage with IAM policies](#)):


- **arn:aws:iam::aws:policy/ReadOnlyAccess**
- **arn:aws:iam::aws:policy/AWSCodeArtifactReadOnlyAccess**

- **arn:aws:iam::aws:policy/AWSCloudShellFullAccess**

Create and attach the following policy to the IAM entity (user, role) you are using to initiate the discovery:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "support:DescribeSeverityLevels",
      "Resource": "*"
    }
  ]
}
```

2. From any page or AWS Region in the AWS Management Console, open the AWS CloudShell to run the account discovery script shown next. Make sure that you're logged into the AWS Management Console with the correct level of permissions outlined in step 1.

 **Note**

Do not change the `--domain-owner 354220221581` and `--region us-west-2` sections from the following script. Copy the script as-is.

```
python3 -m venv awsdiscovery
source ~/awsdiscovery/bin/activate
pip install pip --upgrade
aws codeartifact login --tool pip \
  --repository AwsAccountDiscoveryCli \
  --domain aws-account-discovery-cli \
  --domain-owner 354220221581 \
  --region us-west-2
pip install awsaccountdiscoverycli
```

3. Verify that the installation completed successfully:

```
awsdiscover --version
```

4. Start the collection for the current account:

```
awsdiscover
```

5. Discovery takes more time on large accounts. Once finished, compress the output folder to download the report:

```
tar -czvf DiscoveryReports.tar.gz /home/cloudshell-user/AwsAccountDiscoveryReports/
```

6. Select **Actions** in the top right corner, then choose **Download file**.
7. For the **Individual file path**, specify the following path and then choose **Download**.

/home/cloudshell-user/DiscoveryReports.tar.gz

Important

If the account you want to discover is part of an AWS Organization, **AwsAccountDiscoveryCli** must be called from the Organization's management account or by a member account that is a delegated administrator for an AWS service to collect organization-level information. Otherwise, the organization-level information isn't collected. To learn more about these concepts, see [AWS Organizations terminology and concepts](#)

Other methods of account discovery

Prerequisites if not using the AWS CloudShell method

- The latest AWS CLI configured with read-only permissions, see [Configuring the AWS CLI](#). AMS recommends the following AWS Managed policies:
 - **arn:aws:iam::aws:policy/ReadOnlyAccess**
 - **arn:aws:iam::aws:policy/AWSCodeArtifactReadOnlyAccess**

Create and attach the following policy to the IAM entity (user, role) you are using to initiate the discovery:

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

```
    "Effect": "Allow",
    "Action": "support:DescribeSeverityLevels",
    "Resource": "*"
  }
]
```

- Python 3.6 and later (only required if not using AWS CloudShell. We highly recommend using CloudShell for discovery):
 - Linux: [Python downloads](#)
 - Mac OS X: [Python Releases for Mac OS X](#)
 - Windows: [Python Releases for Windows](#)

Account discovery with Linux in Accelerate

Note

We highly recommend using Discovery of an AWS account from AWS CloudShell, as previously described. Use this method only if AWS CloudShell cannot be used due to IAM permission issues.

To discover an AWS account using Linux, follow these steps:

1. To authenticate to the AMS CodeArtifact repository and install **AwsAccountDiscoveryCli** on Linux, use the following script:

```
python3 -m venv awsdiscovery
source ~/awsdiscovery/bin/activate

pip install pip --upgrade

aws codeartifact login --tool pip --repository AwsAccountDiscoveryCli \
  --domain aws-account-discovery-cli --domain-owner 354220221581 --region us-west-2

pip install awsaccountdiscoverycli
```

2. Verify that the installation completed successfully (by checking that a version exists):

```
awsdiscover --version
```

3. Start the collection for the current account:

```
awsdiscover
```

4. Discovery takes more time on large accounts. After it's finished, you'll see the output location printed in the tool's output on screen.

Account discovery with Windows in Accelerate

Note

We highly recommend using Discovery of an AWS account from AWS CloudShell, as previously described. Use this method only if AWS CloudShell cannot be used due to IAM permission issues.

To discover an AWS account using Windows, follow these steps:

1. To authenticate to the AMS CodeArtifact repository and install **AwsAccountDiscoveryCli** on Windows, use the following script:

```
py -m venv env
.\env\Scripts\activate

pip install pip --upgrade

$CODEARTIFACT_TOKEN = Get-CAAuthorizationToken -Domain aws-account-discovery-cli `
    -DomainOwner 354220221581 `
    -Region us-west-2
pip config set global.index-url https://aws:
$(($CODEARTIFACT_TOKEN.AuthorizationToken)@aws-account-discovery-
cli-354220221581.d.codeartifact.us-west-2.amazonaws.com/pypi/AwsAccountDiscoveryCli/
simple/

pip install awsaccountdiscoverycli
```

2. Verify that the installation completed successfully:

```
awsdiscover --version
```

3. Start the collection for the current account:

```
awsdiscover
```

4. Discovery takes more time on large accounts. Once finished, you'll see the output location.

Account discovery using a Hub account

This process involves creating a read-only IAM role in each account you want to discover (child accounts) from a central (hub) account.

Hub account - IAM configuration

Attach the following policy to the IAM entity (user, role) you are using to initiate the discovery:

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "sts:AssumeRole"
    ],
    "Resource": "arn:aws:iam::*:role/AwsAccountDiscoveryRole"
  }]
}
```

Child account - IAM configuration

1. Create a file: `*trust_policy.json*` (replace HUB_ACCOUNT_ID with the AWS account ID of your hub account):

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Principal": {
      "AWS": "arn:aws:iam::HUB_ACCOUNT_ID:root"
    },
  },
```

```
"Action": "sts:AssumeRole"
}]
}
```

2. Configure your AWS CLI to use the child account. You'll need the following permissions:

```
iam:CreateRole
iam:AttachRolePolicy
```

3. Create a role called the **AwsAccountDiscoveryRole** in your child account with a trust to the hub account:

```
aws iam create-role --role-name AwsAccountDiscoveryRole \
  --assume-role-policy-document file://trust_policy.json
```

4. Attach the **ReadOnlyAccess** policy to the role:

```
aws iam attach-role-policy --role-name AwsAccountDiscoveryRole \
  --policy-arn arn:aws:iam::aws:policy/ReadOnlyAccess
```

Run Discovery from the Hub account

Open the AWS Management Console in the Hub account and open [AWS CloudShell](#).

Configure your AWS CLI to HUB_ACCOUNT_ID

Run the following command with the desired CHILD_ACCOUNT_ID to discover

```
awsdiscover -a CHILD_ACCOUNT_ID
```

Note

After you have completed account discovery of a child account, AMS recommends deleting the **AwsAccountDiscoveryRole** role if you have no further use for it.

CLI reference

Use the **Help** menu of the tool to get the latest information about the available commands.

Step 2. Onboarding management resources in Accelerate

This is an overview of the process of onboarding management resources.

You accept terms

Your cloud services delivery manager (CSDM) guides you through the acceptance process. You need to accept the Terms and Conditions, select AWS Regions, add-ons, and a Service Level Agreement (SLA).

You grant permissions to AMS roles

You need to grant access to AMS processes and to your Cloud Architect. You do this by creating a AWS CloudFormation stack for each role. See [The template to create AMS roles](#) and then [Create `aws_managedservices_onboarding_role` with AWS CloudFormation](#). For more details see [Access management in AMS Accelerate](#).

AMS reviews your configuration

Your Cloud Architect (CA) also looks for possible configuration issues in your account, like Service Control Policies (SCPs), and security findings that might prevent AMS from deploying the tools and resources required by AMS. Your CA works with you to help you remediate findings and remove any blockers to the deployment of AMS tools and resources.

AMS reviews your AWS CloudTrail trail configurations

Your Cloud Architect (CA) will review your CloudTrail trail configurations, and confirm if you want AMS to deploy a global CloudTrail trail, or integrate Accelerate with your CloudTrail account or Organization trail resources. If you choose to have Accelerate integrate with your CloudTrail trail, your CA will guide you through required updates to the configurations for your CloudTrail trail resources.

AMS deploys management resources

The AMS team deploys tools and AWS resources to provide the different services of AMS Accelerate. After it's completed, AMS has built the AWS Managed Services account and AMS notifies you that your account is active.

This concludes the *Onboarding management resources* stage. You can proceed directly to the next step of the onboarding process: [Step 3. Onboarding AMS features with default policies](#).

Note

Now that your account is active, you have the option to perform any of these tasks:

- Create incidents and service requests for AWS infrastructure using the Support Center Console. See [Incident reports, service requests, and billing questions in AMS Accelerate](#).
- See the conformance status in your account of the AWS Config Rules deployed by AMS, [Configuration compliance in Accelerate](#).
- Locate and analyze GuardDuty and Macie (optional) findings. See [GuardDuty](#).
- Access and audit CloudTrail logs
- Track changes in your AMS Accelerate account. See [Tracking changes in your AMS Accelerate accounts](#).
- Use Resource Tagger to create tags. See [Resource Tagger](#).
- Request Patch, Backup, and AWS Config Reports. See [Reports and options](#).

Review and update your configurations to enable AMS Accelerate to use your CloudTrail trail

AMS Accelerate relies on AWS CloudTrail logging in order to manage audits and compliance for all resources in your account. During onboarding, you choose whether Accelerate deploys a CloudTrail trail in your primary AWS Region or uses events generated by your existing CloudTrail account or Organization trail. If your account does not have a trail configured, then Accelerate will deploy a managed CloudTrail trail during onboarding.

Important

CloudTrail log management configuration is only required when you choose to integrate AMS Accelerate with your CloudTrail account or Organization trail.

Review your CloudTrail trail configurations, Amazon S3 bucket policy, and AWS KMS key policy for your CloudTrail events delivery destination with your Cloud Architect (CA)

Before Accelerate can use your CloudTrail trail, you must work with your Cloud Architect (CA) to review and update your configurations to meet Accelerate requirements. If you choose to integrate Accelerate with your CloudTrail Organization trail, then your CA works with you to

update your CloudTrail events delivery destination Amazon S3 bucket and AWS KMS key policies to enable cross-account queries from your Accelerate account. Your Amazon S3 bucket can be in an account that's managed by Accelerate, or an account that you manage. During onboarding, Accelerate validates that queries can be made to your CloudTrail Organization trail events delivery destination, and pauses the onboarding if the queries fail. You work with your CA to correct these configurations so that onboarding can resume.

Review and update your CloudTrail account or Organization trail configurations

The following configurations are required to integrate Accelerate CloudTrail log management your CloudTrail account or Organization trail resources:

- Your CloudTrail trail is configured to log events from all AWS Regions.
- Your CloudTrail trail has [global service events](#) enabled.
- Your CloudTrail account or Organization trail logs all [management events](#), including [read and write events](#), and AWS AWS KMS and Amazon RDS Data API event logging is enabled.
- Your CloudTrail trail has [log file integrity validation](#) enabled.
- The Amazon S3 bucket your CloudTrail trail delivers events to encrypts events using either [SSE-S3](#) or [SSE-KMS](#) encryption.
- The Amazon S3 bucket your CloudTrail trail delivers event to has [server access logging](#) enabled.
- The Amazon S3 bucket your CloudTrail trail delivers event to has a [lifecycle configuration](#) that retains your CloudTrail trail data for at least 18 months.
- The Amazon S3 bucket your CloudTrail trail delivers event to has [Object Ownership](#) set to Bucket owner enforced.
- The Amazon S3 bucket your CloudTrail trail delivers event to is accessible by Accelerate.

Review and update the Amazon S3 bucket policy for your CloudTrail events delivery destination

During onboarding, you work with your Cloud Architect (CA) to add Amazon S3 bucket policy statements to your CloudTrail events delivery destination. To enable your users to query changes in your CloudTrail events delivery destination Amazon S3 bucket from your Accelerate account, you can deploy a uniformly named IAM role in each account in your Organization that Accelerate manages, and add it to the `aws:PrincipalArn` list in all Amazon S3 bucket policy statements. With this configuration, your users can query and analyze your account's CloudTrail Organization trail events in Accelerate using Athena. For more information about how to update an Amazon S3

bucket policy, see [Adding a bucket policy by using the Amazon S3 console](#) in the *Amazon Simple Storage Service User Guide*.

Important

Updating your Amazon S3 bucket policy is required only when Accelerate integrates with a CloudTrail trail that delivers events to a centralized S3 bucket. Accelerate doesn't support integrating with a CloudTrail trail that delivers to a centralized bucket but doesn't have the accounts under an AWS Organization.

Note

Before updating your Amazon S3 bucket policy, replace *red* fields with applicable values:

- *YOUR-S3-BUCKET-NAME* with the name of the Amazon S3 bucket that contains the trail events from your accounts.
- *YOUR-ORGANIZATION-ID* with the ID of the AWS Organization that your accounts are a member of.
- *YOUR-OPTIONAL-S3-LOG-DELIVERY-PREFIX* with your CloudTrail trail's Amazon S3 bucket delivery prefix. For example, `my-bucket-prefix`, that you might have set when [you created your CloudTrail trail](#).

If you haven't configured a Amazon S3 bucket delivery prefix for your trail, then remove "*YOUR-OPTIONAL-S3-LOG-DELIVERY-PREFIX*" and the preceding forward slash (/) from the following Amazon S3 bucket policy statements.

The following three Amazon S3 bucket policy statements grant Accelerate access to retrieve the configurations of and run AWS Athena queries to [analyze the CloudTrail events](#) in your events delivery destination Amazon S3 bucket from your Accelerate account.

```
{
  "Sid": "DONOTDELETE-AMS-ALLOWBUCKETCONFIGAUDIT",
  "Effect": "Allow",
  "Principal": {
    "AWS": "*"
  },
  "Action": [
```

```

    "s3:GetBucketLogging",
    "s3:GetBucketObjectLockConfiguration",
    "s3:GetLifecycleConfiguration",
    "s3:GetEncryptionConfiguration"
  ],
  "Resource": "arn:aws:s3:::YOUR-S3-BUCKET-NAME",
  "Condition": {
    "StringEquals": {
      "aws:PrincipalOrgID": "YOUR-ORGANIZATION-ID"
    },
    "ArnLike": {
      "aws:PrincipalArn": [
        "arn:aws:iam::*:role/ams-access-*"
      ]
    }
  }
},
{
  "Sid": "DONOTDELETE-AMS-ALLOWLISTBUCKET",
  "Effect": "Allow",
  "Principal": {
    "AWS": "*"
  },
  "Action": "s3:ListBucket",
  "Resource": "arn:aws:s3:::YOUR-S3-BUCKET-NAME",
  "Condition": {
    "ForAnyValue:StringEquals": {
      "aws:CalledVia": "athena.amazonaws.com"
    },
    "StringLike": {
      "s3:prefix": "YOUR-OPTIONAL-S3-LOG-DELIVERY-PREFIX/AWSLogs/*"
    },
    "StringEquals": {
      "aws:PrincipalOrgID": "YOUR-ORGANIZATION-ID"
    },
    "ArnLike": {
      "aws:PrincipalArn": [
        "arn:aws:iam::*:role/ams-access-*"
      ]
    }
  }
},
{
  "Sid": "DONOTDELETE-AMS-ALLOWGETOBJECT",

```

```
"Effect": "Allow",
"Principal": {
  "AWS": "*"
},
"Action": "s3:GetObject",
"Resource": "arn:aws:s3:::YOUR-S3-BUCKET-NAME/YOUR-OPTIONAL-S3-LOG-DELIVERY-PREFIX/
AWSLogs/*",
"Condition": {
  "ForAnyValue:StringEquals": {
    "aws:CalledVia": "athena.amazonaws.com"
  },
  "StringEquals": {
    "aws:PrincipalOrgID": "YOUR-ORGANIZATION-ID"
  },
  "ArnLike": {
    "aws:PrincipalArn": [
      "arn:aws:iam::*:role/ams-access-*"
    ]
  }
}
}
```

Review and update the AWS KMS key policy for your CloudTrail events delivery destination

During onboarding, you work with your Cloud Architect (CA) to update the AWS KMS key policy used to encrypt the CloudTrail trail events delivered to your Amazon S3 bucket. Make sure that you append the reference AWS KMS key policy statements to your existing AWS KMS key. This configures Accelerate to integrate with your existing CloudTrail trail event delivery destination Amazon S3 bucket and decrypt events. To enable your users to query changes in your CloudTrail events delivery destination Amazon S3 bucket from your Accelerate account, you can deploy a uniformly named IAM Role in each account in your Organization that Accelerate is managing, and add it to the "aws:PrincipalArn" list. With this configuration, your users can query events.

There are different AWS KMS key policy update scenarios to consider. You might only have a AWS KMS key configured to your CloudTrail trail to encrypt all events, and not have a AWS KMS key that encrypts objects in your Amazon S3 bucket. Or, you might have one AWS KMS key that encrypts events delivered by CloudTrail, and another AWS KMS key that encrypts all objects stored in your Amazon S3 bucket. When you have two AWS KMS keys, you update the AWS KMS key policy for each key to grant Accelerate access to your CloudTrail events. Make sure that you amend the reference AWS KMS key policy statement to your existing AWS KMS key policy before you update

the policy. For more information about how to update a AWS KMS key policy, see [Changing a key policy](#) in the *AWS Key Management Service User Guide*.

Important

You're required to update your AWS KMS key policy only when Accelerate integrates with a CloudTrail trail with log file SSE-KMS encryption enabled.

Note

Before you apply this AWS KMS key policy statement to the AWS KMS key used to encrypt your AWS CloudTrail events delivered to your Amazon S3 bucket, replace the following *red* fields with applicable values:

- *YOUR-ORGANIZATION-ID* with the ID of the AWS Organization your accounts are a member of.

This AWS KMS key policy statement grants Accelerate access to decrypt and query trail events delivered to Amazon S3 bucket from each account in your Organization with access restricted to Athena, used by Accelerate to [query and analyze CloudTrail events](#).

```
{
  "Sid": "DONOTDELETE-AMS-ALLOWTRAILOBJECTDECRYPTION",
  "Effect": "Allow",
  "Principal": {
    "AWS": "*"
  },
  "Action": [
    "kms:Decrypt",
    "kms:DescribeKey"
  ],
  "Resource": "*",
  "Condition": {
    "ForAnyValue:StringEquals": {
      "aws:CalledVia": "athena.amazonaws.com"
    },
    "StringEquals": {
      "aws:PrincipalOrgID": "YOUR-ORGANIZATION-ID"
    }
  },
}
```

```
    "ArnLike": {
      "aws:PrincipalArn": [
        "arn:aws:iam::*:role/ams-access-*"
      ]
    }
  }
}
```

The template to create AMS roles

The following AMS role grants permissions to your AMS cloud architect (CA). The provided AWS CloudFormation template simplifies creating the IAM role, permissions policy, and trust policy. For more information, consult with your CA.

Role Name	Required by	Sample Templates
aws_managedservice s_onboarding_role	AMS personnel during onboarding only	onboarding_role_minimal.zip

Note

After you select and download a sample template (one per role), you will upload these as definitions of AWS CloudFormation stacks in [Create aws_managedservices_onboarding_role with AWS CloudFormation](#).

Create aws_managedservices_onboarding_role with AWS CloudFormation

You can create the AWS Identity and Access Management role, `aws_managedservices_onboarding_role`, with AWS CloudFormation from the AWS Management Console. Or, you can use commands from AWS CloudShell to deploy the role.


Use the AWS Management Console

Note

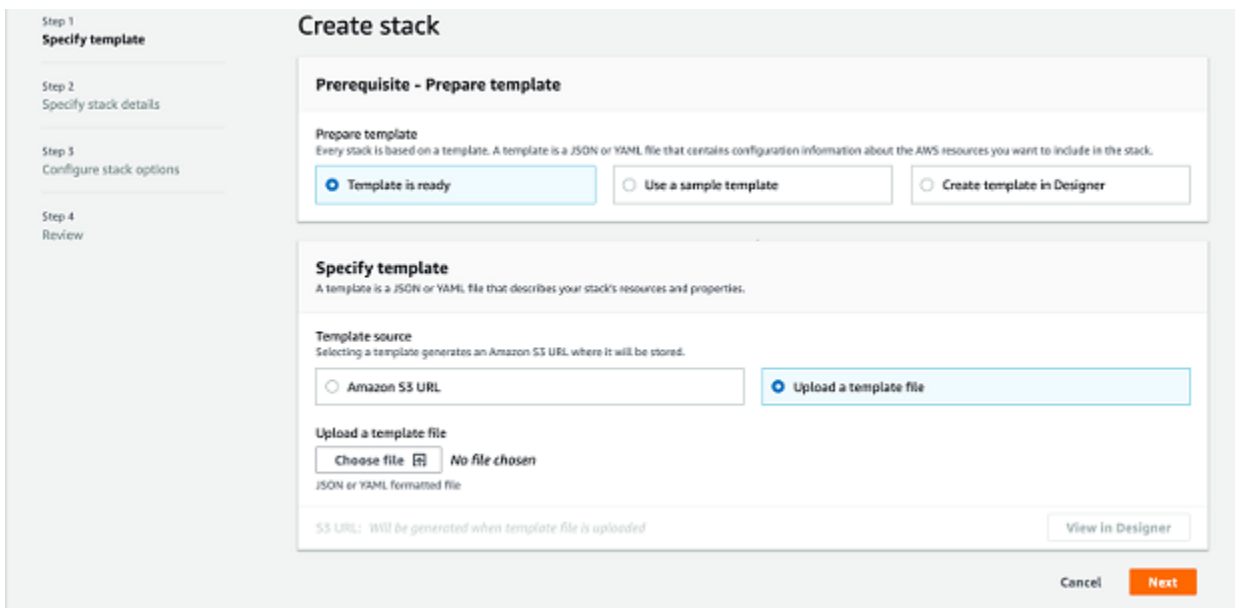
Before starting, have a JSON or YAML file for each role ready to upload. For more information, see [The template to create AMS roles](#).

To create the role from the AWS Management Console, complete the following steps:

1. Sign in to the AWS Management Console and open the AWS CloudFormation console at <https://console.aws.amazon.com/cloudformation>.



2. Choose **Create Stack > With new resources (standard)**. You see the following page.



3. Choose **Upload a template file**, upload the JSON or YAML file of the IAM role, and then choose **Next**. You see the following page.

- Enter the stack name **"ams-onboarding-role"** in the **Stack Name** field. Enter a **DateOfExpiry** using the format **"YYYY-MM-DDT00:00:00Z"** (30 days from the current date is recommended). Continue scrolling down and selecting next until you reach this page:

- Make sure the check box is selected and then select **Create Stack**.
- Make sure the stack was created successfully.

Use commands from AWS CloudShell

To deploy the `aws_managedservices_onboarding_role` IAM role, run the following command in [AWS CloudShell](#):

AWS CLI

```
curl -s "https://docs.aws.amazon.com/en_us/managedservices/latest/accelerate-guide/samples/onboarding_role_minimal.zip" -o "onboarding_role_minimal.zip"
unzip -q -o onboarding_role_minimal.zip
aws cloudformation create-stack \
  --stack-name "aws-managedservices-onboarding-role" \
  --capabilities CAPABILITY_NAMED_IAM \
  --template-body file://onboarding_role_minimal.json \
  --parameters ParameterKey=DateOfExpiry,ParameterValue="`date -d '+30 days' -u '+%Y-%m-%dT%H:%M:%SZ'`"
```

AWS Tools for PowerShell

```
Invoke-WebRequest -Uri 'https://docs.aws.amazon.com/en_us/managedservices/latest/accelerate-guide/samples/onboarding_role_minimal.zip' -OutFile 'onboarding_role_minimal.zip'
Expand-Archive -Path 'onboarding_role_minimal.zip' -DestinationPath . -Force
New-CFNStack `
  -StackName 'aws-managedservices-onboarding-role' `
  -Capability CAPABILITY_NAMED_IAM `
  -TemplateBody (Get-Content 'onboarding_role_minimal.json' -Raw) `
  -Parameter @{ParameterKey = "DateOfExpiry"; ParameterValue = (Get-Date).AddDays(30).ToString('yyyy-MM-ddTHH:mm:ssZ')}
```

After you create the role, work with your Cloud Architect (CA) to complete the [Step 2. Onboarding management resources in Accelerate](#) process. After AMS informs you that your account is active, you're ready to onboard your instances.

Step 3. Onboarding AMS features with default policies

In this stage, you onboard AMS features with default policies. These include adding Amazon EC2 instances and configuring monitoring, backup, AWS Config remediation, and patch (if applicable, AMS Patch Orchestrator is an add-on that you must specifically request) according to your preferences. You can do this yourself, or request that AMS onboards features based on your inputs. To request help from AMS, create a service request and provide all the required inputs to complete the task. Remember that service requests are not resolved immediately.

Note

While an account might use default policies for back-up, patch, or monitoring, resources need to be tagged for the appropriate policy to be effective.

Topics

- [\(Optional\) Quick Start template](#)
- [Onboarding Accelerate monitoring](#)
- [Onboarding EC2 instances to Accelerate](#)
- [Onboarding AWS Backup in Accelerate](#)
- [Onboarding patching in Accelerate](#)
- [Review non-conformance reports in Accelerate](#)

(Optional) Quick Start template

The Quick Start template automates the deployment and configuration of AMS Resource Tagger in Accelerate-enabled AWS accounts. This template saves time and effort compared to manual setup. Use this template as-is to define monitoring, patching, and backup basics in one account. Or, use it as a StackSet to apply settings across Organizations Units to standardise settings for multiple accounts.

You can also use it as a starting point to build your own custom AMS Resource Tagger profile and use snippets from the documentation to create more complex definitions for tagging.


Quick Start template functionality

The Quick Start template completes the following tasks:

- Creates and deploys a configuration version for AMS Resource Tagger.
- Applies tags to Amazon EC2 instances that enable AMS management and creation of monitoring resources.
- (Optional) Applies tags to managed EC2 instances that enable them to be patched on a desired schedule, and creates a Patching Maintenance Window to facilitate the patching.

Warning: By default, instances are rebooted and stopped instances are started automatically for patch installation.

- (Optional) Applies tags to managed EC2 instances that enable them to be backed up as defined in the [default AMS Backup Plan](#).
- (Optional) Applies tags to managed Amazon Relational Database Service (Amazon RDS) resources that enable them to be backed up according to the [default AMS Backup Plan](#). A backup plan also enables Point in time recovery (PITR) for Amazon RDS. If Amazon RDS automated backup isn't enabled, then the database restarts close to the time of the next backup window.



Quick Start template overrides and exclusions

After you use this template to create a Quick Start stack, use the following steps to exclude specific instances from Management/Monitoring, Patching, or Backup:

- **Management and Monitoring:** To exclude an EC2 instance from being managed and monitored by AMS, add this tag to the instance: `ExcludeFromAMSQuickStartMonitoring=true`.
- **Patching:** EC2 instances that are members of an Auto Scaling group, Amazon Elastic Container Service, or Amazon Elastic Kubernetes Service cluster are excluded from Patching by this Quick Start template.

To disable the creation of a Patching Window and patch-related tagging of EC2 instances, set the CloudFormation stack parameter to `EnablePatching=false`.

To exclude an EC2 instance from being targeted by the Quick Start patching window when `EnablePatching=true`, add this tag to the instance: `ExcludeFromAMSQuickStartPatching=true`.

- **Backup:** EC2 instances that are members of an Auto Scaling groups, ECS, or EKS cluster are excluded from Backup by this Quick Start template.

To exclude an EC2 instance from being targeted by the default AMS Backup Plan, when `EnableBackup=true`, add this tag to that instance: `ExcludeFromAMSQuickStartBackup=true`.

Tip

You can tag EC2 instances in bulk. Use [Tag Editor](#) to bulk-select and tag resources in one step in the AWS Management Console.

Quick Start template parameters

This Quick Start template configures Resource Tagger to add the tag `ams:rt:ams-managed=true` to all EC2 instances in the account, excluding instances that you add the `ExcludeFromAMSQuickStartMonitoring=true` tag to. Use the following parameters to control the optional parts of this stack according to your requirements:

CFN Parameter	Value	Effect
<code>EnableBackup</code>	'true' (default)	All AMS-managed EC2 instances (<code>ams:rt:ams-managed=true</code>) are tagged with <code>ams:rt:backup-</code>

CFN Parameter	Value	Effect
		<p><code>orchestrator=true</code> for backup daily at 4AM UTC as per the default AMS Backup Plan.</p> <p>For all RDS instances, the template applies the (<code>ams:rt:backup-orchestrator-data-sensitive=true</code>) backup daily at 4AM UTC as per the default AMS Backup Plan.</p>
	'false'	No instances are targeted for backup by this quick start template.
EnablePatching	'false' (default)	No instances are targeted for patching by this Quick Start template.
	'true'	All AMS-managed instances (<code>ams:rt:ams-managed=true</code>) are tagged with <code>ams:rt:Patch Group=AMSQuickStartPatchWindow</code> and a basic SSM Maintenance Window resource is created to facilitate patching according to the schedule defined in <code>CronExpression</code> .
CronExpression	-	The default value of <code>cron(0 30 19 ? * SAT#2 *)</code> sets the 2nd Saturday of the month at 7:30PM according to the <code>Timezone</code> parameter. Use cron syntax .

CFN Parameter	Value	Effect
Timezone	-	Target instances are patched according to CronExpression . Use IANA format .

Download the Quick Start template

Download the [AMSQuickStart.zip file](#).

Or, run the following command in [AWS CloudShell](#) to deploy the AMSQuickStart.yaml:

AWS Command Line Interface

```
curl -s "https://docs.aws.amazon.com/en_us/managedservices/latest/accelerate-guide/samples/AMSQuickStart.zip" -o "AMSQuickStart.zip"
unzip -q -o AMSQuickStart.zip
for region in region1 region2 ;
do
aws cloudformation create-stack \
  --region $region \
  --stack-name "AMSQuickStart" \
  --template-body file://AMSQuickStart.yaml \
  --parameters \
    ParameterKey=EnableBackup,ParameterValue="true" \
    ParameterKey=EnablePatching,ParameterValue="false" \
    ParameterKey=Timezone,ParameterValue="US/Eastern" \
    ParameterKey=CronExpression,ParameterValue="cron(0 30 19 ? * SAT#2 *)" \
;
done
```

AWS Tools for PowerShell

```
Invoke-WebRequest -Uri 'https://docs.aws.amazon.com/en_us/managedservices/latest/accelerate-guide/samples/AMSQuickStart.zip' -OutFile 'AMSQuickStart.zip'
Expand-Archive -Path 'AMSQuickStart.zip' -DestinationPath . -Force
@('region1', 'region2') | `
ForEach-Object { `
New-CFNStack `
  -Region $_ `
  -StackName 'AMSQuickStart' `
```



```
-TemplateBody (Get-Content 'AMSQuickStart.yaml' -Raw) `
-Parameter @(
  @{ParameterKey = "EnableBackup"; ParameterValue = "true"},
  @{ParameterKey = "EnablePatching"; ParameterValue = "false"},
  @{ParameterKey = "Timezone"; ParameterValue = "US/Eastern"},
  @{ParameterKey = "CronExpression"; ParameterValue = "cron(0 30 19 ? * SAT#2
*)"}
)
}
```

For a description of each parameter, see the preceding section, [Quick Start template parameters](#).

Onboarding Accelerate monitoring

Monitoring is enabled by default for all new resources except Amazon EC2 instances. You can start monitoring your Amazon EC2 instances by tagging your instances.

To onboarding monitoring, first make sure that your configuration monitors the resources you want AMS to monitor, and ignores the resources you want it to ignore.

You can use the following CloudWatch dashboards to explore how many of your resources are targeted by AMS monitoring and tagging, and how many are not. In your account, navigate to the CloudWatch dashboards console, and select one of the following:

- AMS-Alarm-Manager-Reporting-Dashboard
- AMS-Resource-Tagger-Reporting-Dashboard

For a complete description of the dashboard metrics, see:

- [Viewing the number of resources monitored by Alarm Manager in Accelerate](#)
- [Viewing the number of resources managed by Resource Tagger](#)

Onboarding resources to be monitored in Accelerate

To override the default behavior, for example, to disable default monitoring for non-EC2 resources, you need to untag those resources using a custom configuration profile. For more information about tagging for monitoring, see [Monitoring](#).

Monitoring is disabled for EC2 instances until you onboard your instances, which includes tagging your instances using a custom configuration profile. The next section describes EC2 instance onboarding.

Creating a monitoring configuration profile in Accelerate

- For information about using the default configuration, see [Alarm Manager](#).
- For information about using a custom configuration, see [Modifying the Accelerate alarm default configuration](#).

Onboarding EC2 instances to Accelerate

EC2 instances are onboarded to AMS Accelerate through a process called Automated Instance Configuration, which ensures that each instance is writing the correct logs and emitting the correct metrics for AMS to properly manage the instance. You should onboard all of your EC2 instances, unless you specifically want AMS to ignore some. Automated Instance Configuration requires that specific conditions are met that enable AMS to configure the instance (for details see [Prerequisites for automated instance configuration in Accelerate](#)). The most important condition is that the AWS Systems Manager agent (SSM agent) needs to be installed on each Amazon EC2 instance that you want AMS to manage for you. For more information on SSM agent, see [Working with SSM agent](#).

SSM pre-installed in standard AMIs for Accelerate

The SSM agent is already installed on AWS-provided AMIs for the following operating systems.

- Amazon Linux and Amazon Linux 2
- SUSE Linux Enterprise Server (SLES) 12 and 15
- Microsoft Windows Server 2019, 2016, 2012 R2, 2012
- Ubuntu Linux 18.04 and 20.04

If you are using one of these AWS-provided AMIs, see [Tagging instances in Accelerate](#).

Manual SSM installation of SSM in Accelerate

For the following operating systems, or when using a custom AMI, you can manually install the SSM agent. Or, you can use the AMS SSM Agent auto installation feature. To learn more about SSM auto installation, see [SSM Agent automatic installation](#). For instructions on manual installation, select the link for your operating system:

- [CentOS SSM installation](#)
- [Oracle SSM installation](#)
- [Red Hat SSM installation](#)
- [SUSE Linux Enterprise Server SSM installation](#)
- [Windows SSM installation](#)

Tagging instances in Accelerate

After the SSM agent is installed, you must tag your instances. See [Tagging in AMS Accelerate](#).

Automated instance configuration in Accelerate

Once your instance is tagged, AMS performs an **Automated instance configuration**, which includes:

- Record operating system logs and metrics
- Enable remote access for AMS engineers
- Execute remote commands on the instance

These tasks are essential for AMS monitoring, patch, and log services; and for AMS to respond to incidents. For details on setting up **Automated Instance Configuration**, see [Automated instance configuration in AMS Accelerate](#).

After **Automated instance configuration** is complete, you are able to:

- Create incidents and service requests for Amazon EC2 instances and operating systems using the Support Center Console. For more information, see [Incident reports, service requests, and billing questions in AMS Accelerate](#).
- Access and audit Amazon EC2 logs
- Obtain patch reports

Onboarding AWS Backup in Accelerate

To configure backups, you need to create backup policies called *backup plans*. A backup plan specifies which AWS resources to back up, how frequently they need to be backed up, and the

backup retention period. We recommend evaluating your organization's continuity, security, and compliance requirements to determine what backup plans you need.

Opt-in

- Ensure that AWS Backup is enabled for each account, Region, and resource type by following the steps here:

[Getting Started 1: Service Opt-in.](#)

Optionally, [Getting started 2: Create an on-demand backup.](#)

Choose a backup plan

- To choose a backup plan, see [Select an AMS backup plan.](#)

Add resources

Resources are not associated with a backup plan by default. They need to be added to a backup plan.

- To add resources to a backup plan, see [Tag your resources to apply AMS backup plans.](#)
- To enable backup on all resources using tags, see [Backups.](#)

Onboarding patching in Accelerate

You need to configure patching to ensure that your software is up-to-date and meets your compliance policies.

AWS Backup prerequisite: To allow creation of a root volume snapshot during the patching maintenance window, ensure that AWS Backup is enabled for each account and region for the *Amazon EBS* resource type by following the steps here: [Getting started 1: Service Opt-in.](#) (You do not need to continue to *Getting started 2: Create an on-demand backup.*)

When to patch: Patching occurs during a *maintenance window*. You can schedule maintenance windows so that patches are only applied during preset times.

What to patch: You have to associate the Amazon EC2 instances you want to patch with a maintenance window. To associate the instances with a maintenance window, the Amazon EC2 instances must be tagged, and the maintenance window should have those tags as a target.

Which patches to install: Using patch baselines, you set rules to auto-approve certain types of patches, such as operating system or high-severity patches. You can also specify exceptions to your rules, for example, lists of patches that are always approved or rejected.

See [Patching recommendations](#) for guidance with Amazon EC2 patch policies.

- To start configuring patch management, see [Patch management in AMS Accelerate](#)
- To create a custom patch configuration, see [Custom patch baseline](#).

Review non-conformance reports in Accelerate

AMS deploys AWS Config rules that help you identify violations against standards set by the Center for Internet Security (CIS), the National Institute of Standards and Technology (NIST) Cloud Security Framework (CSF). We recommend you to review the non-conformance report with your delivery team to prioritize remediation actions so that your account is baselined to a conformant state.

Step 4. Customize features in Accelerate

In this stage, you have already onboarded monitoring, patching, and backup with default policies. Now you have the opportunity to customize policies to suit your needs.

You can choose to use the default policies for patch, backup or monitoring, or choose a custom policy based on your needs. AMS uses tags to associate resources to operational policies. AMS provides a Resource Tagger that allows you to specify rules on how tags are applied to your AWS resources based on application grouping or other grouping logic. For more information, see [Resource Tagger](#).

The customer-provided tags feature allows you to add and delete custom tags to AMS resources. For more information, see [Customer-provided tags](#).

Topics

- [Customize monitoring in Accelerate](#)
- [Customize backup in Accelerate](#)
- [Customize patching in Accelerate](#)

Customize monitoring in Accelerate

To customize monitoring of your cloud resources based on your application needs:

1. Create a custom monitoring policy. See [Modifying the Accelerate alarm default configuration](#).
2. Apply a custom policy to resources using tags. See [Monitoring](#)
3. Route alerts to the resource owner. See [Tag-based alert notification](#).

You can use the following CloudWatch dashboards to explore how many of your resources are targeted by AMS monitoring and tagging, and how many are not. In your account, navigate to the CloudWatch dashboards console, and select one of the following:

- AMS-Alarm-Manager-Reporting-Dashboard
- AMS-Resource-Tagger-Reporting-Dashboard

For a complete description of the dashboard metrics, see:

- [Viewing the number of resources monitored by Alarm Manager in Accelerate](#)
- [Viewing the number of resources managed by Resource Tagger](#)

Customize backup in Accelerate

You cannot customize AMS default back-up plans. Instead, create a new backup plan based on your application needs, and then attach resources to your custom plan using tags. It is up to you to choose which resources AMS should back up, how often, and for what retention period. We recommend evaluating the continuity, security, and compliance requirements of your organization to determine what backup plans you need.

- To create a backup plan, see [Creating a backup plan](#).
- To assign resources to a backup plan, see [Assigning resources to a backup plan](#).

Customize patching in Accelerate

Patching ensures that your software is up-to-date and meets your compliance policies.

When to patch: Patching occurs during a *maintenance window*. You can schedule maintenance windows so that patches are only applied during preset times.

What to patch: You have to associate the Amazon EC2 instances you want to patch with a maintenance window. To associate the instances with a maintenance window, the Amazon EC2 instances must be tagged, and the maintenance window should have those tags as a target.

Which patches to install: Using patch baselines, you set rules to auto-approve certain types of patches, such as operating system or high-severity patches. You can also specify exceptions to your rules, for example, lists of patches that are always approved or rejected.

- For general patching recommendations, see [Patching recommendations](#).
- To create custom maintenance windows, see [Create a patch maintenance window](#).
- To create custom patch baselines, see [Custom patch baseline](#).
- To route patch alerts to the resource owner, see [Patch monitoring and failure remediation](#).

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 policies to the default AMS policy to restrict what you can see and do in the console.

The AMS Accelerate console has these features:

- **Opening page:** The opening page has information boxes and links to facilitate your access to your 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:
 - **Incidents on your resources:** A button for opening an incident case in AWS Support Center, plus how many incident cases are **Awaiting approval** and require your attention and how many are **Open**
 - **Compliance status:** Links to **Rules** and **Resources** that are noncompliant or compliant
 - **Service requests:** A button for opening a service request case in AWS Support Center, plus how many are **Awaiting approval** and require your attention and how many are **Open**
 - **Account-level security:** Links to details on **Real time threat detection** GuardDuty findings and **Data security and privacy** Macie findings
 - **Quick actions:** Open your **Backup vaults** or **Patch instances** configuration pages
 - **Reports:** Opens the **Reports** page and the default reports, **Daily Backup** and **Daily Patch** and **Monthly Billing**
 - **Configuration:** Ensure your resources are being managed successfully and according to your specifications.
 - **Install SSM agent:** The SSM agent is required
 - **Configure tagging rules:** Opens AMS Resource Tagger
 - **Configure alarms:** Opens AMS CloudWatch alarm configuration
 - **Configure patch schedule:** Opens the AWS Systems Manager console
 - **Configure patch baselines:** Opens the AWS Patch Manager console
 - **Configure backup plans:** Opens the AWS Backup console
- **Feature spotlight:** Information on the latest updates to the console
- **Documentation:** The AWS Managed Services documentation landing page

AMS patterns

An AWS Managed Services (AMS) pattern is a generalized solution that solves for a family of use cases in the AMS managed environment.

As you operate on the AMS platform, AMS cloud architects (CAs) work with you to meet your business and operational requirements. While AMS customers operate in a unique way, we notice

that customers have similar use-cases. In such cases, CAs create general solution templates, or "patterns", that are used in multiple customer environments with minimal configuration and deployment effort.

AMS patterns are built to help deliver features to AMS customers and usually built by the account CA of the customer that requests it.

How AMS patterns work

To request more details about each pattern including cloudformation templates required so you can deploy the pattern, submit an AMS service request with the **Subject** "Request for additional details about pattern *Pattern_Name*" (substitute the pattern that you want) and add your AMS cloud architect (CA) to the **Additional contacts** option.

AMS patterns are classified into two (2) categories:

- **General Use:** Patterns are considered stable as they have been deployed and being utilized by multiple AMS customers
- **Preview Mode:** AMS recommends deploying Preview Mode patterns in your non-production environments for validation, and engaging with your Cloud Architect to discuss the use case before deployment.

Important

AMS patterns do not adhere to your default AMS Service Level Agreements [Service Level Agreements \(SLAs\)](#) and [Service Level Objectives \(SLOs\)](#). Support and updates to the pattern are done on a best-effort basis.

This AWS content is provided subject to the terms of the [AWS Customer Agreement](#) or other written agreement between the customer and either Amazon Web Services, Inc. or Amazon Web Services EMEA SARL ("AWS Europe") or both.

The material embodied in this software is provided to you "as-is" and without warranty of any kind, express, implied or otherwise, including without limitation, any warranty of fitness for a particular purpose.

AMS patterns

AMS patterns.

AMS patterns

Name	Overview	Benefits	Category
Customize CloudWatch Alarm Notifications	Customize CloudWatch alarm notifications to include information from instance tags such as instance name, application ID, and so forth.	Adding contextual information to the alarm notifications will make them more meaningful and provides actionable information.	Monitoring
Disk Usage Reporting	The Disk Usage Reporting pattern collects consumption space of volumes across multiple app accounts and present result as a centralized report in Amazon S3 with Athena table querying capability.	Provide insights into account volumes' actual usage to determine cost saving opportunities.	Cost optimization
Prowler Stack	Runs Prowler checks on accounts using Amazon EC2 where CloudShell can't be used.	Help unblock the Accelerate onboarding (Prowler) in cases where CloudShell can't be used either due to permissions or timeout issues without any impact to their current security posture.	Security
AMS Amazon S3 Replication with Custom Object Keys	Make copies of Amazon S3 objects and retain all	Customize the object keys (folders) during Amazon S3 replication without requiring	Reliability

Name	Overview	Benefits	Category
	<p>metadata and object keys (folders).</p> <p>Strips part of the source object keys, or creates custom destination object keys during replication.</p>	<p>additional scripts to move objects to required folders.</p>	
<p>Amazon EBS Snapshot Deletion</p>	<p>Automation based on Lambda and CloudWatch Events to automate deletion of Amazon EBS snapshots taken outside of AWS Backup, based on retention.</p>	<p>Help purge individual snapshots taken outside of AWS Backup orchestrator, saving added cost over time.</p>	<p>Cost optimization</p>
<p>AMS Amazon RDS Secrets Rotation</p>	<p>Using a CloudFormation template, automatically deploy all required resources (Lambda function, Security groups, elastic network interfaces or ENIs) needed for rotating secrets for supported Amazon RDS databases, Redshift, and DocumentDB.</p>	<p>Automate database secrets rotation, and provide a notification mechanism when rotation failure occurs.</p>	<p>Security</p>

Name	Overview	Benefits	Category
Automated Key Rotation	Automatically rotate access and secret keys for IAM users based, on CloudWatch Events and Lambda.	Easier rotation of access and secret keys for IAM users.	Security
Amazon EBS Volume Snapshot Tagger	Tag all Amazon EBS volumes and snapshots using the tags in the Amazon EC2 instances.	Help categorize and track costs with meaningful, relevant business information making it easier to validate where money is being spent, and enable the use of automation for tagged volumes and snapshots. Highly recommended best practice by the AWS Cost Optimization pillar.	Tagging (cost optimization, Security, incident management and automation)

Automated instance configuration in AMS Accelerate

AMS Accelerate provides an automated instance configuration service. This service ensures that an instance is emitting the correct logs and metrics for AMS to properly manage the instance. Automated instance configuration has its own prerequisites and steps for onboarding, described later in this section.

Topics

- [How automated instance configuration works in Accelerate](#)
- [SSM Agent automatic installation](#)
- [Automated instance configuration changes](#)

How automated instance configuration works in Accelerate

Automated instance configuration enables AMS Accelerate to perform certain configurations on a daily basis on instances that you indicate by adding particular agents and tags.

Prerequisites for automated instance configuration in Accelerate

These conditions must be met to enable AMS Accelerate to perform the previously described automated actions on managed instances.

The SSM Agent is installed

AMS Accelerate automated instance configuration requires that the AWS Systems Manager SSM Agent is installed.

For information on using the AMS SSM Agent auto installation feature see [SSM Agent automatic installation](#).

For information on manually installing the SSM Agent, see the following:

- Linux: [Manually install SSM Agent on Amazon EC2 instances for Linux - AWS Systems Manager](#)
- Windows: [Manually install SSM Agent on Amazon EC2 instances for Windows Server - AWS Systems Manager](#)

The SSM Agent is in the managed state

AMS Accelerate automated instance configuration requires an operational SSM Agent. The SSM Agent must be installed, and the Amazon EC2 instance must be in the managed state. For more information, see the AWS documentation, [Working with SSM Agent](#).

Automated instance configuration setup

Assuming the prerequisites have been met, adding a specific Amazon EC2 instance tag automatically initiates the AMS Accelerate automated instance configuration. Use one of the following methods to add this tag:

1. (Strongly recommended) Use the AMS Accelerate Resource Tagger

To configure the tagging logic for your account, see [How tagging works](#). After tagging is complete, tags and automated instance configuration are handled automatically.

2. Manually add tags

Manually add the following tag to the Amazon EC2 instances:

Key:**ams:rt:ams-managed**, Value:**true**.

Note

The instance configuration service attempts to apply the required AMS configurations once the **ams:rt:ams-managed** tag is applied to the instance. The service asserts the AMS required configurations whenever an instance is started, and when a the AMS daily configuration check occurs.

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

If your accounts are onboarded to AMS Accelerate Patch Add-On, then configure a patch maintenance window (MW) for the instances. If SSM Agent is missing or offline 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 instance configuration changes

The AMS Accelerate instance configuration automation makes the following changes in your account:

1. IAM permissions

Adds the IAM-managed Policies required to grant the instance permission to use the agents installed by AMS Accelerate.

2. Agents

- a. The Amazon CloudWatch Agent is responsible for emitting OS logs and metrics. The instance configuration automation ensures that the CloudWatch agent is installed and running the AMS Accelerate minimum version.
- b. The AWS Systems Manager SSM Agent is responsible for running remote commands on the instance. The instance configuration automation ensures that the SSM Agent is running the AMS Accelerate minimum version.

3. CloudWatch Configuration

- a. To ensure that the required metrics and logs are emitted, AMS Accelerate customizes the CloudWatch configuration. For more information, see the following section, [CloudWatch configuration change details](#).

Automated instance configuration makes changes or additions to your IAM instance profiles and CloudWatch configuration.

IAM permissions change details

Each managed instance must have an instance profile that includes the following managed policies:

- `arn:aws:iam::aws:policy/AmazonSSMManagedInstanceCore`
- `arn:aws:iam::aws:policy/CloudWatchAgentServerPolicy`
- `arn:aws:iam::aws:policy/AMSInstanceProfileBasePolicy`

The first two are AWS-managed policies. The AMS-managed policy is:

AMSInstanceProfileBasePolicy

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "secretsmanager:CreateSecret",
        "secretsmanager:UpdateSecret"
      ],
      "Resource": [
        "arn:aws:secretsmanager:*:*:secret:/ams/byoa/*"
      ],
      "Effect": "Allow"
    },
    {
      "Action": [
        "kms:Encrypt"
      ],
      "Resource": [
        "*"
      ],
      "Effect": "Allow"
    }
  ]
}
```


If your instance already has a profile, but is missing any of these policies, AMS adds the missing policies to your profile. If your instance does not have a profile, AMS attaches the **AMSOSConfigurationCustomerInstanceProfile-<REGION>**, which has all policies required by AMS Accelerate.

CloudWatch configuration change details

Additional detail on the CloudWatch configuration.

- CloudWatch configuration file location on the instance:
 - Windows: %ProgramData%\Amazon\AmazonCloudWatchAgent\amazon-cloudwatch-agent.json
 - Linux: /opt/aws/amazon-cloudwatch-agent/etc/amazon-cloudwatch-agent.d/ams-accelerate-config.json
- CloudWatch configuration file location in Amazon S3:
 - Windows: [https://ams-configuration-artifacts-*REGION_NAME*.s3.*REGION_NAME*.amazonaws.com/configurations/cloudwatch/latest/windows-cloudwatch-config.json](https://ams-configuration-artifacts-<i>REGION_NAME</i>.s3.<i>REGION_NAME</i>.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](https://ams-configuration-artifacts-<i>REGION_NAME</i>.s3.<i>REGION_NAME</i>.amazonaws.com/configurations/cloudwatch/latest/linux-cloudwatch-config.json)
- Metrics collected:
 - Windows:
 - AWS Systems Manager SSM Agent (CPU_Usage)
 - CloudWatch Agent (CPU_Usage)
 - Disk space utilization for all disks (% free space)
 - Memory (% committed bytes in use)
 - Linux:
 - AWS Systems Manager SSM Agent (CPU_Usage)
 - CloudWatch Agent (CPU_Usage)
 - CPU (cpu_usage_idle, cpu_usage_iowait, cpu_usage_user, cpu_usage_system)
 - Disk (used_percent, inodes_used, inodes_total)
 - Diskio (io_time, write_bytes, read_bytes, writes, reads)
 - Mem (mem_used_percent)
 - Swap (swap_used_percent)

- Logs collected:
 - 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

Offboarding AMS Accelerate

AMS Accelerate provides an easy way for you to understand how to operate an enterprise-class operating environment, and the supporting infrastructure operations model, for AWS migration and usage. However, after using AMS Accelerate, you might decide to in-source, or re-assign AWS infrastructure operations responsibilities to other teams. To do this, you must offboard your accounts from the AMS service.

When you offboard an account from AMS Accelerate, AMS transfers all the responsibilities defined in our service description back to you. For example, you won't have the ability to cut incidents or service requests to AMS. Similarly, our operations engineers and automation will no longer have access to your Accelerate accounts, preventing us from remediating health, availability, and security and compliance findings. Your AWS workloads can continue to run in the same accounts that AMS was operating.

The team that will be performing your infrastructure operations services going forward should be included to define what people, tools, and processes will be used after your Accelerate offboarding. AMS leaves some of the AMS tooling such as guardrails and logs to allow the development of a "to be" operating environment and model. Review carefully the following documentation to understand the tools that you can continue using and how to request to offboard an account.

AMS Accelerate offboarding considerations

While preparing to offboard from Accelerate, keep the following considerations in mind.

- **Access:** The `ams-access-management` AWS CloudFormation stack that defines the `ams-access-management` AWS Identity and Access Management role isn't deleted. After offboarding, these resources remain, but are unused by other components that are left behind. You can delete the stack and role at your convenience.
- **Incident management:** Incident management is the process the AMS service uses to respond to your reported incidents. After offboarding, Accelerate no longer detects and responds to incidents, or assists your team in resolving issues. You will not be able to exchange incident and service request communications with Accelerate and Accelerate console access for your Accelerate accounts is deactivated.
- **Monitoring:** Monitoring is the process the AMS service uses to track your resources. During offboarding, AMS removes AMS-specific tools, such as Alarm Manager and Resource Tagger, and any EventBridge event rules and CloudWatch alarms that AMS deployed as part of the AMS

monitoring baseline. Accelerate will no longer respond to alarms or configure new ones after offboarding. For details on Alarm Manager and Resource Tagger, see [Tag-based Alarm Manager](#) and [Resource Tagger](#).

- **Security:** Security management is the process the AMS service uses to protect your resources. After offboarding, you keep your Amazon GuardDuty detector and findings, and any AWS Config rules that you created. AWS Config rules deployed by Accelerate are removed. Accelerate no longer monitors, remediates, or reports on the findings from these tools.
- **Patch management:** Patch management is the process the AMS service uses to update your EC2 instances. After offboarding, 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. You retain the Patch baselines and snapshots created in the past. Additionally, the configuration of your patch maintenance windows remains but the patches are no longer installed by Accelerate.
- **Problem management:** After offboarding, Accelerate no longer performs analysis to identify and investigate problems and to identify the root cause.
- **Designated experts:** After offboarding, your designated Cloud Service Delivery Manager (CSDM) and a Cloud Architect (CA) no longer provide guidance, report about, or drive operational and security excellence for your off-boarded Accelerate accounts.
- **Backup management:** AMS uses backup management to take snapshots of your resources. After offboarding, you can continue using the backup schedules, frequency, and retention periods defined on AWS Backup except for the AMS backup plans; see [Select an AMS backup plan](#). The IAM resources created by AMS Backup Orchestrator are removed, but not the AMS-authored backup vaults and corresponding KMS key. Accelerate no longer monitors the backup jobs or performs restoration actions during incidents.
- **Operations tools:** AMS Accelerate can provide ongoing operations for your workload's infrastructure in AWS. After offboarding an Accelerate account you no longer have access to tools like Resource Tagger to help you tag your resources based on rules, or automated instance configuration to install required agents in your EC2 instances. CloudWatch and SSM Agents on instances are left in place with existing configurations. The `AMSOSConfigurationCustomerInstanceRole` IAM profile and the `AMSInstanceProfileBasePolicy` are detached from your instances and be removed from your Accelerate accounts.
- **Cost optimization:** After offboarding, AMS Resource Scheduler is deleted. AMS Resource Scheduler helps you reduce operational costs by stopping the resources that are not in use and starting them back when their capacity is needed. AMS does not continue providing cost

optimization recommendations. For details on Resource Scheduler, see [Cost optimization with AMS Resource Scheduler](#).

- **Logging and Reporting:** After offboarding, you retain the logs stored as a result of CloudWatch, CloudTrail, and VPC Flow Logs. You can leave the configuration of those services as-is to continue generating logs; however, AMS no longer monitors such configurations. Accelerate no longer provides the monthly service report that summarizes key performance metrics of AMS. You retain the data generated from Self-Service Reporting (SSR) (see [Self-service reporting](#)), but Accelerate does not generate new ones.
- **Automation:** After offboarding, the AMS-curated AWS SSM automation runbooks, and AWS Lambda functions, are no longer available.

Getting offboarding assistance for an Accelerate account

AMS offboards your account after receiving at least 30 days notice through a AMS Account Service Termination Request. The Service Termination Date is 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.

To request off-boarding an account you must:

1. Submit a formal request to offboard the account using a service request. One service request (SR) documenting all of the accounts you want to offboard, or one SR per account.

In the request, provide the list of account IDs to offboard, the reason for offboarding, and any additional considerations.

2. Inform your CSDM about the accounts you want to offboard and request their help executing the offboarding process.

Notification settings in Accelerate

Communications between you and AMS occur for many reasons:

- Events created by monitoring alerts
- Patching service notifications, if you have opted-in to the Patch add on
- Service requests and incident reports

- Occasional important AWS announcements (your CSDM contacts you if any action on your part is required)

All notifications are sent using an email that you provided for patch notifications when you were onboarded. Otherwise, notifications are sent to the default email that you provided to 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.

You can use named lists of contacts for non-resource based notifications, such as alerts based on GuardDuty or AWS Config. For example, you might have a list named `SecurityContacts` and another named `OperationsContacts`. AMS sends alarms and notifications to these lists.

See [AWS Config Control Compliance report](#) for more details.

Tagging in AMS Accelerate

Most Accelerate features (patching, backup, monitoring) use *tags* and *configuration profiles* to decide which of your resources to manage, what actions to apply, and when to apply them. Tags are labels that you apply to resources. Configuration profiles contain rules based on those tags.

Each Accelerate feature has its own tagging requirements. Some features require you to use specific tags, while others allow you to use any of your own.

For information about required tags, see [Customer-managed tags](#).

For information about tags that can be defined customers, see [Customer-provided tags](#)

Contents

- [Tags in AMS Accelerate](#)
 - [What are tags?](#)
 - [How tagging works](#)
 - [Customer-managed tags](#)
 - [Monitoring](#)
 - [Configuring EC2 instances](#)
 - [Patching](#)
 - [Backups](#)
 - [Accelerate-managed tags](#)
 - [Customer-provided tags](#)
- [Tag management tools](#)
 - [Resource Tagger](#)
 - [What is Resource Tagger?](#)
 - [How Resource Tagger works](#)
 - [Resource Tagger Configuration Profiles](#)
 - [Syntax and structure](#)
 - [Resource Tagger use cases in AMS Accelerate](#)
 - [Viewing the tags applied by Resource Tagger](#)
 - [Using Resource Tagger to create tags](#)

- [Preventing Resource Tagger from modifying resources](#)
- [Example configuration profile](#)
- [Merging the default configuration](#)
- [Disabling the default configuration](#)
- [Removing tags applied by Resource Tagger](#)
- [Viewing or making changes to the Resource Tagger configuration](#)
- [Deploying configuration changes](#)
- [Configuring Terraform to ignore Resource Tagger tags](#)
- [Viewing the number of resources managed by Resource Tagger](#)
- [CloudFormation](#)
 - [AWS CloudFormation Use Cases](#)
 - [Tagging an EC2 instance](#)
 - [Tagging an AutoScaling Group \(ASG\)](#)
 - [Deploying a configuration profile](#)
- [Terraform](#)

Tags in AMS Accelerate

Contents

- [What are tags?](#)
- [How tagging works](#)
- [Customer-managed tags](#)
 - [Monitoring](#)
 - [Configuring EC2 instances](#)
 - [Patching](#)
 - [Backups](#)
- [Accelerate-managed tags](#)
- [Customer-provided tags](#)

What are tags?

A tag is a label that you assign to an AWS resource. Each tag consists of a key and an optional value, both of which you define.

You use tags to categorize your AWS resources in different ways, for example, by purpose, owner, or environment. For example, you could define a set of tags for your account's Amazon EC2 instances, which helps you track each instance's owner and stack level.

Tags don't have any semantic meaning to Amazon EC2 and are interpreted strictly as a string of characters.

To learn more, see [Tagging AWS resources](#).

How tagging works

There are multiple ways to apply tags to your resources. You can tag resources directly in the console of each AWS service when you create the resource; use AWS [Tag Editor](#) to add, remove, or edit tags for multiple resources; or use provisioning tools such as AWS CloudFormation [Resource tag](#). AMS Accelerate also provides the AMS Accelerate Resource Tagger that you use to define rules for an automated tag lifecycle manager. For information about using Resource Tagger in AMS Accelerate, see [Resource Tagger](#). AMS Accelerate also provides customer-provided tags to add and remove custom tags to your AMS resources. For more information about customer-provided tags, see [Customer-provided tags](#).

Customer-managed tags

Certain tags are required to trigger various AMS Accelerate actions.

Contents

- [Monitoring](#)
- [Configuring EC2 instances](#)
- [Patching](#)
- [Backups](#)

Monitoring

AMS Accelerate monitors supported resources for health, availability, and reliability. For more information about this service offering, see [Monitoring and event management in AMS Accelerate](#).

AMS Accelerate periodically onboards additional AWS services to baseline monitoring. If you use the Resource Tagger default configuration, these updates are automatically deployed to your accounts, and changes are reflected to the supported resources.

To opt-in to have your Amazon EC2 instances managed by AMS Accelerate, you must apply the following tag via Customization profile in AppConfig; for more information, see [Step 3: Creating a configuration and a configuration profile](#).

Apply the following tag to your resources:

Key	Value
ams:rt:ams-managed	true

For example, you can create a Customized configuration document like this one to apply the tags to all your AMS-supported EC2 resources:

```
{
  "AWS::EC2::Instance": {
    "AllEC2": {
      "Enabled": true,
      "Filter": {
        "Platform": "*"
      },
      "Tags": [
        {
          "Key": "ams:rt:ams-managed",
          "Value": "true"
        }
      ]
    }
  }
}
```

Important

Remember to deploy your configuration changes after you have made them. In SSM AppConfig, you must deploy a new version of the configuration after creating it.

Services other than Amazon EC2 will have default baseline monitoring. In order to **opt out** your resources to be monitored by AMS Accelerate, you can use the customization configuration profile to exclude specific resources or AWS services. This allows you to control which resources should have monitoring tags to deploy baseline alarm definitions. See [Resource Tagger use cases in AMS Accelerate](#).

Using Resource Tagger

The AMS Accelerate Resource Tagger configuration in your account ensures that the following tags are deployed automatically, if you apply this one tag (**ams:rt:ams-managed**).

You will see the following tags being applied to your supported resources for baseline monitoring.

Key	Value	Rule
ams:rt:ams-monitoring-policy	ams-monitored	Applies to all EC2 resources supported by AMS
ams:rt:ams-monitoring-policy-platform	ams-monitored-linux	Applies to all Amazon EC2 instances running Linux OS
ams:rt:ams-monitoring-policy-platform	ams-monitored-windows	Applies to all Amazon EC2 instances running Windows OS

For other supported services

Apply the following tags to your resources, according to the given rules:

Key	Value	Rule
ams:rt:ams-monitoring-policy	ams-monitored	Applies to all resources supported by AMS Accelerate monitoring.
ams:rt:ams-monitoring-with-kms	ams-monitored-with-kms	OpenSearch Domain with KMS

Key	Value	Rule
ams:rt:ams-monitoring-with-master	ams-monitored-with-master	OpenSearch Domain with Dedicated Master Node

If you're not using Resource Tagger

See [Tags without Resource Tagger](#) for help on applying the correct monitoring tags using methods other than using AMS Resource Tagger.

Configuring EC2 instances

AMS Accelerate manages agents on your Amazon EC2 instances, such as the SSM agent and the CloudWatch agent. For more information about this service offering, see [Automated instance configuration in AMS Accelerate](#)

To opt-in to have your Amazon EC2 instances managed by AMS Accelerate, you must apply the following tag to your Amazon EC2 instances:

Key	Value
ams:rt:ams-managed	true

Patching


AMS Accelerate manages the patching of supported resources. For more information about this service offering, see [Patch management in AMS Accelerate](#).

Note

AMS Accelerate patching is an optional add-on service.

You can use any tag **key:value** combination to associate your resources with your patch maintenance windows. AMS Accelerate patch management uses tags to identify which resources should be patched in the default patch cycle. AMS Accelerate provides a default patch cycle when you onboard to patching. To make use of the default patch cycle, add the following tag to your supported resources:

Key	Value
AmsDefaultPatchKey	True

 **Note**

This is the default tag for the default patch cycle.

Backups

AMS Accelerate manages the backing up of supported resources. For more information about this service offering, see [Continuity management in AMS Accelerate](#).

AMS Accelerate backup management uses tags to identify which resources should be automatically backed up (and also provides manual backup capabilities). You can use any tag key:value combination to associate your resources with backup plans. To opt in to automated backups using the **ams-default-backup-plan** AWS Backup plan, you must apply the following tag to your supported resources:

Key	Value
ams:rt:backup-orchestrator	true

 **Note**

During onboarding, AMS Accelerate tags all resources with **ams:rt:backup-orchestrator-onboarding** with value **true** for short interval, short retention snapshots. This is managed by the **ams-onboarding-backup-plan** backup plan. For more information about AMS Accelerate-managed AWS Backup plans, see [Select an AMS backup plan](#).

Accelerate-managed tags

During onboarding to AMS Accelerate, several AWS resources are deployed to your account. So you can identify them, these resources are tagged with the following:

Key	Value
ams:resourceOwner	AMS
ams:resourceOwnerService	A description of which AMS Accelerate service offering this resource comes from, for instance, AMS Deployment, Backup, Controls, Monitoring, Patch, and so forth.
AppId	AMSInfrastructure
AppName	
Environment	

Note

These tags are applied using AWS CloudFormation stack-level tags, and rely on AWS CloudFormation propagating the tags to created resources. For more information, see [Resource tag](#).

Customer-provided tags

What are customer-provided tags?

Customer-provided tags is an AMS Accelerate service feature used to specify rules governing how AMS resources are tagged in your account. With customer-provided tags, you can add user-defined custom tags to AMS resources deployed to your accounts. The customer-provided tags feature is automatically added to the requested accounts when you request to your Cloud Service Delivery Manager (CSDM) using an automated service. Note that you can't override [AMS tags](#). AMS tags start with 'ams:'.

You can define your own [tags](#) (labels) and specify the functionalities of these tags for all of your [AMS Accelerate resources](#). You can provide these tags before you onboard to AMS so that AMS tags and your custom tags are applied during the onboarding process. Or, you can provide tags after onboarding.

How can I add customer-provided tags?

To request the addition of these tags to your resources, [contact your CSDM](#). These tags will be applied to AMS resources in your account.

What is the scope of the tags?

This feature is currently only available for Accelerate customers and in AWS commercial Regions. You can add tags to all accounts that you own or to a specific list of accounts.

Note

These tags apply only to AMS resources and don't affect your own resources.

Tag management tools

Contents

- [Resource Tagger](#)
- [CloudFormation](#)
- [Terraform](#)

Resource Tagger

With Resource Tagger, you can specify rules to govern how AWS resources are tagged in your account. While onboarding an account, AMS Accelerate deploys your tagging policy to ensure resources within your managed accounts are tagged.

Contents

- [What is Resource Tagger?](#)
- [How Resource Tagger works](#)
- [Resource Tagger Configuration Profiles](#)
 - [Syntax and structure](#)
- [Resource Tagger use cases in AMS Accelerate](#)
 - [Viewing the tags applied by Resource Tagger](#)

- [Using Resource Tagger to create tags](#)
- [Preventing Resource Tagger from modifying resources](#)
- [Example configuration profile](#)
- [Merging the default configuration](#)
- [Disabling the default configuration](#)
- [Removing tags applied by Resource Tagger](#)
- [Viewing or making changes to the Resource Tagger configuration](#)
- [Deploying configuration changes](#)
- [Configuring Terraform to ignore Resource Tagger tags](#)
- [Viewing the number of resources managed by Resource Tagger](#)

What is Resource Tagger?

Resource Tagger is an AMS Accelerate service offering you use to specify rules to govern how AWS resources are tagged in your account. It aims to provide you with complete visibility into how your tags are applied to your AWS resources.

Resource Tagger automatically creates, updates, and deletes tags on supported AWS resources, based on the tagging rules you specify in your configuration profiles. For example, you can specify a rule that applies a tag to a collection of Amazon EC2 instances, indicating that they should be managed by AMS Accelerate, which results in the instances being monitored or backed up. You can use tags like this to identify compliance status for the AWS resources based on the defined policy in your AWS AppConfig configuration profiles. For more information, see [AWS AppConfig](#).

AMS Accelerate provides a default managed tagging configuration so you can have your resources monitored by AMS Accelerate. You define which resources should be managed by AMS Accelerate, and the managed tagging rules ensure that the resources having the appropriate tags are monitored by AMS Accelerate.

With Resource Tagger, if you choose, you can override or deactivate the default AMS Accelerate managed tags, provide your own tagging rules to meet your policies, and use other mechanisms, such as Terraform, to avoid drift. You can define the exceptions to scale, based on your operations. For example, you could define policy to apply tags for all Amazon EC2 instances with supported platforms (such as Windows and Linux), and exclude from tagging specific instance IDs.

⚠ Important

Resource Tagger controls all tags in your account that have the **ams:rt:** prefix. Any tags that begin with this prefix are deleted unless they are present in Resource Tagger's configuration rules. To summarize, any tag on supported resources that starts with **ams:rt:** is considered owned by Resource Tagger. If you manually tag something with, for example, **ams:rt:**, that tag would automatically be removed if it wasn't specified in one of the Resource Tagger configuration profiles.

How Resource Tagger works

When your account is onboarded to AMS Accelerate, two JSON configuration documents are deployed to your account in AWS AppConfig. The two documents, called *Configuration profiles*, are **AMSManagedTags**, referred to as the **default configuration profile**, and **CustomerManagedTags**, referred to as the **customization configuration profile**. You use the customization configuration profile to define your own policies and rules for your accounts, and those are not overwritten by AMS Accelerate.

Both profiles reside in the **AMSResourceTagger** application, and in the **AMSInfrastructure** environment. All tags applied by the resource tagger have the key prefix **ams:rt:**.

Customization configuration profile:

The customization configuration profile is initially empty at the time of account onboarding; however, any rules placed in the profile document are enforced, in addition to the rules in the default configuration profile. Any configuration in the customization configuration profile is entirely managed by you, and is not overwritten by AMS Accelerate, except by your request.

You can specify any custom tagging rules you want in the custom configuration profile for the supported AWS resources, and you can also specify modifications to the AMS Accelerate-managed default configuration here, see [Resource Tagger use cases in AMS Accelerate](#).

⚠ Important

If you update this profile, the Resource Tagger automatically enforces the changes across all relevant resources in your AWS account. The changes are enacted automatically, but they may take up to 60 minutes to take effect.

You can update this profile by using the AWS Management Console, or through AWS CLI/SDK tools. For information about updating a customization configuration profile, see the AWS AppConfig user guide: [What Is AWS AppConfig?](#)

Default configuration profile:

The default configuration profile document is internal to AMS Accelerate and it contains AMS Accelerate-supplied default rules that you can't modify or delete permanently. This profile can be updated at any time by AMS Accelerate and made available to you for review; any changes you have made to it are automatically deleted. If you want to modify or disable any of the default configuration rules you use the customization configuration profile, see [Resource Tagger use cases in AMS Accelerate](#).

Resource Tagger Configuration Profiles

Configuration profiles help ensure that tags are applied uniformly to resources throughout the lifetime of the resources.

Syntax and structure

A configuration profile is a JSON object with the following structure:

```
{
  "Options": {
    "ReadOnly": false
  },
  "ResourceType": {
    "ConfigurationID": {
      "Enabled": true,
      "Filter": { ... },
      "Tags": [ ... ]
    },
    "ConfigurationID": {
      ...
    }
  },
  "ResourceType": {
    ...
  }
}
```

Options: (optional) Specify options for how you would like the ResourceTagger to behave. Omitting the block is equivalent to setting all options to their default values. See below for available **Options** settings:

- **ReadOnly:** (optional, defaults to false): Specifies ReadOnly mode for Resource Tagger. Set ReadOnly to true to disable Resource Tagger creating or removing tags on AWS resources. For more information, see [Preventing Resource Tagger from modifying resources](#).

ResourceType: This key must be one of the following supported strings, and represents all configuration related to the resource type indicated:

- AWS::AutoScaling::AutoScalingGroup
- AWS::EC2::Instance
- AWS::EC2::NatGateway
- AWS::EC2::VPNConnection
- AWS::EFS::FileSystem
- AWS::EKS::Cluster
- AWS::ElasticLoadBalancing::LoadBalancer
- AWS::ElasticLoadBalancingV2::LoadBalancer
- AWS::Elasticsearch::Domain
- AWS::FSx::FileSystem
- AWS::OpenSearch::Domain
- AWS::RDS::DBCluster
- AWS::RDS::DBInstance
- AWS::Redshift::Cluster
- AWS::Synthetics::Canary

ConfigurationID: This key must be unique in the profile document, and uniquely names the following block of configuration. If two configuration blocks in the same **ResourceType** block have the same **ConfigurationID**, the one that appears last in the profile takes effect. If you specify a **ConfigurationID** in your customization profile that is the same as one specified in the default document, the configuration block defined in the customization profile takes effect.

⚠ Important

The **ConfigurationID** should *not* overlap with the AMS Accelerate profile; for example, it should not be **AMSMonitoringLinux** or **AMSMonitoringWindows**, otherwise it disables the respective configuration of the **AMSManagedTags** configuration profile.

Enabled (optional, defaults to **true**): Specifies if the configuration block takes effect. Set this to **false** to disable a configuration block. A disabled configuration block has no effect.

Filter: Specifies the resources that the configuration applies to. Each filter object can have any one (but only one) of the following fields:

- **AWS::AutoScaling::AutoScalingGroup:**
 - **AutoScalingGroupName**: The Autoscaling Group name. This field supports wildcard matching.
- **AWS::EC2::Instance:**
 - **InstanceId**: The filter matches an EC2 instance with the specified instance ID. This field supports wildcard matching, so **i-00000*** would match any instance that has an instance ID starting with **i-00000**.
 - **Platform**: The filter matches an EC2 instance with the specified platform. Valid values are **windows**, **linux** or the wildcard ***** (to match any platform).
- **AWS::EC2::NatGateway:**
 - **NatGatewayId**: The ID of the NAT Gateway. This field supports wildcard matching.
 - **State**: The state of the NAT gateway (pending | failed | available | deleting | deleted or wildcard **"*"**)
 - **VpcId**: The ID of the VPC in which the NAT Gateway resides. This field supports wildcard matching.
 - **SubnetId**: The ID of the Subnet in which the NAT Gateway resides. This field supports wildcard matching
- **AWS::EC2::VPNConnection:**
 - **VpnConnectionId**: The ID of the connection. This field supports wildcard matching.
- **AWS::EKS::Cluster:**
 - **ClusterName**: The name of the cluster. This field supports wildcard matching.
- **AWS::ElasticLoadBalancing::LoadBalancer (Classic Load Balancer):**

- **LoadBalancerName:** The LoadBalancer Name. This field supports wildcard matching.
- **Scheme:** Can be either "internet-facing", "internal" or wildcard "*".
- **VPCId:** The VPCId in which the loadbalancer is deployed, can be wildcard "*".
- **AWS::ElasticLoadBalancingV2::LoadBalancer (Application Load Balancer (ALB)):**
 - **LoadBalancerArn:** The LoadBalancer Amazon Resource Name (ARN).
 - **DNSName:** The DNSName of the LoadBalancer. This field supports wildcard matching.
 - **LoadBalancerName:** The LoadBalancer Name. This field supports wildcard matching.
- **AWS::Elasticsearch::Domain:**
 - **DomainId:** The DomainId of the ElasticSearch resource. This field supports wildcard matching.
 - **DomainName:** The DomainName of the ElasticSearch resource. This field supports wildcard matching.
 - **HasMasterNode:** Boolean value of true or false. Matches if the Domain has a dedicated master node.
 - **HasKmsKey:** Boolean value of true or false. Matches if the Domain has a KMS key for encryption at rest.
- **AWS::FSx::FileSystem:**
 - **FileSystemId:** The ID of the FSx filesystem. This field supports wildcard matching.
- **AWS::OpenSearch::Domain:**
 - **DomainId:** The DomainId of the OpenSearch resource. This field supports wildcard matching.
 - **DomainName:** The DomainName of the OpenSearch resource. This field supports wildcard matching.
 - **HasMasterNode:** Boolean; If the Domain has a dedicated master node, this can be set to true.
 - **HasKmsKey:** If the Domain has a KMS key for encryption at rest, this can be set to true.
- **AWS::RDS::DBCluster:**
 - **DBClusterIdentifier:** The filter matches an RDS cluster identifier with the specified identifier. This field does not support wildcard matching, so a cluster identifier must be specified.
 - **Engine:** The engine in use by the RDS Instance. This field supports wildcard matching.
 - **EngineVersion:** The engine version. This field supports wildcard matching.
- **AWS::RDS::DBInstance:**
 - **DBInstanceIdentifier:** The filter matches an RDS instance with the specified instance ID. This field does not support wildcard matching, so an instance identifier must be specified.

- **Engine:** The engine in use by the RDS Instance. This field supports wildcard matching.
- **EngineVersion:** The engine version. This field supports wildcard matching.
- **AWS::Redshift::Cluster:**
 - **ClusterIdentifier:** The Cluster Identifier. This field supports wildcard matching.
- **AWS::Synthetics::Canary:**
 - **CanaryName:** The name of the Synthetics canary.

Other Filter properties:

- **Tag:** The filter applies to any resource that already has the given tag applied. The value for this property must be a JSON object with the following fields:
 - **Key:** Must be an exact string, and specifies that the resources must have a tag with that exact key.
 - **Value:** Specifies the matching value for the tag. Supports wildcards, so a value of **Sample** matches any value that ends with the string **Sample**.
- **Fn::AND:** A JSON array of JSON objects. Each object follows the same rules as the **Filter** configuration block. This specifies that the filter match any resource that matches all of the sub-filters.
- **Fn::OR:** A JSON array of JSON objects. Each object follows the same rules as the **Filter** configuration block. This specifies that the filter match any resource that matches any of the sub-filters.
- **Fn::NOT:** A JSON object that follows the same rules as the **Filter** configuration block. This specifies that the filter explicitly not match any resource that matches the sub-filter. Use this to specify exclusions to your tagging rules.

Tags: The tags to be applied to the matched resources. (See [Tag naming and usage conventions](#).) This field is an array of key-value pairs:

- **Key:** The key of the tag to be applied.
- **Value:** The value of the tag to be applied.

Note

Tags applied by Resource Tagger always have keys that begin with **ams:rt:**. If you don't specify this prefix in your profile, Resource Tagger inserts it for you. This is how Resource Tagger distinguishes the tags it owns and manages from tags used by other tools for other purposes.

Resource Tagger use cases in AMS Accelerate

This section lists commonly performed actions with Resource Tagger.

Topics

- [Viewing the tags applied by Resource Tagger](#)
- [Using Resource Tagger to create tags](#)
- [Preventing Resource Tagger from modifying resources](#)
- [Example configuration profile](#)
- [Merging the default configuration](#)
- [Disabling the default configuration](#)
- [Removing tags applied by Resource Tagger](#)
- [Viewing or making changes to the Resource Tagger configuration](#)
- [Deploying configuration changes](#)
- [Configuring Terraform to ignore Resource Tagger tags](#)
- [Viewing the number of resources managed by Resource Tagger](#)

Viewing the tags applied by Resource Tagger

All tags applied by Resource Tagger have the key prefix **ams:rt:**. For example, the following tag definition results in a tag with key **ams:rt:sampleKey** and value **sampleValue**. All tags with this prefix are treated as being part of Resource Tagger.

```
{
  "Key": "sampleKey",
  "Value": "sampleValue"
```

```
}
```

Important

If you manually create your own tag with the **ams:rt:** prefix, it's considered managed by Resource Tagger. This means that if the resource is managed by Resource Tagger, but the configuration profiles do not indicate that the tag should be applied, then Resource Tagger removes your manually added tag. If you do manually tag resources managed by Resource Tagger, do not use the **ams:rt:** prefix for tag keys.

Using Resource Tagger to create tags

The AMS Accelerate Resource Tagger is a component that is deployed in your account during AMS Accelerate onboarding. Resource Tagger has a configurable set of rules that define how your resources should be tagged, then it enforces those rules, automatically adding and removing tags on resources to ensure they comply with your rules.

If you want to use the Resource Tagger to tag your resources, see [Resource Tagger](#).

The following is an example Resource Tagger configuration snippet that adds the tag **ams:rt:ams-managed** with the value **true** to all Amazon EC2 instances. The **ams:rt:ams-managed** tag opts you in to having your resources monitored by AMS Accelerate.

```
{
  "AWS::EC2::Instance": {
    "SampleConfigurationBlock": {
      "Enabled": true,
      "Filter": {
        "Platform": "*"
      },
      "Tags": [
        {
          "Key": "ams:rt:ams-managed",
          "Value": "true"
        }
      ]
    }
  }
}
```


⚠ Warning

Be careful when specifying the name for your new configuration (SampleConfigurationBlock in the provided example) as you may inadvertently override the AMS-managed configuration with the same name.

Preventing Resource Tagger from modifying resources

Resource Tagger can be set to a read-only mode that prevents it from adding or removing any tags on your resources. This is useful if you want to provide your own tagging mechanism.

When in read-only mode, Resource Tagger still examines the tagging rules that are being specified in the managed and customer configuration profiles, and scans for resources that do not meet these tagging rules. Any non-compliant resources are surfaced with AWS Config. The AWS Config Rules that you can look for have the AMSResourceTagger- prefix. For example the AMSResourceTagger-EC2Instance AWS Config rule evaluates if appropriate tags are created for AWS::EC2::Instance resources based on the configuration profile.

Resource Tagger stops at this point, and does not make any changes to your resources (does not add or remove tags).

You can enable the read-only mode by modifying the customer configuration profile to include the **ReadOnly** key in the **Options** block. For example, the following configuration profile snippet shows how this might look:

```
{
  "Options": {
    "ReadOnly": true
  },
  "AWS::EC2::Instance": {
    [... the rest of your configuration ...]
  }
}
```

Resource Tagger would react to this new configuration as soon as it has finished deploying, and stop adding and removing tags on resources.

Note

To re-enable tag modification, change the **ReadOnly** value to **false**, or remove the key altogether, since the default value is **false**.

For more on the **Options** setting, see [Syntax and structure](#), next.

Example configuration profile

The following example profile document specifies that all Windows EC2 instances that are part of a stack-* CloudFormation stack be managed by AMS Accelerate; however, explicitly excludes a particular EC2 instance with ID i-000000000000000001.

```
{
  "AWS::EC2::Instance": {
    "AMSMonitoringWindows": {
      "Enabled": true,
      "Filter": {
        "Fn::AND": [
          {
            "Platform": "Windows"
          },
          {
            "Tag": {
              "Key": "aws:cloudformation:stack-name",
              "Value": "stack-*"
            }
          },
          {
            "Fn::NOT": {
              "InstanceId": "i-000000000000000001"
            }
          }
        ]
      },
      "Tags": [
        {
          "Key": "ams:rt:ams-managed",
          "Value": "true"
        }
      ]
    }
  }
}
```

```
    }  
  }  
}
```

Warning

Be careful when specifying the name for your new configuration (SampleConfigurationBlock in the provided example) as you may inadvertently override the AMS-managed configuration with the same name.

Merging the default configuration

The default configuration profile is supplied by AMS Accelerate at the time of account onboarding. This profile provides default rules that are deployed in your account.

While you can't modify the default configuration profile, you can provide overrides to the defaults by specifying a configuration block in your customization configuration profile with the same **ConfigurationID** as the default configuration block. If you do this, your configuration block overwrites the default configuration block.

For example, consider the following default configuration document:

```
{  
  "AWS::EC2::Instance": {  
    "AMSMangedBlock1": {  
      "Enabled": true,  
      "Filter": {  
        "Platform": "Windows"  
      },  
      "Tags": [{  
        "Key": "my-tag",  
        "Value": "SampleValueA"  
      }]  
    }  
  }  
}
```

In order to change the tag value applied here from SampleValueA to SampleValueB, and have the tag applied to all instances, not just Windows instances, you would provide the following customization configuration profile:

```
{
  "AWS::EC2::Instance": {
    "AMSMangedBlock1": {
      "Enabled": true,
      "Filter": {
        "Platform": "*"
      },
    },
    "Tags": [{
      "Key": "my-tag",
      "Value": "SampleValueB"
    }]
  }
}
```

Important

Remember to deploy your configuration changes after you have made them; for information, see [Deploying configuration changes](#). In SSM AppConfig, you must deploy a new version of the configuration after creating it.

Disabling the default configuration

You can disable a default configuration rule by adding a configuration block with the same **ConfigurationID** to your customization configuration profile and giving the **Enabled** field for a value of **false**.

For example, if the following configuration were present in the default configuration profile:

```
{
  "AWS::EC2::Instance": {
    "AMSMangedBlock1": {
      "Enabled": true,
      "Filter": {
        "Platform": "Windows"
      },
    },
    "Tags": [{
      "Key": "my-tag",
      "Value": "SampleValueA"
    }]
  }
}
```

```
}  
}  
}
```

You could disable this tagging rule by including the following in your customization configuration profile:

```
{  
  "AWS::EC2::Instance": {  
    "AMSMangedBlock1": {  
      "Enabled": false  
    }  
  }  
}
```

Important

Remember to deploy your configuration changes after you have made them; for information, see [Deploying configuration changes](#). In SSM AppConfig, you must deploy a new version of the configuration after creating it.

Removing tags applied by Resource Tagger

Any tags prefixed with **ams:rt** are removed by Resource Tagger if the tags do not exist in the configuration profiles, or, if they do exist, where the filter doesn't match. This means that you can remove tags applied by Resource Tagger by doing one of the following:

- Modifying the customization configuration section that defines the tag.
- Adding an exception for the specific resources so they no longer match the filter.

For example: if a **Linux** instance has the following tags:

```
"Tags": [{  
  "Key": "ams:rt:MyOwnTag",  
  "Value": true  
}], {  
  "Key": "myTag",  
  "Value": true
```

```
}]
```

And you deploy the following Resource Tagger configuration profile:

```
{
  "AWS::EC2::Instance": {
    "AMSMonitoringWindows": {
      "Enabled": true,
      "Filter": {
        "Platform": "Windows"
      },
      "Tags": [{
        "Key": "ams:rt:ams-managed",
        "Value": "true"
      }]
    }
  }
}
```

Resource Tagger reacts to the new configuration changes, and the only tag on the instance becomes:

```
"Tags": [{
  "Key": "myTag",
  "Value": true
}]
```

Warning

Be careful when specifying the name for your new configuration (SampleConfigurationBlock in the provided example) as you may inadvertently override the AMS-managed configuration with the same name.

Important

Remember to deploy your configuration changes after you have made them; for information, see [Deploying configuration changes](#). In SSM AppConfig, you must deploy a new version of the configuration after creating it.

Viewing or making changes to the Resource Tagger configuration

The two JSON configuration profiles, **AMSManagedTags** and **CustomerManagedTags**, deployed to your account in [AWS AppConfig](#) at onboarding and residing in the `AMSResourceTagger` application, and in the **AMSInfrastructure** environment, can be reviewed through AppConfig's [GetConfiguration](#) API.

The following is an example of this `GetConfiguration` call:

```
aws appconfig get-configuration
--application AMSResourceTagger
--environment AMSInfrastructure
--configuration AMSManagedTags
--client-id ANY_STRING
outfile.json
```

Application: AppConfig logical unit to provide capabilities, for the Resource Tagger, this is `AMSResourceTagger`.

- **Environment:** `AMSInfrastructure`.
- **Configuration:** To view AMS Accelerate default tag definitions, the value is `AMSManagedTags`, while to view customer tag definitions, the value is `CustomerManagedTags`.
- **Client ID:** The unique application instance identifier, this can be any string.
- The tag definitions can then be viewed in the specified output file, in this case, `outfile.json`.

The alarm definitions can then be viewed in the specified output file, in this case, `outfile.json`.

You can see which version of configuration is deployed to your account by viewing the past deployments in the **AMSInfrastructure** environment.

To override tag rules:

Any of the existing tag rules can be overridden by updating the customization profile either with AWS CloudFormation by [Deploying a configuration profile](#) or, or directly using using AppConfig's [CreateHostedConfigurationVersion](#) API. Using the same **ConfigurationID** as a default configuration tag rule overrides the default rule, and applies the custom rule in its place.

To deploy changes made to the **CustomerManagedTags** document:

After you make changes to the customization configuration profile, you must deploy the changes for them. To deploy the new changes, AppConfig's [StartDeployment](#) API must be run using the AWS AppConfig Console or the CLI.

Deploying configuration changes

Once the customization is completed, these changes must be deployed through the AWS AppConfig [StartDeployment](#) API. The following instructions show how to deploy using the AWS CLI. Additionally, you can use the AWS Management Console to make these changes. For information, see [Step 5: Deploying a configuration](#).

```
aws appconfig start-deployment
--application-id <application_id>
--environment-id <environment_id>
--deployment-strategy-id <deployment_strategy_id>
--configuration-profile-id <configuration_profile_id>
--configuration-version 1
```

- **Application ID:** The application ID of the application AMSResourceTagger. Get this with the [ListApplications](#) API call.
- **Environment ID:** The environment ID; get this with the [ListEnvironments](#) API call.
- **Deployment Strategy ID:** The deployment strategy ID; get this with the [ListDeploymentStrategies](#) API call.
- **Configuration Profile ID:** The configuration profile ID of CustomerManagedTags; get this with the [ListConfigurationProfiles](#) API call.
- **Configuration Version:** The version of the configuration profile you intend to deploy.

Important

Resource Tagger applies the tags as specified in the configuration profiles. Any manual modifications you make (with the AWS Management Console, or CloudWatch CLI/SDK) to the resource tags are automatically reverted back, so ensure your changes are defined through Resource Tagger. To know which tags are created by the Resource Tagger, look for tag keys prefixed with `ams:rt:`.

Restrict access to the deployment with the [StartDeployment](#) and the [StopDeployment](#) API actions to trusted users who understand the responsibilities and consequences of deploying a new configuration to your targets.

To learn more about how to use AWS AppConfig features to create and deploy a configuration, see the documentation at [Working with AWS AppConfig](#).

Configuring Terraform to ignore Resource Tagger tags

If you use Terraform to provision your resources, and you want to use Resource Tagger to tag your resources, the Resource Tagger tags may be identified as drift by Terraform.

You can configure Terraform to ignore all Resource Tagger tags using the **lifecycle** configuration block, or the **ignore_tags** global configuration block. For more information, see the Terraform documentation on Resource Tagging at [Resource Tagging](#).

The following example shows how to create a global configuration to ignore all tags that begin with the Resource Tagger tag prefix `ams:rt::`

```
provider "aws" {  
  # ... potentially other configuration ...  
  
  ignore_tags {  
    key_prefixes = ["ams:rt:"]  
  }  
}
```

Viewing the number of resources managed by Resource Tagger

Resource Tagger sends metrics every hour to Amazon CloudWatch, in the `AMS/ResourceTagger` namespace. Metrics are emitted only for resource types supported by Resource Tagger.

Metric Name	Dimensions	Description
ResourceCount	Component, ResourceType	Number of resources (of the specified resource type) deployed in this region. Units: Count

Metric Name	Dimensions	Description
Resources MissingManagedTags	Component, ResourceType	Number of resources (of the specified resource type) that require managed tags, according to the configuration profiles, but have not yet been tagged by Resource Tagger. Units: Count
Unmanaged Resources	Component, ResourceType	Number of resources (of the specified resource type) with no managed tags applied by Resource Tagger. Typically, these resources did not match any Resource Tagger configuration block, or are explicitly excluded from configuration blocks. Units: Count
MatchingResourceCount	Component, ResourceType, ConfigClauseName	Number of resources (of the specified resource type) that match the Resource Tagger configuration block. For a resource to match the configuration block, the block must be enabled and the resource must match the block's filter. Units: Count

These metrics are also viewable as graphs, in the **AMS-Resource-Tagger-Reporting-Dashboard**. To see the dashboard, from the AWS CloudWatch management console, select **AMS-Resource-Tagger-Reporting-Dashboard**. By default, the graphs in this dashboard display the data for the prior 12-hour period.

AMS Accelerate deploys CloudWatch alarms to your account to detect significant increases in the number of unmanaged resources, for example, resources excluded from management by AMS Resource Tagger. AMS Operations will investigate increases in unmanaged resources that exceed: either 3 resources of the same type, or a 50% increase over all resources of the same type. If the change does not appear to be deliberate, AMS Operations may contact you to review the change.

CloudFormation

You can use AWS CloudFormation to apply tags at the stack level (see AWS CloudFormation documentation, [Resource tag](#)) or at the individual resource level (for example, see [Tagging your Amazon EC2 resources](#)).

Important

Some AMS Accelerate service components require tags with the **ams:rt:** prefix. Resource Tagger believes that it owns these tags, and will delete them if no Resource Tagger configuration rules permit them. You always need to deploy a Resource Tagger configuration profile for these tags, even if you are using AWS CloudFormation or Terraform.

AWS CloudFormation Use Cases

This section lists commonly performed actions with AWS CloudFormation.

Topics

- [Tagging an EC2 instance](#)
- [Tagging an AutoScaling Group \(ASG\)](#)
- [Deploying a configuration profile](#)

Tagging an EC2 instance

The following is an example of how you can apply the tag **ams:rt:ams-managed** with the value **true** to an Amazon EC2 instance managed by AWS CloudFormation. The **ams:rt:ams-managed** tag opts you in to having your resources monitored by AMS Accelerate.

```
Type: AWS::EC2::Instance
```

```
Properties:
```

```
  InstanceType: "t3.micro"
```

```
  # ...other properties...
```

```
Tags:
  - Key: "ams:rt:ams-managed"
    Value: "true"
```

Tagging an AutoScaling Group (ASG)

The following is an example of how you can apply the tag **ams:rt:ams-managed** with the value **true** to an Auto Scaling group managed by AWS CloudFormation. Note that the Auto Scaling group will propagate its tags to Amazon EC2 instances that are created by it. The **ams:rt:ams-managed** tag opts you in to having your resources monitored by AMS Accelerate.

```
Type: AWS::AutoScaling::AutoScalingGroup
Properties:
  AutoScalingGroupName: "SampleASG"

# ...other properties...

Tags:
  - Key: "ams:rt:ams-managed"
    Value: "true"
```

Deploying a configuration profile

If you wish to deploy your `CustomerManagedTags` configuration profile using AWS CloudFormation, you can use the following CloudFormation templates. Put your desired JSON configuration in the `AMSResourceTaggerConfigurationVersion.Content` field.

When you deploy the templates in a CloudFormation Stack or Stack Set, the deployment of the `AMSResourceTaggerDeployment` resource will fail if you have not followed the required JSON format for the configuration. See [Syntax and structure](#) for details on the expected format.

For help on deploying these templates as a CloudFormation stack or stack set, see the relevant AWS CloudFormation documentation below:

- [Creating a stack on the AWS CloudFormation console](#)
- [Creating a stack with AWS CLI](#)
- [Creating a stack set](#)

Note

If you deploy a configuration version using one of these templates, and then subsequently delete the CloudFormation stack/stack set, the template configuration version will remain as the current deployed version, and no additional deployment will be made. If you wish to revert back to a default configuration, you will need to either manually deploy an empty configuration (i.e. just {}), or update your stack to an empty configuration, rather than deleting the stack.

JSON

```
{
  "Description": "Custom configuration for the AMS Resource Tagger.",
  "Resources": {
    "AMSResourceTaggerConfigurationVersion": {
      "Type": "AWS::AppConfig::HostedConfigurationVersion",
      "Properties": {
        "ApplicationId": {
          "Fn::ImportValue": "AMS-ResourceTagger-Configuration-ApplicationId"
        },
        "ConfigurationProfileId": {
          "Fn::ImportValue": "AMS-ResourceTagger-Configuration-CustomerManagedTags-ProfileID"
        },
        "Content": "{\"Options\": {\"ReadOnly\": false}}",
        "ContentType": "application/json"
      }
    },
    "AMSResourceTaggerDeployment": {
      "Type": "AWS::AppConfig::Deployment",
      "Properties": {
        "ApplicationId": {
          "Fn::ImportValue": "AMS-ResourceTagger-Configuration-ApplicationId"
        },
        "ConfigurationProfileId": {
          "Fn::ImportValue": "AMS-ResourceTagger-Configuration-CustomerManagedTags-ProfileID"
        },
        "ConfigurationVersion": {
          "Ref": "AMSResourceTaggerConfigurationVersion"
        }
      }
    }
  }
}
```

```
    "DeploymentStrategyId": {
      "Fn::ImportValue": "AMS-ResourceTagger-Configuration-Deployment-StrategyID"
    },
    "EnvironmentId": {
      "Fn::ImportValue": "AMS-ResourceTagger-Configuration-EnvironmentId"
    }
  }
}
}
```

YAML

Description: Custom configuration for the AMS Resource Tagger.

Resources:

AMSResourceTaggerConfigurationVersion:

Type: AWS::AppConfig::HostedConfigurationVersion

Properties:

ApplicationId:

!ImportValue AMS-ResourceTagger-Configuration-ApplicationId

ConfigurationProfileId:

!ImportValue AMS-ResourceTagger-Configuration-CustomerManagedTags-ProfileID

Content: |

```
{
  "Options": {
    "ReadOnly": false
  }
}
```

ContentType: application/json

AMSResourceTaggerDeployment:

Type: AWS::AppConfig::Deployment

Properties:

ApplicationId:

!ImportValue AMS-ResourceTagger-Configuration-ApplicationId

ConfigurationProfileId:

!ImportValue AMS-ResourceTagger-Configuration-CustomerManagedTags-ProfileID

ConfigurationVersion:

!Ref AMSResourceTaggerConfigurationVersion

DeploymentStrategyId:

!ImportValue AMS-ResourceTagger-Configuration-Deployment-StrategyID

EnvironmentId:

!ImportValue AMS-ResourceTagger-Configuration-EnvironmentId

Terraform

If you don't want to use AMS Accelerate Resource Tagger, you can apply your own tags using Terraform. However, if you don't want to use Resource Tagger because of its drift from your Terraform definitions, there is a way for you to use the Resource Tagger and ignore the drift it causes; see [Configuring Terraform to ignore Resource Tagger tags](#).

Important

Some AMS Accelerate service components require tags with the **ams:rt:** prefix. Resource Tagger believes that it owns these tags, and will delete them if no Resource Tagger configuration rules permit them. You always need to deploy a Resource Tagger configuration profile for these tags, even if you are using AWS CloudFormation or Terraform.

The following is an example of how you can apply the tag **ams:rt:ams-managed** with the value **true** to an Amazon EC2 instance managed by Terraform. The **ams:rt:ams-managed** tag opts you in to having your resources monitored by AMS Accelerate.

```
resource "aws_instance" "sample_linux_instance" {
  # ...ami and other properties...

  instance_type = "t3.micro"

  tags = {
    "ams:rt:ams-managed" = "true"
  }
}
```

The following is an example of how you can apply the tag **ams:rt:ams-managed** with the value **true** to an Auto Scaling group managed by Terraform. Note that the Auto Scaling group propagates its tags to the Amazon EC2 instances that are created by it. The **ams:rt:ams-managed** tag opts you in to having your resources monitored by AMS Accelerate.

```
resource "aws_autoscaling_group" "sample_asg" {
  # ...other properties...
```

```
name = "terraform-sample"

tags = {
  "ams:rt:ams-managed" = "true"
}
}
```

For a description of how to manage Terraform-created resource tags, see [Configuring Terraform to ignore Resource Tagger tags](#).

Incident reports, service requests, and billing questions in AMS Accelerate

With AMS Accelerate, you can request help with operational issues and requests at any time through the AWS Support Center in the AWS console. AMS Accelerate 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 Accelerate operations engineers proactively notify you of important alerts and questions using the same mechanisms.

AMS Accelerate provides a range of operational services to help you achieve operational excellence on AWS. To gain a quick understanding of how AMS helps your teams achieve overall operational excellence in AWS Cloud with some of our key operational capabilities including 24x7 helpdesk, proactive monitoring, security, patching, logging and backup, see [AMS Reference Architecture Diagrams](#).

Topics

- [Incident management](#)
- [Service request management](#)
- [Incident report and service request testing](#)
- [Billing questions](#)

Incident management

In AMS Accelerate, you use the **AWS Support Center** in the **AWS Console** to file incident reports. Incidents are AWS service performance issues that impact your managed environment, as determined by AMS Accelerate or you. Incidents identified by the AMS Accelerate 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 Accelerate operations team determines if the event is non-impacting, or an incident (a service interruption or degradation), or a problem (the underlying root cause of one or more incidents).

Note

The AMS Accelerate team also receives incidents created by you programmatically using the [AWS Support API](#) with service code `service-ams-operations-report-incident`.

For information about using AWS Support, see [Getting started with AWS Support](#).

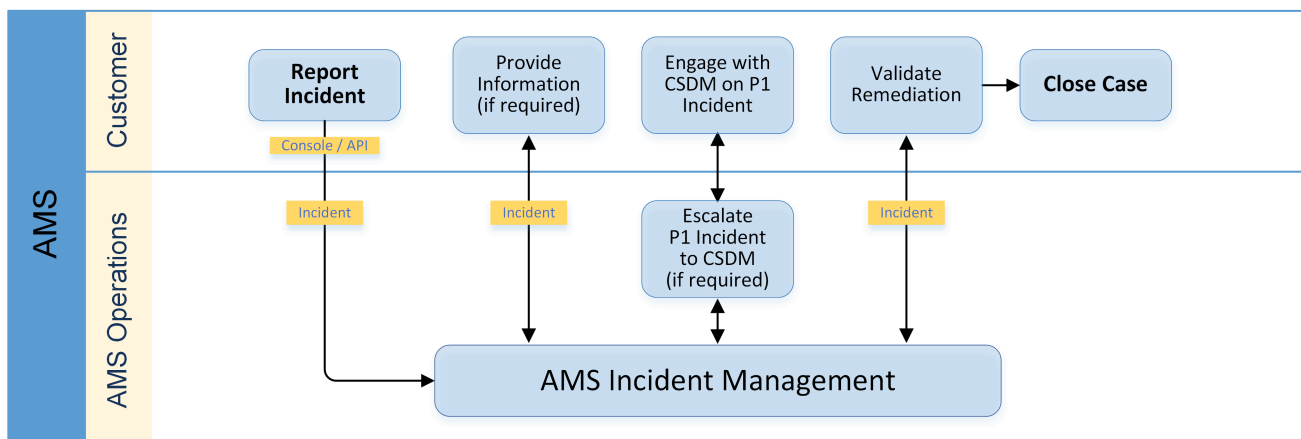
What is incident management?

Incident management is the process AMS uses to record, act on, communicate progress of, and provide notification of, active incidents.

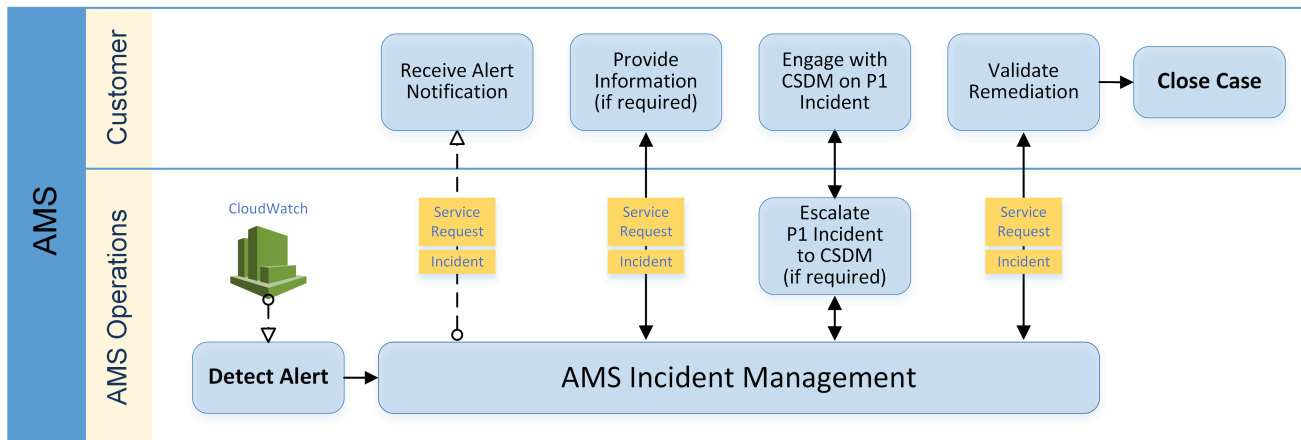
The goal of the incident management process is to ensure that normal operation of your managed service is restored as quickly as possible, the business impact is minimized, and all concerned parties are kept informed.

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.

How incident response and resolution work

AMS Accelerate uses 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 five operations centers around the world with dedicated operators actively watching monitoring dashboards and incident queues.

Our operations engineers use internal incident tracking tools to identify, log, categorize, prioritize, diagnose, resolve, and close incidents; we provide you with updates on all of these activities through AWS Support Center and through the AWS Support API. Our operators leverage a variety of internal AWS support tools to help with all of those activities. These operators are deeply familiar with AMS Accelerate-supported infrastructure and have expert-level technical skills to address identified support issues. In the event our operators need assistance, the Premium Support and AWS service teams are available.

After your incident is received by the AMS Accelerate operations team, we validate the priority and classification working with you if there are any clarifications required. For example, if the incident report is better classified as a service request, it's reclassified and the AMS Accelerate service request team takes over and you're notified. If the incident can be resolved by the receiving operator, steps are taken to quickly resolve the incident. AMS Accelerate operators consult internal documentation for a resolution and, if needed, escalate the incident to other support resources until the incident is resolved. After it's resolved, the AMS Accelerate operations team documents the incident and resolution for future use.

In cases where critical severity incidents are impacting your critical workloads, AMS Accelerate may recommend an infrastructure restore. There is often a trade-off between troubleshooting an issue and simply restoring from a known functional backup, and your risks and impacts from service downtime are the deciding factors. If you have time to devote to troubleshooting issues, AMS Accelerate will assist you, and your cloud service delivery manager (CSDM) may get involved, but if the urgency to restore is high, AMS Accelerate can initiate a restore right away.

Working with incidents

From AWS Support Center, you can perform the following tasks:

- Report and update an incident. To report an AMS Accelerate incident, choose **AMS Operations -- Report Incident** from the **Services** menu.
- Get a list of, and detailed information about, all of your submitted incidents.
- Narrow your search for incidents by status and other filters.
- Add communications and file attachments to your incidents, and add email recipients for case correspondence.
- Initiate a live chat or request a call back on your incident.
- Resolve incidents.
- Rate incident communications.

The following examples describe using Support Center to submit an incident. After it's submitted, the AMS Accelerate team works with you to resolve the incident per the standard AMS Accelerate SLA.

Submitting an incident

To report an incident using Support Center, refer to the support documentation: [Creating a support case](#)

To report an incident using AWS Support center:

1. Click **Create case**. The create incident case page opens.
2. Open the **Technical support issue type** menu and choose **AMS Operations -- Report Incident**. Supply information about your incident and choose **Create**.
3. To be kept informed by email at each step of the incident resolution process, be sure to fill in the **CC Emails** option; if you connect by federation, log in before following the link in the email that AMS Accelerate sends you about the incident.

Note

Make your description as detailed as possible. Include relevant resource information, along with anything else that might help us understand your issue. For example, to troubleshoot performance, include timestamps and logs. For feature requests or general guidance questions, include a description of your environment and purpose. In all cases, follow the Description Guidance that appears on your case submission form.

When you provide as much detail as possible, you increase the chances that your case can be resolved quickly.

You can also use the [AWS Support API](#) with service code `service-ams-operations-report-incident` to report an incident.

Monitoring and updating an incident

You can update, monitor, and review incident reports and service requests, both called *cases*, by using Support Center, or programmatically using the AWS Support API, [DescribeCases](#) operation.

To monitor a case, incident or service request, using AWS Support Center, follow these steps.

1. In the AWS Management console, browse to **Support**.
2. From the left navigation, select **Your support cases**, browse to a case and choose the **Subject** link to open a details page with current status and correspondences.

If you want to use phone or chat at this point, click **Open case in Support Center** to open the case **Create** page in the AWS Support Center, auto-populated with the AMS service type.

When a reported incident or service request case is updated by the Accelerate operations team, you receive an email and a link to the incident in the Support Center so you can respond.

 **Note**

You can't respond to case correspondence by replying to the email.

If there are many cases in the dashboard, you can use the **Filter** option:

- **Subject:** Use this filter to search on keywords in the subject of the case.
 - **Severity:** Use this to filter cases by severity by selecting a severity from the list.
 - **Case type:** Use this to see all cases of a particular case type. Accelerate incidents and service requests appear under the Technical Support Case Type along with any service-specific cases.
 - **Status:** Use this to filter cases by status by selecting a specific status from the list.
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 support API

You can use the [AWS Support API](#) to create incidents and add correspondence with AWS Support staff during investigations into your issues. The AWS Support API models much of the behavior of the [AWS Support Center](#).

For information about you can use this AWS support service, see [Programming the Life of an AWS Support Case](#).

Note

The AMS Accelerate team receives incidents created by you programmatically using the with service code `service-ams-operations-report-incident`.

Responding to an AMS Accelerate-generated incident

AMS Accelerate proactively monitors your resources. For more information, see [Monitoring and event management in AMS Accelerate](#). Sometimes AMS Accelerate identifies and creates an incident, most often to notify you of an event. If action on your part is required to resolve an incident, notification is sent by the AMS Accelerate team to the contact information you have provided for the account. You respond to this notification in the same way as for any other incident—usually through Support Center, though in some cases contact through email or phone is required.

Important

To receive state change notifications for an incident case or service request, enter an email address in the addresses field.

[Watch Akshay's video to learn more \(4:15\)](#)

Service request management

AMS Accelerate uses service request management to record, act on, communicate progress of, and provide notification of active service requests.

The goal of the service request management process is to ensure that your managed service is delivering what you need.

For billing-related queries, create a service request.

Note

The AMS Accelerate team receives service requests created by you programmatically using the [AWS Support API](#) with service code `service-ams-operations-service-request`.

How service request management works

Service requests are handled by the on-call AMS Accelerate operations team.

After your service request is received by the AMS Accelerate operations team, it's reviewed to make sure that it's properly classified as a service request or an incident. If it's reclassified as an incident, the AMS Accelerate incident management process begins and you're sent a notification.

If the AMS Accelerate 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 Accelerate operations, the operator either sends the request to your cloud service delivery manager, so they can communicate with you, or to the appropriate AWS support 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 are satisfied with the outcome.

Note

We recommend you provide a contact email, name, and phone number in all cases to facilitate communications.

Working with service requests

Using the AWS Support Center, you can perform the following tasks:


- Report and update a service request. To an AMS Accelerate service request, choose **AMS Operation -- Service Request** from the **Services** menu.
- Get a list of, and detailed information about, all of your submitted service requests.
- Narrow your search for service requests by status and other filters.
- 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 AWS Management console to create a service request for AMS Accelerate. After it's submitted, the AMS Accelerate team works with you to resolve the request per your AMS SLA.

Creating a service request

To create a service request, follow these steps:

1. From the AMS Accelerate console, browse to [Dashboard](#).
2. Choose **Open a service request**, the **AMS – Service Request** is pre-selected.
3. Choose a **Category**:
 - **Backup related**
 - **Monitoring related**
 - **Other** (non-resource-specific help or ask a how-to question)
 - **Patch related**
 - **Reporting Query** (request AMS-specific report data)
 - **Resource Tagger**
4. Choose a **Severity** (Plus or Premium tiers only):
 - **General Guidance:** Non-critical functions of your business service or application related to AWS or AMS Accelerate resources are impacted. This is the default for most service requests.

- **System Impaired:** A non-production business service or application related to AWS/AMS Accelerate resources is moderately impacted and functioning in a degraded state due to one of the categories listed in step 4.
 - **Production System Impaired:** A production business service or application related to AWS or AMS Accelerate resources is moderately impacted and functioning in a degraded state due to one of the categories mentioned in the previous step.
 - **Production System Down:** Critical functions of a production business service or application related to AWS or AMS Accelerate resources are unavailable. In most cases, we recommend using the incident form instead.
 - **Business Critical System Down (Premium tier only):** Your business is significantly impacted. Critical functions of your application related to AWS or AMS Accelerate resources are unavailable. In most cases, we recommend using the incident form instead.
5. Enter information for:
 - **Subject:** A descriptive title for the service request.
 - **Description:** A comprehensive description of the service request, the systems impacted, and the expected outcome of a resolution.
 6. To add an attachment, choose **Attach files**, browse to the attachment you want, and choose **Open**. To delete the attachment, choose the delete icon 
 7. **Contact us:** The default contact AMS through the web. To select other options:
 - **Preferred contact language:** English is the supported language for AMS Accelerate service requests.
 - **Web:** Your service request is submitted through the web and handled by the AMS operations team.
 - **Chat:** Chat online with an AMS Accelerate operations representative. This option adds you to the chat queue.
 - **Phone:** An AMS operations representative calls you back. Enter your AWS Region, phone number, and extension if applicable.
 - **Additional contacts:** Enter any additional email addresses you want copied on your service request.
 8. Choose **Submit**.

A case details page opens with information on the service request, such as **Type**, **Subject**, **Created**, **ID**, and **Status**. Plus, a **Correspondence** area that includes the description of the request you create.

To open a correspondence area and provide additional details or updates in status, choose **Reply**.

After the service request has been resolved, choose **Resolve Case**.

If there are so many correspondences that they don't all appear on the page, choose **Load More**.

Be sure to rate the service through the 1-5 star rating to let AMS know how we're doing.

Note

If you're going to test service request functionality, we recommend you add a no-action flag to your service request's subject, such as `AMSTestNoOpsActionRequired`. Then you can test without starting the service request resolution process.

The AMS Accelerate team receives service requests created by you programmatically using the [AWS Support API](#) with service code `service-ams-operations-service-request`.

Monitoring and updating a service request

To monitor a case, incident or service request, using AWS Support Center, follow these steps.

1. In the AWS Management console, browse to **Support**.
2. From the left navigation, select **Your support cases**, browse to a case and choose the **Subject** link to open a details page with current status and correspondences.

If you want to use phone or chat at this point, click **Open case in Support Center** to open the case **Create** page in the AWS Support Center, auto-populated with the AMS service type.

When a reported incident or service request case is updated by the Accelerate operations team, you receive an email and a link to the incident in the Support Center so you can respond.

Note

You can't respond to case correspondence by replying to the email.

If there are many cases in the dashboard, you can use the **Filter** option:

- **Subject:** Use this filter to search on keywords in the subject of the case.
 - **Severity:** Use this to filter cases by severity by selecting a severity from the list.
 - **Case type:** Use this to see all cases of a particular case type. Accelerate incidents and service requests appear under the Technical Support Case Type along with any service-specific cases.
 - **Status:** Use this to filter cases by status by selecting a specific status from the list.
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 service requests with the support API

You can use the [AWS Support API](#) to create service requests 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](#).

The AMS team also receives service requests created by you programmatically using the AWS Support API with the service code **service-ams-operations-service-request**.

For more information about how you can use this AWS support service, see [Programming the Life of an AWS Support Case](#).

Responding to an AMS Accelerate-generated service request

AMS Accelerate proactively monitors your resources; for more information, see [Monitoring and event management in AMS Accelerate](#). Sometimes AMS Accelerate creates a service request, or service notification for you, typically if action on your part is required to resolve a service request. In that case, the AMS Accelerate team sends a notification to the contact you have provided for the account. You respond to this service request in the same way as any other case—usually through the Support Center, though in some cases, email or phone correspondence is required.

Important

To receive state change notifications for a service request or incident case, you must have entered an email address in the addresses field. Notifications go only to the email address added to the case when it's created.

The link in the notification email works only if you're using an email server on your AMS Accelerate federated network. Otherwise, you can respond to the correspondence by going to your AMS Accelerate console and using the case details page.

Note

AMS Accelerate sends communications to your primary email address on your AWS account; we recommend adding an alternate operations contact email alias to facilitate the service request/notification management process. This is covered during the AMS Accelerate onboarding process and within the related onboarding documentation.

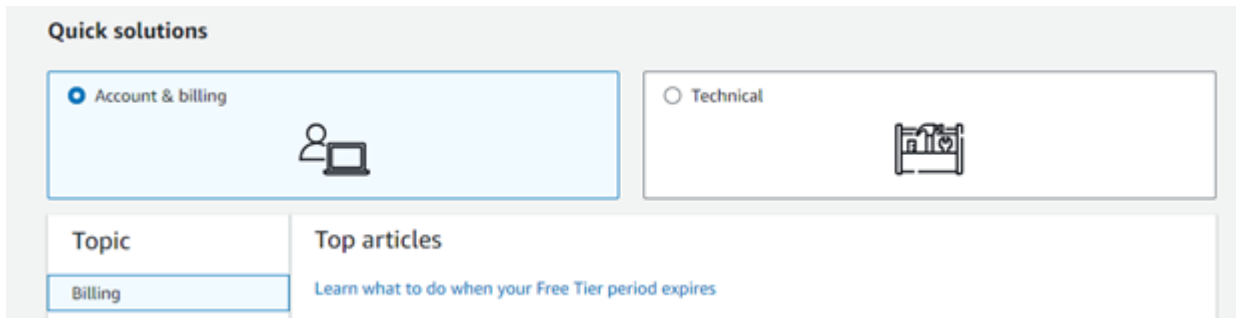
Incident report and service request testing

When testing incident reports or service requests, we ask that you include **AMSTestNoOpsActionRequired** in the subject text. This lets AMS know that the incident or request is only for testing. When AMS operations engineers see it, they will not respond in any way.


Billing questions

To submit a billing-related question, complete the following steps:


1. Open the [AWS Support Center](https://console.aws.amazon.com/support/home#/) 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 components 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	<p>AMS provides ongoing management of AMI building and vending for customers.</p> <p>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 assistance 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.</p>	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	<p>AMS provides ongoing operations of AWS landing zones deployed through AWS Landing Zone Accelerator (LZA).</p> <p>Our engineers handle configuration file changes, AWS Control Tower (CT) environment management (account vending, OU creation, CT guardrails), service control 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.</p>	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	<p>AMS collaborates with you to onboard, implement, and manage the ongoing operations of your SQL Server databases deployed on EC2 instances.</p> <p>Our engineers leverage our operational best practices and automation to free up your database teams by performing tasks such as backup and patching, extending AMS operational support to SQL Server patching to include cluster-aware rolling updates, backup and restore services aligned with our ransomware defense strategy, and monitoring adherence to customer-provided backup and patching controls.</p>	SQL Server customers assisted with offloading 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	<p>AMS frees your container developers by handling the ongoing maintenance and health of your Amazon Elastic Kubernetes Service (Amazon EKS) deployments. AMS performs the end-to-end procedures necessary to update a cluster addressing the components 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.</p>	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 automation 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 protection from known and unmitigated 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.

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

AWS Regions introduced before March 20, 2019 are considered "Original" Regions and are enabled by default. Regions introduced after this date are "Opt-in" Regions and are disabled by default. If your account uses multiple Regions and you onboard AMS Accelerate to an account with an enabled "Opt-in" Region as the default Region, the AMS Reporting feature is only available in that Region. If you do not set a default Region, the last Region you visited is your default Region.

To enable a Region, see [Enabling a Region](#). To set a default Region, see [Choosing a Region](#). For a list of the Opt-in status for each Region, see [Available Regions](#) in the *Amazon Elastic Compute Cloud User Guide*.

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 host management](#)
- [AMS Backup reporting](#)
- [AWS Config Control Compliance report](#)
- [AMS Config Rules Response Configuration report](#)
- [Incidents Prevented report](#)
- [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 host management

Topics

- [SSM Agent Coverage report](#)

SSM Agent Coverage report

AMS SSM Agent 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

Field Name	Definition
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

Field Name	Definition
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
# 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
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

Field Name	Definition
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. Resources are the ones supported by AWS Backup (Aurora, DocumentDB, DynamoDB, 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

Field Name	Definition
perc_coverage	Percentage of supported resources with a backup executed in the last 30 days.

AWS Config Control Compliance report

AWS Managed Services (AMS) AWS Config Control Compliance report provides an in-depth look at resource and config rule compliance of AMS accounts, including:

- Insights on top non-compliant resources in your environment to discover potential threats and misconfigurations.
- Insights on compliance of resources and config rules over time.
- Insights on config rule description, recommended severity of rule, and remediation steps to fix non-compliant resources.

Field	Description
Date	Report date
Customer name	Customer name
AWS account ID	Associated AWS account ID for customer
Source identifier	AWS Config rule unique source identifier
Rule Description	AWS Config rule description
Rule Type	AWS Config rule type
Compliance Flag	AWS Config rule compliance state
Resource Type	AWS resource type
Resource Name	AWS resource name
Severity	Default recommended severity defined by AMS for the AWS config rule

Field	Description
Remediation Category	Associated remediation response category for a config rule
Remediation Description	Remediation action explained to make config rule to be compliant
Customer action	Customer action required to make the config rule to be compliant
Delta metrics report	Changes for compliance of a rule between given 2 dates

AMS Config Rules Response Configuration report

AWS Managed Services (AMS) Config Rules Response configuration Report provides an in-depth look at the current configuration of how you want AMS Accelerate to respond to some findings (non-compliant Config rules). To learn more about how to change the response for config rules, please visit [AMS Accelerate Customized findings responses](#).

This report only shows the configurations that you have changed, and excludes the AMS default configurations which you can find here [AMS Config Rules Library](#). The report provides insights on resource and config rule response configuration of AMS accounts, including:

- The list of AWS accounts for which you have changed the default response for Config Rules.
- The list of tags for which you have associated a response for Config Rules.
- The list of response configurations for each rule, account, and/or tag.
- The list of resources for which you have changed the default response for Config Rules.

Latest Response Configurations Report

Field	Description
Date	Date in which the report was generated
Customer name	Customer name

Field	Description
AWS account ID	The AWS account ID associated with the configuration
Account Name	AWS account name of account level resource group
Finding Type	Type of finding identified. In this case, AWS Config
Source Identifier	AWS Config Rule Unique Source Identifier
Resource Group ID	The Resource Group ID associated with the Response Configuration
Response Action Configured	Action Type Triggered by AMS
SSM Runbook Associated	Remediation Runbook that will be executed, if any
Resource Group Type	Resource Group Type, Can be Account or Tag

Resources with Custom Default Response of Config Rules

Field Name	Definition
Customer Name	Customer name
Date	Date in which the report was generated
AWS Account Name	AWS Account Name
Account ID	Associated AWS Account ID
AMS Config Rule	AWS Config Rule which is targeting the resource and applying with a configuration

Field Name	Definition
Resource ID	The Resource ID in the customer account targeted by the Config Rule
Resource Region	The AWS Region which the configuration is applied
Resource Type	AWS resource type
Resource Group ID	The Resource Group ID associated with the Response Configuration
Resource AMS Flag	If the AWS resource is deployed by AMS, this field is set to true
Trigger Type	The type of response configured for the resource
Compliance Flag	AWS Config rule compliance state

Incidents Prevented report

AWS Managed Services (AMS) Incidents Prevented report lists the alarms that were automatically remediated, preventing a possible incident. To learn more, see [Auto remediation](#).

Field Name	Definition
execution_start_time_utc	Date in which the automation was executed
customer_name	Account customer name
account_name	The name of the account
AwsAccountId	The ID of the AWS account
document_name	The name of the SSM document or automation executed

Field Name	Definition
duration_in_minutes	The length of the automation in minutes
Region	AWS region where the resource is located
automation_execution_id	The ID of the execution
automation_execution_status	The status of the execution

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
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

Field Name	Definition
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; $\text{adjusted_aws_charges} * \text{uplift_percent}$
Instances EC2 RDS Spend	Spend on EC2 and RDS instances
AMS Charges	Total AMS charges for the product; $\text{uplifted_aws_charges} + \text{instance_ec2_rds_spend} + \text{uplifted_ris} + \text{uplifted_sp}$
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

Using custom keys with AWS Glue

If you 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 for AMS Patching](#)
- [Patch details](#)
- [Instances that missed patches](#)

Instance details summary for AMS Patching

This is an informational report that helps identify all the instances onboarded to AMS Patching, 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 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

Console Field Name	Dataset Field Name	Definition
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 infrastructure within customer account) or Customer stack (AMS managed infrastructure that supports customer applications)
Auto Scaling Group Name	instance_asg_name	Name of Auto Scaling Group (ASG) that contains the instance
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

Console Field Name	Dataset Field Name	Definition
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'.

Field Name	Dataset Field Name	Definition
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 infrastructure within customer account) or Customer stack (AMS managed infrastructure that supports customer applications)
Instance Patch Group Type	instance_patch_group_type	<p>DEFAULT: default patch group w/ default maintenance window, determined by AMSDefaultPatchGroup:True tag on the instance</p> <p>CUSTOMER: customer created patch group</p> <p>NOT_ASSIGNED: no patch group assigned</p>
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

Field Name	Dataset Field Name	Definition
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 maintenance 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_next_execution_time
	window_next_exec_mm	Month part of window_next_execution_time
	window_next_exec_D	Day part of window_next_execution_time
	window_next_exec_HHMM	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

Field Name	Dataset Field Name	Definition
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
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

Field Name	Dataset Field Name	Definition
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

Field Name	Dataset Field Name	Definition
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
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

Field Name	Dataset Field Name	Definition
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 responsible for paying the ams charges

Field Name	Dataset Field Name	Definition
Linked Account Id	linked_account_id	The 12 digit id identifying the AMS account that consumes services that generates expenses
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 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; $\text{adjusted_aws_charges} * \text{uplift_percent}$
Instances EC2 RDS Spend	instances_ec2_rds_spend	Spend on EC2 and RDS instances
Reserved Instance Charges	ris_charges	Reserved instance charges

Field Name	Dataset Field Name	Definition
Uplifted Reserved Instance Charges	uplifted_ris	The percentage of reserved instance charges to be charged for AMS; $\text{ris_charges} * \text{uplift_percent}$
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; $\text{sp_charges} * \text{uplift_percent}$
AMS Charges	ams_charges	Total ams charges for the product; $\text{uplifted_aws_charges} + \text{instance_ec2_rds_spend} + \text{uplifted_ris} + \text{uplifted_sp}$
Prorated Minimum Fee	prorated_minimum	The amount we charge to meet the contractual minimum
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
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

Field Name	Dataset Field Name	Definition
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
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

Field Name	Dataset Field Name	Definition
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
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

Field Name	Dataset Field Name	Definition
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
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.
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

Field Name	Dataset Field Name	Definition
		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.
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 example, 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 for AMS Patching	2 Months	2 Years
Patch Details	2 Months	2 Years
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 and AMS operations engineers will help you offboard from SSR. In the ticket please provide the reason for offboarding.

If you are offboarding your account and want to do a cleanup please create an SR through the AMS console and AMS operations engineers will help delete the 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 the offboarding process.

Access management in AMS Accelerate

Access management is how your resources are protected by allowing only authorized and authenticated access. With AMS Accelerate, you're responsible for managing access to your AWS accounts and their underlying resources, such as access management solutions, access policies, and related processes. In order to help you manage your access solution, AMS Accelerate deploys AWS Config rules that detect common IAM misconfigurations, and then deliver remediation notifications. A common IAM misconfiguration is that the root user has access keys. The `iam-root-access-key-check` config rule checks if the root user access key is available and is compliant or if the access key does not exist. For a list of config rules deployed by AMS, see the [AMS AWS Config Rule library](#).

Topics

- [Get access to the Accelerate console](#)
- [Permissions to use AMS features](#)
- [Why and when AMS accesses your account](#)
- [How AMS accesses your account](#)
- [How and when to use the root user account in AMS](#)

Get access to the Accelerate console

When you onboard with Accelerate, you automatically have access to the Accelerate console. You can access the console by searching for **Managed Services** in your AWS management console. The Accelerate console gives you a summarized view into the features you have with Accelerate. This view includes individual components presented on the dashboard and the configuration pages.

Permissions to use AMS features

To allow your users to read and configure AMS Accelerate capabilities, like accessing the AMS Console or configuring backups, you must grant explicit permissions to their IAM roles to perform those actions. The following AWS CloudFormation template contains the policies required to read and configure services associated with AMS so you can assign them to your IAM roles. They are designed to closely align with common job responsibilities in the IT industry, where Administrator or Read-Only permissions are required; however, if you need to grant different permissions to

users, you can edit the policy to include or exclude specific permissions. You can also create your own custom policy.

The template provides two policies. The `AMSAccelerateAdminAccess` policy is meant to be used for setting up and operating the AMS Accelerate components. This policy is typically assumed by an IT admin and grants permissions to configure AMS features such as patching and backups. The `AMSAccelerateReadOnly` grants minimum required permissions for viewing AMS Accelerate-related resources.

```

AWSTemplateFormatVersion: 2010-09-09
Description: AMSAccelerateCustomerAccessPolicies

Resources:
  AMSAccelerateAdminAccess:
    Type: 'AWS::IAM::ManagedPolicy'
    Properties:
      ManagedPolicyName: AMSAccelerateAdminAccess
      Path: /
      PolicyDocument:
        Fn::Sub:
          - |
            {
              "Version": "2012-10-17",
              "Statement": [
                {
                  "Sid": "AmsSelfServiceReport",
                  "Effect": "Allow",
                  "Action": "amsssr:*",
                  "Resource": "*"
                },
                {
                  "Sid": "AmsBackupPolicy",
                  "Effect": "Allow",
                  "Action": "iam:PassRole",
                  "Resource": "arn:aws:iam::${AWS:AccountId}:role/ams-backup-iam-role"
                },
                {
                  "Sid": "AmsChangeRecordKMSPolicy",
                  "Effect": "Allow",
                  "Action": [
                    "kms:Encrypt",
                    "kms:Decrypt",
                    "kms:GenerateDataKey"
                  ]
                }
              ]
            }

```



```

    ],
    "Resource": [
      "arn:aws:kms:${AWS::Region}:${AWS::AccountId}:key/*"
    ],
    "Condition": {
      "ForAnyValue:StringLike": {
        "kms:ResourceAliases": "alias/AMSCloudTrailLogManagement"
      }
    }
  },
  {
    "Sid": "AmsChangeRecordAthenaReadPolicy",
    "Effect": "Allow",
    "Action": [
      "athena:BatchGetNamedQuery",
      "athena:Get*",
      "athena:List*",
      "athena:StartQueryExecution",
      "athena:UpdateWorkGroup",
      "glue:GetDatabase*",
      "glue:GetTable*",
      "s3:GetAccountPublicAccessBlock",
      "s3:ListAccessPoints",
      "s3:ListAllMyBuckets"
    ],
    "Resource": "*"
  },
  {
    "Sid": "AmsChangeRecordS3ReadPolicy",
    "Effect": "Allow",
    "Action": [
      "s3:Get*",
      "s3:List*"
    ],
    "Resource": [
      "arn:aws:s3::ams-a${AWS::AccountId}-athena-results-${AWS::Region}",
      "arn:aws:s3::ams-a${AWS::AccountId}-athena-results-${AWS::Region}/",
      "arn:aws:s3::ams-a${AWS::AccountId}-cloudtrail-${AWS::Region}",
      "arn:aws:s3::ams-a${AWS::AccountId}-cloudtrail-${AWS::Region}/"
    ]
  },
  {
    "Sid": "AmsChangeRecordS3WritePolicy",

```

```
"Effect": "Allow",
"Action": [
  "s3:PutObject",
  "s3:PutObjectLegalHold",
  "s3:PutObjectRetention"
],
"Resource": [
  "arn:aws:s3:::ams-a${AWS::AccountId}-athena-results-${AWS::Region}/*"
]
},
{
  "Sid": "MaciePolicy",
  "Effect": "Allow",
  "Action": [
    "macie2:GetFindingStatistics"
  ],
  "Resource": "*"
},
{
  "Sid": "GuardDutyPolicy",
  "Effect": "Allow",
  "Action": [
    "guardduty:GetFindingsStatistics",
    "guardduty:ListDetectors"
  ],
  "Resource": "*"
},
{
  "Sid": "SupportPolicy",
  "Effect": "Allow",
  "Action": "support:*",
  "Resource": "*"
},
{
  "Sid": "ConfigPolicy",
  "Effect": "Allow",
  "Action": [
    "config:Get*",
    "config:Describe*",
    "config:Deliver*",
    "config:List*",
    "config:StartConfigRulesEvaluation"
  ],
}
```

```

    "Resource": "*"
  },
  {
    "Sid": "AppConfigReadPolicy",
    "Effect": "Allow",
    "Action": [
      "appconfig:List*",
      "appconfig:Get*"
    ],
    "Resource": "*"
  },
  {
    "Sid": "AppConfigPolicy",
    "Effect": "Allow",
    "Action": [
      "appconfig:StartDeployment",
      "appconfig:StopDeployment",
      "appconfig:CreateHostedConfigurationVersion",
      "appconfig:ValidateConfiguration"
    ],
    "Resource": [
      "arn:aws:appconfig:*:${AWS::AccountId}:application/
      ${AMSAlarmManagerConfigurationApplicationId}",
      "arn:aws:appconfig:*:${AWS::AccountId}:application/
      ${AMSAlarmManagerConfigurationApplicationId}/configurationprofile/
      ${AMSAlarmManagerConfigurationCustomerManagedAlarmsProfileID}",
      "arn:aws:appconfig:*:${AWS::AccountId}:application/
      ${AMSAlarmManagerConfigurationApplicationId}/environment/*",
      "arn:aws:appconfig:*:${AWS::AccountId}:application/
      ${AMSResourceTaggerConfigurationApplicationId}",
      "arn:aws:appconfig:*:${AWS::AccountId}:application/
      ${AMSResourceTaggerConfigurationApplicationId}/configurationprofile/
      ${AMSResourceTaggerConfigurationCustomerManagedTagsProfileID}",
      "arn:aws:appconfig:*:${AWS::AccountId}:application/
      ${AMSResourceTaggerConfigurationApplicationId}/environment/*",
      "arn:aws:appconfig:*:${AWS::AccountId}:deploymentstrategy/*"
    ]
  },
  {
    "Sid": "CloudFormationStacksPolicy",
    "Effect": "Allow",
    "Action": [
      "cloudformation:DescribeStacks"
    ],
  },

```

```
"Resource": "*"
},
{
  "Sid": "EC2Policy",
  "Action": [
    "ec2:DescribeInstances"
  ],
  "Effect": "Allow",
  "Resource": "*"
},
{
  "Sid": "SSMPolicy",
  "Effect": "Allow",
  "Action": [
    "ssm:AddTagsToResource",
    "ssm:CancelCommand",
    "ssm:CancelMaintenanceWindowExecution",
    "ssm:CreateAssociation",
    "ssm:CreateAssociationBatch",
    "ssm:CreateMaintenanceWindow",
    "ssm:CreateOpsItem",
    "ssm:CreatePatchBaseline",
    "ssm>DeleteAssociation",
    "ssm>DeleteMaintenanceWindow",
    "ssm>DeletePatchBaseline",
    "ssm:DeregisterPatchBaselineForPatchGroup",
    "ssm:DeregisterTargetFromMaintenanceWindow",
    "ssm:DeregisterTaskFromMaintenanceWindow",
    "ssm:Describe*",
    "ssm:Get*",
    "ssm:List*",
    "ssm:PutConfigurePackageResult",
    "ssm:RegisterDefaultPatchBaseline",
    "ssm:RegisterPatchBaselineForPatchGroup",
    "ssm:RegisterTargetWithMaintenanceWindow",
    "ssm:RegisterTaskWithMaintenanceWindow",
    "ssm:RemoveTagsFromResource",
    "ssm:SendCommand",
    "ssm:StartAssociationsOnce",
    "ssm:StartAutomationExecution",
    "ssm:StartSession",
    "ssm:StopAutomationExecution",
    "ssm:TerminateSession",
    "ssm:UpdateAssociation",
```

```

        "ssm:UpdateAssociationStatus",
        "ssm:UpdateMaintenanceWindow",
        "ssm:UpdateMaintenanceWindowTarget",
        "ssm:UpdateMaintenanceWindowTask",
        "ssm:UpdateOpsItem",
        "ssm:UpdatePatchBaseline"
    ],
    "Resource": "*"
},
{
    "Sid": "AmsPatchRestrictAMSResources",
    "Effect": "Deny",
    "Action": [
        "ssm:DeletePatchBaseline",
        "ssm:UpdatePatchBaseline"
    ],
    "Resource": [
        "arn:aws:ssm:${AWS::Region}:${AWS::AccountId}:patchbaseline/*"
    ],
    "Condition": {
        "StringLike": {
            "aws:ResourceTag/ams:resourceOwner": "*"
        }
    }
},
{
    "Sid": "AmsPatchRestrictAmsTags",
    "Effect": "Deny",
    "Action": [
        "ssm:AddTagsToResource",
        "ssm:RemoveTagsFromResource"
    ],
    "Resource": "*",
    "Condition": {
        "ForAnyValue:StringLike": {
            "aws:TagKeys": [
                "AMS*",
                "Ams*",
                "ams*"
            ]
        }
    }
},
{

```

```
    "Sid": "TagReadPolicy",
    "Effect": "Allow",
    "Action": [
      "tag:GetResources",
      "tag:GetTagKeys"
    ],
    "Resource": "*"
  },
  {
    "Sid": "CloudtrailReadPolicy",
    "Effect": "Allow",
    "Action": [
      "cloudtrail:DescribeTrails",
      "cloudtrail:GetTrailStatus",
      "cloudtrail:LookupEvents"
    ],
    "Resource": "*"
  },
  {
    "Sid": "EventBridgePolicy",
    "Effect": "Allow",
    "Action": [
      "events:Describe*",
      "events:List*",
      "events:TestEventPattern"
    ],
    "Resource": "*"
  },
  {
    "Sid": "IAMReadOnlyPolicy",
    "Action": [
      "iam:ListRoles",
      "iam:GetRole"
    ],
    "Effect": "Allow",
    "Resource": "*"
  },
  {
    "Sid": "AmsResourceSchedulerPassRolePolicy",
    "Effect": "Allow",
    "Action": "iam:PassRole",
    "Resource": "arn:aws:iam::${AWS::AccountId}:role/ams_resource_scheduler_ssm_automation_role",
    "Condition": {
```

```

        "StringEquals": {
            "iam:PassedToService": "ssm.amazonaws.com"
        }
    }
}
]
}
- AMSAlarmManagerConfigurationApplicationId: !ImportValue "AMS-Alarm-Manager-
Configuration-ApplicationId"
    AMSAlarmManagerConfigurationCustomerManagedAlarmsProfileID: !ImportValue
"AMS-Alarm-Manager-Configuration-CustomerManagedAlarms-ProfileID"
    AMSResourceTaggerConfigurationApplicationId: !ImportValue "AMS-
ResourceTagger-Configuration-ApplicationId"
    AMSResourceTaggerConfigurationCustomerManagedTagsProfileID: !ImportValue
"AMS-ResourceTagger-Configuration-CustomerManagedTags-ProfileID"

AMSAccelerateReadOnly:
Type: 'AWS::IAM::ManagedPolicy'
Properties:
    ManagedPolicyName: AMSAccelerateReadOnly
    Path: /
    PolicyDocument: !Sub |
    {
        "Version": "2012-10-17",
        "Statement": [
            {
                "Sid": "AmsSelfServiceReport",
                "Effect": "Allow",
                "Action": "amssrv:*",
                "Resource": "*"
            },
            {
                "Sid": "AmsBackupPolicy",
                "Effect": "Allow",
                "Action": [
                    "backup:Describe*",
                    "backup:Get*",
                    "backup:List*"
                ],
                "Resource": "*"
            },
            {
                "Action": [
                    "rds:DescribeDBSnapshots",

```

```

        "rds:ListTagsForResource",
        "rds:DescribeDBInstances",
        "rds:describeDBSnapshots",
        "rds:describeDBEngineVersions",
        "rds:describeOptionGroups",
        "rds:describeOrderableDBInstanceOptions",
        "rds:describeDBSubnetGroups",
        "rds:DescribeDBClusterSnapshots",
        "rds:DescribeDBClusters",
        "rds:DescribeDBParameterGroups",
        "rds:DescribeDBClusterParameterGroups",
        "rds:DescribeDBInstanceAutomatedBackups"
    ],
    "Effect": "Allow",
    "Resource": "*"
},
{
    "Action": [
        "dynamodb:ListBackups",
        "dynamodb:ListTables"
    ],
    "Effect": "Allow",
    "Resource": "*"
},
{
    "Action": [
        "elasticfilesystem:DescribeFilesystems"
    ],
    "Resource": "arn:aws:elasticfilesystem:*:*:file-system/*",
    "Effect": "Allow"
},
{
    "Action": [
        "ec2:DescribeSnapshots",
        "ec2:DescribeVolumes",
        "ec2:describeAvailabilityZones",
        "ec2:DescribeVpcs",
        "ec2:DescribeAccountAttributes",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeImages",
        "ec2:DescribeSubnets",
        "ec2:DescribePlacementGroups",
        "ec2:DescribeInstances",
        "ec2:DescribeInstanceTypes"
    ]
}

```



```
    ],
    "Effect": "Allow",
    "Resource": "*"
  },
  {
    "Action": [
      "tag:GetTagKeys",
      "tag:GetTagValues",
      "tag:GetResources"
    ],
    "Effect": "Allow",
    "Resource": "*"
  },
  {
    "Effect": "Allow",
    "Action": [
      "storagegateway:DescribeCachediSCSIVolumes",
      "storagegateway:DescribeStorediSCSIVolumes"
    ],
    "Resource": "arn:aws:storagegateway:*:*:gateway/*/volume/*"
  },
  {
    "Effect": "Allow",
    "Action": [
      "storagegateway:ListGateways"
    ],
    "Resource": "arn:aws:storagegateway:*:*:*"
  },
  {
    "Effect": "Allow",
    "Action": [
      "storagegateway:DescribeGatewayInformation",
      "storagegateway:ListVolumes",
      "storagegateway:ListLocalDisks"
    ],
    "Resource": "arn:aws:storagegateway:*:*:gateway/*"
  },
  {
    "Action": [
      "iam:ListRoles",
      "iam:GetRole"
    ],
    "Effect": "Allow",
    "Resource": "*"
  }
```

```

    },
    {
      "Effect": "Allow",
      "Action": "organizations:DescribeOrganization",
      "Resource": "*"
    },
    {
      "Action": "fsx:DescribeBackups",
      "Effect": "Allow",
      "Resource": "arn:aws:fsx:*:*:backup/*"
    },
    {
      "Action": "fsx:DescribeFileSystems",
      "Effect": "Allow",
      "Resource": "arn:aws:fsx:*:*:file-system/*"
    },
    {
      "Action": "ds:DescribeDirectories",
      "Effect": "Allow",
      "Resource": "*"
    },
    {
      "Sid": "AmsChangeRecordKMSPolicy",
      "Effect": "Allow",
      "Action": [
        "kms:Encrypt",
        "kms:Decrypt",
        "kms:GenerateDataKey"
      ],
      "Resource": [
        "arn:aws:kms:${AWS::Region}:${AWS::AccountId}:key/*"
      ],
      "Condition": {
        "ForAnyValue:StringLike": {
          "kms:ResourceAliases": "alias/AMSCloudTrailLogManagement"
        }
      }
    },
    {
      "Sid": "AmsChangeRecordAthenaReadPolicy",
      "Effect": "Allow",
      "Action": [
        "athena:BatchGetNamedQuery",
        "athena:Get*"
      ],

```

```

    "athena:List*",
    "athena:StartQueryExecution",
    "athena:UpdateWorkGroup",
    "glue:GetDatabase*",
    "glue:GetTable*",
    "s3:GetAccountPublicAccessBlock",
    "s3:ListAccessPoints",
    "s3:ListAllMyBuckets"
  ],
  "Resource": "*"
},
{
  "Sid": "AmsChangeRecordS3ReadPolicy",
  "Effect": "Allow",
  "Action": [
    "s3:Get*",
    "s3:List*"
  ],
  "Resource": [
    "arn:aws:s3::ams-a${AWS::AccountId}-athena-results-${AWS::Region}",
    "arn:aws:s3::ams-a${AWS::AccountId}-athena-results-${AWS::Region}/*",
    "arn:aws:s3::ams-a${AWS::AccountId}-cloudtrail-${AWS::Region}",
    "arn:aws:s3::ams-a${AWS::AccountId}-cloudtrail-${AWS::Region}/*"
  ]
},
{
  "Sid": "AmsChangeRecordS3WritePolicy",
  "Effect": "Allow",
  "Action": [
    "s3:PutObject",
    "s3:PutObjectLegalHold",
    "s3:PutObjectRetention"
  ],
  "Resource": [
    "arn:aws:s3::ams-a${AWS::AccountId}-athena-results-${AWS::Region}/*"
  ]
},
{
  "Sid": "MaciePolicy",
  "Effect": "Allow",
  "Action": [
    "macie2:GetFindingStatistics"
  ],
  "Resource": "*"
}

```

```
    },
    {
      "Sid": "GuardDutyReadPolicy",
      "Effect": "Allow",
      "Action": [
        "guarddduty:GetFindingsStatistics",
        "guarddduty:ListDetectors"
      ],
      "Resource": "*"
    },
    {
      "Sid": "SupportReadPolicy",
      "Effect": "Allow",
      "Action": "support:Describe*",
      "Resource": "*"
    },
    {
      "Sid": "ConfigReadPolicy",
      "Effect": "Allow",
      "Action": [
        "config:Get*",
        "config:Describe*",
        "config:List*"
      ],
      "Resource": "*"
    },
    {
      "Sid": "AppConfigReadPolicy",
      "Effect": "Allow",
      "Action": [
        "appconfig:List*",
        "appconfig:Get*"
      ],
      "Resource": "*"
    },
    {
      "Sid": "CloudFormationReadPolicy",
      "Effect": "Allow",
      "Action": [
        "cloudformation:DescribeStacks"
      ],
      "Resource": "*"
    },
    {
```

```
    "Sid": "EC2ReadPolicy",
    "Effect": "Allow",
    "Action": [
      "ec2:DescribeInstances"
    ],
    "Resource": "*"
  },
  {
    "Sid": "SSMReadPolicy",
    "Effect": "Allow",
    "Action": [
      "ssm:Describe*",
      "ssm:Get*",
      "ssm:List*"
    ],
    "Resource": "*"
  },
  {
    "Sid": "TagReadPolicy",
    "Effect": "Allow",
    "Action": [
      "tag:GetResources",
      "tag:GetTagKeys"
    ],
    "Resource": "*"
  },
  {
    "Sid": "CloudtrailReadPolicy",
    "Effect": "Allow",
    "Action": [
      "cloudtrail:DescribeTrails",
      "cloudtrail:GetTrailStatus",
      "cloudtrail:LookupEvents"
    ],
    "Resource": "*"
  },
  {
    "Sid": "EventBridgePolicy",
    "Effect": "Allow",
    "Action": [
      "events:Describe*",
      "events:List*",
      "events:TestEventPattern"
    ],
  },
```

```

    "Resource": "*"
  }
]
}

```

Why and when AMS accesses your account

AMS Accelerate (Accelerate) operators can access your account console and instances, in certain circumstances, for managing your resources. 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
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)
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

AMS operators require the following roles to service your account.

Important

Do not modify or delete these roles.

IAM roles for AMS access to customer accounts

Role Name	Description
ams-access-admin	This role has full administrative access to your account without restrictions. AMS services use this role with restrictive session policies that limit access to deploy AMS infrastructure and operate your account.
ams-access-admin-operations	This role grants AMS operators administrative permissions to operate your account. This role does not grant read, write, or delete permissions to customer content in AWS services commonly used as data stores, such as Amazon Simple Storage Service, Amazon Relational Database Service, Amazon DynamoDB, Amazon Redshift, and Amazon ElastiCache. Only qualified AMS operators who have a strong understanding and background in access management can assume this role. These operators serve as an escalation point for access management issues and access your accounts to troubleshoot AMS operator access issues.
ams-access-operations	This role has permissions to perform administrative tasks in your accounts. This role does not have read, write, or delete permissions to customer content in AWS services commonly used as data stores, such as Amazon Simple Storage Service, Amazon Relational Database Service, Amazon DynamoDB, Amazon

Role Name	Description
	Redshift, and Amazon ElastiCache. Permissions to perform AWS Identity and Access Management write operations are also excluded from this role. AMS Accelerate operations staff and cloud architects (CAs) can assume this role.
ams-access-read-only	This role has read-only access to your account. AMS Accelerate operations staff and cloud architects (CAs) can assume this role. Read permissions to customer content in AWS services commonly used as data stores, such as Amazon S3, Amazon RDS, DynamoDB, Amazon Redshift, and ElastiCache, are not granted this role.
ams-access-security-analyst	This AMS security role has permissions in your AMS account to perform dedicated security alert monitoring and security incident handling. Only a very few select AMS Security individuals can assume this role.
ams-access-security-analyst-read-only	This AMS security role is limited to read-only permissions in your AMS account to perform dedicated security alert monitoring and security incident handling.

Note

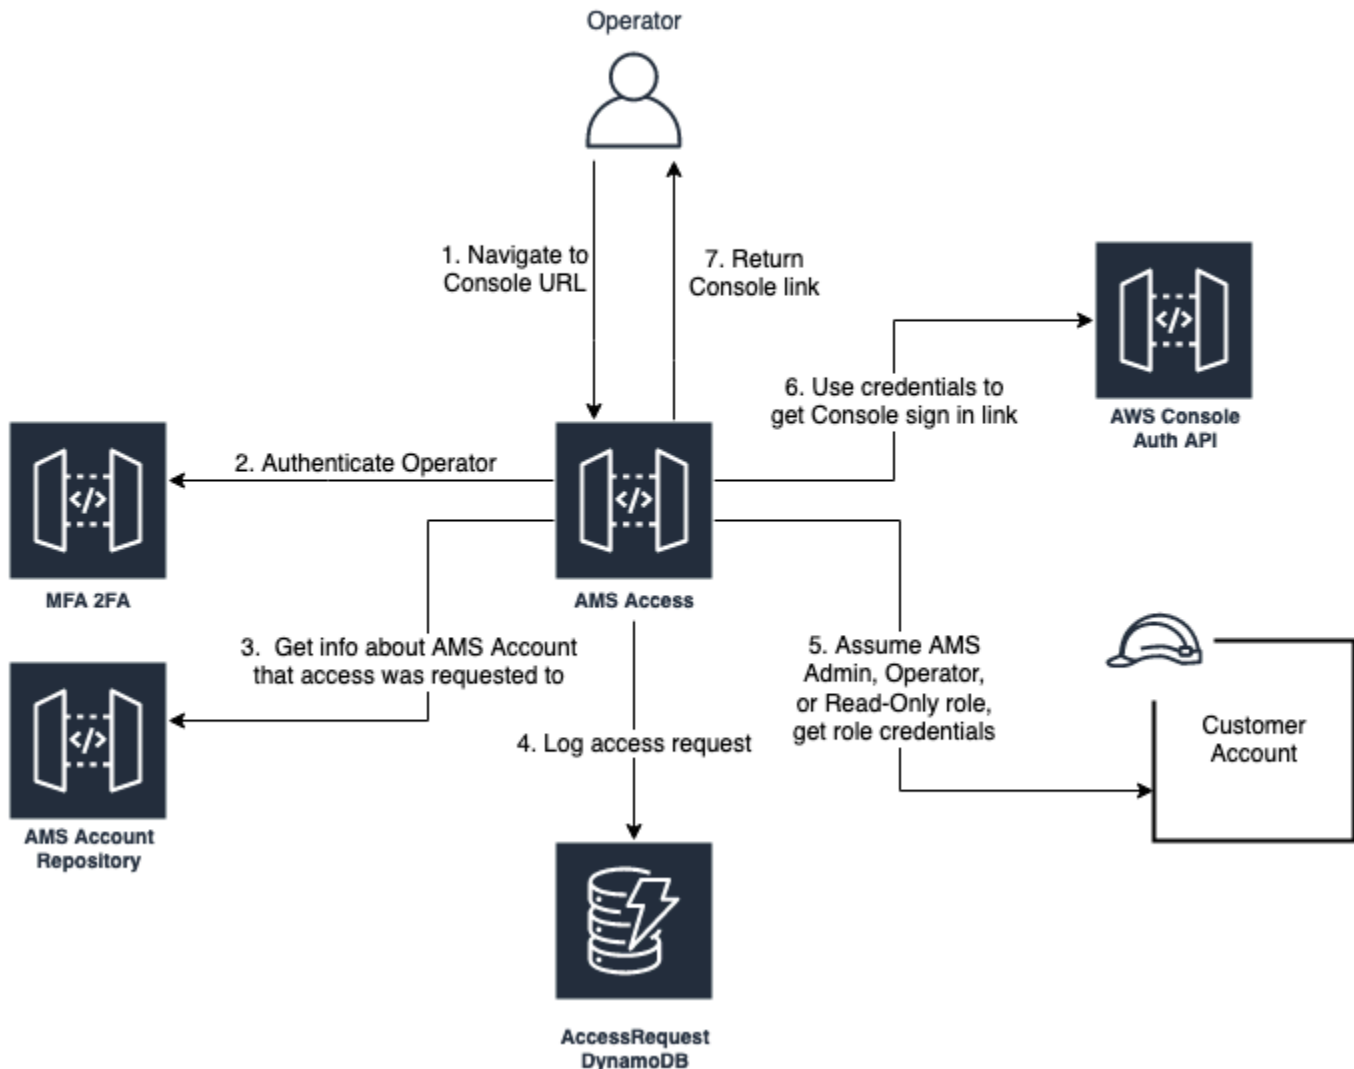
This is the template for the `ams-access-management` role. It is the stack that cloud architects (CAs) manually deploy in your account at onboarding time: [management-role.yaml](#).

This is the template for the different access roles for the different access levels: `ams-access-read-only`, `ams-access-operations`, `ams-access-admin-operations`, `ams-access-admin`: [accelerate-roles.yaml](#).

How AMS accesses your account

AMS Accelerate operators can access your account console and instances, in certain circumstances.

Console Access Method



AMS operators use the internal AMS Accelerate access service to access your accounts in a secured and audited manner. To access your instances, AMS operators use the same internal AMS access service as the broker and, after access is granted, AMS Accelerate operators use SSM session manager to gain access by using session credentials. RDP access for Windows instances is provided by establishing port forwarding to the instance and creating a local user using SSM. The local user credentials are used for RDP access and removed at the end of the session.

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](#).

Root with AMS Accelerate:

AMS does not prohibit you from using your root user account. However, AMS Operations and Security does treat its usage as an issue to investigate and we will reach out to your Security team with every use.

We recommend that you contact your CSDM and CA twenty-four hours in advance, to advise them of the root access work you intend to perform.

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.

Security in AMS Accelerate

AWS Managed Services protects your information assets and helps keep your AWS infrastructure secure by using multiple controls. AMS Accelerate maintains a library of AWS Config Rules and remediation actions to ensure that all your accounts comply with industry standards for security and operational integrity. AWS Config Rules continuously tracks the configuration change among your recorded resources. If a change violates any rule conditions, AMS reports its findings, and allows you to remediate violations automatically or by request, according to the severity of the violation. AWS Config Rules facilitate compliance with standards set by: the Center for Internet Security (CIS), the National Institute of Standards and Technology (NIST) Cloud Security Framework (CSF), the Health Insurance Portability and Accountability Act (HIPAA), and the Payment Card Industry (PCI) Data Security Standard (DSS).

In addition, AMS leverages Amazon GuardDuty to identify potentially unauthorized or malicious activity in your AWS environment. GuardDuty findings are monitored 24x7 by AMS. AMS collaborates with you to understand the impact of the findings and remediations based on best practice recommendations. AMS also uses Amazon Macie to protect your sensitive data such as personal health information (PHI), personally identifiable information (PII) and financial data.

Note

Amazon Macie is an optional service and is not enabled by default.

AMS Accelerate provides a range of operational services to help you achieve operational excellence on AWS. To gain a quick understanding of how AMS helps your teams achieve overall operational excellence in AWS Cloud with some of our key operational capabilities including 24x7 helpdesk, proactive monitoring, security, patching, logging and backup, see [AMS Reference Architecture Diagrams](#).

Topics

- [Using the Log4j SSM Document to Discover Occurrences](#)
- [Infrastructure security](#)
- [Data protection in AMS Accelerate](#)
- [AWS Identity and Access Management in AMS Accelerate](#)
- [Security Incident Response in AMS](#)

- [Security event logging and monitoring](#)
- [Configuration compliance in Accelerate](#)
- [Incident response in Accelerate](#)
- [Resilience in AMS Accelerate](#)
- [Security control for end-of-support operating systems](#)
- [Security best practices](#)
- [Change request security reviews](#)

Using the Log4j SSM Document to Discover Occurrences

AWS Managed Services (AMS) has released a new automation that will assist with searching for the Apache Log4j2 library within ingested workloads. The automation will provide a report of the Process ID of the Java application(s) in which the Log4j2 library is active.

The report includes information about the Java Archives (JAR Files), found within the specified environment that contain the JndiLookup class. AMS recommends upgrading the discovered libraries to the latest available version to mitigate the Remote Code Execution (RCE) identified via CVE-2021-44228. The latest version of the Log4j library can be downloaded directly from Apache at [Download Apache Log4j 2](#).

The document is shared to all the regions onboarded to Accelerate, and can be found in the AWS Systems Manager console under **Documents**. There, select **Shared with me** at the top, then search by Document name prefix equal to `AWSManagedServices-GatherLog4jInformation`. You can use [rate control](#) to run it at scale.

The parameters that `AWSManagedServices-GatherLog4jInformation` gathers are:

- **InstanceId:** (Required) ID of your EC2 instance.
- **S3Bucket:** (Optional) The S3 pre-signed URL or S3 URI (`s3://BUCKET_NAME`) to upload the results to.
- **AutomationAssumeRole:** (Required) The ARN of the role that allows Automation to perform the actions on your behalf.

AWS Managed Services recommends you run this document using rate control. You can set the rate control parameter to be the **InstanceId**, and assign either a list of instances to it, or apply a tag-key combination to target all EC2 instances that have a certain tag. AWS Managed Services also

recommends you provide an S3 bucket to upload the results to, so that you can build a report from the data stored in S3. For an example of how to aggregate the results in S3, see [EC2 Instance Stack | Gather Log4j Information](#).

If you are unable to upgrade the package, follow the guidelines outlined by AWS Security at [Using AWS security services to protect against, detect, and respond to the Log4j vulnerability](#). Additionally, Amazon has developed a Log4j hot patch, that you can run inline with your Java application(s), which will mitigate the vulnerability by removing the JndiLookup class functionality. For more details on the hot patch, please see [Hotpatch for Apache Log4j](#).


If you have any questions about the output of the automation, or how to proceed with additional mitigations, submit a service request.

Infrastructure security

During onboarding, AMS Accelerate deploys the following AWS Config baseline infrastructure and set of rules that AMS Accelerate uses to monitor your accounts:

- **AWS Config service-linked role:** AMS Accelerate deploys the service-linked role named **AWSServiceRoleForConfig**, which is used by AWS Config to query the status of other AWS services. The **AWSServiceRoleForConfig** service-linked role trusts the AWS Config service to assume the role. The permissions policy for the **AWSServiceRoleForConfig** role contains read-only and write-only permissions on AWS Config resources and read-only permissions for resources in other services that AWS Config supports. If you already have a role configured with AWS Config Recorder, AMS Accelerate validates that the existing role has an AWS Config managed-policy attached. If not, AMS Accelerate replaces the role with the service-linked role **AWSServiceRoleForConfig**.
- **AWS Config recorder and delivery channel:** AWS Config uses the configuration recorder to detect changes in your resource configurations and capture these changes as configuration items. AMS Accelerate deploys the configuration recorder in all service AWS Regions, with recording of all resources. AMS Accelerate also creates the config delivery channel, an Amazon S3 bucket, which is used to record changes that occur in your AWS resources; it updates configuration states through the delivery channel. The config recorder and delivery channel are required for AWS Config to work. AMS Accelerate creates the recorder in all AWS Regions, and a delivery channel in a single AWS Region. If you already have a recorder and delivery channel in an AWS Region, AMS Accelerate does not delete the existing AWS Config resources, instead AMS Accelerate utilizes your existing recorder and delivery channel after validating that they are properly configured.

- **AWS Config rules:** AMS Accelerate maintains a library of AWS Config Rules and remediation actions to help you comply with industry standards for security and operational integrity. AWS Config Rules continuously tracks configuration changes among your recorded resources. If a change violates any rule conditions, AMS reports its findings, and allows you to remediate violations automatically or by request, according to the severity of the violation. AWS Config Rules facilitate compliance with standards set by: the Center for Internet Security (CIS), the National Institute of Standards and Technology (NIST) Cloud Security Framework (CSF), the Health Insurance Portability and Accountability Act (HIPAA), and the Payment Card Industry (PCI) Data Security Standard (DSS).
- **AWS Config aggregator authorization:** An aggregator is an AWS Config resource type that collects AWS Config configuration and compliance data from multiple accounts and multiple Regions. AMS Accelerate onboards your account to a config aggregator from which AMS Accelerate aggregates your account's resource configuration information and config compliance data and generates the compliance report. If there are existing aggregators configured in the AMS-owned account, AMS Accelerate deploys an additional aggregator and the existing aggregator is not modified.

 **Note**

The Config aggregator is not set up in your accounts; rather, it is set up in AMS-owned accounts and your account(s) are onboarded to it.

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](#)

For information on reports, see [AWS Config Control Compliance report](#).

Using service-linked roles for AMS Accelerate

AMS Accelerate uses AWS Identity and Access Management (IAM) [service-linked roles](#). A service-linked role (SLR) is a unique type of IAM role that is linked directly to AMS Accelerate. Service-

linked roles are predefined by AMS Accelerate and include all the permissions that the service requires to call other AWS services on your behalf.

A service-linked role makes setting up AMS Accelerate easier because you don't have to manually add the necessary permissions. AMS Accelerate defines the permissions of its service-linked roles, and unless defined otherwise, only AMS Accelerate can assume its roles. The defined permissions include the trust policy and the permissions policy, and that permissions policy cannot be attached to any other IAM entity.

For information about other services that support service-linked roles, see [AWS services that work with IAM](#) and look for the services that have **Yes** in the **Service-linked roles** column. Choose a **Yes** with a link to view the service-linked role documentation for that service.

Deployment toolkit service-linked role for AMS Accelerate

AMS Accelerate uses the service-linked role (SLR) named **AWSServiceRoleForAWSManagedServicesDeploymentToolkit** – this role deploys AMS Accelerate infrastructure into customer accounts.

Note

This policy has recently been updated; for details, see [Accelerate updates to service-linked roles](#).

AMS Accelerate deployment toolkit SLR

The **AWSServiceRoleForAWSManagedServicesDeploymentToolkit** service-linked role trusts the following services to assume the role:

- `deploymenttoolkit.managedservices.amazonaws.com`

The policy named [AWSManagedServicesDeploymentToolkitPolicy](#) allows AMS Accelerate to perform actions on the following resources:

- `arn:aws*s3:::ams-cdktoolkit*`
- `arn:aws*:cloudformation:*:*:stack/ams-cdk-toolkit*`
- `arn:aws:ecr:*:*:repository/ams-cdktoolkit*`

This SLR grants Amazon S3 permissions to create and manage the deployment bucket used by AMS to upload resources, like CloudFormation templates or Lambda asset bundles, into the account for component deployments. This SLR grants CloudFormation permissions to deploy the CloudFormation stack that defines the deployment buckets. For details or to download the policy, see [AWSManagedServices_DeploymentToolkitPolicy](#).

You must configure permissions to allow an IAM entity (such as a user, group, or role) to create, edit, or delete a service-linked role. For more information, see [Service-linked role permissions](#) in the *IAM User Guide*.

Creating an deployment toolkit SLR for AMS Accelerate

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, AMS Accelerate creates the service-linked role for you.

Important

This service-linked role can appear in your account if you were using the AMS Accelerate service before June 09, 2022, when it began supporting service-linked roles, then AMS Accelerate created the `AWSServiceRoleForAWSManagedServicesDeploymentToolkit` 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 Accelerate creates the service-linked role for you again.

Editing an deployment toolkit SLR for AMS Accelerate

AMS Accelerate does not allow you to edit the `AWSServiceRoleForAWSManagedServicesDeploymentToolkit` 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 deployment toolkit SLR for AMS Accelerate

You don't need to manually delete the `AWSServiceRoleForAWSManagedServicesDeploymentToolkit` role. When you Offboard from AMS

in the AWS Management Console, the AWS CLI, or the AWS API, AMS Accelerate 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 Accelerate 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 Accelerate resources used by the `AWSServiceRoleForAWSManagedServicesDeploymentToolkit` service-linked role

Delete `ams-cdk-toolkit` stack from all Regions your account was onboarded to in AMS (you might have to manually empty the S3 buckets first).

To manually delete the service-linked role using IAM

Use the IAM console, the AWS CLI, or the AWS API to delete the `AWSServiceRoleForAWSManagedServicesDeploymentToolkit` service-linked role. For more information, see [Deleting a service-linked role](#) in the *IAM User Guide*.

Detective controls service-linked role for AMS Accelerate

AMS Accelerate uses the service-linked role (SLR) named `AWSServiceRoleForManagedServices_DetectiveControlsConfig` – AWS Managed Services uses this service-linked role to deploy config-recorder, config rules and S3 bucket detective controls..

Attached to the `AWSServiceRoleForManagedServices_DetectiveControlsConfig` service-linked role is the following managed policy:

[AWSManagedServices_DetectiveControlsConfig_ServiceRolePolicy](#). For updates to this policy, see [Accelerate updates to AWS managed policies](#).

Permissions for detective controls SLR for AMS Accelerate

The `AWSServiceRoleForManagedServices_DetectiveControlsConfig` service-linked role trusts the following services to assume the role:

- `detectivecontrols.managedservices.amazonaws.com`

Attached to this role is the

`AWSManagedServices_DetectiveControlsConfig_ServiceRolePolicy` AWS managed policy (see [AWS managed policy](#):

[AWSManagedServices_DetectiveControlsConfig_ServiceRolePolicy](#) The service uses the role to create configure AMS Detective Controls in your account, which requires deployment of resources like s3 buckets, config rules and an aggregator. 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*.

Creating a detective controls SLR for AMS Accelerate

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, AMS Accelerate creates the service-linked role for you.

Important

This service-linked role can appear in your account if you were using the AMS Accelerate service before June 09, 2022, when it began supporting service-linked roles then AMS Accelerate created the `AWSServiceRoleForManagedServices_DetectiveControlsConfig` 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 Accelerate creates the service-linked role for you again.

Editing a detective controls SLR for AMS Accelerate

AMS Accelerate does not allow you to edit the `AWSServiceRoleForManagedServices_DetectiveControlsConfig` 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 a detective controls SLR for AMS Accelerate

You don't need to manually delete the `AWSServiceRoleForManagedServices_DetectiveControlsConfig` role. When you Offboard from AMS in the AWS Management Console, the AWS CLI, or the AWS API, AMS Accelerate 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 Accelerate 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 Accelerate resources used by the `AWSServiceRoleForManagedServices_DetectiveControlsConfig` service-linked role

Delete `ams-detective-controls-config-recorder`, `ams-detective-controls-config-rules-cdk` and `ams-detective-controls-infrastructure-cdk` stacks from all Regions your account was onboarded to in AMS (you might have to manually empty the S3 buckets first).

To manually delete the service-linked role using IAM

Use the IAM console, the AWS CLI, or the AWS API to delete the `AWSServiceRoleForManagedServices_DetectiveControlsConfig` service-linked role. For more information, see [Deleting a service-linked role](#) in the *IAM User Guide*.

Amazon EventBridge rule service-linked role for AMS Accelerate

AMS Accelerate 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 Amazon 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 Accelerate

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 AMS. 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 Accelerate

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, AMS Accelerate creates the service-linked role for you.

Important

This service-linked role can appear in your account if you were using the AMS Accelerate 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 Accelerate creates the service-linked role for you again.

Editing an EventBridge SLR for AMS Accelerate

AMS Accelerate 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 Accelerate

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 Accelerate 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 Accelerate 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 Accelerate 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*.

Contacts service-linked role for AMS Accelerate

AMS Accelerate uses the service-linked role (SLR) named **`AWSServiceRoleForManagedServices_Contacts`** – This role facilitates automated notifications when incidents occur by allowing the service to read the existing tags of the affected resource and retrieve the configured email of the appropriate point of contact.

This is the only service that uses this service-linked role.

Attached to the **`AWSServiceRoleForManagedServices_Contacts`** service-linked role is the following managed policy: [`AWSManagedServices_ContactsServiceRolePolicy`](#). For updates to this policy, see [Accelerate updates to AWS managed policies](#).

Permissions for Contacts SLR for AMS Accelerate

The `AWSManagedServices_Contacts` service-linked role trusts the following services to assume the role:

- `contacts-service.managedservices.amazonaws.com`

Attached to this role is the `AWSManagedServices_ContactsServiceRolePolicy` AWS managed policy (see [AWS managed policy: AWSManagedServices_ContactsServiceRolePolicy](#)). The service uses the role to read the tags on any AWS resource and find the email contained in the tag, of the appropriate point of contact for when incidents occur. This role facilitates automated notifications when incidents occur by allowing AMS to read that tag on an affected resource and retrieve the email. For more information, see [Service-Linked Role Permissions](#) in the *AWS Identity and Access Management User Guide*.

Important

Do not store personally identifiable information (PII) or other confidential or sensitive information in tags. AMS uses tags to provide you with administration services. Tags are not intended to be used for private or sensitive data.

The role permissions policy named `AWSManagedServices_ContactsServiceRolePolicy` allows AMS Accelerate to complete the following actions on the specified resources:

- Action: Allows the Contacts Service to read the tags specifically set up to contain the email for AMS to send incident notifications on any AWS resource.

You can download the JSON `AWSManagedServices_ContactsServiceRolePolicy` in this ZIP: [ContactsServicePolicy.zip](#).

Creating a Contacts SLR for AMS Accelerate

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, AMS Accelerate creates the service-linked role for you.

⚠ Important

This service-linked role can appear in your account if you were using the AMS Accelerate service before February 16, 2023, when it began supporting service-linked roles then AMS Accelerate created the `AWSServiceRoleForManagedServices_Contacts` 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 Accelerate creates the service-linked role for you again.

Editing a Contacts SLR for AMS Accelerate

AMS Accelerate does not allow you to edit the `AWSServiceRoleForManagedServices_Contacts` 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 a Contacts SLR for AMS Accelerate

You don't need to manually delete the `AWSServiceRoleForManagedServices_Contacts` role. When you Offboard from AMS in the AWS Management Console, the AWS CLI, or the AWS API, AMS Accelerate 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 Accelerate 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 Accelerate resources used by the `AWSServiceRoleForManagedServices_Contacts` 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_Contacts` service-linked role. For more information, see [Deleting a service-linked role](#) in the *IAM User Guide*.

Supported regions for AMS Accelerate service-linked roles

AMS Accelerate supports using service-linked roles in all of the regions where the service is available. For more information, see [AWS regions and endpoints](#).

Accelerate updates to service-linked roles

View details about updates to Accelerate service-linked roles since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the Accelerate [Document history](#) page.

Change	Description	Date
Updated policy – Deployment Toolkit	<ul style="list-style-type: none"> These new permissions were added for resource <code>arn:aws:ecr:*:*:repository/ams-cdktoolkit* :</code> <pre>ecr:BatchGetRepositoryScanningConfiguration ecr:PutImageScanningConfiguration</pre> 	April 4, 2024
Updated policy – Deployment Toolkit	<ul style="list-style-type: none"> These new permissions were added for resource <code>arn:aws:cloudformation:*:*:stack/ams-cdk-toolkit* :</code> <pre>cloudformation:DeleteChangeSet cloudformation:DescribeStackEvents cloudformation:GetTemplate cloudformation:TagResource cloudformation:UntagResource</pre> These new permissions were added for resource <code>arn:aws:ecr:*:*:repository/ams-cdktoolkit* :</code> <pre>ecr:CreateRepository</pre> 	May 09, 2023

Change	Description	Date
	<pre> ecr:DeleteLifecyclePolicy ecr:DeleteRepository ecr:DeleteRepositoryPolicy ecr:DescribeRepositories ecr:GetLifecyclePolicy ecr:ListTagsForResource ecr:PutImageTagMutability ecr:PutLifecyclePolicy ecr:SetRepositoryPolicy ecr:TagResource ecr:UntagResource </pre> <ul style="list-style-type: none"> Also, some existing actions with wildcard were scoped down to individual actions: <pre> - s3:DeleteObject* + s3:DeleteObject + s3:DeleteObjectTagging + s3:DeleteObjectVersion + s3:DeleteObjectVersionTagging - s3:GetObject* + s3:GetObject + s3:GetObjectAcl + s3:GetObjectAttributes + s3:GetObjectLegalHold + s3:GetObjectRetention + s3:GetObjectTagging + s3:GetObjectVersion + s3:GetObjectVersionAcl + s3:GetObjectVersionAttributes + s3:GetObjectVersionForReplication + s3:GetObjectVersionTagging + s3:GetObjectVersionTorrent - cloudformation:UpdateTermination* + cloudformation:UpdateTerminationProt ection </pre> 	

Change	Description	Date
Updated policy – Detective Controls	<ul style="list-style-type: none"> The CloudFormation actions have been scoped down further after confirmation with security and access team The Lambda actions have been removed from the policy as they don't impact onboarding/off boarding 	April 10, 2023
Updated policy – Detective Controls	Updated the policy and added the permissions boundary policy.	March 21, 2023
New service-linked role – Contacts SLR	<p>Accelerate added a new service-linked role for the Contacts service.</p> <p>This role facilitates automated notifications when incidents occur by allowing the service to read the existing tags of the affected resource and retrieve the configured email of the appropriate point of contact.</p>	February 16, 2023
New service-linked role – EventBridge	<p>Accelerate added a new service-linked role for an Amazon EventBridge rule.</p> <p>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 Amazon 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.</p>	February 7, 2023

Change	Description	Date
Updated service-linked role – Deployment Toolkit	<p>Accelerate updated AWSServiceRoleForAWSManagedServicesDeploymentToolkit with new S3 permissions.</p> <p>These new permissions were added:</p> <pre>"s3:GetLifecycleConfiguration", "s3:GetBucketLogging", "s3:ListBucket", "s3:GetBucketVersioning", "s3:PutLifecycleConfiguration", "s3:GetBucketLocation", "s3:GetObject*"</pre>	January 30, 2023
Accelerate started tracking changes	Accelerate started tracking changes for its service-linked roles.	November 30, 2022
New service-linked role – Detective Controls	<p>Accelerate added a new service-linked role to deploy Accelerate detective controls.</p> <p>AWS Managed Services uses this service-linked role to deploy config-recorder, config rules and S3 bucket detective controls.</p>	October 13, 2022
New service-linked role – Deployment Toolkit	<p>Accelerate added a new service-linked role to deploy Accelerate infrastructure.</p> <p>this role deploys AMS Accelerate infrastructure into customer accounts.</p>	June 09, 2022

AWS managed policies for AMS Accelerate

An AWS managed policy is a standalone policy that is created and administered by AWS. AWS managed policies are designed to provide permissions for many common use cases so that you can start assigning permissions to users, groups, and roles.

Keep in mind that AWS managed policies might not grant least-privilege permissions for your specific use cases because they're available for all AWS customers to use. We recommend that you reduce permissions further by defining [customer managed policies](#) that are specific to your use cases.

You cannot change the permissions defined in AWS managed policies. If AWS updates the permissions defined in an AWS managed policy, the update affects all principal identities (users, groups, and roles) that the policy is attached to. AWS is most likely to update an AWS managed policy when a new AWS service is launched or new API operations become available for existing services.

For more information, see [AWS managed policies](#) in the *IAM User Guide*.

For a table of changes, see [Accelerate updates to AWS managed policies](#).

AWS managed policy: AWSManagedServices_AlarmManagerPermissionsBoundary

AWS Managed Services (AMS) uses the `AWSManagedServices_AlarmManagerPermissionsBoundary` AWS managed policy. This AWS-managed policy is used in the `AWSManagedServices_AlarmManager_ServiceRolePolicy` to restrict permissions of IAM roles created by `AWSServiceRoleForManagedServices_AlarmManager`.

This policy grants IAM roles created as part of [How Alarm Manager works](#), permissions to perform operations like AWS Config evaluation, AWS Config read to fetch Alarm Manager configuration, and creation of necessary Amazon CloudWatch alarms.

The `AWSManagedServices_AlarmManagerPermissionsBoundary` policy is attached to the `AWSServiceRoleForManagedServices_DetectiveControlsConfig` service-linked role. For updates to this role, see [Accelerate updates to service-linked roles](#).

You can attach this policy to your IAM identities.

Permissions details

This policy includes the following permissions.

- `AWS Config` – Allows permissions to evaluate config rules and select resource configuration.
- `AWS AppConfig` – Allows permissions to fetch AlarmManager configuration.
- `Amazon S3` – Allows permissions to operate AlarmManager buckets and objects.
- `Amazon CloudWatch` – Allows permissions to read and put AlarmManager managed alarms and metrics.

- **AWS Resource Groups and Tags** – Allows permissions to read resource tags.
- **Amazon EC2** – Allows permissions to read Amazon EC2 resources.
- **Amazon Redshift** – Allows permissions to read Redshift instances and clusters.
- **Amazon FSx** – Allows permissions to describe file systems, volumes and resource tags.
- **Amazon CloudWatch Synthetics** – Allows permissions to read Synthetics resources.
- **Amazon Elastic Kubernetes Service** – Allows permissions to describe Amazon EKS cluster.
- **Amazon ElastiCache** – Allows permissions to describe resources.

You can download the policy file in this ZIP: [RecommendedPermissionBoundary.zip](#).

AWS managed policy:

AWSManagedServices_DetectiveControlsConfig_ServiceRolePolicy

AWS Managed Services (AMS) uses the `AWSManagedServices_DetectiveControlsConfig_ServiceRolePolicy` AWS managed policy. This AWS-managed policy is attached to the [AWSServiceRoleForManagedServices_DetectiveControlsConfig service-linked role](#), (see [Detective controls service-linked role for AMS Accelerate](#)). For updates to the `AWSServiceRoleForManagedServices_DetectiveControlsConfig` service-linked role, see [Accelerate updates to service-linked roles](#).

The policy allows the service-linked role to complete actions for you.

You can attach the `AWSManagedServices_DetectiveControlsConfig_ServiceRolePolicy` policy to your IAM entities.

For more information, see [Using service-linked roles for AMS Accelerate](#).

Permissions details

This policy has the following permissions to allow AWS Managed Services Detective Controls to deploy and configure all necessary resources.

- **CloudFormation** – Allows AMS Detective Controls to deploy CloudFormation stacks with resources like s3 buckets, config rules and config-recorder.
- **AWS Config** – Allows AMS Detective Controls to create AMS config rules, configure an aggregator and tag resources.

- Amazon S3 – allows AMS Detective Controls to manage its S3 buckets.

You can download the JSON policy file in this ZIP: [DetectiveControlsConfig_ServiceRolePolicy.zip](#).

AWS managed policy: AWSManagedServicesDeploymentToolkitPolicy

AWS Managed Services (AMS) uses the `AWSManagedServicesDeploymentToolkitPolicy` AWS managed policy. This AWS-managed policy is attached to the [AWSServiceRoleForAWSManagedServicesDeploymentToolkit service-linked role](#), (see [Deployment toolkit service-linked role for AMS Accelerate](#)). The policy allows the service-linked role to complete actions for you. You can't attach this policy to your IAM entities. For more information, see [Using service-linked roles for AMS Accelerate](#).

For updates to the `AWSServiceRoleForManagedServicesDeploymentToolkitPolicy` service-linked role, see [Accelerate updates to service-linked roles](#).

Permissions details

This policy has the following permissions to allow AWS Managed Services Detective Controls to deploy and configure all necessary resources.

- CloudFormation – Allows AMS Deployment Toolkit to deploy CFN stacks with S3 resources required by CDK.
- Amazon S3 – allows AMS Deployment Toolkit to manage its S3 buckets.
- Elastic Container Registry – allows AMS Deployment Toolkit to manage its ECR repository that is used to deploy assets needed by AMS CDK apps.

You can download the JSON policy file in this ZIP:

[AWSManagedServicesDeploymentToolkitPolicy.zip](#).

AWS managed policy: AWSManagedServices_EventsServiceRolePolicy

AWS Managed Services (AMS) uses the `AWSManagedServices_EventsServiceRolePolicy` AWS managed policy. This AWS-managed policy is attached to the [AWSServiceRoleForManagedServices_Events service-linked role](#). The policy allows the service-linked role to complete actions for you. You can't attach this policy to your IAM entities. For more information, see [Using service-linked roles for AMS Accelerate](#).

For updates to the `AWSServiceRoleForManagedServices_Events` service-linked role, see [Accelerate updates to service-linked roles](#).

Permissions details

This policy has the following permissions to allow Amazon EventBridge to deliver alarm state change information from your account to AWS Managed Services.

- `events` – Allows Accelerate to create Amazon 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.

You can download the JSON policy file in this ZIP: [EventsServiceRolePolicy.zip](#).

AWS managed policy: `AWSManagedServices_ContactsServiceRolePolicy`

AWS Managed Services (AMS) uses the `AWSManagedServices_ContactsServiceRolePolicy` AWS managed policy. This AWS-managed policy is attached to the [AWSServiceRoleForManagedServices_Contacts service-linked role](#), (see [Creating a Contacts SLR for AMS Accelerate](#)). The policy allows the AMS Contacts SLR to look at your resource tags, and their values, on AWS resources. You can't attach this policy to your IAM entities. For more information, see [Using service-linked roles for AMS Accelerate](#).

Important

Do not store personally identifiable information (PII) or other confidential or sensitive information in tags. AMS uses tags to provide you with administration services. Tags are not intended to be used for private or sensitive data.

For updates to the `AWSServiceRoleForManagedServices_Contacts` service-linked role, see [Accelerate updates to service-linked roles](#).

Permissions details

This policy has the following permissions to allow the Contacts SLR to read your resource tags to retrieve resource contact information that you have set up ahead of time.

- `IAM` – Allows Contacts service to look at tags on IAM Roles and IAM users.
- `Amazon EC2` – Allows Contacts service to look at tags on Amazon EC2 resources.

- Amazon S3 – Allows Contacts Service to look at tags on Amazon S3 buckets. This action uses a Condition to ensure AMS accesses your bucket tags using the HTTP Authorization header, using the SigV4 signature protocol, and using HTTPS with TLS 1.2 or greater. For more information, see [Authentication Methods](#) and [Amazon S3 Signature Version 4 Authentication Specific Policy Keys](#).
- Tag – Allows Contacts service to look at tags on other AWS resources.
- "iam:ListRoleTags", "iam:ListUserTags", "tag:GetResources", "tag:GetTagKeys", "tag:GetTagValues", "ec2:DescribeTags", "s3:GetBucketTagging"

You can download the JSON policy file in this ZIP: [ContactsServicePolicy.zip](#).

Accelerate updates to AWS managed policies

View details about updates to AWS managed policies for Accelerate since this service began tracking these changes.

Change	Description	Date
Updated policy – Deployment Toolkit	<ul style="list-style-type: none"> • These new permissions were added for resource <code>arn:aws:ecr:*:*:repository/ams-cdktoolkit* :</code> <pre>ecr:BatchGetRepositoryScanningConfiguration ecr:PutImageScanningConfiguration</pre>	April 4, 2024
Updated policy – Deployment Toolkit	<ul style="list-style-type: none"> • These new permissions were added for resource <code>arn:aws:cloudformation:*:*:stack/ams-cdk-toolkit* :</code> <pre>cloudformation:DeleteChangeSet cloudformation:DescribeStackEvents cloudformation:GetTemplate cloudformation:TagResource cloudformation:UntagResource</pre>	May 9, 2023

Change	Description	Date
	<ul style="list-style-type: none"> These new permissions were added for resource <code>arn:aws:ecr:*:*:repository/ams-cdktoolkit*</code> : <div data-bbox="435 380 1219 932" style="border: 1px solid #ccc; border-radius: 10px; padding: 10px; margin: 10px 0;"> <pre>ecr:CreateRepository ecr:DeleteLifecyclePolicy ecr:DeleteRepository ecr:DeleteRepositoryPolicy ecr:DescribeRepositories ecr:GetLifecyclePolicy ecr:ListTagsForResource ecr:PutImageTagMutability ecr:PutLifecyclePolicy ecr:SetRepositoryPolicy ecr:TagResource ecr:UntagResource</pre> </div> Also, some existing actions with wildcard were scoped down to individual actions: <div data-bbox="435 1073 1219 1879" style="border: 1px solid #ccc; border-radius: 10px; padding: 10px; margin: 10px 0;"> <pre>- s3:DeleteObject* + s3:DeleteObject + s3:DeleteObjectTagging + s3:DeleteObjectVersion + s3:DeleteObjectVersionTagging - s3:GetObject* + s3:GetObject + s3:GetObjectAcl + s3:GetObjectAttributes + s3:GetObjectLegalHold + s3:GetObjectRetention + s3:GetObjectTagging + s3:GetObjectVersion + s3:GetObjectVersionAcl + s3:GetObjectVersionAttributes + s3:GetObjectVersionForReplication + s3:GetObjectVersionTagging + s3:GetObjectVersionTorrent</pre> </div> 	

Change	Description	Date
	<pre>- cloudformation:UpdateTermination* + cloudformation:UpdateTerminationProtection</pre>	
Updated policy – Detective Controls	<ul style="list-style-type: none"> The CloudFormation actions have been scoped down further after confirmation with security and access team The Lambda actions have been removed from the policy as they don't impact onboarding/off boarding 	April 10, 2023
Updated policy – Detective Controls	The ListAttachedRolePolicies action is removed from the policy. The action had Resource as wildcard (*). As "list" is a non-mutative action, it is given access over all resources, and the wildcard is disallowed.	March 28, 2023
Updated policy – Detective Controls	Updated the policy and added the permissions boundary policy.	March 21, 2023
New policy – Contacts Service	<p>Accelerate added a new policy to look at your account contact information from your resource tags.</p> <p>Accelerate added a new policy to read your resource tags so that it can retrieve the resource contact information that you have set up ahead of time.</p>	February 16, 2023
New policy – Events Service	<p>Accelerate added a new policy to deliver alarm state change information from your account to AWS Managed Services.</p> <p>Grants IAM roles created as part of How Alarm Manager works permissions to create a required Amazon EventBridge managed rule.</p>	February 07, 2023

Change	Description	Date
Updated policy – Deployment Toolkit	Added S3 permissions to support customer offboarding from Accelerate.	January 30, 2023
New policy – Detective Controls	Allows the service-linked role, Detective controls service-linked role for AMS Accelerate , to complete actions for you to deploy Accelerate detective controls.	December 19, 2022
New policy – Alarm Manager	Accelerate added a new policy to allow permissions to perform alarm manager tasks. Grants IAM roles created as part of How Alarm Manager works permissions to perform operations like AWS Config evaluation, AWS Config read to fetch alarm manager configuration, creation of necessary Amazon CloudWatch alarms.	November 30, 2022
Accelerate started tracking changes	Accelerate started tracking changes for its AWS managed policies.	November 30, 2022
New policy – Deployment Toolkit	Accelerate added this policy for deployment tasks. Grants the service-linked role AWSServiceRoleForAWSManagedServicesDeploymentToolkit permissions to access and update deployment-related Amazon S3 buckets and AWS CloudFormation stacks.	June 09, 2022

Data protection in AMS Accelerate

AMS Accelerate 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 Accelerate assumes responsibility for the initial triage and response to the alarm. Our response processes are based on NIST standards. AMS Accelerate regularly tests its response processes using Security Incident Response Simulation with you to align your workflow with existing customer security response programs.

When AMS Accelerate 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 Accelerate 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 Accelerate 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 Accelerate 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 Accelerate 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](#).

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 AWS 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, see 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 suppression rules

A suppression rule is a set of criteria, consisting of a filter attribute paired with a value, used to filter findings by automatically archiving new findings that match the specified criteria. Suppression rules can be used to filter low-value findings, false positive findings, or known activities, you do not intend to act on, to make it easier to recognize the security threats with the most impact to your environment.

AMS has a defined set of criteria to identify suppression rules for your managed accounts. AMS implements suppression rule to filter false positive findings and reduce frequent unactionable notifications. When a managed account meets this criteria, AMS applies the filters and notifies you with the details of the suppression filter deployed through a Service Request (SR).

You can communicate with AMS via a Service Request (SR) to modify or revert the suppression filters.

Suppressed findings are not sent to AWS Security Hub, Amazon S3, or CloudTrail Events, reducing finding of unactionable data if you consume GuardDuty findings via Security Hub or a third-party SIEM, alerting and ticketing applications.

GuardDuty continues to generate findings even when they match your suppression rules, however, those findings automatically marked as **archived**. The archived finding is stored in GuardDuty for 90-days and can be viewed at any time during that period. You can view suppressed findings in the GuardDuty console by selecting **Archived** from the findings table, or through the GuardDuty API using the [ListFindings](#) API with a **findingCriteria** of **service.archived equal to true**.

Common Use Cases for Suppression Rules:

The following are finding types with common use cases for applying suppression rules, select the finding name to learn more about how to apply a suppression rules for that use case.

- **Recon:EC2/Portscan**: Use a suppression rule to automatically archive findings when using an authorized vulnerability scanner.

- **UnauthorizedAccess:EC2/SSHBruteForce:** Use a suppression rule to automatically archive findings when it is targeted to bastion instances.
- **Recon:EC2/PortProbeUnprotectedPort:** Use a suppression rule to automatically archive findings when it is targeted to intentionally exposed instances.

Data encryption in AMS Accelerate

AMS Accelerate uses several AWS services for data encryption.

Amazon Simple Storage Service 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](#).

AWS Identity and Access Management in AMS Accelerate

AWS Identity and Access Management 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 Accelerate onboarding, you are responsible for creating cross-account IAM administrator roles within each of your managed accounts.

With AMS Accelerate, you're responsible for managing access to your AWS accounts and their underlying resources, such as access management solutions, access policies, and related processes. This means that you manage your user lifecycle, permissions in directory services, and federated authentication system, to access the AWS console or AWS APIs. In order to help you manage your access solution, AMS Accelerate deploys AWS Config rules that detect common IAM misconfigurations, and then deliver remediation notifications. For more information, see [AWS Config Managed Rules](#).

Authenticating with identities in AMS Accelerate

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 permissions 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. AMS uses the following IAM roles:

Role name	Description
Used by (entity): AMS Access Service only	
ams-access-management	Deployed manually by you during onboarding. Assumed only by AMS access to deploy or update access roles. Remains in your account after onboarding for any future updates to the access roles.
Used by (entity): AMS Operations	
ams-access-admin-operations	This role has administrative permissions to operate in accounts, but does not have permissions to read, write, or delete customer content in AWS services commonly used as data stores, such as Amazon Simple Storage Service, Amazon Relational Database Service, Amazon DynamoDB, Amazon Redshift, and Amazon ElastiCache. Only a very few select AMS individuals can assume this role.
ams-access-operations	This AMS Operations role has permissions to perform administrative tasks in your accounts. This role does not have read, write, or delete permissions to customer content in AWS services commonly used as data stores, such as Amazon Simple Storage Service, Amazon Relational Database Service, Amazon DynamoDB, Amazon Redshift, and Amazon ElastiCache. Permissions to perform AWS Identity and Access Management write operations are also excluded from this role.

Role name	Description
ams-access-read-only	This AMS read-only role is limited to read-only permissions in your AMS account. Read permissions to customer content in AWS services commonly used as data stores, such as Amazon S3, Amazon RDS, DynamoDB, Amazon Redshift, and ElastiCache, are not granted by this role.
Used by (entity): AMS Operations and AMS Services	
ams_ssm_automation_role	Assumed by AWS Systems Manager to execute SSM Automation documents within your account.
ams_ssm_automation_role	
Used by (entity): AMS Security	
ams-access-security-analyst	This AMS security role has permissions in your AMS account to perform dedicated security alert monitoring and security incident handling. Only a very few select AMS Security individuals can assume this role.
ams-access-security-analyst-read-only	This AMS security role is limited to read-only permissions in your AMS account to perform dedicated security alert monitoring and security incident handling.
Used by (entity): AWS Services	
ams-access-admin	This AMS admin role has full permissions to operate in accounts without restrictions. Only AMS internal services (with a scoped-down session policy) can assume the admin role.

Role name	Description
ams-opscenter-eventbridge-role	Assumed by Amazon EventBridge to create AWS Systems Manager OpsItems as a part of AMS-specific AWS Config Rules remediation workflow.
AMSOSConfigurationCustomerInstanceRole	This IAM role is applied to your Amazon EC2 instances when AMS OS-Configuration service discovers that the required IAM policies are missing. It allows your Amazon EC2 instances to interact with AWS Systems Manager, Amazon CloudWatch, and Amazon EventBridge services. It also has attached the AMS custom-managed policy to enable RDP access to your Windows instances.
mc-patch-glue-service-role	Assumed by AWS Glue ETL workflow to perform data transformation and prepare it for AMS Patch report generator.
Used by (entity): AMS Service	
ams-alarm-manager-AWSManagedServicesAlarmManagerDe-<8-digit hash>	Assumed by AMS alarm manager infrastructure within your AMS account to perform AWS Config Rules evaluation for a new AWS AppConfig deployment.
ams-alarm-manager-AWSManagedServicesAlarmManagerRe-<8-digit hash>	Assumed by AMS alarm manager remediation infrastructure within your AMS account to allow the creation or deletion of alarms for remediation.
ams-alarm-manager-AWSManagedServicesAlarmManagerSS-<8-digit hash>	Assumed by AWS Systems Manager to invoke the AMS alarm manager remediation service within your AMS account.

Role name	Description
ams-alarm-manager-AWSManagedServicesAlarmManagerTr-<8-digit hash>	Assumed by AMS alarm manager infrastructure within your AWS account to conduct periodic AMS AWS Config Rules evaluation.
ams-alarm-manager-AWSManagedServicesAlarmManagerVa-<8-digit hash>	Assumed by AMS alarm manager infrastructure within your AMS account to ensure that the required alarms exists in the AWS account.
ams-backup-iam-role	This role is used to run AWS Backup within your accounts.
ams-monitoring-AWSManagedServicesLogGroupLimitLamb-<8-digit hash>	Assumed by AMS Logging & Monitoring infrastructure in your AMS account to evaluate Amazon CloudWatch Logs groups limit and compare with the service quotas.
ams-monitoring-AWSManagedServicesRDSMonitoringRDSE-<8-digit hash>	Assumed by AMS Logging & Monitoring infrastructure in your AMS account to forward Amazon RDS events to Amazon CloudWatch Events.
ams-monitoring-AWSManagedServicesRedshiftMonitorin-<8-digit hash>	Assumed by AMS Logging & Monitoring infrastructure in your AMS account to forward Amazon Redshift events (CreateCluster and DeleteCuster) to Amazon CloudWatch Events.
ams-monitoring-infrastruc-AWSManagedServicesMonito-<8-digit hash>	Assumed by AMS Logging & Monitoring infrastructure in your AMS account to publish messages to Amazon Simple Notification Service to validate that the account is reporting all necessary data.
ams-opscenter-role	Assumed by AMS Notification Management system in your AMS account to manage AWS Systems Manager OpsItems related to alerts in your account.

Role name	Description
ams-opsitem-autoexecution-role	Assumed by AMS Notification Management system to handle automated remediation using SSM documents for monitoring alerts related to resources in your account.
ams-patch-infrastructure-ampatchconfigruleroleC1-<8-digit hash>	Assumed by AWS Config to evaluate AMS patch resources and detect drift in its AWS CloudFormation stacks.
ams-patch-infrastructure-ampatchcwruleopsitemams-<8-digit hash>	Assumed by Amazon EventBridge to create AWS Systems Manager OpsItems for patching failures.
ams-patch-infrastructure-ampatchservicebusamspat-<8-digit hash>	Assumed by Amazon EventBridge to send an event to the AMS Patch orchestrator event bus for AWS Systems Manager Maintenance Windows state change notifications.
ams-patch-reporting-infra-ampatchreportingconfigr-<8-digit hash>	Assumed by AWS Config to evaluate AMS Patch reporting resources and detect drift in its AWS CloudFormation stacks.
ams-resource-tagger-AWSManagedServicesResourceTagg-<8-digit hash>	Assumed by AMS Resource Tagger infrastructure within your AMS account to perform AWS Config Rules evaluation upon new AWS AppConfig deployment.
ams-resource-tagger-AWSManagedServicesResourceTagg-<8-digit hash>	Assumed by AMS Resource Tagger infrastructure within your AMS account to validate that required AWS tags exist for the managed resources.
ams-resource-tagger-AWSManagedServicesResourceTagg-<8-digit hash>	Assumed by AWS Systems Manager to invoke AMS Resource Tagger remediation workflow in your AMS account.

Role name	Description
ams-resource-tagger-AWSManagedServicesResourceTagg-<8-digit hash>	Assumed by AMS Resource Tagger remediation infrastructure within your AMS account to create or delete AWS tags for the managed resources.
ams-resource-tagger-AWSManagedServicesResourceTagg-<8-digit hash>	Assumed by AMS Resource Tagger infrastructure within your AWS account to conduct periodic AMS Config Rule evaluation.
ams_os_configuration_event_rule_role-<AWS Region>	Assumed by Amazon EventBridge to forward events from your account to AMS OS-Configuration service EventBus in the correct Region.
mc-patch-reporting-service	Assumed by AMS patch data aggregator and report generator.

Note

This is the template for the `ams-access-management` role. It is the stack that cloud architects (CAs) manually deploy in your account at onboarding time: [management-role.yaml](#).

This is the template for the different access roles for the different access levels: `ams-access-read-only`, `ams-access-operations`, `ams-access-admin-operations`, `ams-access-admin`: [accelerate-roles.yaml](#).

To learn more about AWS Cloud Development Kit (CDK) identifiers, including hashes, see [UniqueIDs](#).

AMS Accelerate feature services assume the **ams-access-admin** role for programmatic access to the account, but with a session policy scoped down for the respective feature service (for example, patch, backup, monitoring, and so forth).

AMS Accelerate follows industry best practices to meet and maintain compliance eligibility. AMS Accelerate access to your account is recorded in CloudTrail and also available for your review

through change tracking. For information about queries that you can use to get this information, see [Tracking changes in your AMS Accelerate accounts](#).

Managing access using policies

Various AMS Accelerate support teams such as Operations Engineers, Cloud Architects, and Cloud Service Delivery Managers (CSDMs), sometimes require access to your accounts in order to respond to service requests and incidents. Their access is governed by an internal AMS access service that enforces controls, such as business justification, service requests, operations items, and support cases. The default access is read-only, and all access is tracked and recorded; see also [Tracking changes in your AMS Accelerate accounts](#).

Validation of IAM resources

The AMS Accelerate access system periodically assumes roles in your accounts (at least every 24 hours) and validates that all of our IAM resources are as expected.

In order to protect your accounts, AMS Accelerate has a "canary" that monitors and alarms on the presence and status of the IAM roles, as well as their attached policies, mentioned above. Periodically, the canary assumes the **ams-access-read-only** role and initiates CloudFormation and IAM API calls against your accounts. The canary evaluates the status of the AMS Accelerate access roles to make sure they are always unmodified and up-to-date. This activity creates CloudTrail logs in the account.

The AWS Security Token Service (AWS STS) session name of the canary is **AMS-Access-Roles-Auditor-`{uuid4()}`** as seen in CloudTrail and the following API calls occur:

- Cloud Formation API Calls: `describe_stacks()`
- IAM API Calls:
 - `get_role()`
 - `list_attached_role_policies()`
 - `list_role_policies()`
 - `get_policy()`
 - `get_policy_version()`
 - `get_role_policy()`

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	CI	RA

Function	Customer	AMS
contact details to use during Security Events and Security Escalations		
Maintain Customer internal contact details, escalation 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
Training		
Provide customer with documentation to support AMS during Incident Response Process	I	RA
Practice shared responsibility 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

Function	Customer	AMS
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, configurations, policies and access management	RAC	I

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
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
Define, deploy, and maintain AMS standard detection services (for example, Amazon GuardDuty, Amazon Macie, and AWS Config)	CI	RA
Record AWS infrastructure change logs	RA	I

Function	Customer	AMS
Enable and configure logging, monitoring to enable event management for the application	RAI	C
Implement and maintain an allow-list, deny-list, and custom detections on AWS security services	RA	RCI
Implement and maintain allow-list, deny-list, and custom detections on other security services supported by AMS	RACI	-
Security event reporting		
Notify AMS of a suspicious activity or an active security investigation	RA	CI
Security event and incident notification to the customer	CI	RA
Notification of planned event that might trigger Security Incident Response process	RA	I

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		

Function	Customer	AMS
Perform initial response for any security alert generated by a detection source	I	RAC
False/True positive assessment based on available data	RAI	RAC
Generate a snap shot of affected instances to be shared with the customer if needed	I	RA
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 investigations	CI	R
Engage third party vendors during investigation (for example, for EPS antimalware investigation and engaging with TrendMicro support team)	RACI	-

Function	Customer	AMS
Share investigation logs to customers during an investigation	AI	R
Communication		
Alert and notifications from AMS detection sources for managed resources	I	RA
Alert and notifications on application security events	RA	I
Engage customer leadership and C-Level executives during a security incident investigation	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 applicati	RA	CI

Function	Customer	AMS
on specific configuration or response activity)		
Define the containment strategy based on the security incident and the affected resource	ACI	RA
Encryption and secure storage of point in time backups of affected systems	RACI	C
Contain AWS Resources including EC2 Instance, Network and IAM containment.	ACI	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		
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	ACI	RA
Eradicate and harden AWS resources including Amazon EC2 instances, network, and IAM eradication.	ACI	RA
Eradicate and harden applications and workloads (through application specific	RA	-

Function	Customer	AMS
configuration or response activity)		

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

Function	Customer	AMS
Back up customer application, APP configuration, and deployment settings	RA	-
Review backup plans to fully restore customer applications and workloads post-incident	RA	-
Restore applications and customer workloads (through application specific restoration steps)	RA	-

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, Network Traffic, and Amazon internal Threat Intelligence feeds.

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. 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.

Phase 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

Accounts enrolled in AMS Accelerate are configured with a baseline deployment of CloudWatch [Events](#) and [Alarms](#) that have been optimized to reduce noise and to identify indications of a true incident. AMS Accelerate also employs GuardDuty for account monitoring; see [GuardDuty](#) for more detail.

Configuration compliance in Accelerate

AMS Accelerate helps you configure your resources to high standards for security and operational integrity, and comply with the following industry standards:

- Center for Internet Security (CIS)
- National Institute of Standards and Technology (NIST) Cloud Security Framework (CSF)
- Health Insurance Portability and Accountability Act (HIPAA)
- Payment Card Industry (PCI) Data Security Standard (DSS)

We do this by deploying our entire compliance AWS Config rule set to your account, see [AMS Config Rule library](#). An AWS Config rule represents desired configurations for a resource and is evaluated against configuration changes on the settings of your AWS resources. Any configuration change triggers a large number of rules to test compliance. For example, suppose you create an Amazon S3 bucket, and configure it to be publicly readable, in violation of NIST standards. The [ams-nist-cis-s3-bucket-public-read-prohibited rule](#) detects the violation and labels your S3 bucket **Noncompliant** in your Configuration Report. Because this rule belongs to the **Auto Incident** remediation category, it immediately creates a Incident Report, alerting you to the issue. Other more severe rule violations might cause AMS to automatically remediate the issue. See [Responses to violations in Accelerate](#).

⚠ Important

If you want us to do more, for example, if you want AMS to remediate a violation for you, regardless of its remediation category, submit a Service Request that asks AMS to remediate the noncompliant resources for you. In the Service Request, include a comment such as "As part of the AMS config rule remediation, please remediate non-complaint resources *RESOURCE_ARNS_OR_IDS*, config rule *CONFIG_RULE_NAME* in the account" and add the required inputs to remediate the violation.

If you want us to do less, for example, if you don't want us to take action on a particular S3 bucket that requires public access by design, you can create exceptions, see [Creating rule exceptions in Accelerate](#).

AMS Config Rule library

Accelerate deploys a library of AMS config rules to protect your account. These config rules begin with `ams-`. You can view rules within your account, and their compliance state, from either the AWS Config console, AWS CLI, or the AWS Config API. For general information about using AWS Config, see [Viewing Configuration Compliance](#).

📘 Note

For opt-in AWS Regions, and gov cloud Regions, we only deploy a subset of the config rules due to Region restrictions. Check the rule availability in Regions by checking the link associated to the identifier in the AMS Accelerate config rules table.

You cannot remove any of the deployed AMS Config Rules at this time.

Table of Rules

Download as [ams_config_rules.zip](#).

AMS Configuration Rules

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-guardduty-	GuardDuty	Periodic	Remediate	CIS: CIS.13,CIS.14; NIST-CSF: PR.DS-1;

Rule Name	Service	Trigger	Action	Frameworks
enabled-centralized				HIPAA: 164.312(a)(2)(iv),164.312(e)(2)(ii); PCI: 2.2,3.4,8.2.1;
ams-nist-cis-vpc-flow-logs-enabled	VPC	Periodic	Remediate	CIS: CIS.6; NIST-CSF: DE.AE-1,DE.AE-3,PR.DS-5,PR.PT-1; HIPAA: 164.308(a)(3)(ii)(A),164.312(b); PCI: 2.2,10.1,10.3.2,10.3.3,10.3.4,10.3.5,10.3.6;
ams-eks-secrets-encrypted	EKS	Periodic	Incident	CIS: NA; NIST-CSF: NA; HIPAA: NA; PCI: NA;
ams-eks-endpoint-no-public-access	EKS	Periodic	Incident	CIS: NA; NIST-CSF: NA; HIPAA: NA; PCI: NA;
ams-nist-cis-vpc-default-security-group-closed	VPC	Config Changes	Incident	CIS: CIS.11,CIS.12,CIS.9; NIST-CSF: DE.AE-1,PR.AC-3,PR.AC-5,PR.PT-4; HIPAA: 164.312(e)(1); PCI: 1.2,1.3,2.1,2.2,1.2.1,1.3.1,1.3.2,2.2.2;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-iam-password-policy	IAM	Periodic	Incident	CIS: NA; NIST-CSF: PR.AC-1,PR.AC-4; HIPAA: 164.308(a)(3)(i),164.308(a)(3)(ii)(A),164.308(a)(3)(ii)(B),164.308(a)(4)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(B),164.308(a)(4)(ii)(C),164.312(a)(1); PCI: 7.1.2,7.1.3,7.2.1,7.2.2;
ams-nist-cis-iam-root-access-key-check	IAM	Periodic	Incident	CIS: CIS.16,CIS.4; NIST-CSF: PR.AC-1,PR.AC-4,PR.PT-3; HIPAA: 164.308(a)(3)(i),164.308(a)(3)(ii)(A),164.308(a)(3)(ii)(B),164.308(a)(4)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(B),164.308(a)(4)(ii)(C),164.312(a)(1); PCI: 2.2,7.1.2,7.1.3,7.2.1,7.2.2;
ams-nist-cis-iam-user-mfa-enabled	IAM	Periodic	Incident	CIS: CIS.16; NIST-CSF: PR.AC-1,PR.AC-4; HIPAA: 164.308(a)(3)(i),164.308(a)(3)(ii)(A),164.308(a)(3)(ii)(B),164.308(a)(4)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(B),164.308(a)(4)(ii)(C),164.312(a)(1); PCI: 2.2,7.1.2,7.1.3,7.2.1,7.2.2;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-restricted-ssh	Security Groups	Config Changes	Incident	CIS: CIS.16; NIST-CSF: PR.AC-1,PR.AC-4; HIPAA: 164.308(a)(3)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(B),164.308(a)(4)(ii)(C),164.312(a)(1); PCI: 2.2,7.2.1,8.1.4;
ams-nist-cis-restricted-com-mon-ports	Security Groups	Config Changes	Incident	CIS: CIS.11,CIS.12,CIS.9; NIST-CSF: DE.AE-1,P R.AC-3,PR.AC-5,PR.PT-4; HIPAA: 164.308(a)(3)(i),164.308(a)(3)(ii)(B),164.308(a)(4)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(B),164.308(a)(4)(ii)(C),164.312(a)(1),164.312(e)(1); PCI: 1.2,1.3,2.2,1.2.1,1.3.1,1.3.2,2.2.2;
ams-nist-cis-s3-account-level-public-access-blocks	S3	Config Changes	Incident	CIS: CIS.9,CIS.12,CIS.14; NIST-CSF: PR.AC-3,PR.AC-4,PR.AC-5,PR.DS-5,PR.PT-3,PR.PT-4; HIPAA: 164.308(a)(3)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(C),164.312(a)(1),164.312(e)(1); PCI: 1.2,1.2.1,1.3,1.3.1,1.3.2,1.3.4,1.3.6,2.2,2.2.2;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-s3-bucket-public-read-prohibited	S3	Config Changes	Incident	CIS: CIS.12,CIS.14,CIS.9; NIST-CSF: PR.AC-3,P R.AC-4,PR.AC-5,PR.DS-5,PR.PT-3,PR.PT-4; HIPAA: 164.308(a)(3)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(C),164.312(a)(1),164.312(e)(1); PCI: 1.2,1.3,2.2,1.2.1,1.3.1,1.3.2,1.3.4,1.3.6,2.2.2;
ams-nist-cis-s3-bucket-public-write-prohibited	S3	Config Changes	Incident	CIS: CIS.12,CIS.14,CIS.9; NIST-CSF: PR.AC-3,P R.AC-4,PR.AC-5,PR.DS-5,PR.PT-3,PR.PT-4; HIPAA: 164.308(a)(3)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(C),164.312(a)(1),164.312(e)(1); PCI: 1.2,1.3,2.2,1.2.1,1.3.1,1.3.2,1.3.4,1.3.6,2.2.2;
ams-nist-cis-s3-bucket-server-side-encryption-enabled	S3	Config Changes	Incident	CIS: CIS.13,CIS.14; NIST-CSF: PR.DS-1; HIPAA: 164.312(a)(2)(iv),164.312(c)(2),164.312(e)(2)(ii); PCI: 2.2,3.4,10.5,8.2.1;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-securityhub-enabled	Security Hub	Periodic	Incident	CIS: CIS.3,CIS.4,CIS.6, CIS.12,CIS.16,CIS.19; NIST-CSF: PR.DS-5,P R.PT-1; HIPAA: 164.312(b); PCI: NA;
ams-nist-cis-ec2-instance-managed-by-systems-manager	EC2	Config Changes	Report	CIS: CIS.2,CIS.5; NIST-CSF: ID.AM-2,PR.IP-1; HIPAA: 164.308(a)(5)(ii)(B); PCI: 2.4;
ams-nist-cis-cloudtrail-enabled	CloudTrail	Periodic	Report	CIS: CIS.16,CIS.6; NIST-CSF: DE.AE-1,DE.AE-3,PR.DS-5,PR.MA-2,PR.PT-1; HIPAA: 164.308(a)(3)(ii)(A),164.308(a)(5)(ii)(C),164.312(b); PCI: 10.1,10.2.1,10.2.2,10.2.3,10.2.4,10.2.5,10.2.6,10.2.7,10.3.1,10.3.2,10.3.3,10.3.4,10.3.5,10.3.6;
ams-nist-cis-access-keys-rotated	IAM	Periodic	Report	CIS: CIS.16; NIST-CSF: PR.AC-1; HIPAA: 164.308(a)(4)(ii)(B); PCI: 2.2;
ams-nist-cis-acm-certificate-expiration-check	Certificate Manager	Config Changes	Report	CIS: CIS.13,CIS.14; NIST-CSF: PR.AC-5,PR.PT-4; HIPAA: NA; PCI: 4.1;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-alb-http-to-https-redirection-check	ALB	Periodic	Report	CIS: CIS.13,CIS.14; NIST-CSF: PR.DS-2; HIPAA: 164.312(a)(2)(iv), 164.312(e)(1),164.312(e)(2)(i),164.312(e)(2)(ii); PCI: 2.3,4.1,8.2.1;
ams-nist-cis-api-gw-cache-enabled-and-encrypted	API Gateway	Config Changes	Report	CIS: CIS.13,CIS.14; NIST-CSF: PR.DS-1; HIPAA: 164.312(a)(2)(iv), 164.312(e)(2)(ii); PCI: 3.4;
ams-nist-cis-api-gw-execution-logging-enabled	API Gateway	Config Changes	Report	CIS: CIS.6; NIST-CSF: DE.AE-1,DE.AE-3,PR.PT-1; HIPAA: 164.312(b); PCI: 10.1,10.3.1,10.3.2,10.3.3,10.3.4,10.3.5,10.3.6,10.5.4;
ams-nist-autoscaling-group-elb-healthcheck-required	ELB	Config Changes	Report	CIS: NA; NIST-CSF: PR.PT-1,PR.PT-5; HIPAA: 164.312(b); PCI: 2.2;
ams-nist-cis-cloud-trail-encryption-enabled	CloudTrail	Periodic	Report	CIS: CIS.13,CIS.14; NIST-CSF: PR.DS-1; HIPAA: 164.312(a)(2)(iv),164.312(e)(2)(ii); PCI: 2.2,3.4,10.5;
ams-nist-cis-cloud-trail-log-file-validation-enabled	CloudTrail	Periodic	Report	CIS: CIS.6; NIST-CSF: PR.DS-6; HIPAA: 164.312(c)(1),164.312(c)(2); PCI: 2.2,10.5,11.5,10.5.2,10.5.5;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-cloudtrail-s3-dataevents-enabled	CloudTrail	Periodic	Report	CIS: CIS.6; NIST-CSF: DE.AE-1,DE.AE-3,PR.DS-5,PR.PT-1; HIPAA: 164.308(a)(3)(ii)(A),164.312(b); PCI: 2.2,10.1,10.2.1,10.2.2,10.2.3,10.2.5,10.3.1,10.3.2,10.3.3,10.3.4,10.3.5,10.3.6;
ams-nist-cis-cloudwatch-alarm-action-check	CloudWatch	Config Changes	Report	CIS: CIS.13,CIS.14; NIST-CSF: NA; HIPAA: 164.312(a)(2)(iv),164.312(e)(2)(ii); PCI: 3.4;
ams-nist-cis-cloudwatch-log-group-encrypted	CloudWatch	Periodic	Report	CIS: CIS.13,CIS.14; NIST-CSF: NA; HIPAA: 164.312(a)(2)(iv),164.312(e)(2)(ii); PCI: 3.4;
ams-nist-cis-codebuild-project-envvar-awscred-check	CodeBuild	Config Changes	Report	CIS: CIS.18; NIST-CSF: PR.DS-5; HIPAA: 164.308(a)(3)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(C),164.312(a)(1); PCI: 8.2.1;
ams-nist-cis-codebuild-project-source-repo-url-check	CodeBuild	Config Changes	Report	CIS: CIS.18; NIST-CSF: PR.DS-5; HIPAA: 164.308(a)(3)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(C),164.312(a)(1); PCI: 8.2.1;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-db-instance-backup-enabled	RDS	Config Changes	Report	CIS: CIS.10; NIST-CSF: ID.BE-5,PR.DS-4,PR.IP-4,PR.PT-5,RC.RP-1; HIPAA: 164.308(a)(7)(i),164.308(a)(7)(ii)(A),164.308(a)(7)(ii)(B); PCI: NA;
ams-nist-cis-dms-replication-not-public	DMS	Periodic	Report	CIS: CIS.12,CIS.14,CIS.9; NIST-CSF: PR.AC-3,PR.AC-4,PR.AC-5,PR.DS-5,PR.PT-3,PR.PT-4; HIPAA: 164.308(a)(3)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(C),164.312(a)(1),164.312(e)(1); PCI: 1.2,1.3,1.2.1,1.3.1,1.3.2,1.3.4,1.3.6,2.2.2;
ams-nist-dynamodb-autoscaling-enabled	DynamoDB	Periodic	Report	CIS: NA; NIST-CSF: ID.BE-5,PR.DS-4,PR.PT-5,RC.RP-1; HIPAA: 164.308(a)(7)(i),164.308(a)(7)(ii)(C); PCI: NA;
ams-nist-cis-dynamodb-pitr-enabled	DynamoDB	Periodic	Report	CIS: CIS.10; NIST-CSF: ID.BE-5,PR.DS-4,PR.IP-4,PR.PT-5,RC.RP-1; HIPAA: 164.308(a)(7)(i),164.308(a)(7)(ii)(A),164.308(a)(7)(ii)(B); PCI: NA;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-dynamodb-throughput-limit-check	DynamoDB	Periodic	Report	CIS: NA; NIST-CSF: NA; HIPAA: 164.312(b); PCI: NA;
ams-nist-efs-optimized-instance	EBS	Config Changes	Report	CIS: NA; NIST-CSF: NA; HIPAA: 164.308(a)(7)(i); PCI: NA;
ams-nist-cis-efs-snapshot-public-restorable-check	EBS	Periodic	Report	CIS: CIS.12,CIS.14,CIS.9; NIST-CSF: PR.AC-3,PR.AC-4,PR.AC-5,PR.DS-5,PR.PT-3,PR.PT-4; HIPAA: 164.308(a)(3)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(C),164.312(a)(1),164.312(e)(1); PCI: 1.2,1.3,1.2.1,1.3.1,1.3.2,1.3.4,1.3.6,2.2.2;
ams-nist-ec2-instance-detail-monitoring-enabled	EC2	Config Changes	Report	CIS: NA; NIST-CSF: DE.AE-1,PR.PT-1; HIPAA: 164.312(b); PCI: NA;
ams-nist-cis-ec2-instance-no-public-ip	EC2	Config Changes	Report	CIS: CIS.12,CIS.14,CIS.9; NIST-CSF: PR.AC-3,PR.AC-4,PR.AC-5,PR.PT-3,PR.PT-4; HIPAA: 164.308(a)(3)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(C),164.312(a)(1),164.312(e)(1); PCI: 1.2,1.3,1.2.1,1.3.1,1.3.2,1.3.4,1.3.6,2.2.2;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-ec2-managedinstances-association-compliance-status-check	EC2	Config Changes	Report	CIS: CIS.12,CIS.9; NIST-CSF: PR.AC-3,PR.AC-4,PR.AC-5,PR.PT-3,PR.PT-4; HIPAA: 164.308(a)(3)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(C),164.312(a)(1),164.312(e)(1); PCI: 1.2,1.3,1.2.1,1.3.1,1.3.2,1.3.4,1.3.6,2.2.2;
ams-nist-cis-ec2-managedinstances-patch-compliance-status-check	EC2	Config Changes	Report	CIS: CIS.2,CIS.5; NIST-CSF: ID.AM-2,PR.IP-1; HIPAA: 164.308(a)(5)(ii)(B); PCI: 6.2;
ams-nist-cis-ec2-stopped-instance	EC2	Periodic	Report	CIS: CIS.2; NIST-CSF: ID.AM-2,PR.IP-1; HIPAA: NA; PCI: NA;
ams-nist-cis-ec2-volume-inuse-check	EC2	Config Changes	Report	CIS: CIS.2; NIST-CSF: PR.IP-1; HIPAA: NA; PCI: NA;
ams-nist-cis-efs-encrypted-check	EFS	Periodic	Report	CIS: CIS.13,CIS.14; NIST-CSF: PR.DS-1; HIPAA: 164.312(a)(2)(iv),164.312(e)(2)(ii); PCI: 3.4,8.2.1;
ams-nist-cis-eip-attached	EC2	Config Changes	Report	CIS: CIS.13,CIS.14; NIST-CSF: PR.DS-1; HIPAA: 164.312(a)(2)(iv),164.312(e)(2)(ii); PCI: 3.4,8.2.1;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-elasticache-redis-cluster-automatic-backup-check	ElastiCache	Periodic	Report	CIS: CIS.10; NIST-CSF: ID.BE-5,PR.DS-4,PR.IP-4,PR.PT-5,RC.RP-1; HIPAA: 164.308(a)(7)(i),164.308(a)(7)(ii)(A),164.308(a)(7)(ii)(B); PCI: NA;
ams-nist-cis-opensearch-encrypted-at-rest	OpenSearch	Periodic	Report	CIS: CIS.14,CIS.13; NIST-CSF: PR.DS-1; HIPAA: 164.312(a)(2)(iv),164.312(e)(2)(ii); PCI: 3.4,8.2.1;
ams-nist-cis-opensearch-in-vpc-only	OpenSearch	Periodic	Report	CIS: CIS.13,CIS.14; NIST-CSF: PR.DS-1; HIPAA: 164.312(a)(2)(iv),164.312(e)(2)(ii); PCI: 3.4,8.2.1;
ams-nist-cis-elb-acm-certificate-required	Certificate Manager	Config Changes	Report	CIS: CIS.12,CIS.9; NIST-CSF: PR.AC-3,PR.AC-4,PR.AC-5,PR.DS-5,PR.PT-3,PR.PT-4; HIPAA: 164.308(a)(3)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(C),164.312(a)(1),164.312(e)(1); PCI: 1.2,1.3,1.2.1,1.3.1,1.3.2,1.3.4,1.3.6,2.2.2;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-elb-deletion-protection-enabled	ELB	Config Changes	Report	CIS: CIS.13,CIS.14; NIST-CSF: PR.DS-2; HIPAA: 164.312(a)(2)(iv), 164.312(e)(1),164.312(e)(2)(i),164.312(e)(2)(ii); PCI: 4.1,8.2.1;
ams-nist-cis-elb-logging-enabled	ELB	Config Changes	Report	CIS: CIS.6; NIST-CSF: DE.AE-1,DE.AE-3,PR.PT-1; HIPAA: 164.312(b); PCI: 10.1,10.3.1,10.3.2,10.3.3,10.3.4,10.3.5,10.3.6,10.5.4;
ams-nist-cis-emr-kerberos-enabled	EMR	Periodic	Report	CIS: CIS.6; NIST-CSF: DE.AE-1,DE.AE-3,PR.PT-1; HIPAA: 164.312(b); PCI: 10.1,10.3.1,10.3.2,10.3.3,10.3.4,10.3.5,10.3.6,10.5.4;
ams-nist-cis-emr-master-no-public-ip	EMR	Periodic	Report	CIS: CIS.14,CIS.16; NIST-CSF: PR.AC-1,PR.AC-4,PR.AC-6; HIPAA: 164.308(a)(3)(i),164.308(a)(3)(ii)(A),164.308(a)(3)(ii)(B),164.308(a)(4)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(B),164.308(a)(4)(ii)(C),164.312(a)(1); PCI: 7.2.1;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-encrypted-volumes	EBS	Config Changes	Report	CIS: CIS.12,CIS.9; NIST-CSF: PR.AC-3,PR.AC-4,PR.AC-5,PR.PT-3,PR.PT-4; HIPAA: 164.308(a)(3)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(C),164.312(a)(1),164.312(e)(1); PCI: 1.2,1.3,1.2.1,1.3.1,1.3.2,1.3.4,1.3.6,2.2.2;
ams-nist-cis-guardduty-non-archived-findings	GuardDuty	Periodic	Report	CIS: CIS.12,CIS.13,CIS.16,CIS.19,CIS.3,CIS.4,CIS.6,CIS.8; NIST-CSF: DE.AE-2,DE.AE-3,DE.CM-4,DE.DP-5,ID.RA-1,ID.RA-3,PR.DS-5,PR.PT-1; HIPAA: 164.308(a)(5)(ii)(C),164.308(a)(6)(ii),164.312(b); PCI: 6.1,11.4,5.1.2;
ams-nist-iam-group-has-users-check	IAM	Config Changes	Report	CIS: NA; NIST-CSF: PR.AC-4,PR.AC-1; HIPAA: 164.308(a)(3)(i),164.308(a)(3)(ii)(A),164.308(a)(3)(ii)(B),164.308(a)(4)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(B),164.308(a)(4)(ii)(C),164.312(a)(1); PCI: 7.1.2,7.1.3,7.2.1,7.2.2;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-iam-policy-no-statements-with-admin-access	IAM	Config Changes	Report	CIS: CIS.16; NIST-CSF: PR.AC-6,PR.AC-7; HIPAA: 164.308(a)(4)(ii)(B),164.308(a)(5)(ii)(D),164.312(d); PCI: 8.2.3,8.2.4,8.2.5;
ams-nist-cis-iam-user-group-membership-check	IAM	Config Changes	Report	CIS: CIS.16,CIS.4; NIST-CSF: PR.AC-1,PR.AC-4,PR.PT-3; HIPAA: 164.308(a)(3)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(B),164.308(a)(4)(ii)(C),164.312(a)(1),164.312(a)(2)(i); PCI: 2.2,7.1.2,7.2.1,8.1.1;
ams-nist-cis-iam-user-no-policies-check	IAM	Config Changes	Report	CIS: CIS.16; NIST-CSF: PR.AC-1,PR.AC-7; HIPAA: 164.308(a)(4)(ii)(B),164.312(d); PCI: 8.3;
ams-nist-cis-iam-user-unused-credentials-check	IAM	Periodic	Report	CIS: CIS.16; NIST-CSF: PR.AC-1,PR.AC-4,PR.PT-3; HIPAA: 164.308(a)(3)(i),164.308(a)(3)(ii)(A),164.308(a)(3)(ii)(B),164.308(a)(4)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(B),164.308(a)(4)(ii)(C),164.312(a)(1); PCI: 2.2,7.1.2,7.1.3,7.2.1,7.2.2;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-ec2-instances-in-vpc	EC2	Config Changes	Report	CIS: CIS.11,CIS.12,CIS.9; NIST-CSF: DE.AE-1,PR.AC-3,PR.AC-5,PR.PT-4; HIPAA: 164.308(a)(3)(i),164.308(a)(3)(ii)(B),164.308(a)(4)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(B),164.308(a)(4)(ii)(C),164.312(a)(1),164.312(e)(1); PCI: 1.2,1.3,2.2,1.2.1,1.3.1,1.3.2,2.2.2;
ams-nist-cis-internet-gateway-authorized-vpc-only	Internet Gateway	Periodic	Report	CIS: CIS.9,CIS.12; NIST-CSF: NA; HIPAA: NA; PCI: NA;
ams-nist-cis-kms-cmk-not-scheduled-for-deletion	KMS	Periodic	Report	CIS: CIS.13,CIS.14; NIST-CSF: PR.DS-1; HIPAA: NA; PCI: 3.5,3.6;
ams-nist-lambda-concurrency-check	Lambda	Config Changes	Report	CIS: NA; NIST-CSF: NA; HIPAA: 164.312(b); PCI: NA;
ams-nist-lambda-dlq-check	Lambda	Config Changes	Report	CIS: NA; NIST-CSF: NA; HIPAA: 164.312(b); PCI: NA;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-lambda-function-public-access-prohibited	Lambda	Config Changes	Report	CIS: CIS.12,CIS.9; NIST-CSF: PR.AC-3,P R.AC-4,PR.AC-5,PR. DS-5,PR.PT-3,PR.PT-4; HIPAA: 164.308(a)(3) (i),164.308(a)(4)(ii)(A) ,164.308(a)(4)(ii)(C),164. 312(a)(1),164.312(e)(1); PCI: 1.2,1.3,1.2.1,1.3. 1,1.3.2,1.3.4,2.2.2;
ams-nist-cis-lambda-inside-vpc	Lambda	Config Changes	Report	CIS: CIS.12,CIS.9; NIST- CSF: PR.AC-3,PR.AC-4,PR .AC-5,PR.PT-3,PR.PT-4; HIPAA: 164.308(a)(3) (i),164.308(a)(4)(ii)(A) ,164.308(a)(4)(ii)(C),164. 312(a)(1),164.312(e)(1); PCI: 1.2,1.3,1.2.1,1.3. 1,1.3.2,1.3.4,2.2.2;
ams-nist-cis-mfa-enabled-for-iam-console-access	IAM	Periodic	Report	CIS: CIS.16; NIST- CSF: PR.AC-7; HIPAA: 164.312(d); PCI: 2.2,8.3;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-multi-region-cloudtrail-enabled	CloudTrail	Periodic	Report	CIS: CIS.6; NIST-CSF: DE.AE-1,DE.AE-3,PR.DS-5,PR.MA-2,PR.PT-1; HIPAA: 164.308(a)(3)(ii)(A),164.312(b); PCI: 2.2,10.1,10.2.1,10.2.2,10.2.3,10.2.4,10.2.5,10.2.6,10.2.7,10.3.1,10.3.2,10.3.3,10.3.4,10.3.5,10.3.6;
ams-nist-rds-enhanced-monitoring-enabled	RDS	Config Changes	Report	CIS: NA; NIST-CSF: PR.PT-1; HIPAA: 164.312(b); PCI: NA;
ams-nist-cis-rds-instance-public-access-check	RDS	Config Changes	Report	CIS: CIS.12,CIS.14,CIS.9; NIST-CSF: PR.AC-3,PR.AC-4,PR.AC-5,PR.DS-5,PR.PT-3,PR.PT-4; HIPAA: 164.308(a)(3)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(C),164.312(a)(1),164.312(e)(1); PCI: 1.2,1.3,1.2.1,1.3.1,1.3.2,1.3.4,1.3.6,2.2.2;
ams-nist-rds-multi-az-support	RDS	Config Changes	Report	CIS: NA; NIST-CSF: ID.BE-5,PR.DS-4,PR.PT-5,RC.RP-1; HIPAA: 164.308(a)(7)(i),164.308(a)(7)(ii)(C); PCI: NA;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-rds-snapshots-public-prohibited	RDS	Config Changes	Report	CIS: CIS.12,CIS.14,CIS.9; NIST-CSF: PR.AC-3,PR.AC-4,PR.AC-5,PR.DS-5,PR.PT-3,PR.PT-4; HIPAA: 164.308(a)(3)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(C),164.312(a)(1),164.312(e)(1); PCI: 1.2,1.3,1.2.1,1.3.1,1.3.2,1.3.4,1.3.6,2.2.2;
ams-nist-cis-rds-storage-encrypted	RDS	Config Changes	Report	CIS: CIS.13,CIS.5,CIS.6; NIST-CSF: DE.AE-1,DE.AE-3,PR.DS-1,PR.PT-1; HIPAA: 164.312(a)(2)(iv),164.312(b),164.312(e)(2)(ii); PCI: 3.4,10.1,10.2.1,10.2.2,10.2.3,10.2.4,10.2.5,10.3.1,10.3.2,10.3.3,10.3.4,10.3.5,10.3.6,8.2.1;
ams-nist-cis-redshift-cluster-configuration-check	RedShift	Config Changes	Report	CIS: CIS.6,CIS.13,CIS.5; NIST-CSF: DE.AE-1,DE.AE-3,PR.DS-1,PR.PT-1; HIPAA: 164.312(a)(2)(iv),164.312(b),164.312(e)(2)(ii); PCI: 3.4,8.2.1,10.1,10.2.1,10.2.2,10.2.3,10.2.4,10.2.5,10.3.1,10.3.2,10.3.3,10.3.4,10.3.5,10.3.6;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-redshift-cluster-public-access-check	RedShift	Config Changes	Report	CIS: CIS.12,CIS.14,CIS.9; NIST-CSF: PR.AC-3,PR.AC-4,PR.AC-5,PR.DS-5,PR.PT-3,PR.PT-4; HIPAA: 164.308(a)(3)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(C),164.312(a)(1),164.312(e)(1); PCI: 1.2,1.3,1.2.1,1.3.1,1.3.2,1.3.4,1.3.6,2.2.2;
ams-nist-cis-redshift-require-tls-ssl	RedShift	Periodic	Report	CIS: CIS.13,CIS.14; NIST-CSF: PR.DS-2; HIPAA: 164.312(a)(2)(iv),164.312(e)(1),164.312(e)(2)(i),164.312(e)(2)(ii); PCI: 2.3,4.1;
ams-nist-cis-root-account-hardware-mfa-enabled	IAM	Periodic	Report	CIS: CIS.16,CIS.4; NIST-CSF: PR.AC-7; HIPAA: 164.312(d); PCI: 2.2,8.3;
ams-nist-cis-root-account-mfa-enabled	IAM	Periodic	Report	CIS: CIS.16,CIS.4; NIST-CSF: PR.AC-7; HIPAA: 164.312(d); PCI: 2.2,8.3;
ams-nist-cis-s3-bucket-default-lock-enabled	S3	Config Changes	Report	CIS: CIS.14,CIS.13; NIST-CSF: ID.BE-5,PR.PT-5,RC.RP-1; HIPAA: NA; PCI: NA;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-s3-bucket-logging-enabled	S3	Config Changes	Report	CIS: CIS.6; NIST-CSF: DE.AE-1,DE.AE-3,PR.DS-5,PR.PT-1; HIPAA: 164.308(a)(3)(ii)(A),164.312(b); PCI: 2.2,10.1,10.2.1,10.2.2,10.2.3,10.2.4,10.2.5,10.2.7,10.3.1,10.3.2,10.3.3,10.3.4,10.3.5,10.3.6;
ams-nist-cis-s3-bucket-replication-enabled	S3	Config Changes	Report	CIS: CIS.10; NIST-CSF: ID.BE-5,PR.DS-4,PR.IP-4,PR.PT-5,RC.RP-1; HIPAA: 164.308(a)(7)(i),164.308(a)(7)(ii)(A),164.308(a)(7)(ii)(B); PCI: 2.2,10.5.3;
ams-nist-cis-s3-bucket-ssl-requests-only	S3	Config Changes	Report	CIS: CIS.13,CIS.14; NIST-CSF: PR.DS-2; HIPAA: 164.312(a)(2)(iv),164.312(c)(2),164.312(e)(1),164.312(e)(2)(i),164.312(e)(2)(ii); PCI: 2.2,4.1,8.2.1;
ams-nist-cis-s3-bucket-versioning-enabled	S3	Periodic	Report	CIS: CIS.10; NIST-CSF: ID.BE-5,PR.DS-4,PR.DS-6,PR.IP-4,PR.PT-5,RC.RP-1; HIPAA: 164.308(a)(7)(i),164.308(a)(7)(ii)(A),164.308(a)(7)(ii)(B),164.312(c)(1),164.312(c)(2); PCI: 10.5.3;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-sagemaker-endpoint-configuration-kms-key-configured	SageMaker	Periodic	Report	CIS: CIS.13,CIS.14; NIST-CSF: PR.DS-1; HIPAA: 164.312(a)(2)(iv),164.312(e)(2)(ii); PCI: 3.4,8.2.1;
ams-nist-cis-sagemaker-notebook-instance-kms-key-configured	SageMaker	Periodic	Report	CIS: CIS.13,CIS.14; NIST-CSF: PR.DS-1; HIPAA: 164.312(a)(2)(iv),164.312(e)(2)(ii); PCI: 3.4,8.2.1;
ams-nist-cis-sagemaker-notebook-no-direct-internet-access	SageMaker	Periodic	Report	CIS: CIS.12,CIS.9; NIST-CSF: PR.AC-3,P R.AC-4,PR.AC-5,PR. DS-5,PR.PT-3,PR.PT-4; HIPAA: 164.308(a)(3)(i),164.308(a)(4)(ii)(A),164.308(a)(4)(ii)(C),164.312(a)(1),164.312(e)(1); PCI: 1.2,1.3,1.2.1,1.3.1,1.3.2,1.3.4,1.3.6,2.2.2;
ams-nist-cis-secretsmanager-rotation-enabled-check	Secrets Manager	Config Changes	Report	CIS: CIS.16; NIST-CSF: PR.AC-1; HIPAA: 164.308(a)(4)(ii)(B); PCI: NA;

Rule Name	Service	Trigger	Action	Frameworks
ams-nist-cis-secretsmanager-schedule-rotation-success-check	Secrets Manager	Config Changes	Report	CIS: CIS.16; NIST-CSF: PR.AC-1; HIPAA: 164.308(a)(4)(ii)(B); PCI: NA;
ams-nist-cis-sns-encrypted-kms	SNS	Config Changes	Report	CIS: CIS.13,CIS.14; NIST-CSF: PR.DS-1; HIPAA: 164.312(a)(2)(iv),164.312(e)(2)(ii); PCI: 8.2.1;
ams-nist-cis-vpc-sg-open-only-to-authorized-ports	VPC	Config Changes	Report	CIS: CIS.11,CIS.12,CIS.9; NIST-CSF: DE.AE-1,PR.AC-3,PR.AC-5,PR.PT-4; HIPAA: 164.312(e)(1); PCI: 1.2,1.3,1.2.1,1.3.1,1.3.2,2.2.2;
ams-nist-vpc-vpn-2-tunnels-up	VPC	Config Changes	Report	CIS: NA; NIST-CSF: ID.BE-5,PR.DS-4,PR.PT-5,RC.RP-1; HIPAA: 164.308(a)(7)(i); PCI: NA;
ams-cis-ec2-ebs-encryption-by-default	EC2	Periodic	Report	CIS: CIS.13,CIS.14; NIST-CSF: PR.DS-1; HIPAA: 164.312(a)(2)(iv),164.312(e)(2)(ii); PCI: 2.2,3.4,8.2.1;
ams-cis-rds-snapshot-encrypted	RDS	Config Changes	Report	CIS: CIS.13,CIS.14; NIST-CSF: PR.DS-1; HIPAA: 164.312(a)(2)(iv),164.312(e)(2)(ii); PCI: 3.4,8.2.1;

Rule Name	Service	Trigger	Action	Frameworks
ams-cis-redshift-cluster-maintenance-settings-check	RedShift	Config Changes	Report	CIS: CIS.5; NIST-CSF: PR.DS-4,PR.IP-1,PR.IP-4; HIPAA: 164.308(a)(5)(ii)(A),164.308(a)(7)(ii)(A); PCI: 6.2;

Responses to violations in Accelerate

All Config Rule violations appear in your Configuration Report. This is a universal response. Depending on the *Remediation Category* (severity) of the rule, AMS might take additional actions, summarized in the following table. For details on how to customize the *Action Code* for certain rules, see [Customized findings responses](#).

Remediation Actions

Action Code	AMS Actions
Report	1. Add to Config Report
Incident	1. Add to Config Report 2. Automatic incident report in Accelerate
Remediate	1. Add to Config Report 2. Automatic incident report in Accelerate 3. Automatic remediation in Accelerate

Requesting Additional Help

Note

AMS can remediate any violation for you, regardless of its remediation category. To request help, submit a Service Request, and indicate which resources you want AMS to remediate with a comment such as "As part of the AMS config rule remediation, please remediate non-complaint resources *RESOURCE_ARNS_OR_IDS*resource ARNs/IDs>, config

rule `CONFIG_RULE_NAME` in the account" and add the required inputs to remediate the violation.

AMS Accelerate has a library of AWS Systems Manager automation documents and runbooks to assist in remediating noncompliant resources.

Add to Config Report

AMS generates a Config Report that tracks the compliance status of all rules and resources in your account. You can request the report from your CSDM. You can also review compliance status from the AWS Config console, AWS CLI, or AWS Config API. Your Config Report includes:

- The top, noncompliant resources in your environment, to discover potential threats and misconfigurations
- Compliance of resources and config rules over time
- Config rule descriptions, severity of rules, and recommended remediation steps to fix noncompliant resources

When any resource goes into a noncompliant state, the resource status (and rule status) becomes **Noncompliant** in your Config Report. If the rule belongs to the **Config Report Only** remediation category, by default, AMS takes no further action. You can always create a Service Request to request additional help or remediation from AMS.

For more details, see [AWS Config Reporting](#).

Automatic incident report in Accelerate

For moderately severe rule violations, AMS automatically creates an Incident Report to notify you that a resource has gone into a noncompliant state, and asks which actions you would like to be performed. You have the following options when responding to an incident:

- Request that AMS remediate the noncompliant resources listed in the incident. Then, we attempt to remediate the noncompliant resource, and notify you once the underlying incident has been resolved.
- You can resolve the noncompliant item manually in the console or through your automated deployment system (for example, CI/CD Pipeline template updates); then, you can resolve

the incident. The noncompliant resource is re-evaluated as per the rule's schedule and, if the resource is evaluated as noncompliant, a new incident report is created.

- You can choose to not resolve the noncompliant resource and simply resolve the incident. If you update the configuration of the resource later, AWS Config will trigger a re-evaluation and you will again be alerted to evaluate the noncompliance of that resource.

Automatic remediation in Accelerate

The most critical rules belong to the **Auto Remediate** category. Noncompliance with these rules may strongly impact the security and availability of your accounts. When a resource violates one of these rules:

1. AMS automatically notifies you with an Incident Report.
2. AMS starts an automated remediation using our automated SSM documents.
3. AMS updates the Incident Report with success or failure of the automated remediation.
4. If automated remediation failed, an AMS engineer investigates the issue.

Creating rule exceptions in Accelerate

The AWS Config Rules resource exception feature allows you to suppress reporting of specific, noncompliant resources for a specific rules.

Note

The exempted resources still show up as **Noncompliant** in your AWS Config Service console. The exempted resources appear with a special flag in Config Reports (`resource_exception:True`). Your CSDMs can filter out those resources according to that column when generating reports.

If you have resources that you know are not compliant, you can eliminate a specific resource for a specific config rule in their Config Reports. To do this:

Submit a service request to Accelerate against your account, with a list of the config rules and resources that to be exempted from report. You must provide an explicit business justification (such as, no need to report that *resource_name_1* and *resource_name_2* are not backed up because

we do not want them backed up). For help submitting an Accelerate service request, see [Creating a service request](#).

Paste into the request the following inputs (for every resource add a separate block with all the required fields, as shown), and then submit:

```
[
  {
    "resource_name": "resource_name_1",
    "config_rule_name": "config_rule_name_1",
    "business_justification": "REASON_TO_EXEMPT_RESOURCE",
    "resource_type": "resource_type"
  },
  {
    "resource_name": "resource_name_2",
    "config_rule_name": "config_rule_name_2",
    "business_justification": "REASON_TO_EXEMPT_RESOURCE",
    "resource_type": "resource_type"
  }
]
```

Customized findings responses

You can choose how you want AMS Accelerate to respond to some findings (non-compliant Config rules). You can configure AMS to respond to findings it by remediating the finding, asking for your approval to remediate, or just reporting to you in your next Monthly Business Review (MBR). You can change the default responses for AMS Accelerate Config rules. To see the rules, go to [Configuration Compliance > Table of rules](#) or download the rules table as a ZIP file [ams_config_rules.zip](#).

Changing default responses helps you to increase the security and compliance state of your account by allowing more findings to be remediated. When you remediate more findings, you have fewer cases that need to wait for a manual review and approval. The extensive library of AMS remediation runbooks constantly fixes non-compliant resources and you are contacted only when required.

Customized responses are used only with new resources or existing resources with new events. For example, a resource that became non-compliant after a change. This is because older resources tend to require a deeper inspection before remediation and it's easier to enforce the resource

remediation as they are created or changed. To request remediation of the finding for any resource at any time, submit a service request.

Requesting a change in the default responses

Cloud Architects (CAs) work with you during on-boarding to collect your preferences. CAs then setup the initial configuration on internal AMS systems. After onboarded, create a Service Request to request updates to your configurations. You can request configuration updates as many times as needed. Please note that Operations only updates configurations for the account in which the service request was created. If you need to update multiple accounts at the same time, contact your Cloud Architect. Your CA will ask you to cut a service request with your preferences for audit purposes.

Changing default responses for your findings and accounts

You always need a response preference for each account and finding. AMS provides a default response (see [Configuration Compliance](#)), so this configuration is optional. You can change the default responses for each finding to the following options:

- Remediate: AMS manually or automatically remediates the finding. AMS reviews the remediation and lets you know if it fails.
- Request approval: AMS creates an outbound case to notify you about the finding. Use this option when you want to review the finding before approving its remediation or exempting it. AMS then executes the action you prefer.
- No action (report only): AMS takes no action to remediate or escalate the finding. The findings might still appear on the console and reports presented during MBRs.

Note

You can't change the configuration of rules that must be remediated by AMS. For example, enabling Amazon GuardDuty and VPC Flow Logs.

Changing default responses by resources

You can further configure the response to specific resources using tags. You can use your pre-existing tags or tag resources using Resource Tagger. For details, see [Resource Tagger](#). Configuration for resources with tags take precedence over the default action for the finding. When

a resource has multiple tags with different associated configurations, AMS can't run customized remediations. Instead, AMS sends you an outbound Service Request to inform you of the situation. For example, for the [s3-bucket-server-side-encryption-enabled](#) finding you can:

- Change the response to 'remediate' unencrypted S3 buckets with the tag key value pair "Regulated: True"
- Change the response to 'no action' when unencrypted S3 buckets has the tags "Regulated: False", and
- Change the default response of unencrypted S3 buckets to be 'ask for approval'. This applies for all S3 buckets that don't have the tags "Regulated: True" or "Regulated: False"

You can also add the input required to run custom finding response. For example, for remediations that require an encryption key, you can provide your key IDs to AMS. You can change the input parameters of the remediation runbooks, but AMS doesn't support integration with custom runbooks. For a description of AMS remediation runbooks in the Config Report, see [AWS Config Control Compliance report](#).

Incident response in Accelerate

Upon receiving an alert, the AMS team uses automated and manual remediations to bring the resources back to a healthy state. If remediation fails, AMS starts the incident management process to collaborate with your team. You can change the baselines by updating the default configuration in a configuration file.

Incident response and onboarding in AMS Accelerate

During onboarding, AMS Accelerate suppresses automatic incident creation for your existing noncompliant resources; instead, your Cloud Service Deliver Manager (CSDM) provides you with a report that contains all the noncompliance rules and resources for your review. After you have identified the rules that you want AMS to remediate, create a service request in the AWS Support Center console indicating those rule and resources. The following Service Request template is an example of a customer request to AMS to manually remediate noncompliant resources. If AMS has additional questions, we work with you in the Service Request to gather the information required.

```
Hello,  
Please remediate the following resources for the Config Rule "ENCRYPTED_VOLUMES".  
Resource List:
```

```
"Vol-12345678"
```

```
"Vol-87654312"
```

```
Thank you
```

After the onboarding process is completed, AMS Accelerate automatically creates an incident for each noncompliant resource for the rules marked as Automatic Incident.

Resilience in AMS Accelerate

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](#).

For information about AMS Accelerate continuity management, see [Continuity management in AMS Accelerate](#).

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.

Security best practices

AMS Accelerate uses conformance packs, which provide a general-purpose compliance framework designed to enable you to create security, operational, or cost-optimization governance checks using managed or custom AWS Config Rules and AWS Config remediation actions. For information on how to best configure these conformance packs, see AWS Config's [Operational Best Practices for NIST CSF](#) and [Operational Best Practices for CIS Top 20](#).

Change request security reviews

The AWS Managed Services change request review process ensures that AMS performs a security review of the requested changes as they are implemented on your behalf in your account.

[AMS Accelerate technical standards](#) define the minimum security criteria, configurations, and processes to establish the baseline security of your accounts. When AMS implements the requested changes, we follow these standards.

AMS evaluates all change requests against the AMS technical standards. Any change that might lower your account's security posture by deviating from the technical standards goes through a security review process. During this process, relevant risk is highlighted by AMS and reviewed and approved by your authorised risk approver to balance security and business needs.

Customer Security Risk Management process

The AMS Accelerate Customer Security Risk Management (CSRM) process helps to clearly identify and communicate risks to the right owners. This process minimizes the security risks in your environment and reduces ongoing operational overhead for identified risks.

By default, when someone from your organization requests that AMS implement a change to your managed environment, AMS reviews the change to determine if the request falls outside of the technical standards, which might alter the security posture of your account. If there is a high or very high security risk, then the change review is accepted or rejected by your authorized security personnel. Requested changes are also evaluated for adverse effects on AMS's ability to operate the account. If the review finds possible adverse impacts, then additional reviews and approvals are required within AMS.

You can opt-out from the approval based workflow in the CSRM process for high or very high risks. To change the CSRM option for specific accounts from **Standard CSRM** to **Notification Only**, work with your Cloud Service Delivery Managers to create a one-time risk acceptance. If you choose

to proceed with the **Notification Only** option, then AMS implements the requested changes regardless of the risk category. And, AMS sends a risk notification to your authorized risk approvers instead of seeking approval prior to the change implementation. Speak with your Cloud Architects or Cloud Service Delivery Managers for more information about the AMS CSRM process, how to change the default CSRM option when onboarding new AMS accounts, or how to update existing accounts.

Note

AMS strongly recommends that you use the default option of **Standard CSRM** in all of your accounts.

AMS Accelerate technical standards

The following are Accelerate technical standards categories:

ID	Category
AMS-STD-X002	AWS Identity and Access Management
AMS-STD-X003	Network Security
AMS-STD-X004	Penetration Testing
AMS-STD-X005	Amazon GuardDuty
AMS-STD-X007	Logging

Standard controls in AMS Accelerate

The following are the standard controls in AMS:

AMS-STD-X002 - AWS Identity and Access Management (IAM)

ID	Technical standard
1.0	Timeout Duration

ID	Technical standard
1.1	A federated user default timeout session is one hour and may be increased to up to four hours.
1.2	RDP session timeout for Microsoft Windows Server is set to 15 minutes and can be extended based on use case.
2.0	AWS Root Account Usage
2.1	If there is a root account usage for any reason, Amazon GuardDuty must be configured to generate relevant findings.
2.2	Access keys for root account must not be created.
3.0	Users Creation and Modification
3.1	IAM users/roles with programmatic access and with read only permissions can be created without any time-limited policy. However, the permission to allow the reading of objects (for example, S3:GetObject) in all the Amazon Simple Storage Service buckets in the account are not permitted.
3.1.1	IAM human users for console access and with read only permissions can be created with the time bound policy (up to 180 days) while the removal/renewal/extension of the time bound policy will result in the risk notification. However, the permission to allow the reading of objects (for example, S3:GetObject) in all the S3 buckets in the account are not permitted.

ID	Technical standard
3.2	IAM users and roles for console and programmatic access with any infrastructure-mutating permissions (write, permission management, or tagging) in the customer account must not be created without risk acceptance. However, S3 object-level write permissions are allowed without risk acceptance as long as the specific buckets are in the scope.
3.3	On Microsoft Windows Servers, only Microsoft group Managed Service Account (gMSA) must be created.
4.0	Policies, Actions, and APIs
4.4	A policy must not provide administrator access with a statement that is equivalent to "Effect": "Allow" with "Action": "*" over "Resource": "*" without risk acceptance.
4.6	API calls against KMS key policies for AMS infrastructure keys in the customer IAM policies must not be permitted.
4.8	Actions, which changes to the AMS infrastructure DNS records in Amazon Route 53 must not be permitted.
4.9	IAM human users with console access created after following the due process, must not have any policies attached directly except trust policy, assume role, and time limited policy.

ID	Technical standard
4.10	Amazon EC2 instance profiles with read access to a specific secret or namespace in AWS Secrets Manager within the same account can be created.
4.12	IAM policy must not include any action which includes action Allow logs:DeleteLogGroup and logs:DeleteLogStream on any AMS Amazon CloudWatch log group.
4.13	Permissions to create multi-region keys must not be permitted.
4.14	Access to S3 bucket ARN which are not yet created in the customer accounts can be provided by restricting the access to the buckets to the customer accounts through specifying the account number using service-specific S3 condition key s3:ResourceAccount.
4.16	SQL Workbench related full permissions can be granted to roles/users to work on Amazon Redshift databases.
4.17	Any AWS CloudShell permissions can be granted to customer roles as an alternative of CLI.
4.18	An IAM role with AWS service as a trusted principal also needs to be inline with the IAM technical standards.
4.19	Service Linked Roles (SLRs) are not subject to AMS IAM technical standards, as they are built and maintained by IAM Service Team.

ID	Technical standard
4.20	IAM policy should not allow reading of objects (for example, S3:GetObject) in all the S3 buckets in the account.
4.21	All the IAM permissions for resource type "savingsplan" can be granted to customers.
4.22	AMS engineer is not permitted to copy or move customer data (files, S3 objects, databases etc) manually in any of the data storage services like Amazon S3, Amazon Relational Database Service, Amazon DynamoDB, and so on, or in the OS file system.
6.0	Cross Account Policies
6.1	IAM roles trust policies between AMS accounts that belong to the same customer as per customer records, can be configured.
6.2	IAM roles trust policies between AMS and non-AMS accounts must be configured only if the 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).
6.3	IAM roles trust policies between AMS accounts and third-party accounts must not be configured without risk acceptance.
6.4	Cross-account policies to access any customer-managed CMKs between AMS accounts of the same customer can be configured.

ID	Technical standard
6.5	Cross-account policies to access any KMS key within a non-AMS account by an AMS account can be configured.
6.6	Cross-account policies to access any KMS key within an AMS account by a third-party account must not be permitted without risk acceptance.
6.6.1	Cross-account policies to access any KMS key within an AMS account by a non-AMS account can be configured only if the non-AMS account is owned by the same AMS customer.
6.7	Cross-account policies to access any S3 bucket data or resources where data can be stored (such as Amazon RDS, Amazon DynamoDB, or Amazon Redshift) between AMS accounts of the same customer can be configured.
6.8	Cross-account policies to access any S3 bucket data or resources where data can be stored (such as Amazon RDS, Amazon DynamoDB, or Amazon Redshift) in a non-AMS account from an AMS account with read-only access can be configured.

ID	Technical standard
6.9	Cross-account policies to access any S3 bucket data or resources where data can be stored (such as Amazon RDS, Amazon DynamoDB, or Amazon Redshift) with write permissions from AMS to a non-AMS account (or a non-AMS to AMS account) must be configured only if the 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).
6.10	Cross-account policies to access any S3 bucket data or resources where data can be stored (such as Amazon RDS, Amazon DynamoDB, or Amazon Redshift) in a third-party account from an AMS account with read only access can be configured.
6.11	Cross-account policies to access any S3 bucket data or resources where data can be stored (such as Amazon RDS, Amazon DynamoDB, or Amazon Redshift) in a third-party account from an AMS account with write access must not be configured.
6.12	Cross-account policies from third-party accounts to access an AMS customer S3 bucket or resources where data can be stored (such as Amazon RDS, Amazon DynamoDB, or Amazon Redshift) must not be configured without risk acceptance.
7.0	User Groups

ID	Technical standard
7.1	IAM groups with readonly and non mutative permissions are permitted.
8.0	Resource-based policies
8.4	AMS infrastructure resources should be protected from management by unauthorized identities by the attachment of resource based policies.
8.2	Customer resources should be configured with least-privilege resource-based policies, unless the customer explicitly specifies a different policy.

AMS-STD-X003 - Network Security

The following is the standard control for X003 - Network Security:

ID	Technical standard
	Networking
1.0	<i>Reserved for future control</i>
2.0	Elastic IP on EC2 instances is permitted
3.0	AMS control plane and by extension in data plane TLS 1.2+ must be used.
5.0	A security group must not have source as 0.0.0.0/0 in the inbound rule if it is not attached to a load balancer as per 9.0
6.0	S3 bucket or objects must not be made public without risk acceptance.

ID	Technical standard
7.0	Servers management access on ports SSH/22 or SSH/2222 (Not SFTP/2222), TELNET/23, RDP/3389, WinRM/5985-5986, VNC/ 5900-5901 TS/CITRIX/1494 or 1604, LDAP/389 or 636 and RPC/135, NETBIOS/137-139 must not be permitted from outside the VPC through security groups.
8.0	Database management access on ports (MySQL/3306, PostgreSQL/5432, Oracle/1521, MSSQL/1433) or on custom port must not be permitted from public IPs not routed to VPC over DX, VPC-peer, or VPN via security group.
9.0	Direct applications access over port HTTP/80, HTTPS/8443 and HTTPS/443 from the Internet is permitted only to load balancers, but not to any compute resources directly e.g. EC2 instances, ECS/EKS/Fargate containers, etc.
10.0	Applications access over port HTTP/80 and HTTPS/443 from customer private IP range can be permitted.
11.0	Any changes to the security groups which controls the access to the AMS infrastructure must not be permitted without risk acceptance.
12.0	AMS Security will refer the standards every time a security group is requested to be attached to an instance.

ID	Technical standard
14.0	Cross account association of private hosted zones with VPCs from AMS to non-AMS account (or non-AMS to AMS account) must be configured only if non-AMS account is owned by the same AMS customer (by confirming that they are under the same AWS Organization account or by matching the email domain with the customer's company name) using internal tools.
15.0	VPC peering connections between accounts that belong to the same customer can be permitted.
16.0	AMS base AMIs can be shared with non-AMS account as long as both accounts are owned by the same customer (by confirming that they are under the same AWS Organizations account or by matching the email domain with the customer's company name) using internal tools.
17.0	FTP port 21 must not be configured in any of the security group without a risk acceptance.
18.0	Cross account network connectivity via transit gateway is permitted as long as all the accounts are owned by the customer.
19.0	Making a private subnet to public is not permitted
20.0	VPC peering connections with a third party accounts (not owned by the customer) must not be permitted.

ID	Technical standard
21.0	Transit Gateway attachment with a third party account (not owned by the customer) must not be permitted.
22.0	Any network traffic required for AMS to provide the services to customers must not be blocked at the customer network egress point.
23.0	Inbound ICMP request to Amazon EC2 from the customer infra will require risk notification.
24.0	Inbound request from public IPs routed to Amazon VPC over DX, VPC-peer, or VPN via security group is allowed.
25.0	Inbound request from public IPs not routed to Amazon VPC over DX, VPC-peer, or VPN via security group would require a risk acceptance.
26.0	Outbound ICMP request from Amazon EC2 to any destination is allowed.

AMS-STD-X004 - Penetration Testing

The following is the standard control for X004 - Penetration Testing

1. AMS doesn't support pentest infrastructure. It's the customer's responsibility. For example, Kali is not a AMS supported distribution of Linux.
2. Customers need to adhere to [Penetration Testing](#).
3. AMS to be pre-notified 24hrs in advance in the case when the customer would like to perform infrastructure penetration testing within accounts.
4. AMS will provision customer pentesting infrastructure per customer requirements explicitly stated in the change request or service request by the customer.

5. Identity management for customer pentesting infrastructure is the responsibility of the customer.

AMS-STD-X005 - GuardDuty

The following is the standard control for X005 - GuardDuty

1. GuardDuty must be enabled in all the customer accounts at all times.
2. GuardDuty alerts must be stored within the same account or any other managed account under the same organization.
3. Trusted IP list feature of GuardDuty must not be used. Instead auto-archiving can be used as an alternative, which is useful for audit purposes.

AMS-STD-X007 - Logging

The following is the standard control for X007 - Logging

ID	Technical standard
1.0	Log types
1.1	OS Logs: All the hosts must log at minimum host authentication events, access events for all uses of elevated privileges and access events for all changes to access and privilege configuration including success and failure both.
1.2	AWS CloudTrail: CloudTrail management event logging must be enabled and configured to deliver logs to an S3 bucket.
1.3	VPC Flow Logs: All the network traffic logs must be logged via VPC Flow Logs and must be sent to S3 bucket and can optionally be sent to CloudWatch Logs.

ID	Technical standard
1.4	Amazon S3 Server Access Logging: AMS mandated S3 buckets that store logs must have server access logging enabled.
1.5	AWS Config Snapshots: AWS Config must record configuration changes for all supported resources in all the regions and deliver the configuration snapshot files to S3 buckets at least once per day.
1.7	Application Logs: Customers are empowered to enable logging in their applications and store in CloudWatch Logs log group or an S3 bucket.
1.8	S3 Object level logging: Customers are empowered to enable object level logging in their S3 buckets.
1.9	Service Logging: Customers are empowered to enable and forward logs for SSPS services like any core services.
1.10	Elastic Load Balancing(Classic/Application Load Balancer/Network Load Balancer) Logs: Access and error log entries must be stored in the AMS 2.0-managed S3 buckets.
2.0	Access control
2.3	AMS-mandated S3 buckets that store logs must not allow third party accounts users as principles in the bucket policies.
2.4	Logs from CloudWatch Logs log groups must not be deleted without explicit approval from the customer authorised security contact.

ID	Technical standard
3.0	Logs retention
3.1	AMS-mandated CloudWatch Logs log groups must have a minimum retention period of 90 days on the logs.
3.2	AMS-mandated S3 buckets that stores the logs must have a minimum retention period of 18 months on the logs.
3.3	AWS Backup snapshots should be available with minimum retention of 31 days on the supported resources.
4.0	Encryption
4.1	Encryption must be enabled in all S3 buckets required by AMS Teams that stores logs.
4.2	Any log forwarding from customer accounts to any other account must be encrypted.
5.0	Integrity
5.1	The log file integrity mechanism must be enabled. That means configure “Log file validation” in the AWS CloudTrail trails required by AMS teams.
6.0	Logs forwarding
6.1	Any log can be forwarded from one AMS account to another AMS account of the same customer.

ID	Technical standard
6.2	Any log can be forwarded from AMS to non-AMS account only if 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.

Changes that introduce high or very high security risks in your environment

The following changes introduce high or very high security risk in your environment:

AWS Identity and Access Management

- High_Risk-IAM-001: Create access keys for root account
- High_Risk-IAM-002: SCP policy modification to allow additional access
- High_Risk-IAM-003: SCP policy modification that could break AMS infrastructure
- High_Risk-IAM-004: Creation of a role/user with infrastructure mutating permissions (write, permission management or tagging) in customer account
- High_Risk-IAM-005: IAM roles trust policies between AMS accounts and third-party accounts (not owned by the customer)
- High_Risk-IAM-006: Cross-account policies to access any KMS key from an AMS account by a third-party account)
- High_Risk-IAM-007: Cross-account policies from third-party accounts to access an AMS customer S3 bucket or resources where data can be stored (such as Amazon RDS, Amazon DynamoDB, or Amazon Redshift)
- High_Risk-IAM-008: Assign the IAM permissions with any infrastructure mutating permission in customer account
- High_Risk-IAM-009: Allow listing and reading on all the S3 buckets in the account

Network security

- High_Risk-NET-001: Open OS management ports SSH/22 or SSH/2222 (Not SFTP/2222), TELNET/23, RDP/3389, WinRM/5985-5986, VNC/ 5900-5901 TS/CITRIX/1494 or 1604, LDAP/389 or 636 and NETBIOS/137-139 from the internet
- High_Risk-NET-002: Open database management ports MySQL/3306, PostgreSQL/5432, Oracle/1521, MSSQL/1433 or any management customer port from the internet
- High_Risk-NET-003: Open application ports HTTP/80, HTTPS/8443 and HTTPS/443 on any compute resources directly. For example, EC2 instances, ECS/EKS/Fargate containers, and so on from the internet
- High_Risk-NET-004: Any changes to the security groups which controls the access to the AMS infrastructure
- High_Risk-NET-006: VPC peering with the third-party account (not owned by the customer)
- High_Risk-NET-007: Adding customer firewall as egress point for all the AMS traffic
- High_Risk-NET-008: Transit Gateway attachment with the third-party account is not allowed
- High_Risk-S3-001: Provision or enable public access in the S3 bucket

Logging

- High_Risk-LOG-001: Disable CloudTrail.
- High_Risk-LOG-002: Disable VPC Flow Logs.
- High_Risk-LOG-003: Log forwarding via any method (S3 event notification, SIEM agent pull, SIEM agent push etc) from an AMS managed account to third party account (not owned by customer)
- High_Risk-LOG-004: Use non-AMS trail for CloudTrail

Miscellaneous

- High_Risk-ENC-001: Disable encryption in any resource if it is enabled

Monitoring and event management in AMS Accelerate

The AMS Accelerate monitoring system monitors your AWS resources for failures, performance degradation, and security issues.

As a managed account, AMS Accelerate configures and deploys alarms for applicable AWS resources, monitors these resources, and performs remediation when needed.

The AMS Accelerate monitoring system relies on internal tools, such as Resource Tagger and Alarm Manager, and leverages native AWS services, such as [AWS AppConfig](#), Amazon CloudWatch (CloudWatch), Amazon EventBridge(formerly known as CloudWatch), Amazon GuardDuty, Amazon Macie, and AWS Health.

AMS Accelerate provides a range of operational services to help you achieve operational excellence on AWS. To gain a quick understanding of how AMS helps your teams achieve overall operational excellence in AWS Cloud with some of our key operational capabilities including 24x7 helpdesk, proactive monitoring, security, patching, logging and backup, see [AMS Reference Architecture Diagrams](#).

Topics

- [What is monitoring?](#)
- [How monitoring works](#)
- [Alerts from baseline monitoring in AMS](#)
- [Alarm Manager](#)
- [AMS automatic remediation of alerts](#)
- [Using Amazon EventBridge Managed Rules in AMS](#)

What is monitoring?

AMS Accelerate monitoring provides these benefits:

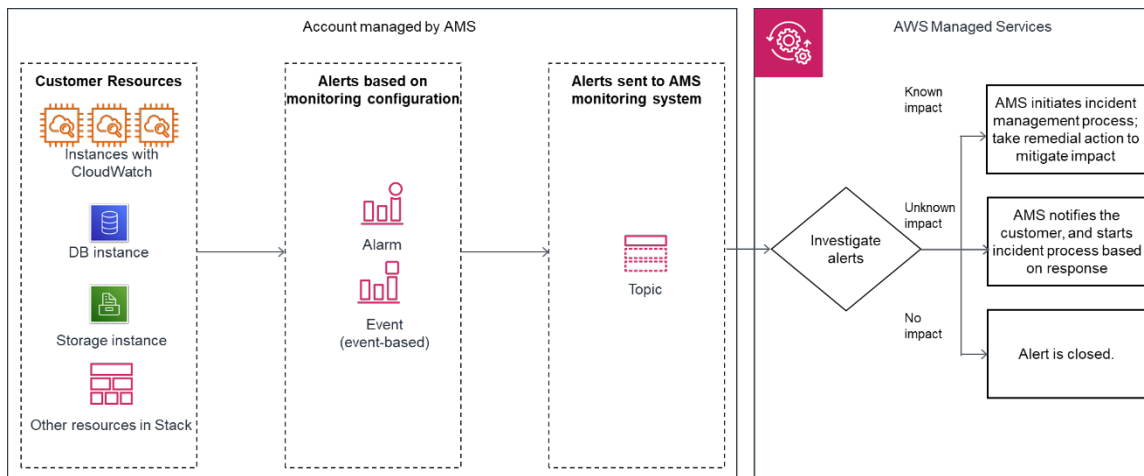
- A default configuration that creates, manages, and deploys policies across your managed account for all or supported AWS resources that you select.
- A monitoring baseline so that you have a default level of protection, even if you don't configure any other monitoring for your managed accounts. For more information, see [Alerts from baseline monitoring in AMS](#).

- The ability to customize the baseline resource alarms to meet your requirements.
- Automatic remediation of alerts by AMS Operations, when possible, to prevent or reduce the impact to 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](#).
- Visibility 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. For more information, see [Incident management](#).
- 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 the 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.

How monitoring works

See the following graphics on monitoring architecture in AWS Managed Services (AMS).

The following diagram depicts the **AMS Accelerate** monitoring architecture.



After your resources are tagged based on the policy defined using Resource tagger, and alarm definitions are deployed, the following diagram depicts the AMS monitoring architecture.

- **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:** 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.

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.

Alerts from baseline monitoring in AMS

Learn about AMS Accelerate monitoring defaults; for more information, see [Monitoring and event management in AMS Accelerate](#).

The following table shows what is monitored and the default alerting thresholds. You can change the alerting thresholds with a custom configuration document, or submit a service request. For instructions on changing your custom alarm configuration, see [Changing the Accelerate alarm configuration](#). 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, [Alarm Manager](#).

Amazon CloudWatch provides extended retention of metrics. For more information, see [CloudWatch Limits](#).

Note

AMS Accelerate 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 Accelerate updates the baseline monitoring in existing accounts on a periodic basis and you may experience a delay before the updates are in place.

Alerts from baseline monitoring

Resource	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).

ALB instance	ApplicationLoadBalancerErrorCount (HTTPCode_ELB_5XX_Count/RequestCount)*100	CloudWatch alarm on excess number of HTTP 5XX response codes generated by the Loadbalancer.
--------------	--	---

Resource	Alert name and trigger condition	Notes
	sum > 15% for 1 min, 5 consecutive times.	
	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
ALB target	TargetConnectionErrorCount (HTTPCode_Target_5XX_Count/ RequestCount)*100 sum > 15% for 1 min, 5 consecutive times.	CloudWatch alarm on excess number of HTTP 5XX response codes generated by a target.
	ApplicationLoadBalancerTargetGroupErrorCount sum > 0% for 1 min, 5 consecutive times.	CloudWatch alarm if number of connections were unsuccessfully established between the load balancer and the registered instances.
Aurora	Average CPU utilization > 90% for 20 mins, 5 consecutive times.	CloudWatch Alarm.
AWS Backup	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.

Resource	Alert name and trigger condition	Notes
EC2 instance - all OSs	CPUUtilization*	<p>CloudWatch alarm. High CPU utilization is an indicator of a change in application state such as deadlocks, infinite loops, malicious attacks, and other anomalies.</p> <p>This is a Direct-Customer-Alerts alarm.</p>
	> 95% for 5 mins, 6 consecutive times.	
StatusCheckFailed		
> 0% for 5 minute , 3 consecutive times.		
EC2 instance - Linux	Minimum mem_used_percent	
	>= 95% for 5 minutes, 6 consecutive times.	
	Average swap_used_percent	
EC2 instance - Windows	>= 95% for 5 minutes, 6 consecutive times.	
	Maximum disk_used_percent	
EC2 instance - Linux	>= 95% for 5 minutes, 6 consecutive times.	
	Maximum LogicalDisk % Free Space	
EC2 instance - Windows	<= 5% for 5 minutes, 6 consecutive times.	
	AMSEFSBurstCreditBalanceExhausted.	
EFS	BurstCreditBalance less than 1000 for fifteen minutes.	

Resource	Alert name and trigger condition	Notes
	AMSEFSClientConnectionsLimit. ClientConnections > 24,000 for fifteen minutes.	CloudWatch alarm on the ClientConnections of the EFS file system.
	AMSEFSThroughputUtilizationLimit. EFS Throughput Utilization > 80% for one hour.	CloudWatch alarm on the Throughput Utilization of the EFS file system.
	AMSEFSPercentIOLimit. PercentIOLimit > 95 for seventy five minutes.	CloudWatch alarm on the PercentIO Limit of the EFS file system.
ELB instance	SpilloverCountBackendConnections > 1 for 1 minute , 15 consecutive times.	CloudWatch alarm if an excess number of requests that were rejected because the surge queue is full.
	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.
	SurgeQueueLength > 100 for 1 minute, 15 consecutive times.	CloudWatch alarm if an excess number of requests are pending routing.
Amazon FSx ONTAP	AMFSXONTAPThroughputUtilization. FSX:ONTAP IOPS Utilization > 80% for two hours.	CloudWatch alarm on the IOPS utilization limit of the FSX:ONTAP instance.

Resource	Alert name and trigger condition	Notes
	AMSFSXONTAPIOPSUtilization. FSX:ONTAP Throughput Utilization > 80% for two hours.	CloudWatch alarm on the volume capacity utilization limit of the FSX:ONTAP volume.
	AMSFSXONTAPVolumelnodeUtilization. FSX:ONTAP Inode Utilization > 80% for two hours.	CloudWatch alarm on the file capacity utilization limit of the FSX:ONTAP volume.
Amazon FSx Windows	AMSFSXWindowsThroughputUtilization. FSX:Windows Throughput Utilization > 80% for two hours.	CloudWatch alarm on the IOPS utilization limit of the FSX:Windows instance.
	AMSFSXWindowsIOPSUtilization. FSX:Windows IOPS Utilization > 80% for two hours.	CloudWatch alarm on the IOPS utilization limit of the FSX:Windows instance.
GuardDuty Service	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	Alert name and trigger condition	Notes
Health	AWS Personal Health Dashboard	<p>Notifications are sent when there are changes in the status of AWS Personal Health Dashboard (AWS Health) events.</p> <p>Service event example: Scheduled EC2 instance store retirement.</p> <p>These Health events are not monitored:</p> <ul style="list-style-type: none">• AWS_VPN_SINGLE_TUNNEL_NOTIFICATION• AWS_ELASTICACHE_UPDATE_COMPLETED• AWS_STORAGEGATEWAY_SOFTWARE_UPDATE_AVAILABLE• AWS_BILLING_NOTIFICATION• AWS_VPN_REDUNDANCY_LOSS
IAM	EC2 IAM Instance Profile does not exist. The instance profile is missing.	For instructions on replacing an EC2 IAM instance profile, see the IAM documentation at Replace IAM role .

Resource	Alert name and trigger condition	Notes
	<p>EC2 IAM Instance Profile has too many policies.</p> <p>The IAM instance profile has 10 policies and additional policies cannot be added.</p>	<ul style="list-style-type: none"> • Modify the AWS Service Quota for IAM to increase number of managed policies per role to 20. For information about service quotas, see Viewing service quotas. • Lower the managed policy count below the current IAM quota by removing unnecessary managed policies for the IAM Role associated with these instances. Be sure to keep AMS required policies. • Lower the managed policy count below the current IAM quota by consolidating policies for the IAM Role associated with these instances. Be sure to keep AMS required policies. <p>For AMS required policies, see the <i>AMS Accelerate User Guide</i>: IAM permissions change details.</p>
Macie	<p>Newly generated alerts and updates to existing alerts.</p> <p>Macie finds any changes in the findings. These changes include newly generated findings or subsequent occurrences of existing findings.</p>	<p>Amazon Macie alert. For a list of supported Amazon Macie alert types, see Analyzing Amazon Macie findings. Note that Macie is not enabled for all accounts.</p>
NATGateways	<p>PacketsDropCount : Alarm if packetsdroppcount is > 0 over 15 minutes period</p>	<p>A value greater than zero may indicate an ongoing transient issue with the NAT gateway.</p>

Resource	Alert name and trigger condition	Notes
	ErrorPortAllocation : Alarm if NAT Gateways could not allocate port for over 15 minutes evaluation period	The number of times the NAT gateway could not allocate a source port. A value greater than Zero indicates that too many concurrent connections are open..
OpenSearch cluster	ClusterStatus red maximum is >= 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 Red Cluster Status .
	KMSKeyError >= 1 for 1 minute, 1 consecutive time.	CloudWatch alarm. At least one primary shard and its replicas are not allocated to a node. To learn more, see Encryption of Data at Rest for Amazon OpenSearch Service .
	KMSKeyInaccessible >= 1 for 1 minute, 1 consecutive time.	
OpenSearch domain	ClusterStatus yellow maximum is >= 1 for 1 minute, 1 consecutive time.	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.	A node in your cluster is down to 20 GiB of free storage space. To learn more, see Lack of Available Storage Space .
	ClusterIndexWritesBlocked >= 1 for 5 minutes, 1 consecutive time.	The cluster is blocking write requests. To learn more, see ClusterBlockException .

Resource	Alert name and trigger condition	Notes
	<p>Nodes</p> <p>minimum < x for 1 day, 1 consecutive time.</p>	<p>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.</p>
	<p>CPUUtilization</p> <p>average >= 80% for 15 minutes, 3 consecutive times.</p>	<p>100% CPU utilization isn't uncommon, but sustained high averages are problematic. Consider right-sizing an existing instance types or adding instances.</p>
	<p>JVMMemoryPressure</p> <p>maximum >= 80% for 5 minutes, 3 consecutive times.</p>	<p>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.</p>
	<p>MasterCPUUtilization</p> <p>average >= 50% for 15 minutes, 3 consecutive times.</p>	<p>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.</p>
	<p>MasterJVMMemoryPressure</p> <p>maximum >= 80% for 15 minutes, 1 consecutive time.</p>	

Resource	Alert name and trigger condition	Notes
OpenSearch instance	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. To learn more, see Red Cluster Status .
RDS	Average CPU utilization $> 75\%$ for 15 mins, 2 consecutive times.	CloudWatch alarms.
	Sum of DiskQueueDepth $> 75\%$ for 1 mins, 15 consecutive times.	
	Average FreeStorageSpace $< 1,073,741,824$ bytes for 5 mins, 2 consecutive times.	
	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 .
	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.	RDS-EVENT-0031, see details at Amazon RDS Event Categories and Event Messages .
RDS -0034 failover not attempted. RDS is not attempting a requested failover because a failover recently occurred on the DB instance.	RDS-EVENT-0034, see details at Amazon RDS Event Categories and Event Messages .	

Resource	Alert name and trigger condition	Notes
	<p>RDS - 0035 DB instance invalid parameters</p> <p>For example, MySQL could not start because a memory-related parameter is set too high for this instance class, so your action would be to modify the memory parameter and reboot the DB instance.</p>	<p>RDS-EVENT-0035, see details at Amazon RDS Event Categories and Event Messages.</p>
	<p>Invalid subnet IDs DB instance</p> <p>The DB instance is in an incompatible network. Some of the specified subnet IDs are invalid or do not exist.</p>	<p>Service event. RDS-EVENT-0036, see details at Amazon RDS Event Categories and Event Messages.</p>
	<p>RDS-0045 DB instance read replica error</p> <p>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.</p>	<p>RDS-EVENT-0045, see details at Amazon RDS Event Categories and Event Messages.</p>
	<p>RDS-0057 Error create statspack user account</p> <p>Replication on the Read Replica was ended.</p>	<p>Service event. RDS-EVENT-0057, see details at Amazon RDS Event Categories and Event Messages.</p>
	<p>RDS-0058 DB instance read replication ended</p> <p>Error while creating Statspack user account PERFSTAT. Drop the account before adding the Statspack option.</p>	<p>Service event. RDS-EVENT-0058, see details at Amazon RDS Event Categories and Event Messages.</p>

Resource	Alert name and trigger condition	Notes
	<p>DB instance recovery start</p> <p>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>[, ...])</p>	<p>Service event. RDS-EVENT-0066 see details at Amazon RDS Event Categories and Event Messages.</p>
	<p>A failover for the DB cluster has failed.</p>	<p>RDS-EVENT-0069, see details at Amazon RDS Event Categories and Event Messages.</p>
	<p>Invalid permissions recovery S3 bucket</p> <p>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.</p>	<p>Service event. RDS-EVENT-0081 see details at Amazon RDS Event Categories and Event Messages.</p>
	<p>Aurora was unable to copy backup data from an Amazon S3 bucket.</p>	<p>RDS-EVENT-0082, see details at Amazon RDS Event Categories and Event Messages.</p>
	<p>Low storage alert when the DB instance has consumed more than 90% of its allocated storage.</p>	<p>Service event. RDS-EVENT-0089 see details at Amazon RDS Event Categories and Event Messages.</p>
	<p>Notification service when scaling failed for the Aurora Serverless DB cluster.</p>	<p>Service event. RDS-EVENT-0143 see details at Amazon RDS Event Categories and Event Messages.</p>

Resource	Alert name and trigger condition	Notes
	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 .
	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 .
RedShift cluster	The health of the cluster when not in maintenance mode < 1 for 5 min	For more information, see Monitoring Amazon Redshift using CloudWatch metrics .

Resource	Alert name and trigger condition	Notes
Site-to-Site VPN	VPNTunnelDown TunnelState <= 0 for 1 min, 20 consecutive times.	TunnelState is 0 when both tunnels are down, .5 when one tunnel is up, and 1.0 when both tunnels are up.
Systems Manager Agent	EC2 Instances Not Managed by Systems Manager SSM agent is not installed. SSM agent is installed on the instance, but the agent service is not running. SSM agent has no network route to the AWS Systems Manager service.	There are additional conditions that cause disruption the Systems Manager Agent; for more information, see Troubleshooting managed node availability .

For information on remediation efforts, see [AMS automatic remediation of alerts](#).

[Watch Andrew's video to learn more \(7:03\)](#)

Alarm Manager

AMS Accelerate applies alarms to your AWS resources using the tag-based Alarm Manager to implement a baseline monitoring strategy and ensure that all your AWS resources are monitored and protected. By integrating with the tag-based Alarm Manager, you can customize the configuration of your AWS resources based on their type, platform, and other tags, to ensure the resources are monitored. Alarm Manager is deployed to your Accelerate account during onboarding.

How Alarm Manager works

When your account is onboarded to AMS Accelerate, two JSON documents, called configuration profiles, are deployed in your account in [AWS AppConfig](#). Both profile documents reside in the Alarm Manager application and in the AMS Accelerate infrastructure environment.

The two configuration profiles are named **AMSManagedAlarms** (the default configuration profile) and **CustomerManagedAlarms** (the customization configuration profile).

- Default configuration profile:

- The configuration found in this profile contains the default configuration that AMS Accelerate deploys in all customer accounts. This configuration contains the default AMS Accelerate monitoring policy, which you should not modify because AMS Accelerate can update this profile at any time, erasing any changes you have made.
- If you want to modify or disable any of these definitions, see [Modifying the Accelerate alarm default configuration](#) and [Disabling the default Accelerate alarm configuration](#).
- Customization configuration profile:
 - Any configuration in this profile is entirely managed by you; AMS Accelerate does not overwrite this profile, unless you explicitly request it.
 - You can specify any custom alarm definitions you want in this profile, and you can also specify modifications to the AMS Accelerate-managed default configuration. For more information, see [Modifying the Accelerate alarm default configuration](#) and [Disabling the default Accelerate alarm configuration](#).
 - If you update this profile, Alarm Manager automatically enforces your changes across all relevant resources in your AWS account. Note that while your changes are enacted automatically, they may take up to 60 minutes to take effect.
 - You can update this profile using the AWS Management Console or AWS CLI/SDK tools. See the [AWS AppConfig User Guide](#) for instructions about updating a configuration.
 - The customization profile is initially empty; however, any alarm definitions placed in the profile document are enforced, in addition to the default configuration.

All CloudWatch alarms created by the Alarm Manager contain the tag key **ams:alarm-manager:managed** and tag value **true**. This is to ensure that the Alarm Manager manages only those alarms that it creates, and won't interfere with any of your own alarms. You can see these tags using the Amazon CloudWatch [ListTagsForResource](#) API.

Important

If custom alarm definitions and default alarm definitions are specified with the same ConfigurationID (see [Configuration profile: monitoring](#)), the custom definitions take priority over default rules.

Getting started with Accelerate Alarm Manager

By default, when you onboard with AMS Accelerate, your configuration is deployed to AWS AppConfig, defining an alarm baseline for your resources. The alarm definitions are applied only to resources with the **ams:rt:*** tags. We recommend that these tags be applied using the [Resource Tagger](#): you set up a basic Resource Tagger configuration in order to let AMS Accelerate know which resources you want managed.

Use Resource Tagger to apply the tag key **ams:rt:ams-managed** with tag value **true** to any resources you want AMS Accelerate to monitor.

The following is an example Resource Tagger customization profile that you can use to opt in to monitoring for all of your Amazon EC2 instances. For general information, see [Resource Tagger](#).

```
{
  "AWS::EC2::Instance": {
    "AMSManageAllEC2Instances": {
      "Enabled": true,
      "Filter": {
        "InstanceId": "*"
      },
      "Tags": [
        {
          "Key": "ams:rt:ams-managed",
          "Value": "true"
        }
      ]
    }
  }
}
```

For information about how to apply this Resource Tagger configuration, see [Viewing or making changes to the Resource Tagger configuration](#).

Accelerate Alarm Manager tags

By default, when you onboard with AMS Accelerate, your configuration is deployed to AWS AppConfig, defining an alarm baseline for your resources. The alarm definitions are applied only to resources with the **ams:rt:*** tags. We recommend that these tags be applied using the [Resource Tagger](#): you set up a basic Resource Tagger configuration in order to let AMS Accelerate know which resources you want managed.

Use Resource Tagger to apply the tag key **ams:rt:ams-managed** with tag value **true** to any resources you want AMS Accelerate to monitor.

Topics

- [Tags using Resource Tagger](#)
- [Tags without Resource Tagger](#)
- [Tags using AWS CloudFormation](#)
- [Tags using Terraform](#)

Tags using Resource Tagger

The tag-based Alarm Manager manages the lifecycle of per-resource CloudWatch alarms; however, it requires that the managed resources have specific tags defined by AMS Accelerate. To use the Resource Tagger to apply the default set of AMS-managed alarms to both Linux and Windows based instances, follow these steps.

1. Browse to the [AppConfig](#) console within your account.
2. Select the ResourceTagger application.
3. Select the **Configuration profiles** tab, and then select **CustomerManagedTags**.
4. Click **Create** to create a new profile.
5. Select **JSON** and define your configuration. For more examples of filter and platform definition, see [Resource Tagger](#).

```
{
  "AWS::EC2::Instance": {
    "MonitorAllInstances": {
      "Enabled": true,
      "Filter": {
        "Platform": "*"
      },
      "Tags": [
        {
          "Key": "ams:rt:ams-managed",
          "Value": "true"
        }
      ]
    }
  }
}
```



```
}
```

6. Click **Create hosted configuration version**.
7. Click **Start deployment**.
8. Define the following deployment details:

```
Environment: AMSInfrastructure Hosted configuration version: <Select the version  
that you have just created>  
Deployment Strategy: AMSNoBakeDeployment
```

9. Click **Start deployment**.

Your instances become tagged with `"ams:rt:ams-managed": "true"` which ensures that additional `"ams:rt:ams-monitoring-policy": "ams-monitored"` and `"ams:rt:ams-monitoring-policy-platform": "ams-monitored-linux"` are applied to the instances. These tags then result in the appropriate alarms being created for the instance. For more information about this process, see [Monitoring](#).

[Watch Himanshu's video to learn more \(11:04\)](#)

Tags without Resource Tagger

The tag-based Alarm Manager manages the lifecycle of per-resource CloudWatch alarms; however, it requires that the managed resources have specific tags defined by AMS Accelerate. AMS Accelerate provides a default configuration profile that assumes that your tags have been applied by Resource Tagger.

If you want to use an alternate method of applying tags to your resources, such as AWS CloudFormation or Terraform, and not Resource Tagger, you need to disable the Resource Tagger so that it doesn't apply tags to your resources and compete with your chosen tagging method. For instructions on changing your custom Resource Tagger configuration profile to enable read-only mode, see [Preventing Resource Tagger from modifying resources](#).

After the Resource Tagger has been set to read-only mode, and the configuration profile is deployed, use your chosen tagging method to apply tags to your resources according to the following guidelines:

Resource type	Tag key	Tag value
All supported resources (described in this table)	ams:rt:ams-monitoring-policy	ams-monitored
EC2 instances (Linux)	ams:rt:ams-monitoring-policy-platform	ams-monitored-linux
EC2 instances (Windows)	ams:rt:ams-monitoring-policy-platform	ams-monitored-windows
OpenSearch Domain with KMS	ams:rt:ams-monitoring-with-kms	ams-monitored-with-kms
OpenSearch Domain with Dedicated Master Node	ams:rt:ams-monitoring-with-master	ams-monitored-with-master

Resources that have these tag keys and values are managed by the AMS Accelerate Alarm Manager.

Tags using AWS CloudFormation

Note

Make sure you have set Resource Tagger to read-only mode first before applying tags using AWS CloudFormation, otherwise Resource Tagger may modify the tags based on the configuration profile. For information on setting Resource Tagger to read-only mode, and guidelines on providing your own tags, see [Tags without Resource Tagger](#).

To apply tags using AWS CloudFormation, you can apply tags at the stack level (see [CloudFormation Resource Tags](#)) or, at the individual resource level, (for example, see [Creating EC2 Instance Tags](#)).

The following is an example of how you can apply AMS Accelerate alarm management tags to an Amazon EC2 instance managed by AWS CloudFormation:

```
Type: AWS::EC2::Instance
Properties:
```

```
InstanceType: "t3.micro"

# ...other properties...

Tags:
  - Key: "aws:rt:ams-monitoring-policy"
    Value: "ams-monitored"
  - Key: "aws:rt:ams-monitoring-policy-platform"
    Value: "ams-monitored-linux"
```

The following is an example of how you can apply AMS Accelerate alarm management tags to an Auto Scaling group managed by AWS CloudFormation. Note that the Auto Scaling group will propagate its tags to Amazon EC2 instances that are created by it:

```
Type: AWS::AutoScaling::AutoScalingGroup
Properties:
  AutoScalingGroupName: "TestASG"

# ...other properties...

Tags:
  - Key: "aws:rt:ams-monitoring-policy"
    Value: "ams-monitored"
  - Key: "aws:rt:ams-monitoring-policy-platform"
    Value: "ams-monitored-linux"
```

Tags using Terraform

Note

Make sure you have set Resource Tagger to read-only mode first before applying tags using AWS CloudFormation, otherwise Resource Tagger may modify the tags based on the configuration profile. For information on setting Resource Tagger to read-only mode, and guidelines on providing your own tags, see [Tags without Resource Tagger](#).

For a description of how to manage resource tags using Terraform, see the Terraform documentation [Resource Tagging](#).

The following is an example of how you can apply AMS Accelerate alarm management tags to an Amazon EC2 instance managed by Terraform.

```
resource "aws_instance" "test_linux_instance" {
  # ...ami and other properties...

  instance_type = "t3.micro"

  tags = {
    "aws:rt:ams-monitoring-policy" = "ams-monitored"
    "aws:rt:ams-monitoring-policy-platform" = "ams-monitored-linux"
  }
}
```

The following is an example of how you can apply AMS alarm management tags to an Auto Scaling group managed by Terraform. Note that the Auto Scaling group propagates its tags to EC2 instances that are created by it:

```
resource "aws_autoscaling_group" "test_asg" {
  name = "terraform-test"
  # ...other properties...

  tags = {
    "aws:rt:ams-monitoring-policy" = "ams-monitored"
    "aws:rt:ams-monitoring-policy-platform" = "ams-monitored-linux"
  }
}
```

Accelerate Alarm Manager configuration profiles

When your account is onboarded to AMS Accelerate, two JSON documents, called configuration profiles, are deployed in your account with AWS AppConfig (see [What is AWS AppConfig](#)). Both profile documents reside in the Alarm Manager application and in the AMS Accelerate infrastructure environment.

Topics

- [Configuration profile: monitoring](#)
- [Configuration profile: pseudoparameter substitution](#)
- [Accelerate alarm configuration examples](#)
- [Viewing your Accelerate Alarm Manager configuration](#)
- [Changing the Accelerate alarm configuration](#)

- [Modifying the Accelerate alarm default configuration](#)
- [Deploying Accelerate alarm configuration changes](#)
- [Rolling back Accelerate alarm changes](#)
- [Disabling the default Accelerate alarm configuration](#)

Configuration profile: monitoring

Both the default configuration profile document and the customization configuration profile document follow the same structure:

```
{
  "<ResourceType>": {
    "<ConfigurationID>": {
      "Enabled": true,
      "Tag": {
        "Key": "...",
        "Value": "..."
      },
      "AlarmDefinition": {
        ...
      }
    },
    "<ConfigurationID>": {
      ...
    }
  },
  "<ResourceType>": {
    ...
  }
}
```

- **ResourceType:** This key must be one of the following supported strings. The configuration within this JSON object will relate only to the specified AWS resource type. Supported resource types:

```
AWS::EC2::Instance
AWS::EC2::Instance::Disk
AWS::RDS::DBInstance
AWS::RDS::DBCluster
AWS::Elasticsearch::Domain
AWS::OpenSearch::Domain
```

```
AWS::Redshift::Cluster
AWS::ElasticLoadBalancingV2::LoadBalancer
AWS::ElasticLoadBalancingV2::LoadBalancer::TargetGroup
AWS::ElasticLoadBalancing::LoadBalancer
AWS::FSx::FileSystem::ONTAP
AWS::FSx::FileSystem::ONTAP::Volume
AWS::FSx::FileSystem::Windows
AWS::EFS::FileSystem
AWS::EC2::NatGateway
AWS::EC2::VPNConnection
```

- **ConfigurationID:** This key must be unique in the profile, and uniquely names the following block of configuration. If two configuration blocks in the same **ResourceType** block have the same **ConfigurationID**, the one that appears latest in the profile takes effect. If you specify a **ConfigurationID** in your customization profile that is the same as one specified in the default profile, the configuration block defined in the customization profile takes effect.
- **Enabled:** (optional, default=true) Specifies if the configuration block will take effect. Set this to false to disable a configuration block. A disabled configuration block behaves as if it's not present in the profile.
- **Tag:** Specifies the tag that this alarm definition applies to. Any resource (of the appropriate resource type) that has this tag key and value will have a CloudWatch alarm created with the given definition. This field is a JSON object with the following fields:
 - **Key:** The key of the tag to match. Keep in mind that if you're using Resource Tagger to apply the tags to the resource, the key for the tag will always begin with **ams:rt:**.
 - **Value:** The value of the tag to match.
- **AlarmDefinition:** Defines the alarm to be created. This is a JSON object whose fields are passed as is to the CloudWatch PutMetricAlarm API call (with the exception of pseudoparameters; for more information, see [Configuration profile: pseudoparameter substitution](#)). For information about what fields are required, see the [PutMetricAlarm](#) documentation.

OR

CompositeAlarmDefinition: Defines a composite alarm to be created. When you create a composite alarm, you specify a rule expression for the alarm that takes into account the alarm state of other alarms that you have created. This is a JSON object whose fields are passed as-is to the CloudWatchPutCompositeAlarm. The composite alarm goes into ALARM state only

if all conditions of the rule are met. The alarms specified in a composite alarm's rule expression can include metric alarms and other composite alarms. For information about what fields are required, see the [PutCompositeAlarm](#) documentation.

Both options provide the following fields:

- **AlarmName:** Specifies the name of the alarm you want to create for the resource. This field has all of the same rules as specified in the [PutMetricAlarm](#) documentation; however, since the alarm name must be unique in a Region, the Alarm Manager has one additional requirement: you must specify the unique identifier pseudoparameter in the name of the alarm (otherwise, Alarm Manager appends the unique identifier of the resource to the front of the alarm name). For example, for the **AWS::EC2::Instance** resource type, you must specify `${EC2::InstanceId}` in the alarm name, or it's implicitly added at the start of the alarm name. For the list of identifiers, see [Configuration profile: pseudoparameter substitution](#).

All other fields are as specified in the [PutMetricAlarm](#) or the [PutCompositeAlarm](#) documentation.

- **AlarmRule:** Specifies which other alarms are to be evaluated to determine this composite alarm's state. For each alarm that you reference, they have to be either exist in CloudWatch or specified in Alarm Manager configuration profile in your account.

Important

You can specify either **AlarmDefinition** or **CompositeAlarmDefinition** in your Alarm Manager configuration document, But they both can't be used at the same time.

In the following example, the system creates an alarm when two specified metric alarms exceeds its threshold:

```
{
  "AWS::EC2::Instance": {
    "LinuxResourceAlarm": {
      "Enabled": true,
      "Tag": {
        "Key": "ams:rt:mylinuxinstance",
        "Value": "true"
      }
    },
  },
}
```

```
"CompositeAlarmDefinition": {
  "AlarmName": "${EC2::InstanceId} Resource Usage High",
  "AlarmDescription": "Alarm when a linux EC2 instance is using too much CPU and
too much Disk",
  "AlarmRule": "ALARM(\"${EC2::InstanceId}: Disk Usage Too High -
${EC2::Disk::UUID}\") AND ALARM(\"${EC2::InstanceId}: CPU Too High\")"
}
}
}
```

Important

When Alarm Manager is not able to create or delete an alarm due to broke configuration, it sends the notification to the **Direct-Customer-Alerts** SNS topic. This alarm is called **AlarmDependencyError**.

We highly recommend that you have confirmed your subscription to this SNS topic. To receive messages published to [a topic](#), you must subscribe [an endpoint](#) to the topic. For details, see [Step 1: Create a topic](#).

Note

Many of the AMS Accelerate-provided baseline alarm definitions list the SNS topic, **MMS-Topic**, as a target. This is for use in the AMS Accelerate monitoring service, and is the transport mechanism for your alarm notifications to get to AMS Accelerate. Do not specify **MMS-Topic** as the target for any alarms other than those provided in the baseline (and overrides of the same), as the service ignores unknown alarms. It **does not** result in AMS Accelerate acting on your custom alarms.

Configuration profile: pseudoparameter substitution

In either of the configuration profiles, you can specify pseudoparameters that are substituted in place as follows:

- Global - anywhere in the profile:
 - `${AWS::AccountId}`: Replaced with your AWS account ID

- `${AWS::Partition}`: Replaced with the partition of the AWS Region the resource is in (this is 'aws' for most Regions); for more information, see the entry for partition in the [ARN reference](#).
- `${AWS::Region}`: Replaced with the Region name of the Region that your resource is deployed to (for example us-east-1)
- In an **AWS::EC2::Instance** resource type block:
 - `${EC2::InstanceId}`: (**identifier**) replaced by the instance ID of your Amazon EC2 instance.
 - `${EC2::InstanceName}`: replaced by the name of your Amazon EC2 instance.
- In an **AWS::EC2::Instance::Disk** resource type block:
 - `${EC2::InstanceId}`: (**identifier**) Replaced by the instance ID of your Amazon EC2 instance.
 - `${EC2::InstanceName}`: Replaced by the name of your Amazon EC2 instance.
 - `${EC2::Disk::Device}`: (**identifier**) Replaced by the name of the disk. (Linux only, on instances managed by the [CloudWatch Agent](#)).
 - `${EC2::Disk::FSType}`: (**identifier**) Replaced by the file system type of the disk. (Linux only, on instances managed by the [CloudWatchAgent](#)).
 - `${EC2::Disk::Path}`: (**identifier**) Replaced by the disk path. On Linux, this is the mount point of the disk (for example, /), while in Windows this is the drive label (for example, c:/) (only on an instance managed by the [CloudWatch Agent](#)).
 - `${EC2::Disk::UUID}`: (**identifier**) Replaced by a generated UUID that uniquely identifies the disk, this must be specified in the name of the alarm, as an alarm under `AWS::EC2::Instance::Disk` resource type will create one alarm per volume. Specifying `${EC2::Disk::UUID}` will maintain uniqueness of alarm names.
- In an **AWS::EKS::Cluster** resource type block:
 - `${EKS::ClusterName}`: (**identifier**) replaced by the name of your EKS Cluster.
- In an **AWS::OpenSearch::Domain** resource type block:
 - `${OpenSearch::DomainName}`: (**identifier**) replaced by the name of your EKS Domain.
- In an **AWS::ElasticLoadBalancing::LoadBalancer** resource type block:
 - `${ElasticLoadBalancing::LoadBalancer::Name}`: (**identifier**) replaced by the name of your V1 Load Balancer.
- In an **AWS::ElasticLoadBalancingV2::LoadBalancer** resource type block:
 - `${ElasticLoadBalancingV2::LoadBalancer::Arn}`: (**identifier**) replaced by the ARN of your V2 Load Balancer.

- `${ElasticLoadBalancingV2::LoadBalancer::Name}`: (**identifier**) replaced by the name of your V2 Load Balancer.
- `${ElasticLoadBalancingV2::LoadBalancer::FullName}`: (**identifier**) replaced by the full name of your V2 Load Balancer.
- In an **AWS::ElasticLoadBalancingV2::LoadBalancer::TargetGroup** resource type block:
 - `${ElasticLoadBalancingV2::TargetGroup::FullName}`: (**identifier**) replaced by the target group name of your V2 Load Balancer.
 - `${ElasticLoadBalancingV2::TargetGroup::UUID}`: (**identifier**) replaced by a generated UUID for your V2 Load Balancer.
- In an **AWS::EC2::NatGateway** resource type block:
 - `${NatGateway::NatGatewayId}`: (**identifier**) replaced by the NAT Gateway ID.
- In an **AWS::RDS::DBInstance** resource type block:
 - `${RDS::DBInstanceIdentifier}`: (**identifier**) replaced by your RDS DB instance identifier.
- In an **AWS::RDS::DBCluster** resource type block:
 - `${RDS::DBClusterIdentifier}`: (**identifier**) replaced by your RDS DB cluster identifier.
- In an **AWS::Redshift::Cluster** resource type block:
 - `${Redshift::ClusterIdentifier}`: (**identifier**) replaced by your Redshift cluster identifier.
- In an **AWS::Synthetics::Canary** resource type block:
 - `${Synthetics::CanaryName}`: (**identifier**) replaced by the name of your CloudWatch Synthetics canary.
- In an **AWS::EC2::VPNConnection** resource type block:
 - `${AWS::EC2::VpnConnectionId}`: (**identifier**) replaced by your VPN ID.
- In an **AWS::EFS::FileSystem** resource type block:
 - `${EFS::FileSystemId}`: (**identifier**) Replaced by the file system ID of your EFS file system.
- In an **AWS::FSx::FileSystem::ONTAP** resource type block:
 - `${FSx::FileSystemId}`: (**identifier**) Replaced by the file system ID of your FSX filesystem.
 - `${FSx::FileSystem::Throughput}`: Replaced by the throughput of your FSX file system.
 - `${FSx::FileSystem::Iops}`: Replaced by the IOPS of the FSX file system.
- In an **AWS::FSx::FileSystem::ONTAP::Volume** resource type block:
 - `${FSx::FileSystemId}`: (**identifier**) Replaced by the file system ID of your FSX file system.
 - `${FSx::ONTAP::VolumeId}`: (**identifier**) Replaced by the volume ID.

- In an **AWS::FSx::FileSystem::Windows** resource type block:
 - `${FSx::FileSystemId}`: (**identifier**) Replaced by the file system ID of your FSX file system.
 - `${FSx::FileSystem::Throughput}`: Replaced by the throughput of your FSX file system.

Note

All parameters marked with **identifier** are used as a prefix for the name of created alarms, unless you specify that identifier in the alarm name.

Accelerate alarm configuration examples

In the following example, the system creates an alarm for each disk attached to the matching Linux instance.

```
{
  "AWS::EC2::Instance::Disk": {
    "LinuxDiskAlarm": {
      "Tag": {
        "Key": "ams:rt:mylinuxinstance",
        "Value": "true"
      },
      "AlarmDefinition": {
        "MetricName": "disk_used_percent",
        "Namespace": "CWAgent",
        "Dimensions": [
          {
            "Name": "InstanceId",
            "Value": "${EC2::InstanceId}"
          },
          {
            "Name": "device",
            "Value": "${EC2::Disk::Device}"
          },
          {
            "Name": "fstype",
            "Value": "${EC2::Disk::FSType}"
          },
          {
            "Name": "path",
            "Value": "${EC2::Disk::Path}"
          }
        ]
      }
    }
  }
}
```

```

        }
      ],
      "AlarmName": "${EC2::InstanceId}: Disk Usage Too High -
${EC2::Disk::UUID}"
      ...
    }
  }
}

```

In the following example, the system creates an alarm for each disk attached to the matching Windows instance.

```

{
  "AWS::EC2::Instance::Disk": {
    "WindowsDiskAlarm": {
      "Tag": {
        "Key": "ams:rt:mywindowsinstance",
        "Value": "true"
      },
      "AlarmDefinition": {
        "MetricName": "LogicalDisk % Free Space",
        "Namespace": "CWAgent",
        "Dimensions": [
          {
            "Name": "InstanceId",
            "Value": "${EC2::InstanceId}"
          },
          {
            "Name": "objectname",
            "Value": "LogicalDisk"
          },
          {
            "Name": "instance",
            "Value": "${EC2::Disk::Path}"
          }
        ],
        "AlarmName": "${EC2::InstanceId}: Disk Usage Too High -
${EC2::Disk::UUID}"
        ...
      }
    }
  }
}

```

```
}
```

Viewing your Accelerate Alarm Manager configuration

Both the `AMSMangedAlarms` and `CustomerManagedAlarms` can be reviewed in AppConfig with [GetConfiguration](#).

The following is an example of the `GetConfiguration` call:

```
aws appconfig get-configuration --application AMSAlarmManager --environment
  AMSInfrastructure --configuration AMSManagedAlarms --client-id
  any-string outfile.json
```

- **Application:** this is AppConfig's logical unit to provide capabilities; for the Alarm Manager, this is `AMSAlarmManager`
- **Environment:** this is the `AMSInfrastructure` environment
- **Configuration:** to view AMS Accelerate baseline alarms, the value is `AMSMangedAlarms`; to view customer alarm definitions, the configuration is `CustomerManagedAlarms`
- **Client ID:** this is a unique application instance identifier, which can be any string
- The alarm definitions can be viewed in the specified output file, which in this case is `outfile.json`

You can see which version of configuration is deployed to your account by viewing the past deployments in the `AMSInfrastructure` environment.

Changing the Accelerate alarm configuration

To add or update new alarm definitions, you can either deploy configuration document [Using AWS CloudFormation to deploy configuration changes](#), or invoke the [CreateHostedConfigurationVersion](#) API.

This is a Linux command line command that generates the parameter value in base64, which is what the AppConfig CLI command expects. For information, see the AWS CLI documentation, [Binary/Blob \(binary large object\)](#).

As an example:

```
aws appconfig create-hosted-configuration-version --application-id application-id --
  configuration-profile-id configuration-profile-id --content base64-string
```

```
--content-type application/json
```

- **Application ID:** ID of the application AMS AlarmManager; you can find this out with the [ListApplications](#) API call.
- **Configuration Profile ID:** ID of the configuration CustomerManagedAlarms; you can find this out with the [ListConfigurationProfiles](#) API call.
- **Content:** Base64 string of the content, to be created by creating a document and encoding it in base64: `cat alarms-v2.json | base64` (see [Binary/Blob \(binary large object\)](#)).

Content Type: MIME type, `application/json` because alarm definitions are written in JSON.

Important

Restrict access to the [StartDeployment](#) and [StopDeployment](#) API actions to trusted users who understand the responsibilities and consequences of deploying a new configuration to your targets.

To learn more about how to use AWS AppConfig features to create and deploy a configuration, see [Working with AWS AppConfig](#).

Modifying the Accelerate alarm default configuration

While you can't modify the default configuration profile, you can provide overrides to the defaults by specifying a configuration block in your customization profile with the same **ConfigurationID** as the default configuration block. If you do this, your whole configuration block overwrites the default configuration block for which tagging configuration to apply.

For example, consider the following default configuration profile:

```
{
  "AWS::EC2::Instance": {
    "AMSManagedBlock1": {
      "Enabled": true,
      "Tag": {
        "Key": "ams:rt:ams-monitoring-policy",
        "Value": "ams-monitored"
      },
    },
    "AlarmDefinition": {
```

```

    "AlarmName": "${EC2::InstanceId}: AMS Default Alarm",
    "Namespace": "AWS/EC2",
    "MetricName": "CPUUtilization",
    "Dimensions": [
      {
        "Name": "InstanceId",
        "Value": "${EC2::InstanceId}"
      }
    ],
    "Threshold": 5,
    ...
  }
}
}
}
}

```

In order to change the threshold of this alarm to 10, **you must provide the entire alarm definition**, not only the parts you want to change. For example, you might provide the following customization profile:

```

{
  "AWS::EC2::Instance": {
    "AMSMangedBlock1": {
      "Enabled": true,
      "Tag": {
        "Key": "ams:rt:ams-monitoring-policy",
        "Value": "ams-monitored"
      }
    },
    "AlarmDefinition": {
      "AlarmName": "${EC2::InstanceId}: AMS Default Alarm",
      "Namespace": "AWS/EC2",
      "MetricName": "CPUUtilization",
      "Dimensions": [
        {
          "Name": "InstanceId",
          "Value": "${EC2::InstanceId}"
        }
      ],
      "Threshold": 10,
      ...
    }
  }
}
}
}
}

```

```
}
```

Important

Remember to deploy your configuration changes after you have made them. In SSM AppConfig, you must deploy a new version of the configuration after creating it.

Deploying Accelerate alarm configuration changes

After you finish a customization, you need to deploy it, either with AppConfig or AWS CloudFormation.

Topics

- [Using AppConfig to deploy Accelerate alarm configuration changes](#)
- [Using AWS CloudFormation to deploy configuration changes](#)

Using AppConfig to deploy Accelerate alarm configuration changes

Once the customization is completed, use AppConfig to deploy your changes with [StartDeployment](#).

```
aws appconfig start-deployment --application-id application_id
--environment-id environment_id --deployment-strategy-id
deployment_strategy_id --configuration-profile-id configuration_profile_id --
configuration-version 1
```

- **Application ID:** ID of the application `AMSAAlarmManager`, you can find this with the [ListApplications](#) API call.
- **Environment ID:** You can find this with the [ListEnvironments](#) API call.
- **Deployment Strategy ID:** You can find this with the [ListDeploymentStrategies](#) API call.
- **Configuration Profile ID:** ID of `CustomerManagedAlarms`; you can find this with the [ListConfigurationProfiles](#) API call.
- **Configuration Version:** The version of the configuration profile to be deployed.

⚠ Important

Alarm Manager applies the alarm definitions as specified in the configuration profiles. Any manual modifications you make with the AWS Management Console or CloudWatch CLI/SDK to the CloudWatch alarms is automatically reverted back, so make sure your changes are defined through Alarm Manager. To understand which alarms are created by the Alarm Manager, you can look for the `ams:alarm-manager:managed` tag with value `true`. Restrict access to the [StartDeployment](#) and [StopDeployment](#) API actions to trusted users who understand the responsibilities and consequences of deploying a new configuration to your targets.

To learn more about how to use AWS AppConfig features to create and deploy a configuration, see the [AWS AppConfig documentation](#).

Using AWS CloudFormation to deploy configuration changes

If you wish to deploy your `CustomerManagedAlarms` configuration profile using AWS CloudFormation, you can use the following CloudFormation templates. Put your desired JSON configuration in the `AMSAlarmManagerConfigurationVersion.Content` field.

When you deploy the templates in a CloudFormation Stack or Stack Set, the deployment of the `AMSResourceTaggerDeployment` resource will fail if you have not followed the required JSON format for the configuration. See [Configuration profile: monitoring](#) for details on the expected format.

For help on deploying these templates as a CloudFormation stack or stack set, see the relevant AWS CloudFormation documentation below:

- [Creating a stack on the AWS CloudFormation console](#)
- [Creating a stack with AWS CLI](#)
- [Creating a stack set](#)

📘 Note

If you deploy a configuration version using one of these templates, and then subsequently delete the CloudFormation stack/stack set, the template configuration version will remain as the current deployed version, and no additional deployment will be made. If you wish to

revert back to a default configuration, you will need to either manually deploy an empty configuration (i.e. just {}), or update your stack to an empty configuration, rather than deleting the stack.

JSON

```
{
  "Description": "Custom configuration for the AMS Alarm Manager.",
  "Resources": {
    "AMSAlarmManagerConfigurationVersion": {
      "Type": "AWS::AppConfig::HostedConfigurationVersion",
      "Properties": {
        "ApplicationId": {
          "Fn::ImportValue": "AMS-Alarm-Manager-Configuration-ApplicationId"
        },
        "ConfigurationProfileId": {
          "Fn::ImportValue": "AMS-Alarm-Manager-Configuration-CustomerManagedAlarms-ProfileID"
        },
        "Content": "{}",
        "ContentType": "application/json"
      }
    },
    "AMSAlarmManagerDeployment": {
      "Type": "AWS::AppConfig::Deployment",
      "Properties": {
        "ApplicationId": {
          "Fn::ImportValue": "AMS-Alarm-Manager-Configuration-ApplicationId"
        },
        "ConfigurationProfileId": {
          "Fn::ImportValue": "AMS-Alarm-Manager-Configuration-CustomerManagedAlarms-ProfileID"
        },
        "ConfigurationVersion": {
          "Ref": "AMSAlarmManagerConfigurationVersion"
        },
        "DeploymentStrategyId": {
          "Fn::ImportValue": "AMS-Alarm-Manager-Configuration-Deployment-StrategyID"
        },
        "EnvironmentId": {
          "Fn::ImportValue": "AMS-Alarm-Manager-Configuration-EnvironmentId"
        }
      }
    }
  }
}
```

```

    }
  }
}
}

```

YAML

Description: Custom configuration for the AMS Alarm Manager.

Resources:

AMSAlarmManagerConfigurationVersion:

Type: AWS::AppConfig::HostedConfigurationVersion

Properties:

ApplicationId:

!ImportValue AMS-Alarm-Manager-Configuration-ApplicationId

ConfigurationProfileId:

!ImportValue AMS-Alarm-Manager-Configuration-CustomerManagedAlarms-ProfileID

Content: |

```
{
```

```
}
```

ContentType: application/json

AMSAlarmManagerDeployment:

Type: AWS::AppConfig::Deployment

Properties:

ApplicationId:

!ImportValue AMS-Alarm-Manager-Configuration-ApplicationId

ConfigurationProfileId:

!ImportValue AMS-Alarm-Manager-Configuration-CustomerManagedAlarms-ProfileID

ConfigurationVersion:

!Ref AMSAlarmManagerConfigurationVersion

DeploymentStrategyId:

!ImportValue AMS-Alarm-Manager-Configuration-Deployment-StrategyID

EnvironmentId:

!ImportValue AMS-Alarm-Manager-Configuration-EnvironmentId

Rolling back Accelerate alarm changes

You can roll back alarm definitions through the same deployment mechanism by specifying a previous configuration profile version and running [StartDeployment](#).

Disabling the default Accelerate alarm configuration

AMS Accelerate provides the default configuration profile in your account based on the baseline alarms. However, this default configuration can be disabled by overriding any of the alarm definitions. You can disable a default configuration rule by overriding the **ConfigurationID** of the rule in your customization configuration profile and specifying the **enabled** field with a value of **false**.

For example, if the following configuration was present in the default configuration profile:

```
{
  "AWS::EC2::Instance": {
    "AMSMangedBlock1": {
      "Enabled": true,
      "Tag": {
        "Key": "ams:rt:ams-monitoring-policy",
        "Value": "ams-monitored"
      },
      "AlarmDefinition": {
        ...
      }
    }
  }
}
```

You could disable this tagging rule by including the following in your customization configuration profile:

```
{
  "AWS::EC2::Instance": {
    "AMSMangedBlock1": {
      "Enabled": false
    }
  }
}
```

To make these changes, the [CreateHostedConfigurationVersion](#) API must be called with the JSON profile document (see [Changing the Accelerate alarm configuration](#)) and subsequently must be deployed (see [Deploying Accelerate alarm configuration changes](#)). Note that when you create the new configuration version, you must also include any previously created custom alarms that you want in the JSON profile document.

⚠ Important

When AMS Accelerate updates the default configuration profile, it's not calibrated against your configured custom alarms, so review changes to the default alarms when you're overriding them in your customization configuration profile.

Creating additional CloudWatch alarms in Accelerate

You can create additional CloudWatch alarms for AMS Accelerate using custom CloudWatch metrics and alarms for Amazon EC2 instances.

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 CloudWatch monitoring scripts for Linux Amazon EC2 instances demonstrate how to produce and consume custom CloudWatch metrics. These sample Perl scripts comprise a fully functional example that reports memory, swap, and disk space utilization metrics for a Linux instance.

⚠ Important

AMS Accelerate does not monitor CloudWatch alarms created by you.

Viewing the number of resources monitored by Alarm Manager in Accelerate

Alarm Manager sends metrics every hour to Amazon CloudWatch, in the AMS/AlarmManager namespace. Metrics are emitted only for resource types supported by Alarm Manager.

Metric Name	Dimensions	Description
ResourceCount	Component, ResourceType	Number of resources (of the specified resource type) deployed in this Region. Units: Count

Metric Name	Dimensions	Description
Resources MissingManagedAlarms	Component, ResourceType	Number of resources (of the specified resource type) that require managed alarms, but Alarm Manager has not applied the alarms yet. Units: Count
Unmanaged Resources	Component, ResourceType	Number of resources (of the specified resource type) that do not have any managed alarms applied to them by Alarm Manager. Typically, these resources did not match any Alarm Manager configuration block, or are explicitly excluded from configuration blocks. Units: Count
MatchingResourceCount	Component, ResourceType, ConfigClauseName	Number of resources (of the specified resource type) that match the Alarm Manager configuration block. For a resource to match the configuration block, the block must be enabled, and the resource must have same tags specified in the configuration block. Units: Count

These metrics are also viewable as graphs, in the **AMS-Alarm-Manager-Reporting-Dashboard**. To see the dashboard, from the AWS CloudWatch management console, select **AMS-Alarm-Manager-Reporting-Dashboard**. By default, the graphs in this dashboard display the data for the prior 12-hour period.

AMS Accelerate deploys CloudWatch alarms to your account to detect significant increases in the number of unmanaged resources, for example, resources excluded from management by AMS Alarm Manager. AMS Operations will investigate increases in unmanaged resources that exceed: either 3 resources of the same type, or a 50% increase over all resources of the same type. If the change does not appear to be deliberate, AMS Operations may contact you to review the change.

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.
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.

Alert name	Description	Remediation
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.

Using Amazon EventBridge Managed Rules in AMS

AMS Accelerate 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 Accelerate 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
AmsAccess RolesRule	This rule listens for modifications in specific AMS Accelerate roles and policies.	<pre>{ "source": ["aws.iam"], "detail-type": ["AWS API Call via CloudTrail"], "detail": { "eventName": ["DeleteRole", "DeletePolicy", "CreatePolicyVersion", "AttachRolePolicy", "DetachRolePolicy"], "requestParameters": {</pre>

Rule Name	Description	Definition
		<pre> "\$or": [{ "roleName": ["ams-access-admin", "ams-access-admin-operations", "ams-access-operations", "ams-access-read-only", "ams-access-security-analyst", "ams-access-security-analyst- read-only"] }, { "policyArn": ["arn:*:iam::*:policy/ams-ac cess-allow-pass-role", "arn:*:iam::*:policy/ams-ac cess-deny-cloudshell-policy", "arn:*:iam::*:policy/ams-ac cess-deny-operations-policy", "arn:*:iam::*:policy/ams-ac cess-deny-update-iam-policy", "arn:*:iam::*:policy/ams-ac cess-ssr-policy", "arn:*:iam::*:policy/ams-ac cess-security-analyst-read-only-policy", "arn:*:iam::*:policy/ams-ac cess-security-analyst-policy", "arn:*:iam::*:policy/ams-ac cess-security-analyst-extended-policy", "arn:*:iam::*:policy/ams-ac cess-admin-policy", "arn:*:iam::*:policy/ams-ac cess-admin-operations-policy"] },] },] }, } </pre>

Rule Name	Description	Definition
AMSCoreRule	This rule forwards AWS Config and Amazon CloudWatch events to AMS Config remediation and AMS monitoring services respectfully. The AWS Config events create and resolve AWS Systems Manager OpsItems. The Amazon CloudWatch events monitor CloudWatch Alarms.	<pre data-bbox="597 226 1507 625">{ { "source": ["aws.config", "aws.cloudwatch"], "detail-type": ["Config Rules Compliance Change", "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.

Continuity management in AMS Accelerate

AMS leverages AWS Backup to centralize and automate the backing up of your data across AWS services. AMS backup plans provide best practices for various use cases; however, you are welcome to continue to use your existing backup plans. After you onboard to AMS backup management, AMS provides backup reports, and AMS experts continuously monitor your backup tasks to ensure you have a reliable backup solution.

To learn more, see [AWS Backup: How It Works](#) and [Supported AWS resources and third-party applications](#).

AMS Accelerate provides a range of operational services to help you achieve operational excellence on AWS. To gain a quick understanding of how AMS helps your teams achieve overall operational excellence in AWS Cloud with some of our key operational capabilities including 24x7 helpdesk, proactive monitoring, security, patching, logging and backup, see [AMS Reference Architecture Diagrams](#).

[Watch Carl's video to learn more \(9:29\)](#)

Topics

- [How continuity management works](#)
- [Select an AMS backup plan](#)
- [Tag your resources to apply AMS backup plans](#)
- [View backups in AMS vaults](#)
- [AMS backup monitoring and reporting](#)

How continuity management works

AMS backup plans define how frequently your data is backed up and the retention policy for your backups. AMS backup vaults keep your backup data organized. Once a resource is associated with a backup plan, [compatible resources](#) are incrementally backed up. The first backup is a full copy and subsequent backups capture incremental changes. Depending on the resource and AMS backup plan selected, [Point-In-Time-Recovery \(PITR\)](#) allows you to rewind your resources by selecting a time for your recovery. To get started with AMS Backup Management, simply select an AMS backup plan and tag your resources.

Note

Ensure that AWS Backup is enabled for each account, AWS Region, and resource type by following the steps here: [Getting Started 1: Service Opt-in](#).

You do not need to continue to *Getting started 2: Create on on-demand backup*.

Related Topics from AWS Backup

- [Working with backups \(Create, Edit, Copy, Restore, Delete\)](#)
- [Create an on-demand backup](#)
- [Creating backup copies across AWS Regions](#)
- [AWS Backup Supported Services](#)
- [Point-In-Time Recovery](#)
- [AWS Backup Features](#)

Select an AMS backup plan

AMS provides three different backup plans with a fourth backup plan to minimize cost during onboarding. To select an AMS backup plan for each supported resource, tag the resource with the plan's associated tag. As you onboard to Accelerate, AMS will work with you to identify the backup plan that best fits your needs.

Important

Do not edit your AMS default backup plans as your changes might be lost. Instead, create new plans for your custom configurations. For more information, see [Creating a backup plan](#).

Default AMS backup plan

AWS Backup "continuous backup" is not enabled for this backup plan; for details, see [Restoring to a specified time using Point-In-Time Recovery \(PITR\)](#).

TAG key: `ams:rt:backup-orchestrator`

TAG value: true

Default AMS backup plan	Start Time	Retention
hourly backup	N/A	N/A
daily backup	daily 4:00 UTC	7 days
weekly backup	Saturday, 2:00 UTC	4 weeks
monthly backup	1st of the month, 2:00 UTC	26 weeks
yearly backup	Jan 1st, 2:00 UTC	2 years

Enhanced backup plan

AWS Backup "continuous backup" is enabled with maximum retention (31 days) on supported resources; for details, see [Restoring to a specified time using Point-In-Time Recovery \(PITR\)](#) and [Supported services and applications for Point-In-Time Recovery \(PITR\)](#).

TAG key: `ams:rt:backup-orchestrator-enhanced`

TAG value: true

Enhanced 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 backup plan

AWS Backup "continuous backup" is enabled with maximum retention (31 days) on supported resources; for details, see [Restoring to a specified time using Point-In-Time Recovery \(PITR\)](#) and [Supported services and applications for Point-In-Time Recovery \(PITR\)](#).

TAG key: `ams:rt:backup-orchestrator-data-sensitive`

TAG value: `true`

Data Sensitive backup plan	Start Time	Retention
hourly backup	every hour	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 Accelerate onboarding backup plan

AWS Backup "continuous backup" is not enabled for this backup plan; for details, see [Restoring to a specified time using Point-In-Time Recovery \(PITR\)](#).

TAG key: `ams:rt:backup-orchestrator-onboarding`

TAG value: `true`

AMS Accelerate onboarding backup plan	Start Time	Retention
hourly backup	every hour	2 weeks
daily backup	N/A	N/A
weekly backup	N/A	N/A

AMS Accelerate onboarding backup plan	Start Time	Retention
monthly backup	N/A	N/A
yearly backup	N/A	N/A

Related AWS Backup Topics

- [Creating a backup plan](#)
- [Point-In-Time-Recovery \(PITR\)](#) enables continuous backups of supported resources and allows you to select a specific time for your recovery. For a list of supported resources, see [Feature availability by resource](#).

Tag your resources to apply AMS backup plans

To assign resources to an AMS backup plan, tag the resource with the plan's tag key-value pair. You can use AMS Resource Tagger to apply the AMS backup plans to a subset of your resources or for all supported resources in your account. If you want to use an alternate method to apply tags to your resources, such as AWS CloudFormation or Terraform, then turn off Resource Tagger so that it doesn't compete with your chosen tagging method. For more information, see [Preventing Resource Tagger from modifying resources](#).

The following example demonstrates how you can use Resource Tagger to apply the default AMS backup plan on Amazon Elastic Compute Cloud instances in your account. For this backup plan, apply the tag key `ams:rt:backup-orchestrator` and value `true`. To use a different backup plan, change the key to match your desired backup plan's tag key. To learn about AMS Resource Tagger and understand how to integrate the following referenced profile with the current (already configured) profile in your Accelerate account, see [Resource Tagger](#).

1. Open the AWS AppConfig console at <https://console.aws.amazon.com/systems-manager/appconfig>.
2. Choose the **ResourceTagger** application.
3. Choose the **Configuration profiles** tab, then choose **CustomerManagedTags**
4. Choose **Create** to create a new profile.
5. Choose **JSON**, and then copy and paste the following JSON object:

```
{
  "AWS::EC2::Instance": {
    "AccelerateBackupPlan": {
      "Enabled": true,
      "Filter": {
        "Fn::AND": [
          {
            "Platform": "*"
          }
        ]
      },
      "Tags": [
        {
          "Key": "ams:rt:backup-orchestrator",
          "Value": "true"
        }
      ]
    }
  }
}
```

6. Choose **Create hosted configuration version**.
7. Choose **Start deployment**.
8. Define the following deployment details:

```
Environment: AMSInfrastructure
Hosted configuration version: Select the version that you have just created.
Deployment Strategy: AMSNoBakeDeployment
```

9. Choose **Start deployment**. Resource Tagger tags your instances `ams:rt:backup-orchestrator: true`, ensuring that your instances are backed up in accordance with the default AMS backup plan.

View backups in AMS 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

ams-automated-backups

ams-automated-enhanced-backups

ams-automated-data-sensitive-backups

ams-onboarding-backups

AMS Backup Plan Tag Key

ams:rt:backup-orchestrator

ams:rt:backup-orchestrator-enhanced

ams:rt:backup-orchestrator-data-sensitive

ams:rt:backup-orchestrator-onboarding

Other AMS Vaults

AMS Vault Name

ams-manual-backups

ams-patch-backups (optional)

ams-custom-backups

Description

This vault contains manually started backups created by the `AWSManagedServices-StartBackupJob` SSM Automation document.

This vault is created for customers with the AMS patch add-on and contains pre-patch backups created by AMS patch automations before patching.

This is the recommended vault for backups created outside of AMS backup plans.

Related AWS Backup Topics

- [View Backups by Resource](#)
- [Working with backups](#)

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.

Patch management in AMS Accelerate

Important

Accelerate Patch reporting periodically deploys an AWS Glue resource-based policy. Be advised that AMS updates to the patching system overwrite existing AWS Glue resource-based policies.

You can use the AMS Accelerate patching system, Patch Add-On, to patch your instances with security-related and other types of updates. Accelerate Patch Add-On is a feature that provides tag-based patching for AMS instances. It 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. The AMS Accelerate Patch Add-On is an onboarding option, if you did not obtain it during onboarding your Accelerate account, contact your cloud service delivery manager (CSDM) to get it.

AMS Accelerate patch management uses the Systems Manager patch baseline functionality to control the definition of the patches that are applied on an instance. The patch baseline contains the list of patches that are pre-approved; for example, all security patches. The compliance of the instance is measured against the patch baseline associated to it. AMS Accelerate, by default, installs all patches available to keep the instance up to date.

Note

AMS Accelerate applies only operating system (OS) patches. For example, for Windows, only Windows updates are applied, not Microsoft updates.

For information on reports, see [AMS host management](#).

AMS Accelerate provides a range of operational services to help you achieve operational excellence on AWS. To gain a quick understanding of how AMS helps your teams achieve overall operational excellence in AWS Cloud with some of our key operational capabilities including 24x7 helpdesk, proactive monitoring, security, patching, logging and backup, see [AMS Reference Architecture Diagrams](#).

Topics

- [Patching recommendations](#)
- [Create a patch maintenance window](#)
- [AMS Accelerate patch baseline](#)
- [Creating an IAM role for on-demand patching](#)
- [Patch monitoring and failure remediation](#)

Patching recommendations

If you are involved in application or infrastructure operations, you understand the importance of an operating system (OS) patching solution that is flexible and scalable enough to meet the varied requirements from your application teams. In a typical organization, some application teams use an architecture that involves immutable instances whereas others deploy their applications on mutable instances.

For more information on AWS Prescriptive Guidance for patching, see [Automated patching for mutable instances in the hybrid cloud using AWS Systems Manager](#).

Note

Accelerate Patch Add-On is a feature that provides tag-based patching for AMS instances. It 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. The AMS Accelerate Patch Add-On is an onboarding option, if you did not obtain it during onboarding your Accelerate account, contact your cloud service delivery manager (CSDM) to get it.

Patch responsibility recommendations

The patching process for persistent instances should involve the following teams and actions:

- **The application (DevOps) teams** define the patch groups for their servers based on application environment, OS type, or other criteria. They also define the maintenance windows specific to each patch group. This information should be stored on tags attached to the instances. Recommended tag names are 'Patch Group' and 'Maintenance Window'. During each patch cycle, the application teams prepare for patching, test the application after patching, and troubleshoot any issues with their applications and OS during patching.

- **The security operations team** defines the patch baselines for various OS types that are used by the application teams, and make the patches available through Systems Manager Patch Manager.
- **The automated patching solution** runs on a regular basis and deploys the patches defined in the patch baselines, based on the user-defined patch groups and maintenance windows.
- **The governance and compliance teams** define patching guidelines and exception processes & mechanisms.

For more information, see [Patching solution design for mutable EC2 instances](#).

Guidance for application teams

- Review and become familiar with creating and managing maintenance windows; see [AWS Systems Manager Maintenance Windows](#) and [Create an SSM Maintenance window for patching](#) to learn more. Understanding the general structure and use of maintenance windows helps you understand what information to provide if you are not the person creating them.
- For High Availability (HA) setups, plan to have one maintenance window per availability zone and per environment (Dev/Test/Prod). This will ensure continued availability during patching.
- Recommended Maintenance Window duration is 4 hours with a 1-hour cutoff, plus 1 additional hour per 50 instances
- Patch Dev and Test versions with enough time between each to allow you to identify any potential issues prior to Production patching.
- Automate common pre- and post-patching tasks via SSM automation and run them as maintenance window tasks. Note that for post-patching tasks you must ensure that there is sufficient time allotted, as tasks will not launch once the cutoff is reached.
- Become familiar with Patch Baselines and their features—particularly around auto-approval delays for patch severity types that can be used to ensure that only the patches that were applied in Dev/Test get applied in Production at a later date. See [About patch baselines](#) for details.

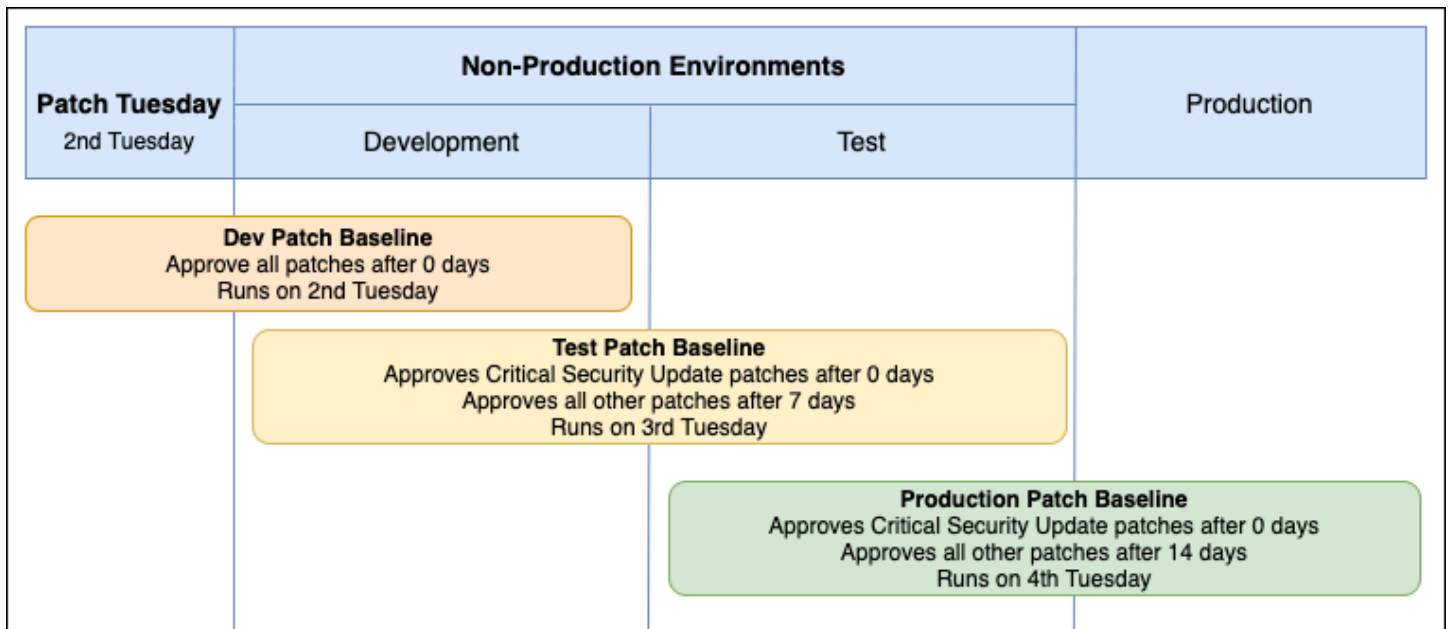
Guidance for security operations teams

- Review and become familiar with patch baselines. Patch approval is handled in an automated fashion and has different rule options. See [About patch baselines](#) for more information.
- Discuss needs around patching Dev/Test/Prod with application teams and develop multiple baselines to accommodate these needs.

Guidance for governance and compliance teams

- Patching should be an "Opt Out" function. A default maintenance window and automated tagging should exist to ensure nothing goes unpatched. AMS Resource Tagger can help with this; discuss this option with your cloud architect (CA) or cloud service delivery manager (CSDM) for guidance on implementation.
- Requests for exemption from patching should require documentation justifying the exemption. A Chief Information Security Officer (CISO) or other approval officer should approve or deny the request.
- Patching compliance should be reviewed on a regular schedule via the Patch Manager console, Security Hub, or a vulnerability scanner.

Example design for high availability Windows application



Overview:

- One Maintenance Window per AZ.
- One Set of Maintenance Windows per Environment.
- One Patch Baseline per Environment:
 - Dev: Approve all severity and classification after 0 days.

- Test: Approve critical security update patches after 0 days and all other severity and classifications after 7 days.
- Prod: Approve critical security update patches after 0 days and all other severity and classifications after 14 days.

CloudFormation Scripts:

These scripts are setup to build out the maintenance windows, baselines, and patching tasks for a two availability zone Windows HA EC2 application using the baseline approval settings described above.

- Windows Dev CFN Stack Example: [HA-Patching-Dev-Stack.json](#)
- Windows Test CFN Stack Example: [HA-Patching-Test-Stack.json](#)
- Windows Prod CFN Stack Example: [HA-Patching-Prod-Stack.json](#)

Patch recommendations FAQs

Q: How do I handle unscheduled patching for "0" day exploits?

A: SSM supports a **Patch Now** feature that uses the current default baseline for the instance's OS. AMS deploys a default set of Patch Baselines that approves all patches after 0 days. However, when using the **Patch Now** feature, a pre-patch snapshot is not taken, as this command runs the AWS-RunPatchBaseline SSM document. We recommend that you take a manual backup prior to patching.

Q: Does AMS support patching for instances in Auto-Scaling Groups (ASGs)?

A: No. At this time, ASG patching is not supported for Accelerate customers.

Q: Are there any limitations for Maintenance Windows to keep in mind?

A: Yes, there are a few limitations you should be aware of.

- Maintenance Windows per Account: 50
- Tasks per Maintenance Window: 20
- Maximum number of concurrent automations per Maintenance Window: 20
- Maximum number of concurrent Maintenance Windows: 5

For a full list of default SSM limits, see [AWS Systems Manager endpoints and quotas](#).

Create a patch maintenance window

During a maintenance window, the system creates a snapshot of each instance's root volume. If needed, AMS Operations engineers can use the snapshot to restore the instance's root volume.

Note

You must have the Accelerate Patch Add-On in order to use the Accelerate tag-based patching feature options, contact your cloud service delivery manager (CSDM) to get it.

Topics

- [Create a recurring "Patch Tuesday" maintenance window from the AMS console \(recommended\)](#)
- [Create a patch maintenance window using AWS CloudFormation](#)
- [Create a maintenance window from the Systems Manager console](#)
- [Create a maintenance window with the Systems Manager command line interface \(CLI\)](#)

Create a recurring "Patch Tuesday" maintenance window from the AMS console (recommended)

The AMS patch maintenance window console is for AMS Accelerate customers with the Accelerate Patch Add-On only. The patch maintenance windows created by this console do not function properly if you do not have the Patch Add-On. The Patch Add-On is an onboarding option, if you did not obtain it during onboarding your Accelerate account, contact your cloud service delivery manager (CSDM) to get it.

Microsoft releases patches for its operating systems on the second Tuesday of each month, also known as Patch Tuesday. It is common to schedule patching for both Windows and Linux instances relative to Patch Tuesday. To schedule recurring patch maintenance windows on the first or second weekends after Patch Tuesday, visit the AMS console and follow these steps:

1. Provide a name for your patch maintenance window.
2. [optional] Provide a description for the patch maintenance window.
3. Select a day relative to Patch Tuesday.

4. Enter a time for the patch maintenance window to start in hh:mm. For example, midnight is **00:00** and 11pm is **23:00**. Then select a timezone.
5. [optional] Change the duration to suit your needs. AMS recommends a four hour minimum duration.
6. Enter a patch tag key and value for the target. For information, see [What are tags?](#)
7. [optional] Expand the optional parameters to adjust concurrency, error rate, and maintenance window cut-off.
 1. Concurrency controls how many target instances are patching at the same time. For example, a 50% concurrency for 10 target instances will patch no more than 5 instances at a time, while 100% concurrency will patch all 10 at once.
 2. Error rate controls the tolerance for errors before patching is suspended. For example, an error rate of 100% for 10 target instances will patch all instances regardless of how many fail, while a 50% error rate will suspend patching once 5 instances have failed to patch. AMS recommends a 100% error rate.
 3. Patch maintenance window cutoff prevents breach of the patch maintenance window by suspending the start of new patching activities the specified hours before the end of the patch maintenance window. For example a cutoff of 1 hour (recommended), ceases new patch activities 1 hour before the end of the patch maintenance window.

Important

Verify the next execution time.

Visit the [SSM Maintenance Window console](#), search for your newly created patch maintenance window, and verify the next execution time. If you have any questions or need to edit your patch maintenance window, create a service request to talk with an AMS patch expert

To schedule a CRON-based patch maintenance window using CloudFormation, see [Create a patch maintenance window using AWS CloudFormation](#).

Create a patch maintenance window using AWS CloudFormation

These AWS CloudFormation templates are for use by AMS Accelerate customers with the Accelerate Patch Add-On only. These templates and patch maintenance windows created by the templates

do not function properly if you do not have the Patch Add-On. The Patch Add-On is an onboarding option, if you did not obtain it during onboarding your Accelerate account, contact your cloud service delivery manager (CSDM) to get it.

To create an AMS Accelerate patch maintenance window using AWS CloudFormation, first log into your Accelerate account and select the AWS Region where your target instances reside. Then follow these steps on the <https://console.aws.amazon.com/cloudformation>:

1. Select one of two custom Accelerate patching CloudFormation templates.
 - Patch Tuesday Scheduling: Microsoft releases patches for its operating systems on the second Tuesday of each month, also known as Patch Tuesday, to schedule patch maintenance windows on the first or second weekends after Patch Tuesday: Once logged into the Accelerate console, use this link [PatchTuesdayScheduling CloudFormation template](#).
 - CRON Scheduling: To create patch maintenance windows using CRON to define the start day, use this link [CRONScheduling CloudFormation template](#). Remember that Systems Manager CRON numbers days 1-7 (for details on Systems Manager CRON, see [Reference: Cron and rate expressions for Systems Manager](#)).

Choosing one of these links causes the template to load automatically on the CloudFormation console. Then click **Next**.

2. On the **Specify stack details** page (step 2 of the Create Stack pages), enter a stack name and template parameters (default parameters shown are AMS recommended defaults, select day and times for your use case). When finished, click **Next**.
3. Configure Stack Options (Optional). For information on the options, see [Setting AWS CloudFormation stack options](#). When finished, click **Next**.
4. Review your stack values (Optional). For information on reviewing stack details to estimate costs, see [Reviewing your stack and estimating stack cost](#). When ready, click **Create stack**.

The stack may take up to a minute to create. Once the stack is created successfully, your patch maintenance window runs at the specified time. You can make changes to your patch maintenance window by creating and executing a CloudFormation change set (recommended) (for details on doing this, see [Creating stacks using changesets](#)), or by updating the patch maintenance window on the Systems Manager **Maintenance window** console (<https://console.aws.amazon.com/systems-manager/maintenance-windows>).

[Watch Namrata's video to learn more \(5:41\)](#)

Create a maintenance window from the Systems Manager console

To create an AMS Accelerate maintenance window from the Systems Manager console, follow these steps:

1. In the left navigation bar in the **Change management** area, click **Maintenance Windows**, and then click **Create Maintenance Window** at the top right of the screen. Fill in the form. For details on any of the options, see [Create a maintenance window \(console\)](#). When finished, click **Create maintenance window**.

The maintenance windows list page opens.

2. Select the newly created maintenance window.

The maintenance window details page opens.

3. Go to the **Targets** tab and choose **Register target**.

The **Register target** page opens.

4. Add your Accelerate target. For information on targets, see [Assign targets to a maintenance window \(console\)](#). When finished, click **Register target**. Make note of the target as you need it later.

The maintenance windows details page reopens on the **Targets** tab with a list including the new target.

5. On the **Tasks** tab of the maintenance window details page, choose **Register Task**, and then pick **Register Automation task** from the drop-down list. Fill in the form. Accelerate notes:
 - Provide a meaningful task name. For example: AcceleratePatch.
 - In the **Automation document** area click in the search box, choose **Owner**, then **Shared documents**.
 - Select the automation document by clicking in the search box and choosing **Document name prefix --> Equals** and then typing: **AWSManagedServices-PatchInstance**. Then select the **AWSManagedServices-PatchInstance** document by selecting its radio button.
 - Under document version, choose **Default version at runtime**.
 - In the **Targets** section:
 - Set **Target by:** to **Selecting registered target groups**.

- In the list of targets, select the target you registered in the **Targets** tab.
- In the **Input parameters** section, fill in the form.
 - **InstanceId**: `{{TARGET_ID}}`
 - **StartInactiveInstances**: `True` to start the instances if they are stopped during the patch maintenance window.

 **Note**

The **InstanceId** parameter value is case sensitive and the **StartInactiveInstances** parameter value can be either `True` or `False`.
Stopped instances cannot be started when targeted by tags. For more information, see [No Invocations to Execute](#).

- In the **Rate control** section, choose percentages. AMS Accelerate recommends **100%** for **Concurrency** and **100%** for **Error threshold** to attempt to patch all instances simultaneously, regardless of automation errors. If you wish to patch half your targets at a time, for example, to keep a half of the target instances behind a load balance running, set **Concurrency** to 50%.
- In the **IAM service role** section, choose **Use a custom service role**, then choose the **ams_ssm_automation_role**.

Click **Register Automation task**.

The patching maintenance window is created. Under the **Description** tab, you can see the **Next execution time**.

Create a maintenance window with the Systems Manager command line interface (CLI)

To create an AMS Accelerate maintenance window with the command line interface:

1. Follow the SSM [Tutorial: Create and configure a maintenance window \(AWS CLI\)](#) . For each step of the tutorial, here are sample CLI commands for patching.


```
--name "AutomationExample" \  
--description "Sample Description" \  
--task-type=AUTOMATION
```

AMS Accelerate patch baseline

A patch baseline defines which patches are approved for installation on your instances. You can specify approved or rejected patches one by one. You can also create auto-approval rules to specify that certain types of updates (for example, critical updates) should be automatically approved. The rejected list overrides both the rules and the approve list.

Default patch baseline

When you onboard to AMS Accelerate patching, the default patch baselines are overridden by the AMS Accelerate default patch baselines for the following operating systems.

- **Windows**
- **Amazon Linux 1**
- **Amazon Linux 2**
- **CentOS**
- **Suse**
- **Rhel**

Important

Default patch baselines are managed by AMS. Do not edit default patch baselines as your changes may be lost. Instead, create a custom patch baseline. See [Custom patch baseline](#)

Note

The AMS Accelerate patch baselines defined as **product = *** mean that all patches are applied to the instance of all security and classifications.

Custom patch baseline

To use a custom patch baseline with AMS Accelerate, first ensure you have a patch group, and then create the custom baseline.

For more informations, see:

- [Working with patch groups](#)
- [Creating a custom patch baseline \(Windows\)](#)
- [Creating a custom patch baseline \(Linux\)](#)
- [Updating or deleting a custom patch baseline \(console\)](#)

Creating an IAM role for on-demand patching

After your account is onboarded to AMS Accelerate patching, AMS Accelerate deploys a managed policy, **amspatchmanagedpolicy**. This policy contains the required permissions for on-demand patching using the AMS automation document `AWSManagedServices-PatchInstance`. To use this automation document, the account administrator creates a IAM role for users. Follow these steps:

Create a role using the AWS Management Console:

1. Sign in to the AWS Management Console and open the [IAM console](#).
2. In the navigation pane of the console, choose **Roles**, then **Create role**.
3. Choose the **Another AWS account** role type.
4. For **Account ID**, enter the AWS account ID that you want to grant access to your resources.

The administrator of the specified account can grant permission to assume this role to any IAM user in that account. To do this, the administrator attaches a policy to the user, or a group, that grants permission for the **sts:AssumeRole** action. That policy must specify the role's Amazon Resource Name (ARN) as the resource. Note the following:

- If you are granting permissions to users from an account that you do not control, and the users will assume this role programmatically, then choose **Require external ID**. The external ID can be any word or number that is agreed upon between you and the administrator of the third-party account. This option automatically adds a condition to the trust policy that enables the user to assume the role only if the request includes the correct **sts:ExternalID**. For

more information, see [How to use an external ID when granting access to your AWS resources to a third party](#).

- If you want to restrict the role to users who sign in with multi-factor authentication (MFA), choose **Require MFA**. This adds a condition to the role's trust policy that checks for an MFA sign-in. A user who wants to assume the role must sign in with a temporary one-time password from a configured MFA device. Users without MFA authentication can't assume the role. For more information about MFA, see [Using multi-factor authentication \(MFA\) in AWS](#).

5. Choose **Next: Permissions**.

IAM includes a list of policies in the account. Under **Add Permissions**, enter **amspatchmanagedpolicy** in the filter box and select the checkbox for this permissions policy. Click **Next**.

6. Under **Role details**, enter a Role name such as PatchRole, add a description for the role (recommended), also add tags to help you identify this role. Role names aren't case sensitive, but must be unique within the AWS account. When finished, click **Create Role**.

Note

Role names can't be edited after they've been created.

Patch monitoring and failure remediation

AMS Accelerate Patch add-on monitors patching and remediates failures.

Patch notification

AMS Accelerate sends notifications for the patching maintenance windows configured, see [Create a patch maintenance window](#). The system sends CloudWatch Events as advance notice of each upcoming maintenance window and a CloudWatch Event at the end of each maintenance window. The advance notices are sent four days and one hour before the maintenance window. The advance notice CloudWatch Events are sent with the following schema:

```
{
  "version": "0",
  "id": "37004d81-458d-2cef-fe1c-8afa8af30406",
  "detail-type": "AMS Patch Window Execution State Change",
  "source": "aws.managedservices",
```

```

"account": "145917996532",
"time": "2021-05-20T02:00:00Z",
"region": "us-east-1",
"resources": [
  "arn:aws:ssm:us-east-1:123456789012:maintenancewindow/mw-00000001235",
  "arn:aws:ec2:us-east-1:123456789012:instance/i-0000000aaaaaaaa",
  "arn:aws:ec2:us-east-1:123456789012:instance/i-0000000aaaaaaaaab"
],
"detail": {
  "State": "PREEMPTIVE",
  "StartTime": "2021-05-24T02:00:00.000000",
  "WindowArn": "arn:aws:ssm:us-east-1:123456789012:maintenancewindow/
mw-00000001235",
  "Results": "[{\"instanceId\": \"i-0000000aaaaaaaa\"}, {\"instanceId\":
\\\"i-0000000aaaaaaaaab\\\"}]"
}
}

```

This table describes the advance notice event schema.

Patch notification details

Property name	Description	Sample values
State	The state of the patching maintenance window	PREEMPTIVE – The patching window scheduled to begin soon
Status	The status of the patching maintenance window	SUCCESS – All instances were patch without failure FAILED – At least one instance has failed to patch
StartTime	The start time, in ISO format, of the patching maintenance window	2021-02-03T22:14:05.814308
WindowArn	The unique identifier of the patching maintenance window.	arn:aws:ssm:us-east-1:123456789012:maint

Property name	Description	Sample values
		enancewindow/mw-0000001235
Results	The list of instances that will be targeted by the patch window	InstanceId – the instance ID targeted

The CloudWatch Events can be used to trigger a CloudWatch rule to notify you whenever a patching maintenance window advance notice is sent. In this case, you would configure the CloudWatch rule with the following:

```
{
  "source": [
    "aws.managedservices"
  ],
  "detail-type": ["AMS Patch Window Execution State Change"],
  "detail": {
    "State": ["PREEMPTIVE"]
  }
}
```

The window end CloudWatch Event is sent with the following schema:

```
{
  "version": "0",
  "id": "0f25add5-44a9-0702-d2bc-bd2102affefe",
  "detail-type": "AMS Patch Window Execution State Change",
  "source": "aws.managedservices",
  "account": "123456789012",
  "time": "2021-02-03T22:14:06Z",
  "region": "us-east-1",
  "resources": [
    "arn:aws:ssm:us-east-1:123456789012:maintenancewindow/mw-00000001235",
    "arn:aws:ec2:us-east-1:123456789012:instance/i-0000000aaaaaaaa",
    "arn:aws:ec2:us-east-1:123456789012:instance/i-0000000aaaaaaaaab"
  ],
  "detail": {
    "State": "[COMPLETED]",
    "Status": "SUCCESS",
    "StartTime": "2021-02-03T22:12:00.814308",
    "EndTime": "2021-02-03T22:14:05.814309",
  }
}
```

```

    "WindowArn": "arn:aws:ssm:us-east-1:123456789012:maintenancewindow/
mw-00000001235",
    "WindowExecutionId": "e32088eb-c05f-4c63-b766-6866e163c818",
    "Results": "[{\"instanceId\": \"i-00000000aaaaaaaaaa\", \"status\":
\\\"Success\\\", \"missing_critical_patch_count\": 0, \"missing_total_patch_count
\\\": 0} }, {\"instanceId\": \"i-00000000aaaaaaaaab\", \"status\": Success},
\\\"missing_critical_patch_count\": 0, \"missing_total_patch_count\": 0}]"
  }
}

```

Patch window end details

Property name	Description	Sample values
State	The state of the patching maintenance window	COMPLETED – The patching window is finished
Status	The status of the patching maintenance window	SUCCESS – All instances were patch without failure FAILED – At least one instance has failed to patch
StartTime	The start time, in ISO format, of the patching maintenance window	2021-02-03T22:14:05.814308
EndTime	The end time, in ISO format, of the patching maintenance window	2021-02-03T23:14:05.814308
WindowArn	The unique identifier of the patching maintenance window.	arn:aws:ssm:us-east-1:123456789012:maintenancewindow/mw-0000001235
WindowExecutionId	The window execution ID, which can be seen from the SSM Maintenance Window Console	e32088eb-c05f-4c63-b766-6866e163c818

Property name	Description	Sample values
Results	The list of instances that will be targeted by the patch window	Instanceid – the instance ID targeted status – the instance patch status missing_critical_patch_count - the count of critical patches missing on the instance missing_total_patch_count - the count of total patches missing on the instance

Note

Patch failure alerts aren't created for instances that have unsupported operating systems, or that are stopped during the maintenance window.

Patch remediation

AWS Managed Services (AMS) manages patching and includes patch failure remediation. When patch fails, AMS Operations is alerted and they proceed to remediate by following AWS and AMS best practices to address the issue.

AMS creates an SSM OpsItem in the account with the following title: **AWS Managed Services – Patch Instance failure for instance <instance-id>**. AMS Operation will resolve the OpsItem if the situation can be corrected without your intervention. If they need your intervention, a service request is created and notice sent to the account owner to collaborate and address the issue. Upon remediation, the OpsItem is resolved.

If no action is taken, the patching maintenance window runs on the next cycle and attempts to patch the instance again.

Note

Patch failure alerts aren't created for instances that have unsupported operating systems, or that are stopped during the maintenance window.

Monitoring and Incident Management for Amazon EKS

Monitoring and Incident Management for Amazon EKS monitors your Amazon EKS resources for failures, performance degradation, and security issues. AMS Accelerate configures and deploys Amazon Managed Service for Prometheus alert manager rules, monitors the alerts, and then performs incident management when these alerts are triggered. Monitoring and Incident Management for EKS relies on AMS Alarm Manager and leverages native AWS services, such as [Amazon Managed Service for Prometheus](#), [Amazon Managed Grafana](#), [Amazon GuardDuty](#), [AWS Lambda](#), and [AWS Config](#).

Note

Monitoring and Incident Management for EKS doesn't support AWS GovCloud (US), Windows nodes, or Windows containers.

What is Monitoring and Incident Management for Amazon EKS?

Monitoring and Incident Management for EKS provides the following:

- A default configuration that creates, manages, and deploys monitors and policies across your managed account for Amazon EKS clusters that you select.
- A monitoring baseline to allow your Amazon EKS workloads to have increased availability, even if you don't configure any other monitoring for your Amazon EKS clusters. For more information, see [Baseline alerts](#).
- Notifications that are generated by the baseline monitoring configured for your Amazon EKS cluster. These notifications are known as alerts. Alerts are generated when there are imminent, on-going, receding, or potential failures, performance degradation, or security issues. Examples of alerts include a Prometheus alert, an event, or a finding from an AWS service, such as Amazon GuardDuty.
- Alert investigation with guidance on appropriate remediation actions that you can take. For more information, see [Incident reports and service requests in AMS Accelerate](#).
- Remediation of alerts and incidents by AMS operations, when possible and with your approval, to prevent or reduce the impact to your applications. For more information, see [Incident reports and service requests in AMS Accelerate](#).

- Optional predefined Amazon Managed Grafana dashboards that provide visibility into resource utilization, performance, health of CoreDNS, active alerts, and previously resolved alerts. If you configure Amazon Managed Grafana using the AMS-provided template, then you can open the Amazon Managed Grafana console to view metrics and alerts for your Amazon EKS cluster.

How Monitoring and Incident Management for Amazon EKS works

Generation: As part of onboarding Monitoring and Incident Management for EKS, AMS configures baseline monitoring for the Amazon EKS clusters that you selected in your managed account. AMS uses a combination of Amazon Managed Service for Prometheus alert manager rules and Amazon CloudWatch event rules to configure baseline monitoring. An AMS-configured Prometheus server in your cluster scrapes and remote-writes your Prometheus metrics to an Amazon Managed Service for Prometheus endpoint in the same Region. The baseline monitoring configuration generates an alert when a Prometheus alert manager rule is triggered or a CloudWatch event is generated.

Aggregation: AMS sends all alerts that your resources generate to the AMS monitoring system by directing them to an Amazon Simple Notification Service topic that's managed by AMS.

Processing and impact analysis: AMS analyzes the alerts and then processes them based on their potential for impact. AMS classifies the alerts as follows:

- **Alerts with known customer impact:** For these alerts, AMS creates a new incident report using the [incident management](#) process.
- **Alerts with uncertain customer impact:** For these alerts, AMS sends an incident report. In many cases, these alerts ask you to verify the impact before AMS can take action. For such alerts, AMS sends an [alert notification](#) with the details and checks whether the alert needs a mitigating action. AMS provides options for mitigating actions in the notification. If your reply confirms that the alert is an incident, AMS then triggers the creation of a new incident report and initiates the 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. Also, the corresponding alert is marked as resolved.
- **Alerts with no customer impact:** If, after evaluation, AMS determines that the alert doesn't have any customer impact, the alert is closed.

AMS responsibility matrix (RACI)

The AMS responsible, accountable, consulted, and informed, or RACI matrix assigns the primary responsibility to either the customer or AMS for a variety of activities. The following table provides an overview of the responsibilities of customer and AMS for activities in an application that uses Monitoring and Incident Management for Amazon EKS.

- **R** stands for the responsible party that does the work to achieve the task.
- **A** stands for the accountable party.
- **C** stands for consulted; the party whose opinions are sought, typically as subject matter experts; and with whom there is bilateral communication.
- **I** stands for informed; the party which is informed on progress, often only on completion of the task or deliverable.

Activity	Customer	AMS
Discovery for AMS requirements	I	R
Enable AMS permissions (RBAC) for cluster access	R	C
Install Amazon EC2 Systems Manager Agent on worker nodes if it isn't already present	R	C
Deploy AMS on-cluster components, such as Prometheus, Prometheus Node Exporter, and kube-state-metrics in an AMS namespace, as needed.	C	R
Provision Amazon Managed Service for Prometheus in the AMS control plane	I	R

Activity	Customer	AMS
Configure Prometheus alert manager in the AMS control plane	I	R
Provide Amazon Managed Grafana template and assist with configuration	C	R
Enable GuardDuty EKS Audit Log Monitoring	C	R
Enable Amazon EKS control plane logging	I	R
Monitor Amazon EKS control plane health and performance	I	R
Monitor Amazon EKS cluster health and performance (cluster, node, workload, pod, API Server and CoreDNS)	I	R
Triage alerts and provide incident response for Amazon EKS	I	R
Run diagnostic commands during incidents	I	R
Analyze logs during incidents (control plane and pod logs)	I	R
Incident response for AWS network issues	I	R

Activity	Customer	AMS
Respond to GuardDuty EKS Audit Log Monitoring findings	I	R
Provide customer guidance on actions to remediate incidents when possible	I	R

Baseline alerts

After verifying the alerts, AMS enables the following alerts for Amazon EKS and then engages in monitoring and incident management for your selected Amazon EKS clusters. The response time Service Level Agreements (SLAs) and Service Level Objectives (SLOs) are dependent on your selected account Service Tier (Plus, Premium). For more information, see [Incident reports and service requests in AMS Accelerate](#).

Alerts and actions

The following table lists the Amazon EKS alerts and respective actions that AMS takes:

Alert	Action
Container OOM killed	AMS investigates whether the OOM kill is caused because of reaching container limit or memory limit overcommit, and then advises you on corrective actions.
Pod Job Failed	AMS investigates why the Kubernetes job or corresponding cron job is failing, and then advises you on corrective actions.
StatefulSet Down	AMS determines why pods aren't ready by reviewing error messages in pod events and error log snippets in pod logs, and then advises you on corrective actions.

Alert	Action
HPA Scaling Ability	AMS determines which Kubernetes Horizontal Pod Autoscaler (HPA) is unable to scale pods for its subsequent workload resource, such as a Deployment or StatefulSet.
HPA Metric Availability	AMS determines why HPA can't collect metrics, such as metrics related to server configuration issues or RBAC authorization issues.
Pod Not Ready	AMS investigates affected pod(s) for details, reviews pod logs for related errors and events, and then advises you on corrective actions.
Pod Crash Looping	AMS investigates the reasons for pod not starting, such as insufficient resources, a file locked by another container, database locked by another container, service dependencies failing, DNS issues for external services, and misconfigurations.
Daemonset Mis-scheduled	AMS determines the reasons for a Daemonset pod not being scheduled on a node, such as the taint and toleration configurations.
Kubernetes API Errors	AMS analyzes control plane logs to determine the volume and types of errors that are causing this alert, and identifies any resource contention issues for master node or etcd autoscaling groups. If the API server doesn't recover, AMS engages the Amazon EKS service team.

Alert	Action
Kubernetes API Latency	AMS analyzes control plane logs to determine the volume and types of errors that are causing latency and identifies any resource contention issues for master node or etcd auto-scaling groups. If the API server doesn't recover, AMS engages the Amazon EKS service team.
Kubernetes Client Cert Expiring	AMS sends this notification to inform you that your cluster certificate will expire in 24 hours.
Node Not Ready	AMS investigates the node conditions and events, such as network issues, that prevent kubelet access to the API server.
Node High CPU	AMS determines whether one or more pods are consuming an unusually high amount of CPU. Then, AMS verifies with you that your requests, limits, and pod activity are as expected.
Node OOM Kill Detected	AMS determines if the OOM kill is caused because of reaching the container limit or node overcommit. If the application activity is normal, AMS advises you on requests and limits for overcommits and revising pod limits.
Node Contrack Limit	AMS advises you on the recommended contrack value per core. Kubernetes nodes set the contrack max value proportional to the total memory capacity of the node. High load applications, especially on smaller nodes, can easily exceed the contrack max value, resulting in connection resets and timeouts.

Alert	Action
Node Clock Not in Sync	AMS determines whether Network Time Protocol (NTP) is installed and functioning properly.
Pod High CPU	AMS investigates pod logs to determine the pod tasks that consume a high amount of CPU.
Pod High Memory	AMS investigates pod logs to determine the pod tasks that consume a high amount of memory.
CoreDNS Down	This is a critical alert that indicates that the domain name resolution for internal or external cluster services stopped. AMS checks the status of the CoreDNS pods, verifies CoreDNS configuration, verifies DNS endpoints that point to CoreDNS pods, verifies CoreDNS limits, and with your approval, enables CoreDNS debug logging.
CoreDNS Errors	This alert might signal an issue with an application or a misconfiguration. AMS checks the status of the CoreDNS pods, verifies CoreDNS configuration, verifies DNS endpoints that point to CoreDNS pods, verifies CoreDNS limits, and with your approval, enables CoreDNS debug logging.
CoreDNS Latency	This alert Signals that CoreDNS might be overloaded. AMS checks the status of CoreDNS pods, verifies CoreDNS configuration, verifies DNS endpoints that point to the CoreDNS pods, verifies CoreDNS limits, and with your approval, enables CoreDNS debug logging.


Alert	Action
CoreDNS Forwarding Latency	When CoreDNS isn't the authoritative server or doesn't have a cache entry for a domain name, CoreDNS forwards the DNS request to an upstream DNS server. This alert signals that CoreDNS might be overloaded or there might be an issue with an upstream DNS server. AMS checks the status of CoreDNS pods, verifies CoreDNS configuration, verifies DNS endpoints that point to CoreDNS pods, verifies CoreDNS limits, and with your approval, enables CoreDNS debug logging.
CoreDNS Forwarding Error	When CoreDNS isn't the authoritative server or doesn't have a cache entry for a domain name, CoreDNS forwards the DNS request to an upstream DNS server. This alert signals a possible misconfiguration or an issue with an upstream DNS server. AMS checks the status of CoreDNS pods, verifies CoreDNS configuration, verifies DNS endpoints that point to CoreDNS pods, verifies CoreDNS limits, and with your approval, enables CoreDNS debug logging.

Requirements

- **Supported Kubernetes versions:** See [Amazon EKS Kubernetes versions](#) in the Amazon EKS User Guide.
- **Node types:** Amazon EKS managed nodes are supported. Windows nodes and containers aren't supported.
- **Kubernetes cluster access:** AMS requires system:masters RBAC cluster role and cluster user.
- **SSM Agent on Amazon EC2 nodes:** Both Bottle Rocket and Amazon EKS AMIs have SSM Agent pre-installed. Be sure that SSM Agent is installed on your custom AMIs and Amazon EC2 nodes.

- **Service Quotas** For more information, see the service quotas for [Amazon Managed Service for Prometheus](#) and [Amazon Managed Grafana](#).
- **Supported AWS Regions:**

Region name	Region
US East (Ohio)	us-east-2
US East (N. Virginia)	us-east-1
US West (Oregon)	us-west-2
Asia Pacific (Singapore)	ap-southeast-1
Asia Pacific (Sydney)	ap-southeast-2
Europe (Ireland)	eu-west-1
Europe (London)	eu-west-2

 **Note**

Amazon EKS clusters must reside in a single Amazon VPC per AWS Region.

Onboarding

1. **Enable Amazon EKS cost optimization tags:** See [Tagging your resources for billing](#) in the Amazon EKS User Guide.
2. **Initiate onboarding of Monitoring and Incident Management for EKS:** Contact your Cloud Service Delivery Manager (CSDM) with account IDs and cluster names to onboard.
3. **Validate requirements:** Your Cloud Architect (CA) validates that all [requirements](#) are met before onboarding begins.
4. **Update Kubernetes role-based access control (RBAC):** AMS shares the `eksctl` commands to implement these changes. You can review these changes and then deploy. You must deploy RBAC updates so that AMS has permissions to run commands on your behalf. These updates

include mapping the AMS IAM role to a Kubernetes user, creating a new Kubernetes cluster role for AMS, and binding the AMS Kubernetes cluster role to the user.

5. **Deploy cluster components:** AMS deploys the following components in an AMS-managed namespace on your cluster:
 - Prometheus server
 - Prometheus Node exporter (not applicable for AWS Fargate)
 - kube-state-metrics
6. **Perform Prometheus configuration updates:** AMS configures Prometheus to enable remote-write for metrics.
7. **(Optional) Configure dashboards:** Your CA helps you configure Amazon Managed Grafana dashboards in your account.

Important

After your cluster is onboarded, AMS suppresses alerts while we review existing issues with your cluster. After existing issues are addressed, AMS alerts are enabled for the cluster.

Offboarding

Notify your Cloud Service Delivery Manager (CSDM) with account IDs and cluster names to start the offboarding process. After you offboard, alert processing, metric storage, and metric querying are suspended and metrics are deleted in accordance with the default [Amazon Managed Service for Prometheus data retention policies](#).

AMS performs the following offboarding steps:

1. AMS disables alerts that are sent to you and AMS Operations.
2. AMS removes the Prometheus instance from your Amazon EKS cluster.
3. AMS removes other AWS resources that are installed in your account, such as IAM roles and AWS Config rules.

After these steps are completed, you must complete the following offboarding steps:

1. Use `eksctl` to remove the Kubernetes RBAC permissions from the `aws-auth` ConfigMap.

2. If you previously installed it, remove the Amazon Managed Grafana instance that you configured to connect to AMS.

Cost optimization with AMS Resource Scheduler

The AMS Resource Scheduler on AWS solution helps you reduce your AWS and AMS costs by stopping resources that are not in use and starting resources when capacity is needed. For example, you can use AMS Resource Scheduler on AWS in a development environment to automatically stop instances outside of business hours every day. If you leave all of your instances running at full utilization, this solution can result in reducing the instance utilization, which reduces overall cost based on the schedules you configure.

Use AWS Managed Services (AMS) Resource Scheduler to schedule the automatic start and stop of Auto Scaling groups, Amazon EC2 instances, and Amazon 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 [AWS Instance Scheduler](#) but contains additional features and customizations specific to AMS customer needs. The customization includes support for scheduling Auto Scaling groups, CloudWatch alarm suppressor for Elastic Load Balancing alarms, support for multiple AWS Systems Manager maintenance windows for Amazon EC2, a cost savings estimator, and operational support from AMS.

AMS Resource Scheduler uses periods and schedules. Periods define the times the resource should run, such as start time, end time, and days of the month. Schedules contain your defined periods, along with additional configurations—SSM maintenance window, timezone, hibernate, etc—and specify when resources should run. You can configure these periods and schedules using AMS-provided AWS Systems Manager automation runbooks. Each schedule must contain at least one period that defines the time(s) the instance should run. A schedule can contain more than one period. When more than one period is used in a schedule, the Instance Scheduler applies the appropriate start action when at least one of the period rules is true. For more details on schedule and periods, see [Solutions Components of AWS Instance Scheduler](#).

AMS Resource Scheduler uses AWS resource tags to associate a schedule to one or more resources in order to target them for scheduled start and stop actions. You tag your resources with the tag key (default is `Schedule`) configured in the Scheduler with the schedule name as the value. You configure the same tag key as the cost allocation tag in AWS Cost Explorer for the cost estimator feature of Scheduler to track and report on cost savings.

AMS Resource Scheduler is an opt in feature that you can enable per account.

Using resources with AMS Resource Scheduler

Amazon EC2

- Amazon EC2 instances that are part of an Auto Scaling group aren't processed individually and skipped by AMS Resource Scheduler, even if they are tagged.
- If the target instance root volume is encrypted with a AWS 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, you can add the permission by updating the CloudFormation stack, `ams-resource-scheduler`, with a list of CMK as value to the `UseCMK` parameter (use one or more CMK Key ARNs in the format `arn:partition:kms:region:account-id:key/key-id` instead of a KMS alias).
- If your Amazon EC2 instances are configured with specific software, or vendor licences managed by AWS License Manager, Resource Scheduler needs permissions to the specific AWS License Manager licences to be able to start the instance. You can grant Resource Scheduler the necessary permissions by adding the list of ARN(s) of the AWS License Manager licences to the `License manager license for EC2 instance` parameter of the CloudFormation stack (`ams-resource-scheduler`).

Amazon EC2 Auto Scaling

- 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.

Amazon RDS

- 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.

Note

The Resource Scheduler doesn't validate that a resource is started or stopped. It makes the API call for the relevant service and moves on. If the API call fails, it logs the error for investigation.

AMS Resource Scheduler does not support AWS Backup window. If you map an AWS Backup-enabled RDS instance with Resource Scheduler schedule, for the backup to work as expected, the backup window must lie within the running window of the schedule.

Onboarding AMS Resource Scheduler

Your account is not automatically onboarded to AMS Resource Scheduler when your account is onboarded to the AMS Accelerate operations plan. However, as part of account onboarding to the AMS Accelerate operation plan, or anytime after, you can request your Cloud Service Delivery Manager (CSDM) to onboard the account to AMS Resource Scheduler. Once your CSDM onboards the account, a CloudFormation stack containing AMS Resource Scheduler resources with default configuration, is automatically provisioned into your account.

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.

Customizing AMS Resource Scheduler

When onboarded, AMS Resource Scheduler is deployed as a CloudFormation stack, with name `ams-resource-scheduler`, in the primary AWS region for your AMS Accelerate account. You can configure the properties of AMS Resource Scheduler based on your preferences through CloudFormation stack parameters and performing a stack update. For information on updating CloudFormation stacks, see [Updating stacks directly](#).

We recommend you customize the following properties and leave the rest at default for optimal functionality.

- **Tag name:** The name of the tag that Resource Scheduler will use to associate instance schedules with resources. The default value is `Schedule`.
- **Service(s) to schedule:** 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 time zone:** Specify the default time zone for the Resource Scheduler to use. The default value is UTC.
- **CMK for encrypted EBS volumes:** A comma-separated list of Amazon KMS Customer Managed Key (CMK) ARNs that Resource Scheduler can be granted permissions to.
- **License manager license for EC2 instance:** 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 stack via stack parameters are preserved.

We strongly recommend against making any customization directly to any of the component resource of AMS Resource Scheduler. Doing so impacts Resource Scheduler functionality and AMS's ability to keep it up to date.

Using AMS Resource Scheduler

How to use AMS Resource Scheduler periods in AMS Accelerate accounts.

Use the following set of AWS Systems Manager automation runbooks to administer the required schedule and period in AMS Resource Scheduler.

Note

These SSM automation runbooks are available in the primary AWS Region of your account.

- `AWSManagedServices-AddOrUpdatePeriod`
- `AWSManagedServices-AddOrUpdateSchedule`
- `AWSManagedServices-DeleteScheduleOrPeriod`
- `AWSManagedServices-DescribeScheduleOrPeriods`
- `AWSManagedServices-EnableOrDisableAMSResourceScheduler`

Additionally, AMS provisions an AWS Identity and Access Management role, `ams_resource_scheduler_ssm_automation_role`, that AWS Systems Manager requires and assumes in order to use the runbooks. The IAM role is scoped down with a least privilege inline policy granting SSM permissions required for the functionality of the runbooks.

Prerequisites

Perform the following steps before you begin using the SSM automation runbook and AMS Resource Scheduler.

Attach the following policy to the appropriate IAM entity (user, group or role) that you want to allow to use the automation runbooks to administer the schedule and period in AMS Resource Scheduler. *The policy is not required if your IAM entity has Administrator or PowerUser permission in your account.*

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowPassingResourceSchedulerRole",
      "Effect": "Allow",
```

```

        "Action": "iam:PassRole",
        "Resource": "arn:aws:iam::ACCOUNT_ID:role/
ams_resource_scheduler_ssm_automation_role",
        "Condition": {
            "StringEquals": {
                "iam:PassedToService": "ssm.amazonaws.com"
            }
        }
    },
    {
        "Sid": "ListAndDescribeAutomationExecutions",
        "Effect": "Allow",
        "Action": [
            "ssm:GetAutomationExecution",
            "ssm:DescribeAutomationStepExecutions"
        ],
        "Resource": "arn:aws:ssm:*:ACCOUNT_ID:automation-execution/*"
    },
    {
        "Sid": "ListAndDescribeResourceSchedulerSSMDocuments",
        "Effect": "Allow",
        "Action": [
            "ssm:ListDocumentVersions",
            "ssm:DescribeDocument",
            "ssm:ListDocumentMetadataHistory",
            "ssm:DescribeDocumentParameters",
            "ssm:GetDocument",
            "ssm:DescribeDocumentPermission"
        ],
        "Resource": [
            "arn:aws:ssm::*:document/AWSManagedServices-AddOrUpdatePeriod",
            "arn:aws:ssm::*:document/AWSManagedServices-AddOrUpdateSchedule",
            "arn:aws:ssm::*:document/AWSManagedServices-DeleteScheduleOrPeriod",
            "arn:aws:ssm::*:document/AWSManagedServices-DescribeScheduleOrPeriods",
            "arn:aws:ssm::*:document/AWSManagedServices-
EnableOrDisableAMSResourceScheduler"
        ]
    },
    {
        "Sid": "AllowExecutionOfResourceSchedulerSSMDocuments",
        "Effect": "Allow",
        "Action": [
            "ssm:StartAutomationExecution"
        ],
    },

```

```

    "Resource": [
      "arn:aws:ssm::*:automation-definition/AWSManagedServices-
AddOrUpdatePeriod:*",
      "arn:aws:ssm::*:automation-definition/AWSManagedServices-
AddOrUpdateSchedule:*",
      "arn:aws:ssm::*:automation-definition/AWSManagedServices-
DeleteScheduleOrPeriod:*",
      "arn:aws:ssm::*:automation-definition/AWSManagedServices-
DescribeScheduleOrPeriods:*",
      "arn:aws:ssm::*:automation-definition/AWSManagedServices-
EnableOrDisableAMSResourceScheduler:*"
    ]
  },
  {
    "Sid": "AllowListingAllDocuments",
    "Effect": "Allow",
    "Action": "ssm:ListDocuments",
    "Resource": "*"
  },
  {
    "Sid": "AllowListingAllSSMExecutions",
    "Effect": "Allow",
    "Action": "ssm:DescribeAutomationExecutions",
    "Resource": "*"
  },
  {
    "Sid": "AllowListingIAMRolesForStartingExecutionViaConsole",
    "Effect": "Allow",
    "Action": "iam:ListRoles",
    "Resource": "*"
  }
]
}

```

You can run the automation either from AWS Systems Manager console or using the AWS CLI. If using the AWS CLI, you may need to install and configure it or the AWS tools for PowerShell, if you haven't already. For information, see [Install or upgrade AWS command line tools](#).

[Watch Navish's video to learn more \(4:52\)](#)

Working with periods and schedules in AWS Managed Services Resource Scheduler

AMS Resource Scheduler and periods in AMS Accelerate accounts.

Adding or updating periods in AMS Resource Scheduler

Add or update a Resource Scheduler period in AMS accounts.

Data you'll need:

- **Action:** The type of operation to perform. Use "add" if you want to add a period or "update" if you want to update an existing period.
- **Name:** The name of the period. You must specify a unique value if you are adding a new period.
- **ams_resource_scheduler_ssm_automation_role:** The ARN of the AWS Identity and Access Management (IAM) role that allows the runbook to add or update the period on your behalf. Specify the role as `ams_resource_scheduler_ssm_automation_role`.
- **Description** (Optional): A meaningful description for the period.
- **BeginTime** (Optional): The time in HH:MM format when you want to start the resources.
- **EndTime** (Optional): The time in HH:MM format when you want to stop the resources.
- **Months** (Optional): A comma-delimited list of months or a hyphenated range of months during which the resources should run.
- **MonthDays** (Optional): A comma-delimited list of days of the month or a hyphenated range of days during which the resources should run.
- **WeekDays** (Optional): A comma-delimited list of the days of the week or a range of days of the week during which the resources should run.

How to do it:

- View the document at [AWSManagedServices-AddOrUpdatePeriod](#) (you may have to choose your onboarded Region).

Specify requirements in the **Input parameters** section, then choose **Execute**. After the operation completes, view results in the **Output** tab.

- AWS CLI:

Run the following command to start an automation. Replace *placeholders* with your own information.

```
aws ssm start-automation-execution --document-name "AWSManagedServices-AddOrUpdatePeriod" --document-version "\$DEFAULT"
  --parameters '{"Action":["add" or "update"], "Name":["NAME"],
  "Description":["DESCRIPTION"],"BeginTime":["TIME"], "EndTime":["TIME"],
  "Months":["MONTH"],"MonthDays":["DAY"], "WeekDays":["DAY"],
  "ams_resource_scheduler_ssm_automation_role" : ["arn:aws:iam::ACCOUNTID:role/ams_resource_scheduler_ssm_automation_role"] }' --region $AWS_REGION
```

Example:

The following example shows how you can add a new period using the AWS Systems Manager console. We have named the period **Period-Name** and configured it to cover 9AM-6PM from Mon-Fri for first 15 days of every month.

1. View the AWS Systems Manager automation document at [AWSManagedServices-AddOrUpdatePeriod](#) (you may have to choose your onboarded Region).

The screenshot shows the AWS Systems Manager console interface for the document 'AWSManagedServices-AddOrUpdatePeriod'. The breadcrumb navigation is 'AWS Systems Manager > Documents > AWSManagedServices-AddOrUpdatePeriod'. The document title is 'AWSManagedServices-AddOrUpdatePeriod' with buttons for 'Delete', 'Actions', and 'Execute automation'. Below the title are tabs for 'Description', 'Content', 'Versions', and 'Details'. The 'Description' tab is active, showing 'Document version 8 (Default)'. A 'Document description' section contains a table with the following data:

Platform	Created	Owner	Target type
Windows, Linux, MacOS	Mon, 28 Feb 2022 10:12:22 GMT	Amazon	/

Below the table, the status is 'Active' with a green checkmark icon. The description text reads: 'Add or Update period definition for AMS Resource Scheduler.'

2. Provide values for the parameters.

Input parameters	
<p>Action (Required) Choose what action.</p> <input type="text" value="add"/>	<p>Name (Required) The name used to identify the period. This name must be unique.</p> <input type="text" value="PeriodName"/>
<p>Description (Optional) The description of the period rule.</p> <input type="text" value="Description for the Period to be added"/>	<p>BeginTime (Optional) The time, in HH:MM format, that the instance will start.</p> <input type="text" value="09:00"/>
<p>EndTime (Optional) The time, in HH:MM format, that the instance will stop.</p> <input type="text" value="18:00"/>	<p>Months (Optional) Enter a comma-delimited list of months, or a hyphenated range of months, during which the instance will run.</p> <input type="text" value="jan-dec"/>
<p>MonthDays (Optional) Enter a comma-delimited list of days of the month, or a hyphenated range of days, during which the instance will run.</p> <input type="text" value="1-15"/>	<p>WeekDays (Optional) Enter a comma-delimited list of days of the week, or a range of days of the week, during which the instance will run.</p> <input type="text" value="0-4"/>
<p>AutomationAssumeRole (Required) The ARN of the role that allows Automation to perform the actions on your behalf.</p> <input type="text" value="ams_resource_scheduler_ssm_automation_role"/>	

3. Click **Execute** and wait for automation to complete.

Adding or updating schedules in AMS Resource Scheduler

Add or update a Resource Scheduler schedule in AMS Accelerate accounts.

Data you'll need:

- **Action:** The type of operation to perform. Use "add" if you want to add a schedule or "update" if you want to update an existing schedule.
- **Name:** The name of the schedule. You must specify a unique value if you are adding a new schedule.
- **ams_resource_scheduler_ssm_automation_role:** The ARN of the AWS Identity and Access Management (IAM) role that allows the runbook to add or update the schedule on your behalf. Specify the role `ams_resource_scheduler_ssm_automation_role`.
- **Description (Optional):** A meaningful description for the schedule.
- **Schedules (Optional):** Specify a comma-delimited list of periods that are to be used with this schedule. Each period must have already been created.
- **RetainRunning (Optional):** Specify "true" to prevent Resource Scheduler from stopping a running resource at the end of a running period if the resource was manually started before the beginning of the running period. By default, Resource Scheduler stops the resource.

- **StopNewInstances** (Optional): Specify "false" to prevent Resource Scheduler from stopping a resource the first time it is tagged if it is running outside of the running period. By default, Resource Scheduler stops the resource.
- **SSMMaintenanceWindow** (Optional): Specify a comma-delimited list of AWS Systems Manager (SSM) maintenance windows that you want to add as running periods for the schedule. You must also specify the "UseMaintenanceWindow" property to "true."
- **TimeZone** (Optional): Specify the time zone that you want Resource Scheduler to use. By default, Resource Scheduler uses UTC.
- **UseMaintenanceWindow** (Optional): Specify "true" if you want to Resource Scheduler to consider Amazon Relational Database Service (RDS) maintenance window as a running period to an Amazon RDS instance schedule, or to add AWS Systems Manager (SSM) maintenance windows as a running period to an Amazon EC2 instance schedule.
- **UseMetrics** (Optional): Specify "true" to enable CloudWatch metrics at the schedule level and "false" to disable CloudWatch metrics. Specifying this property overrides the CloudWatch metrics setting set at the stack level.

How to do it:

- View the document at [AWSManagedServices-AddOrUpdateSchedule](#) (you may have to choose your onboarded Region).

Specify requirements in the **Input parameters** section, then choose **Execute**. After the operation completes, view results in the **Output** tab.

- AWS CLI:

Run the following command to start an automation. Replace *placeholders* with your own information.

```
aws ssm start-automation-execution --document-name "AWSManagedServices-AddOrUpdateSchedule" --document-version "\$DEFAULT" --parameters '{"Action":["add" or "update"], "Name":["NAME"], "Description":["DESCRIPTION"], "Hibernate":["true or false"], "Enforced":["true or false"], "OverrideStatus":["running or stopped"], "Periods":["PERIOD-A, PERIOD-B"], "RetainRunning":["true or false"], "StopNewInstances":["true or false"], "SSMMaintenanceWindow":["WINDOW-NAME"], "TimeZone":["TIMEZONE"], "UseMaintenanceWindow":["true or false"], "UseMetrics":["true or false"],
```

```
"ams_resource_scheduler_ssm_automation_role" : ["arn:aws:iam::ACCOUNTID:role/ams_resource_scheduler_ssm_automation_role"] }' --region $AWS_REGION
```

Example:

The following example shows how to add a schedule for AMS Resource Scheduler. In this example we add a schedule with name CustomSchedule using CustomPeriod.

1. View the AWS Systems Manager automation document at [AWSManagedServices-AddOrUpdateSchedule](#) (you may have to choose your onboarded Region).

AWSManagedServices-AddOrUpdateSchedule
Delete
Actions ▾
Execute automation

Description
Content
Versions
Details

Document version
10 (Default)

Document description

Platform	Created	Owner	Target type
Windows, Linux, MacOS	Mon, 28 Feb 2022 10:12:22 GMT	Amazon	/

Status
✔ Active

Add or update schedule definition for AMS Resource Scheduler.

2. Provide values for the parameters.

<div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p>Action <small>(Required) Choose what action.</small></p> <input type="text" value="add"/> </div> <div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p>Description <small>(Optional) The description of the schedule.</small></p> <input type="text" value="String"/> </div> <div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p>Enforced <small>(Optional) Whether to enforce the schedule. When this field is set to true, the scheduler will stop a running instance if it is manually started outside of the running period or it will start an instance if it is stopped manually during the running period.</small></p> <input type="text" value="true"/> </div> <div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p>Periods <small>(Optional) Comma separated one or more name of the periods that are used in this schedule. The name must match with the existing configured periods.</small></p> <input type="text" value="CustomPeriod"/> </div> <div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p>StopNewInstances <small>(Optional) Whether to stop an instance the first time it is tagged if it is running outside of the running period. By default, this field is set to true.</small></p> <input type="text"/> </div> <div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p>TimeZone <small>(Optional) The time zone the schedule will use. If no time zone is specified, the default time zone (UTC) is used.</small></p> <input type="text"/> </div> <div style="border: 1px solid #ccc; padding: 5px;"> <p>UseMetrics <small>(Optional) Whether to enable CloudWatch metrics at the schedule level. This field overwrites the CloudWatch metrics setting you specified at deployment.</small></p> <input type="text" value="true"/> </div>	<div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p>Name <small>(Required) The name used to identify the schedule. This name must be unique.</small></p> <input type="text" value="CustomSchedule"/> </div> <div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p>Hibernate <small>(Optional) Whether to hibernate EC2 instances that are enabled for hibernation and meets hibernation requirements.</small></p> <input type="text"/> </div> <div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p>OverrideStatus <small>(Optional) Whether to override the current schedule action. If set to running, the instance will be started but not stopped until it is manually stopped. Similarly when set to stopped, the instance will be stopped but not started automatically until manually started.</small></p> <input type="text"/> </div> <div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p>RetainRunning <small>(Optional) Whether to prevent the solution from stopping an instance at the end of a running period if the instance was manually started before the beginning of the period.</small></p> <input type="text"/> </div> <div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p>SSMMaintenanceWindow <small>(Optional) Whether to add an AWS Systems Manager maintenance window as a running period. Enter the name of a maintenance window. To use this field, you must also set the UseMaintenanceWindow parameter to true.</small></p> <input type="text" value="String"/> </div> <div style="border: 1px solid #ccc; padding: 5px; margin-bottom: 5px;"> <p>UseMaintenanceWindow <small>(Optional) Whether to add an Amazon RDS maintenance window as a running period to an Amazon RDS instance schedule, or to add an AWS Systems Manager maintenance window as a running period to an Amazon EC2 instance schedule.</small></p> <input type="text" value="false"/> </div> <div style="border: 1px solid #ccc; padding: 5px;"> <p>AutomationAssumeRole <small>(Required) The ARN of the role that allows Automation to perform the actions on your behalf.</small></p> <input type="text" value="ams_resource_scheduler_ssm_automation_role"/> </div>
--	---

3. Click **Execute** and wait for automation to complete.

Deleting periods or schedules in AMS Resource Scheduler

In order to delete Resource Scheduler periods or schedules in AMS Accelerate accounts, you need the following data:

- **ConfigurationType:** The type of configuration you want to delete. Use "period" if you want to delete periods or "schedule" if you want to delete schedules.
- **Name:** The name of the schedule or period that you want to delete.
- **ams_resource_scheduler_ssm_automation_role:** The ARN of the AWS Identity and Access Management (IAM) role that allows the runbook to delete schedules or periods on your behalf. Specify the role `ams_resource_scheduler_ssm_automation_role`.

How to do it:

- View the document at [AWSManagedServices-DeleteScheduleOrPeriod](#) (you may have to choose your onboarded Region).

Specify requirements in the **Input parameters** section, then choose **Execute**. After the operation completes, view results in the **Output** tab.

- AWS CLI:

Run the following command to start an automation. Replace *placeholders* with your own information.

```
aws ssm start-automation-execution --document-name "AWSManagedServices-DeleteScheduleOrPeriod" --document-version "\$DEFAULT" --parameters '{"ConfigurationType":["period" or "schedule"],"Name":["NAME"],"ams_resource_scheduler_ssm_automation_role":["arn:aws:iam::ACCOUNTID:role/ams_resource_scheduler_ssm_automation_role"]}' --region $AWS_REGION
```

Example:

The following example shows how you can delete a period using the AWS Systems Manager console.

1. View the AWS Systems Manager automation document at [AWSManagedServices-DeleteScheduleOrPeriod](#) (you may have to choose your onboarded Region).

The screenshot shows the AWS Systems Manager console interface for the automation document `AWSManagedServices-DeleteScheduleOrPeriod`. The breadcrumb trail is `AWS Systems Manager > Documents > AWSManagedServices-DeleteScheduleOrPeriod`. The document title is `AWSManagedServices-DeleteScheduleOrPeriod`. There are buttons for `Delete`, `Actions` (with a dropdown arrow), and `Execute automation`. Below the title, there are tabs for `Description` (selected), `Content`, `Versions`, and `Details`. The document version is `8 (Default)`. The `Document description` section contains a table with the following data:

Platform	Created	Owner	Target type
Windows, Linux, MacOS	Mon, 28 Feb 2022 12:19:40 GMT	Amazon	/

Below the table, the `Status` is `Active` (indicated by a green checkmark icon). The description text reads: `Delete a schedule or period configuration of AMS Resource Scheduler.`

2. Provide values for the parameters.

The screenshot shows the `Input parameters` section of the automation document. It contains three required parameters:

- `ConfigurationType`: (Required) Whether to delete schedule or period. The value is `period`.
- `AutomationAssumeRole`: (Required) The ARN of the role that allows Automation to perform the actions on your behalf. The value is `ams_resource_scheduler_ssm_automation_role`.
- `Name`: (Required) Name of the schedule or period to delete. The value is `PeriodName`.

3. Click **Execute** and wait for automation to complete.

Describing periods or schedules in AMS Resource Scheduler

In order to describe (view a details on) a Resource Scheduler period or schedule in AMS Accelerate accounts, you need the following data:

- **ConfigurationType**: The type of configuration you want to describe. Use "period" if you want to describe all periods or "schedule" if you want to describe all schedules.
- **ams_resource_scheduler_ssm_automation_role**: The ARN of the AWS Identity and Access Management (IAM) role that allows the runbook to describe schedules or periods on your behalf. Specify the role `ams_resource_scheduler_ssm_automation_role`.

How to do it:

- View the document at [AWSManagedServices-DescribeScheduleOrPeriods](#) (you may have to choose your onboarded Region):
 1. Specify requirements in the **Input parameters** section and then choose **Execute**.
 2. After the operation completes, view results in the **Output** tab.
- AWS CLI:
 1. Run the following command to start an automation. Replace *placeholders* with your own information.

```
aws ssm start-automation-execution --document-name "AWSManagedServices-DescribeScheduleOrPeriods" --document-version "\$DEFAULT"
    --parameters '{"ConfigurationType":["period"
    or "schedule"],"ams_resource_scheduler_ssm_automation_role":
["arn:aws:iam::ACCOUNTID:role/ams_resource_scheduler_ssm_automation_role"]}'
    --region $AWS_REGION
```

Example:

The following example shows how you can describe a period using the AWS Systems Manager console.

1. View the AWS Systems Manager automation document at [AWSManagedServices-DescribeScheduleOrPeriods](#) (you may have to choose your onboarded Region).

The screenshot shows the AWS Systems Manager console interface for the automation document 'AWSManagedServices-DescribeScheduleOrPeriods'. At the top right, there are buttons for 'Delete', 'Actions', and 'Execute automation'. Below the title, there are tabs for 'Description', 'Content', 'Versions', and 'Details'. The 'Description' tab is selected. The document version is '8 (Default)'. The 'Document description' section contains a table with the following data:

Platform	Created	Owner	Target type
Windows, Linux, MacOS	Tue, 01 Mar 2022 00:15:13 GMT	Amazon	/

Below the table, the status is shown as 'Active' with a green checkmark icon. At the bottom, there is a brief description: 'Describe schedule or period configurations for AMS Resource Scheduler.'

2. Provide values for the parameters.

Input parameters	
ConfigurationType (Required) Whether to list config, periods or schedules.	AutomationAssumeRole (Required) The ARN of the role that allows Automation to perform the actions on your behalf.
periods	ams_resource_scheduler_ssm_automation_role

3. Click **Execute** and wait for automation to complete.

Tagging resources for AMS Resource Scheduler

Tagging resources for AMS Resource Scheduler.

Once you add schedules and periods to AMS Resource Schedule, you need to tag your resources with the Resource Scheduler tag name as the tag key, or the your customized one, and the schedule name as the tag value. For details on how to tag your resources in your AMS Accelerate account, see [Tagging in AMS Accelerate](#).

Note

If Resource Tagger is used to tag the resources, the default Tag key for Resource Scheduler must be customized to have the prefix 'ams:rs:' as all tags applied by the resource tagger have the key prefix 'ams:rs:'. Otherwise, the resources tagged with resource tagger will not be managed by Resource Scheduler. To know more about customizing the default Tag key for Resource Scheduler, see [Customizing AMS Resource Scheduler](#).

Cost estimator in AMS Resource Scheduler

In order to track cost savings, AMS Resource Scheduler features a component that hourly calculates the estimated cost savings for Amazon EC2 and Amazon 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 $24 \times 0.125 = \$3$. 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.

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.

Note

Cost Savings Estimator do 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](#).

Alarm suppressor in AMS Resource Scheduler

AMS Resource Scheduler comes with a CloudWatch Alarm suppressor, that is deployed as separate Lambda function named `AMSAlarmSuppressor` that suppresses alarms for instances that are behind an Elastic Load Balancing, Application Load Balancer, or Network Load Balancer. The function runs every 5 minutes, retrieves all alarms present in the account and groups them based on namespace; for example, `AWS/ELB`, `AWS/ApplicationELB`, `AWS/NetworkELB`. For each

group of alarms the suppressor finds the load balancer name and/or target group (for ALB/NLB) from alarm dimensions, finds the instances that are registered with the load balancer and/or target group, and checks the instance state to discover if the instances are scheduled by AMS Resource Scheduler. If instances are scheduled by Resource Scheduler, and are stopped by Resource Scheduler, the suppressor then marks the alarms to disable them. If at least one instance in the registered instance list is running, suppressor marks corresponding alarms to enable the alarms that are marked for enabling, and disable alarms that are marked for disabling. Logs for this are stored in the `/aws/lambda/AMSAAlarmSuppressor` log group.

Log management in AMS Accelerate

AMS Accelerate configures supported AWS services to collect logs. These logs are used by AMS Accelerate to ensure compliance and auditing of resources within your account.

AMS Accelerate provides a range of operational services to help you achieve operational excellence on AWS. To gain a quick understanding of how AMS helps your teams achieve overall operational excellence in AWS Cloud with some of our key operational capabilities including 24x7 helpdesk, proactive monitoring, security, patching, logging and backup, see [AMS Reference Architecture Diagrams](#).

Topics

- [Log management — AWS CloudTrail](#)
- [Log management — Amazon EC2](#)
- [Log management — Amazon VPC Flow Logs](#)

Log management — AWS CloudTrail

[AWS CloudTrail](#) is a service that is used for account governance: compliance, operational auditing, and risk auditing. With CloudTrail, you can log, continuously monitor, and retain account activity related to actions across your AWS infrastructure.

AMS Accelerate relies on AWS CloudTrail logging in order to manage audits and compliance for all resources in your account. During onboarding, you choose whether Accelerate deploys a CloudTrail multi-region trail in your primary AWS Region or uses events generated by your account or Organization trail. If your account does not have a trail configured, then Accelerate will deploy a CloudTrail multi-region trail during onboarding. If you choose to integrate Accelerate with your CloudTrail trail, work with your Cloud Architect (CA) to review and configure your trail resources to [Accelerate required configurations](#), and enable Accelerate to [use Athena to query and analyze events](#).

AMS Accelerate creates an Amazon S3 bucket for an Accelerate deployed CloudTrail trail as the events delivery destination and uses AWS Key Management Service (AWS KMS) encryption. Your trail events are accessed by AMS Accelerate operators for investigation and diagnosis purposes. If the account already has an existing CloudTrail trail enabled, this trail is in addition to that, if you chose to have Accelerate deploy an Accelerate managed trail during onboarding.

AMS Accelerate deploys AWS Config rules to ensure that your CloudTrail account trails, including an Accelerate deployed CloudTrail trail are correctly set up and encrypted. To learn more, see [AWS Config](#). These are the rules used, presented as links to the AWS documentation describing them:

- [multi-region-cloudtrail-enabled](#). Checks that AMS Accelerate CloudTrail is properly set up with the correct configurations.
- [cloud-trail-encryption-enabled](#). Checks that AWS CloudTrail is configured to use the server-side encryption (SSE) with AWS KMS customer master key (CMK) encryption.
- [cloud-trail-log-file-validation-enabled](#). When enabled, checks that AWS CloudTrail creates a signed digest file with logs. We strongly recommend that you enable file validation on all trails.
- [s3-bucket-default-lock-enabled](#). When enabled, checks that the Amazon S3 bucket has lock enabled.
- [s3-bucket-logging-enabled](#). When enabled, checks whether logging is enabled for Amazon S3 buckets.

AMS Accelerate uses AWS KMS to encrypt the logged events for an Accelerate deployed CloudTrail trail in your account. This key is controlled by, and is accessible to, the account administrators, AMS Accelerate operators, and CloudTrail. For more information about AWS KMS, see [AWS Key Management Service features](#) product documentation.

Accessing and auditing CloudTrail logs

CloudTrail logs for an AMS Accelerate deployed CloudTrail trail are stored in an Amazon S3 bucket within your account. Trail data stored in the Amazon S3 bucket is encrypted using a AWS KMS key created when CloudTrail resources are provisioned.

Amazon S3 buckets leverage a naming pattern of **ams-aaws account id-cloudtrail-AWS Region**, (example: **ams-a123456789-cloudtrail-us-east-1a**) and all the events are stored with the **AWS/CloudTrail** prefix. All access to the primary bucket is logged and the log objects are encrypted and versioned for auditing purposes.

For more information about tracking changes and querying the logs, see [Tracking changes in your AMS Accelerate accounts](#).

Protecting and retaining CloudTrail logs

AMS Accelerate enables Amazon S3 object locking with Governance Mode for an Accelerate deployed CloudTrail trail to ensure that users can't overwrite or delete an object version or alter its lock settings without special permissions. For more information, see [Amazon S3 object locking](#).

By default, all logs in this bucket are kept indefinitely. If you want to change the retention period, you can submit a service request through the [AWS Support Center](#) to set up a different retention policy.

Accessing Amazon EC2 logs

You can access Amazon EC2 instance logs by using the AWS Management Console. Logs produced by instances and AWS services are available in CloudWatch Logs, which is available in each account managed by AMS Accelerate. For information about accessing your logs, see the [CloudWatch Logs documentation](#).

Retaining Amazon EC2 logs

Amazon EC2 instance logs are kept indefinitely, by default. If you want to change the retention period, you can submit a service request through the [AWS Support Center](#) to set up a different retention policy.

Log management — Amazon EC2

AMS Accelerate installs the CloudWatch agent on all Amazon EC2 instances that you have identified as AMS Accelerate-managed. This agent sends system-level logs to Amazon CloudWatch Logs. For information, see [What are Amazon CloudWatch Logs?](#)

The following log files are sent to CloudWatch Logs, into a log group of the same name as the log. Within each log group, a log stream is created for each Amazon EC2 instance, named according to the Amazon EC2 instance ID.

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/maillog
- /var/log/messages
- /var/log/secure
- /var/log/spooler
- /var/log/yum.log
- /var/log/zypper.log

For more information, see [Manually Create or Edit the CloudWatch Agent Configuration File](#).

Windows

- C:\\ProgramData\\Amazon\\SSM\\Logs\\amazon-ssm-agent.log
- C:\\ProgramData\\Amazon\\SSM\\Logs\\amazon-cloudwatch-agent.log
- C:\\ProgramData\\Amazon\\SSM\\Logs\\errors.log
- C:\\cfn\\log\\cfn-init.log

For more information, see [Quick Start: Enable Your Amazon EC2 Instances Running Windows Server 2016 to Send Logs to CloudWatch Logs Using the CloudWatch Logs Agent](#).

Log management — Amazon VPC Flow Logs

[VPC Flow Logs](#) is a feature that captures information about the IP traffic going to and from network interfaces in your VPC. Flow log data can be published to Amazon CloudWatch logs or Amazon S3. Flow log data collection does not affect network throughput or latency. You can create or delete flow logs without any impact to network performance.

Flow logs can help you with a number of tasks, such as:

- Diagnosing overly restrictive Security Group rules
- Monitoring traffic that reaches your instance
- Determining the direction of the traffic to and from the network interfaces

You do not have to enable VPC flow logs for each newly created VPC in Accelerate accounts. AMS will automatically detect whether a VPC has a flow log using the [ams-nist-cis-vpc-flow-logs-enabled](#) Config rule. If VPC flow logs are not enabled, AMS will automatically remediate it by creating a VPC flow log with [custom fields](#). Having these additional fields will enable AMS and customers to better monitor VPC traffic, understand network dependencies, troubleshoot network connectivity issues, and identify network threats.

For information on viewing and searching flow logs, see [Work with flow logs](#).

Tracking changes in your AMS Accelerate accounts

Topics

- [Viewing your change records](#)
- [Default queries](#)
- [Change record permissions](#)

AWS Managed Services helps you track changes made by the AMS Accelerate Operations team and AMS Accelerate automation by providing a queryable interface using the [Amazon Athena](#) (Athena) console and AMS Accelerate log management.

Athena is an interactive query service you can use to analyze data in Amazon S3 by using standard Structured Query Language (SQL) (see [SQL Reference for Amazon Athena](#)). Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run. AMS Accelerate creates Athena tables with daily partitions over CloudTrail logs, and provides queries on your primary AWS Region and within the **ams-change-record** workgroup. You can choose any of the default queries and run them as needed. To learn more about Athena workgroups, see [How Workgroups Work](#).

Note

Only Accelerate can query CloudTrail events for your Accelerate account using Athena when Accelerate [is integrated with your CloudTrail Organization trail](#), unless your Organization administrator deployed an IAM Role for using Athena to query and analyze CloudTrail events in your account, during onboarding.

Using change record, you can easily answer questions like:

- Who (AMS Accelerate Systems or AMS Accelerate Operators) has accessed your account
- What changes have been made by AMS Accelerate in your account
- When did AMS Accelerate perform changes in your account
- Where to go to view changes made in your account
- Why AMS Accelerate needed to make the changes in your account
- How to modify queries to get answers to all those questions for any non-AMS changes too

Viewing your change records

To use Athena queries, sign in to the AWS Management console and navigate to the Athena console in your primary AWS Region.

Note

If you see the **Amazon Athena Get Started** page while performing any of the steps, click **Get Started**. This might appear for you even if your Change Record infrastructure is already in place.

1. Choose **Workgroup** from the upper navigation panel in the Athena console.
2. Choose the **ams-change-record** workgroup, and then click **Switch Workgroup**.
3. Choose **ams-change-record-database** from the **Database Combo** box. The **ams-change-record-database** includes the **ams-change-record-table** table.
4. Choose **Saved Queries** from the upper navigation panel.
5. The **Saved Queries** window shows a list of queries that AMS Accelerate provides, which you can run. Choose the query you want to run from the **Saved Queries** list. For example, **ams_session_accesses_v1** query.

For the full list of preset AMS Accelerate queries, see [Default queries](#).

6. Adjust the **datetime** filter in the query editor box as needed; by default, the query only checks changes from the last day.
7. Choose **Run query**.

Default queries

AMS Accelerate provides several default queries you can use within the Athena console; they are listed in the following table.

Note

- All queries accept **datetime range** as an optional filter; all the queries run over the last 24 hours, by default. For expected input, see the following subsection, [Modifying the datetime filter in queries](#).

- Parameter inputs that you can or need to change are shown in the query as `<PARAMETER_NAME>` with angular braces. Replace the placeholder **and** the angular braces with your parameter value.
- All filters are optional. In the queries, some optional filters are commented out with a double dash (--) at the start of the line. All queries will run without them, with default parameters. If you want to specify parameter values for these optional filters, remove the double dash (--) at the start of the line and replace the parameter as you want.
- All queries return IAM PrincipalId and IAM SessionId in the outputs
- The calculated cost for running a query depends on how many CloudTrail logs are generated for the account. To calculate the cost, use the [AWS Athena Pricing Calculator](#).

Canned queries

Purpose/Description	Inputs	Outputs
Query name: <code>ams_access_session_query_v1</code>		
<p>Tracking AMS Accelerate access sessions</p> <p>Provides information about a specific AMS Accelerate access session. The query accepts the IAM Principal ID as an optional filter and returns event time, business need for accessing the account, requester, and so on.</p> <p>You can filter on a specific IAM Principal ID by uncommenting the line and replacing the placeholder <code>IAM</code></p>	<p>(Optional) IAM PrincipalId : The IAM Principal identifier of the resource that is trying to access. The format is <code>UNIQUE_IDENTIFIER :RESOURCE_NAME</code> . For details see unique identifiers. You can run the query without this filter to determine the exact IAM PrincipalId the you want to filter with.</p>	<ul style="list-style-type: none"> • EventTime: Time of gaining the access • EventName: AWS Event name (AssumeRole) • EventRegion: AWS Region that gets the request • EventId: CloudTrail Event ID • BusinessNeed Type: Business reason type to access the account. Allowed values are: SupportCase, OpsItem, Issue, Text. • BusinessNeed: Business need to access the account. For example, Support Case ID, Ops Item ID, and so forth. • Requester: Operator ID that accesses the account, or Automation system that access the account.

Purpose/Description	Inputs	Outputs
<p><i>PrincipalId</i> with a specific ID in the query editor.</p> <p>You can also list non-AMS access sessions by removing the useragent filter line in the WHERE clause of the query.</p>		<ul style="list-style-type: none"> RequestAccessType: Requester type (System, OpsConsole, OpsAPI, Unset)

Query name: ams_events_query_v1

<p>Track all mutating actions done by AMS Accelerate</p> <p>Returns all write actions done on the account using that AMS Accelerate role filter.</p> <p>You can also track mutating actions done by non-AMS roles by removing the useridentity.arn filter lines from the WHERE clause of the query.</p>	<p>(Optional)</p> <p>Only datetime range. See Modifying the datetime filter in queries.</p>	<ul style="list-style-type: none"> AccountId: AWS Account ID RoleArn: RoleArn for the requester EventTime: Time of gaining the access EventName: AWS Event name (AssumeRole) EventRegion: AWS Region that gets the request EventId: CloudTrail Event ID RequestParameters : Request parameters for the request ResponseElements: Response elements for the response. UserAgent: AWS CloudTrail User Agent
--	--	--

Query name: ams_instance_access_sessions_query_v1

Purpose/Description	Inputs	Outputs
<p>Track instance accesses by AMS Accelerate</p> <p>Returns a list of AMS Accelerate instance accesses; every record includes event time, event Region, instance ID, IAM Principal ID, IAM Session ID, SSM Session ID. You can use the IAM Principal ID to get more details on the business need for accessing the instance by using the <code>ams_access_sessions_query_v1</code> Athena query. You can use the SSM Session ID to get more details on the instance access session, including the start and end time of the session, log details, and using the AWS Session Manager console in the instance's AWS Region.</p> <p>Users can also list non-AMS instance accesses by removing the userid filter line in the WHERE clause of the query.</p>	<p>Only <code>datetime</code> range. See Modifying the datetime filter in queries.</p>	<ul style="list-style-type: none"> • InstanceId: Instance ID • SSMSessionId: SSM Session ID • RoleArn: RoleArn for the requester • EventTime: Time of gaining the access • EventName: AWS Event name (AssumeRole) • EventRegion: AWS Region that gets the request • EventId: CloudTrail Event ID

Purpose/Description	Inputs	Outputs
Query name: ams_privilege_escalation_events_query_v1		

Purpose/Description	Inputs	Outputs
<p>Track permission (escalation) events for AMS and non-AMS users</p> <p>Provides a list of events that can directly or potentially lead to a privilege escalation. The query accepts ActionedBy as an optional filter and returns EventName, EventId, EventTime, and so forth. All fields associated with the event are also returned. Fields are blank if not applicable for that event. The ActionedBy filter is disabled, by default; to enable it, remove "--" from that line.</p> <p>By default, the ActionedBy filter is disabled (it will show privilege escalation events from all users). To show events for a particular user or role, remove the double dash (--) from the useridentity filter line in the WHERE clause and replace the placeholder <i>ACTIONEDBY_PUT_USER_NAME_HERE</i> with an</p>	<p>(Optional) ACTIONEDBY_PUT_USER_NAME : Username for the actionedBy user. This can be an IAM user or role. For example, ams-access-admin.</p> <p>(Optional) datetime range. See Modifying the datetime filter in queries.</p>	<ul style="list-style-type: none"> • AccountId: Account Id • ActionedBy: ActionedBy Username • EventTime: Time of gaining the access • EventName: AWS Event name (AssumeRole). • EventRegion: AWS Region that gets the request • EventId: CloudTrail Event ID

Purpose/Description	Inputs	Outputs
IAM user or role name. You can run the query without the filter to determine the exact user you want to filter with.		

Query name: `ams_resource_events_query_v1`

<p>Track write events for specific resources AMS or non-AMS</p> <p>Provides a list of events done on a specific resource. The query accepts resource ID as part of the filters (replace placeholder <i>RESOURCE_INFO</i> in the WHERE clause of the query), and returns all write actions done on that resource.</p>	<p>(Required) RESOURCE_INFO : The resource identifier, can be an ID for any AWS resource in the account. Do not confuse this with resource ARNs. For example, an instance ID for an EC2 instance, table name for a DynamoDB table, logGroupName for a CloudWatch Log, etc.</p> <p>(Optional) datetime range. See Modifying the datetime filter in queries.</p>	<ul style="list-style-type: none"> • AccountId: Account Id • ActionedBy: ActionedBy Username • EventTime: Time of gaining the access • EventName: AWS Event name (AssumeRole). • EventRegion: AWS Region that gets the request • EventId: CloudTrail Event ID
--	--	---

Query name: `ams_session_events_query_v1`

Purpose/Description	Inputs	Outputs
<p>Track write actions performed by AMS Accelerate during specific session</p> <p>Provides a list of events done on a specific session. The query accepts IAM Principal ID as part of the filters (replace the placeholder <i>PRINCIPAL_ID</i> in the WHERE clause of the query), and returns all write actions done on that resource.</p>	<p>(Required) <i>PRINCIPAL_ID</i> : Principal ID for the session. The format is <i>UNIQUE_ID ENTIIFIER :RESOURCE_NAME</i> . For details see unique identifiers. You can run the query "ams_session_ids_by_requester_v1" to get list of IAM Principal IDs for a requester. You can also run the query without this filter to determine the exact IAM PrincipalId you want to filter with.</p> <p>(Optional) datetime range. See Modifying the datetime filter in queries.</p>	<ul style="list-style-type: none"> • AccountId: Account Id • ActionedBy: ActionedBy Username • EventTime: Time of gaining the access • EventName: AWS Event name (AssumeRole) • EventRegion: AWS Region that gets the request • EventId: CloudTrail Event ID

Query name: ams_session_ids_by_requester_v1

Purpose/Description	Inputs	Outputs
<p>Track IAM Principal/ Session IDs for a specific requester.</p> <p>The query accepts "requester" (replace the placeholder <i>Requester</i> in the WHERE clause of the query), and returns all IAM Principal Ids by that requester during the specified time range.</p>	<p>(Required) <code>Requester</code> : Operator ID that accesses the account (for example: alias of an operator), or Automation system that access the account (for example: OsConfiguration, AlarmManager, etc.).</p> <p>(Optional) <code>datetime</code> range. See Modifying the datetime filter in queries.</p>	<ul style="list-style-type: none"> • IAM PrincipalId - IAM Principal Id of the session. The format is <i>UNIQUE_IDENTIFIER :RESOURCE_NAME</i> . For details see unique identifiers. You can run the query without this filter to determine the exact IAM PrincipalId you want to filter with. • IAM SessionId - IAM Session Id for the access session • EventTime: Time of gaining the access

Modifying the datetime filter in queries

All queries accept **datetime** range as an optional filter. All the queries run over the last one day by default.

The format used for the **datetime** field is yyyy/MM/dd (for example: 2021/01/01). Remember that it only stores the date and not the entire timestamp. For the entire timestamp, use the field **eventtime**, which stores the timestamp in the ISO 8601 format yyyy-MM-ddTHH:mm:ssZ (for example: 2021-01-01T23:59:59Z). However, since the table is [partitioned](#) on the datetime field, you'll need to pass in both the datetime and eventtime filter to the query. See the following examples.

Note

To see all the accepted ways you can modify the range, see the latest [Presto function documentation](#) based on the Athena engine version currently used for the **Date and Time Functions and Operators** to see all the accepted ways you can modify the range.

Date Level: Last 1 day or last 24 hours (Default) example: If the `CURRENT_DATE='2021/01/01'`, the filter will subtract one day from the current date and format it as `datetime > '2020/12/31'`

```
datetime > date_format(date_add('day', - 1, CURRENT_DATE), '%Y/%m/%d')
```

Date Level: Last 2 months example:

```
datetime > date_format(date_add('month', - 2, CURRENT_DATE), '%Y/%m/%d')
```

Date Level: Between 2 dates example:

```
datetime > '2021/01/01'  
AND  
datetime < '2021/01/10'
```

Timestamp Level: Last 12 hours example:

Partition data scanned to last 1 day and then filter all events within the last 12 hours

```
datetime > date_format(date_add('day', - 1, CURRENT_DATE), '%Y/%m/%d')  
AND  
eventtime > date_format(date_add('hour', - 12, CURRENT_TIMESTAMP), '%Y-%m-%dT%H:  
%i:%sZ')
```

Timestamp Level: Between 2 timestamps example:

Get events between Jan 1, 2021 12:00PM and Jan 10, 2021 3:00PM.

```
datetime > '2021/01/01' AND datetime < '2021/01/10'  
AND  
eventtime > '2021-01-01T12:00:00Z' AND eventtime < '2021-01-10T15:00:00Z'
```

Change record permissions

The following permissions are needed to run change record queries:

- **Athena**
 - `athena:GetWorkGroup`

- `athena:StartQueryExecution`
- `athena:ListDataCatalogs`
- `athena:GetQueryExecution`
- `athena:GetQueryResults`
- `athena:BatchGetNamedQuery`
- `athena:ListWorkGroups`
- `athena:UpdateWorkGroup`
- `athena:GetNamedQuery`
- `athena:ListQueryExecutions`
- `athena:ListNamedQueries`
- **AWS KMS**
 - `kms:Decrypt`
 - AWS KMS key ID of `AMSCloudTrailLogManagement`, or your AWS KMS key ID(s), if Accelerate is using your CloudTrail trail events Amazon S3 bucket data store using [SSE-KMS](#) encryption.
- **AWS Glue**
 - `glue:GetDatabase`
 - `glue:GetTables`
 - `glue:GetDatabases`
 - `glue:GetTable`
- **Amazon S3 read access**
 - Amazon S3 bucket CloudTrail datastore: `ams-aAccountId-cloudtrail-primary region`, or your Amazon S3 bucket name, CloudTrail trail events Amazon S3 bucket data store.
- **Amazon S3 write access**
 - Athena events query results Amazon S3 bucket: `ams-aAccountIdathena-results-primary region`

AWS Systems Manager in Accelerate

Topics

- [Available AMS Accelerate SSM documents](#)
- [AMS Accelerate SSM document versions](#)
- [Systems Manager pricing](#)

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 Amazon. 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, please visit [Best practices for shared SSM documents](#).

Available AMS Accelerate SSM documents

AMS Accelerate SSM documents are available exclusively to AMS Accelerate customers, and are used to automate operational workflow to operate your account.

To see the available AMS Accelerate SSM documents from the AWS Management Console:

1. Open the Systems Manager console 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 Accelerate SSM document versions

SSM documents support versioning. AMS Accelerate 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 Accelerate 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 Accelerate 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

There is no cost associated with AMS Accelerate 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](#).

Document history

The following table describes the important changes to the documentation since the last release of the AMS Accelerate guide.

Latest documentation update: April 4, 2024

Change	Description	Date
Updated: AWS managed policies.	Updated the AWSManagedServicesDeploymentToolkitPolicy with new ECR permissions. See Accelerate updates to service-linked roles .	April 4, 2024
Updated Resource Tagger Configuration Profiles section.	Added <code>AWS::EFS::FileSystem</code> to the ResourceType list. For more information, see Resource Tagger Configuration Profiles .	March 21, 2024
Updated (Optional) Quick Start template section.	Added commands to download the template from AWS CloudShell For more information, see Download the Quick Start template .	March 21, 2024
Updated Create aws_managedservices_onboarding_role with AWS CloudFormation section.	Added commands to create the role from AWS CloudShell For more information, see Create aws_managedservices_onboarding_role with AWS CloudFormation .	March 21, 2024
Updated Configuration profile: monitoring section.	Added the following items to the ResourceType list. <ul style="list-style-type: none"> • <code>AWS::Elasticsearch::Domain</code> • <code>AWS::FSx::FileSystem::ONTAP</code> • <code>AWS::FSx::FileSystem::ONTAP::Volume</code> 	March 21, 2024

Change	Description	Date
	<ul style="list-style-type: none"> • AWS::FSx::FileSystem::Windows • AAWS::EFS::FileSystem <p>For more information, see Configuration profile: monitoring.</p>	
Updated Configuration profile: pseudoparameter substitution section.	<p>Added Amazon EFS and Amazon FSx pseudoparameter substitutions.</p> <p>For more information, see Configuration profile: pseudoparameter substitution.</p>	March 21, 2024
Updated How service request management works section.	<p>Added clarification on how AMS handles service requests that contain a feature request or a bug.</p> <p>For more information, see How service request management works.</p>	March 21, 2024
Updated Incident reports and service requests in Accelerate section.	<p>Changed the topic title to Incident reports, service requests, and billing questions in Accelerate.</p> <p>Added a new section, Billing questions.</p> <p>For more information, see Billing questions.</p>	March 21, 2024
Added new section to Features in the Service description topic.	<p>Added a new section, Service request management under AMS Accelerate features.</p> <p>For more information, see AMS Accelerate features.</p>	March 21, 2024

Change	Description	Date
Updated self-service reporting Weekly Incident report	<p>Added the following columns to the Weekly Incident report so that you can filter for incidents based on quarter, month, week, or day that the incident was created or resolved.</p> <ul style="list-style-type: none">• Created Day• Created Week• Created Quarter• Resolved Day• Resolved Week• Resolved Month• Resolved Quarter <p>For more information, see Weekly Incident report.</p>	March 11, 2024
Improvements for AMS Accelerate CloudTrail trail onboarding	<p>Improvements for AMS Accelerate CloudTrail trail onboarding:</p> <ul style="list-style-type: none">• Collect all bucket policies in a single block• Remove the second AWS Organization ID in the policy statements• Clarify customer environment requirements <p>For more information, see Review and update your configurations to enable AMS Accelerate to use your CloudTrail trail.</p>	February 23, 2024

Change	Description	Date
Updated: Account onboarding process.	Restructured the Account onboarding process section to make the steps more clear. Also added an optional Quick Start template for onboarding features. See (Optional) Quick Start template .	February 22, 2024
Updated: Offboarding AMS Accelerate.	Updated the AMS Accelerate offboarding considerations section to indicate that the <code>ams-access-management</code> CloudFormation stack and <code>ams-access-management</code> IAM role aren't deleted by the offboarding process. See AMS Accelerate offboarding considerations .	February 22, 2024
Updated: Configuration compliance in Accelerate.	Changed "Incident Report" to "Service Request" where applicable to avoid confusion on these terms. See Configuration compliance in Accelerate .	February 22, 2024
Updated: Account discovery in Accelerate.	Reorganized Account discovery in Accelerate to better group prerequisites with the relevant section. See Step 1. Account discovery in Accelerate .	February 22, 2024
Renamed: AMS Patch reporting to AMS host management.	Renamed AMS Patch reporting to AMS host management and renamed the report, Patch Details report, to SSM Agent Coverage report. See AMS host management .	February 22, 2024

Change	Description	Date
Updated Operations on Demand Catalog	Updated the Operations on Demand catalog of offerings table to remove references to "health" in Amazon EKS cluster maintenance . See Requesting AMS Operations On Demand .	February 22, 2024
Updated AMS Event Router	Updated the AMSCoreRule in the AMS Event Router section. See Using Amazon EventBridge Managed Rules in AMS .	February 22, 2024
Updated Supported Operating Systems.	Updated Supported Operating Systems to include SUSE Linux Enterprise Server 15 SP5. See Supported configurations .	February 22, 2024
Updated EC2 volume usage remediation automation	Updated the EC2 volume usage remediation automation section with correct capacity expansion schedule. See EC2 volume usage remediation automation .	February 22, 2024
Updated: Review and update your configurations to enable Accelerate to use your CloudTrail trail	Updated the AMS Accelerate Organization CloudTrail S3 bucket policy section. See Review and update your configurations to enable AMS Accelerate to use your CloudTrail trail	February 15, 2024
Added new feature: SSM Agent auto installation	Added a new section for SSM Agent auto installation See SSM Agent automatic installation .	January 26, 2024

Change	Description	Date
Updated: Supported configurations	<p>Added information regarding the supported versions of AWS Control Tower</p> <p>See Supported configurations.</p>	January 26, 2024
Updated: AMS Patch reporting.	<p>Removed three sections from AMS Patch reporting:</p> <ul style="list-style-type: none"> • Patch Instance Details Summary report • Patch Details report • Instances that Missed Patches report <p>See AMS host management.</p>	December 22, 2023
Updated: Accelerate onboarding prerequisites.	<p>Updated the support plans required to onboard AMS Accelerate.</p> <p>See Accelerate onboarding prerequisites.</p>	December 15, 2023
Updated: Create a patch maintenance window.	<p>Removed Default patch cycle section as this feature is deprecated.</p> <p>See Create a patch maintenance window.</p>	December 13, 2023
Updated: Notification settings in Accelerate.	<p>Clarified what email is used for notifications.</p> <p>See Notification settings in Accelerate for more information.</p>	December 12, 2023
Updated: AMSAccelerateCustomerAccessPolicies template.	<p>Updated the AMSAccelerateCustomerAccessPolicies template to correct a syntax error.</p> <p>See Permissions to use AMS features for more information.</p>	December 12, 2023

Change	Description	Date
Added: Change request security reviews	<p>Added a new section Change request security reviews under Security Management.</p> <p>See Change request security reviews for more information.</p>	December 11, 2023
Updated: resource_inventory.xlsx	<p>Updated the resource_inventory.xlsx to include Security Analyst roles.</p> <p>See Resource Inventory for more information.</p>	November 17, 2023
Updated: ams-access-admin-operations role description	<p>Updated ams-access-admin-operations description.</p> <p>See Why and when AMS accesses your account and Authenticating with identities in AMS Accelerate for more information.</p>	November 17, 2023
Updated: AMS Accelerate offboarding considerations	<p>Updated Security section to clarify what is available from Amazon GuardDuty and AWS Config rules after offboarding.</p> <p>See AMS Accelerate offboarding considerations for more information.</p>	November 17, 2023
Added: Monitoring and Incident Management for Amazon EKS	<p>Monitoring and Incident Management for Amazon EKS monitors your Amazon EKS resources for failures, performance degradation, and security issues.</p> <p>See Monitoring and Incident Management for Amazon EKS for more information.</p>	November 14, 2023
Updated: Tagging	<p>Added information on customer-provided tagging.</p> <p>See Customer-provided tags for more information.</p>	November 7, 2023

Change	Description	Date
Updated: Resource Tagger Configuration Profiles	Added AWS::AutoScaling::AutoScalingGroup, AWS::EKS::Cluster, AWS::Elasticsearch::Domain, and AWS::FSx::FileSystem to the Filter section. See Resource Tagger Configuration Profiles for more information.	October 27, 2023
Updated: Service Description	Added Ubuntu 22.04 to Supported Operating Systems. See Service description	September 29, 2023
Updated: AMS Accelerate Onboarding Prerequisites	Added a note to AMS Accelerate VPC endpoints to include CloudFormation template. See Accelerate onboarding prerequisites .	September 29, 2023
Updated: Detect	Removed endpoint protection type from the AMS Accelerate security response. See Detect .	September 29, 2023
Updated: Alerts from Baseline Monitoring in AMS	Added AWS Outposts to the Alerts from Baseline Monitoring table. See Detect monitoring-default-metrics .	September 29, 2023
Updated: Create aws_managedservices_onboarding_role role with AWS CloudFormation	Updated screenshot for Specify Stack Details. See Create aws_managedservices_onboarding_role with AWS CloudFormation .	September 29, 2023
Updated: Amazon EventBridge Managed Rules deployed by AMS Accelerate	Added new AMS Accelerate Amazon EventBridge Managed Rule AMSCoreRule . Updated AMS Accelerate Amazon EventBridge Managed Rule AMSAccessRolesRule to add a new role. See Amazon EventBridge Managed Rules deployed by AMS for more information.	September 19, 2023

Change	Description	Date
Updated: Alarm Manager Configuration Profiles	Added AWS Outposts pseudoparameter substitution identifiers. See Monitoring and event management in AMS Accelerate .	September 11, 2023
Updated: Resource Tagger Configuration Profiles	Added AWS Outposts resource type. See Configuration profile: pseudoparameter substitution .	September 11, 2023
Updated: Supported services	Added Amazon Elastic File System to the Services monitored by CloudWatch alarms section. See Service description for more information.	September 6, 2023
Updated: Patch monitoring and failure remediation	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 Accelerate for more information.	September 6, 2023
Updated: Clarified Response to malware events runbook	Clarified Response to malware events runbook for Security Incident response. See Security Incident Response in AMS for more information.	September 6, 2023
Updated: Connecting your Accelerate account with Transit Gateway	Clarified steps for Connecting a new Accelerate account VPC to the AMS Multi-Account Landing Zone network (creating a TGW VPC attachment): See Connecting your Accelerate account with Transit Gateway for more information.	September 5, 2023

Change	Description	Date
Updated: Alerts from baseline monitoring in AMS	Removed reference to two deprecated alarms AMSReadLatencyAlarm and AMSWriteLatencyAlarm. See Alerts from baseline monitoring in AMS for more information.	September 5, 2023
Added: AMS Event Router	Added documentation for AMS Event Router. See Using Amazon EventBridge Managed Rules in AMS for more information.	September 5, 2023
Updated: List of Alarm Manager pseudoparameters.	Updated the list of Alarm Manager pseudoparameters. EC2 instance name parameter was added to EC2 instance and EC2 disk alarm configurations. See Configuration profile: pseudoparameter substitution for more information.	August 29, 2023
Added: AMS Access Offboarding	Added consideration when offboarding AMS Access. See AMS Accelerate offboarding considerations .	August 24, 2023
Added: AMS Security Incident Response	Added documentation for using AMS Security Incident Response. See Security Incident Response in AMS .	August 18, 2023
Updated: AMS Accelerate access roles	Corrected a typo in the role names. See AWS Identity and Access Management in AMS Accelerate .	August 10, 2023
Updated: Policy statements	Replaced hardcoded role names with wildcards. See Review and update your configurations to enable AMS Accelerate to use your CloudTrail trail .	August 10, 2023

Change	Description	Date
Updated: List of monitored services with EFS alerts.	Updated the list of monitoring services with new EFS alerts for AMS baseline monitoring. 4 new EFS alert types were added. See Alerts from baseline monitoring in AMS for more information.	August 03, 2023
Updated: Accelerate resource inventory table	Removed ams-backup-config-rule-stack and related resources. See Resource Inventory .	July 18, 2023
Updated: AMS Accelerate access roles	Removed roles ams-backup-config-rule-st-amsBackupAlertConfigRule- <code><8-digit hash></code> and ams-backup-config-rule-st-amsBackupPlanConfigRuleH- <code><8-digit hash></code> . See AWS Identity and Access Management in AMS Accelerate .	July 18, 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 for more information.	June 19, 2023
New: AMS Accelerate access roles	Added new access roles for AMS Security.	June 16, 2023
New: AMS Accelerate CloudTrail log management can now use customer CloudTrail trails.	Updated Accelerate supported options for CloudTrail log management, including Accelerate deployed trail or integration with customer managed CloudTrail account or Organization trail. See Review and update your configurations to enable AMS Accelerate to use your CloudTrail trail for more information.	June 09, 2023

Change	Description	Date
Updated: AMS Accelerate Config Rules Response Configuration Report.	Updated on-request reporting for AWS Config Rules Response Configuration Report. See Accelerate updates to on-request reporting. See AMS Config Rules Response Configuration report .	May 26, 2023
Updated: Service Billing Start Date policy.	Updated definitions of Billing Start Date in AMS key terms .	May 15, 2023
Updated: AWS managed policies.	Updated the AWSManagedServicesDeploymentToolkitPolicy with new CFN and ECR permissions, and scoped down existing actions with wildcards. See Accelerate updates to service-linked roles. See Accelerate updates to service-linked roles .	May 09, 2023
Updated: Access role policy links.	The access roles can now be downloaded directly from Accelerate S3 bucket locations. See Why and when AMS accesses your account and AWS Identity and Access Management in AMS Accelerate .	
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 .	April 13, 2023
Updated: List of Alerts.	Removed CloudTrail references. See Log management in AMS Accelerate .	April 13, 2023

Change	Description	Date
Updated: List of Alerts.	Added three new SSM agent alerts. See Alerts from baseline monitoring in AMS .	April 13, 2023
Updated: Accelerate Prerequisites.	Clarified that Accelerate requires one of four AWS Support plans be in place and excludes the Developer plan. See Accelerate onboarding prerequisites .	April 13, 2023
Updated: Accelerate service-linked role policy.	The Contacts Service policy zip file has been updated. See AWS managed policies for AMS Accelerate .	April 13, 2023
Updated: AMS Resource Scheduler.	Incorrect role name, AWSManagedServices-DescribeScheduleOrPeriod, corrected to AWSManagedServices-DescribeScheduleOrPeriods. See Cost optimization with AMS Resource Scheduler .	April 13, 2023
Updated: AWS managed policies.	Updated Customized findings responses with instructions for updating custom responses in single or multiple accounts.	April 13, 2023
Updated: Resource Tagger	Added warnings about "specifying the name for your new configuration (SampleConfigurationBlock in the provided example) as you may inadvertently override the AMS-managed configuration with the same name". See Resource Tagger use cases in AMS Accelerate .	March 16, 2023
Updated: Patch RACI	Several updates and clarifications to the RACI for patching. See Service description .	March 16, 2023

Change	Description	Date
Updated: Actions in deployment toolkit SLR JSON	Updated policy and actions. See: Using service-linked roles for AMS Accelerate .	March 16, 2023
Updated: Auto remediation	Removed LVM support for EC2 volume automation. See: AMS automatic remediation of alerts .	March 16, 2023
Updated: Accelerate Onboarding.	Clarified use of roles, especially the minimal role The template to create AMS roles .	March 16, 2023
Updated: Self-service reporting.	Daily backup reports now support primary and secondary regions. Both are reported in the Resource Region field of Daily backup report .	March 16, 2023
Updated: Patching guidance	Added a warning not to customize the default patching baselines, which are managed by AMS. Instead, create a new custom patching baseline. See: Default patch baseline and Custom patch baseline .	March 16, 2023
Updated Service Termination policy.	Updated definitions of Service Termination and Service Termination Date in AMS key terms . Termination notices must be issued by the 20th day of the month prior to your last full month.	March 16, 2023
Updated: AWS managed policies.	Clarified policy name: Contacts service-linked role for AMS Accelerate .	Feb 16, 2023
New: AWS managed policies.	Added policy: Contacts service-linked role for AMS Accelerate .	Feb 16, 2023
Updated: Configuration compliance.	Fixed a misspelled word in: Configuration compliance in Accelerate .	Feb 16, 2023

Change	Description	Date
New Content: Unsupported OSES	Added information on what services AMS provides for unsupported operating systems (OSes), see Capabilities for unsupported operating systems in Accelerate .	Feb 16, 2023
Updated: Create patch windows	Added a link for using CloudShell to Create a maintenance window with the Systems Manager command line interface (CLI) .	Feb 16, 2023
Updated Content: Onboarding management resources	Updated the zipped JSON templates in The template to create AMS roles .	Feb 16, 2023
New Content: Configuration Compliance	Added a new topic: Customized findings responses .	Feb 16, 2023
New: AWS managed policies.	Added policy: Amazon EventBridge rule service-linked role for AMS Accelerate .	Feb 07, 2023
Updated: AWS managed policies.	Updated the AWSManagedServices DeploymentToolkitPolicy with new S3 permissions. See Accelerate updates to service-linked roles .	Jan 30, 2023
New opt-in region: CPT.	AMS Accelerate is now available in the Capetown (CPT) opt-in region. To opt in, see Managing AWS Regions .	Jan 12, 2023
Updated: Service Description.	Added FSx services monitored by CloudWatch alarms to Service description .	Jan 12, 2023
Updated: Monitoring default metrics.	Added 6 FSx alarms to Alerts from baseline monitoring in AMS .	Jan 12, 2023
Updated: AMS patterns.	Added <i>Customize Cloudwatch Alarm Notifications</i> to AMS patterns .	Jan 12, 2023

Change	Description	Date
Updated: Onboarding management resources.	Updated the table of templates, adding a row for <code>ams-onboarding-ssm-execution-role</code> in The template to create AMS roles .	Jan 12, 2023
Updated: Configuration compliance.	Additional details for requesting custom remediations (in the Important box) on Configuration compliance in Accelerate .	Jan 12, 2023
Updated: Service-linked-role permissions.	Removed older or duplicated permissions. See Using service-linked roles for AMS Accelerate .	Dec 15, 2022
Updated: Patch management, maintenance windows.	Added guidance to console instructions, step 5, for creating a maintenance window. See Create a maintenance window from the Systems Manager console .	Dec 15, 2022
New: Patch management section.	Added a section for Patch Tuesday maintenance windows. See Create a recurring "Patch Tuesday" maintenance window from the AMS console (recommended) .	Dec 15, 2022
Updated: AMS Resource Scheduler.	Updated the AWS CloudFormation stack name. See Using resources with AMS Resource Scheduler .	Dec 15, 2022
Updated: Tag your resources for backup.	Added guidance for using AMS Resource Tagger. See Tag your resources to apply AMS backup plans .	Dec 15, 2022
Updated: Select a backup plan.	Indicated which plans offer continuous backup. See Select an AMS backup plan .	Dec 15, 2022
Updated: AMS Resource Scheduler.	Updated the AWS CLI example for deleting a period or schedule. See Working with periods and schedules in AWS Managed Services Resource Scheduler .	Dec 15, 2022

Change	Description	Date
Updated: AWS managed policies.	Added the <code>AWSManagedServicesDeploymentToolkitPolicy</code> . See AWS managed policies for AMS Accelerate .	Dec 15, 2022
New: Added section describing the AMS new service-linked role, <code>AWSServiceRoleForManagedServices_DetectiveControlsConfig</code> .	Added GovCloud regions and permissions. See Detective controls service-linked role for AMS Accelerate .	Dec 15, 2022
New: AWS-managed policy	Added section describing how the AWS-managed policy, <code>AWSManagedServices_AlarmManagerPermissionsBoundary</code> , is used in the service-linked role policy, <code>AWSManagedServices_AlarmManager_ServiceRolePolicy</code> , to restrict permissions of IAM roles created by the service-linked role <code>AWSServiceRoleForManagedServices_AlarmManager</code> . See AWS managed policies for AMS Accelerate .	Dec 15, 2022
Updated: Operations on Demand.	Added offerings: <i>SQL Server on EC2 Operations</i> and <i>AMI Building and Vending</i> . See Operations On Demand .	Nov 10, 2022
Updated: Monitoring and event management.	Updated explanation of service notifications and incident reports. See How monitoring works .	Nov 10, 2022
Updated: Service-linked role regions	Added GovCloud regions and permissions. See Using service-linked roles for AMS Accelerate .	Nov 10, 2022
New: Service-linked role.	Added new role: <code>AWSServiceRoleForAMSDetectiveControls</code> . See Detective controls service-linked role for AMS Accelerate .	Nov 10, 2022

Change	Description	Date
Updated: Access management.	Updated subsections with improved instructions. See Access management in AMS Accelerate .	Nov 10, 2022
Updated: Service description.	Updated <i>AMS Patterns</i> in the RACI matrix. See Service description .	Nov 10, 2022
Updated: AMS patterns.	Customers are responsible for pattern deployments. See AMS patterns .	Nov 10, 2022
Updated: Offboarding.	Added details of what happens to specific Backup and Monitoring resources during offboarding. See Offboarding AMS Accelerate .	Nov 10, 2022
Updated: Patch management.	Updated and shortened guidance regarding IAM policies. See Creating an IAM role for on-demand patching .	Nov 10, 2022
New: links to architecture diagrams.	Added links to AMS Reference Architecture Diagrams to various topics. For example, see Monitoring and event management in AMS Accelerate .	Nov 10, 2022
New: Operations on Demand offering	Added "Landing Zone Accelerator Operations". See Operations On Demand .	Oct 13, 2022
Update: Monitoring management. Alerts generate incident reports, not service requests	How monitoring works .	Oct 13, 2022
New: Creating a patch maintenance window with an Accelerate-custom CFN template	AWS CloudFormation patch window configuration templates. See Create a patch maintenance window .	Sept 15, 2022

Change	Description	Date
Updated: Offboarding	Emphasized that backup plans in Accelerate no longer work after offboarding. See Offboarding AMS Accelerate .	Sept 15, 2022
Updated: CloudWatch configuration change details	Corrected a mistake in the Windows and Linux examples. See CloudWatch configuration change details .	Sept 15, 2022
Updated: Using AMS Resource Scheduler	Added guidance about Cost Allocation Tags. See Cost estimator in AMS Resource Scheduler .	September 15, 2022
Updated: AMS Config Rule Library	Added two <code>ams-eks-</code> config rules to the Table of Rules. See AMS Config Rule library .	September 15, 2022
Updated: Backup Management	Removed the misleading label PITR (point-in-time-recovery) from backup plan titles and descriptions. See Select an AMS backup plan .	September 15, 2022
Updated: Accelerate Service Description	Updated descriptions of config rules and canaries. See Service description .	September 15, 2022
Updated: Service Description, Supported Configurations	Removed end-of-service date for Windows 2008 R2. Accelerate does not support Windows 2008. See Supported configurations .	August 11, 2022
Updated: Service Description, Roles and Responsibilities	Updated the RACI table. Removed ELB access logs from the last row of the Networking section. We do not enable ELB access logs for Accelerate customers. See Roles and responsibilities .	August 11, 2022
Updated: Configuration Compliance	Corrected a typo in the Table of Rules, Frameworks column. NIST-CSF was incorrectly listed as NIST-CIS. See Configuration compliance in Accelerate .	August 11, 2022

Change	Description	Date
New: Accelerate Offboarding	Considerations and process for offboarding. See Offboarding AMS Accelerate .	August 11, 2022
Updated: List of pre-installed SSM agents' operating systems	Added "Ubuntu Linux 18.04 and 20.04" to the list. See Onboarding EC2 instances to Accelerate .	August 11, 2022
New: Resource Scheduler	Use AMS Resource Scheduler to cost optimize by stopping and starting resources only as needed. See Cost optimization with AMS Resource Scheduler .	July 14, 2022
Updated: Service Description for Resource Scheduler	Several sections of the service description were updated for the new Resource Scheduler offering. See Service description .	July 14, 2022
New: AMS Patterns	AMS offers pattern templates, a generalized solution that solves for a family of use cases in the AMS managed environment. First pattern on offer: AMS patterns .	July 14, 2022
New: Cost optimization note	Added a note explaining how costs can increase with resource usage. See Resource Inventory .	July 14, 2022
Updated: AMS Config Rules	Reorganized the tables in the AMS Config Rule library . The HTML table has fewer columns, to make it easier to read at a glance. The downloadable spreadsheet has additional columns to allow sorting and filtering.	July 14, 2022

Change	Description	Date
Updated: Access Management	Updated the sample CloudFormation template in Permissions to use AMS features . The AMSAccelerateAdminAccess policy now includes the AmsResourceSchedulerPassRolePolicy and IAMReadOnlyPolicy policies.	July 14, 2022
Updated: Self-Service Reporting	Added instructions for encrypting AWS Glue metadata with KMS keys. See box labeled Important on Self-service reporting .	July 14, 2022
Updated: AMS baseline monitoring	Added DeleteRecoveryPoint backup alert. Alerts from baseline monitoring in AMS	July 14, 2022
Updated: Supported operating systems	Added End of Support date for Amazon Linux 2. Service description	July 14, 2022
Updated: AMS Reporting	Added note about Opt-in Regions. Reports and options	July 14, 2022
Resource Scheduler	Added information about onboarding and using AMS Resource Scheduler to assist in cost optimization by scheduling resource stop and start times. Also, updated the Accelerate service description to include mention of Resource Scheduler. Additionally, updated the Amazon Linux 2 supported end of support date to 2024. See Cost optimization with AMS Resource Scheduler and Service description	June 30, 2022
New alarm	Added a AWS Backup alarm. Alerts from baseline monitoring in AMS	June 21, 2022
New content	Added the service-linked role content. Using service-linked roles for AMS Accelerate	June 16, 2022

Change	Description	Date
	AWS Network Firewall Operations added to Operations on Demand (OOD) catalog of offerings. Operations On Demand	June 16, 2022
	Added problem management feature description. Service description	June 16, 2022
	Added note about the set of config rules that does not support in particular opt-in regions. Configuration compliance in Accelerate	June 16, 2022
Updated content	Configuration compliance. "AMS Config Rule library" -> "Table of rules", was updated and removed to ZIP only. Configuration compliance in Accelerate	June 16, 2022
	Removed escalation emails. Escalation path	June 16, 2022
	Moved topic list to below opening paragraphs. What is AMS Accelerate?	June 16, 2022
	Updated the auto remediation content. AMS automatic remediation of alerts	June 16, 2022
Updated content: Service Description	Added EKS to the list of services monitored by AMS Config Rules in Supported services . Updated monitoring description in RACI table in Roles and responsibilities .	May 12, 2022
Updated content: Configuration Compliance	Added EKS-related config rules. See Configuration compliance in Accelerate .	May 12, 2022

Change	Description	Date
Updated content: Getting Started, Account Discovery	Added a newer version of the <code>AwsAccountDiscoveryCli</code> script (in the <i>Account Discovery Changelog zip file</i>) in Step 1. Account discovery in Accelerate .	May 12, 2022
Updated content: Monitoring, default metrics	Updated trigger conditions for ALB-related metrics. See Alerts from baseline monitoring in AMS .	May 12, 2022
Updated content: Patching onboarding	Added an explicit patching prerequisite: you need to opt-in to EBS. See Onboarding patching in Accelerate .	May 12, 2022
Updated content: Accelerate resource inventory table	Changed <code>ams-detective-controls-config-rules-cdk</code> rules, added rules for <code>ams-detective-controls-recorder-cdk</code> and <code>ams-detective-controls-infrastructure-cdk</code> . See Resource Inventory .	April 14, 2022
Updated content: Configuration Compliance	Introduction to industry standards, config rules, and types of responses. Emphasizes that customers do not choose individual config rules or responses. Configuration compliance in Accelerate .	April 14, 2022
Updated content: Service Description	Moved the existing Scope of Changes section under Roles and Responsibilities. See Roles and responsibilities .	April 14, 2022
Updated content: Tagging and Monitoring	Added <code>AWS::Synthetics:Canary</code> to lists of allowed Resource Types for tagging and monitoring. See Resource Tagger Configuration Profiles and Configuration profile: pseudoparameter substitution .	April 14, 2022
Updated content: Accelerate Prerequisites	Added SSM-required bucket permissions to Amazon EC2 Systems Manager in Accelerate .	April 14, 2022

Change	Description	Date
New content: Patching and Monitoring	Added sample code to use Cloudformation to deploy tagging and monitoring configurations. See Deploying a configuration profile and Using AWS CloudFormation to deploy configuration changes .	March 10, 2022
Updated content: Patch maintenance console	Reordered steps in Create a maintenance window from the Systems Manager console to match the console interface.	March 10, 2022
Updated content: Patch maintenance CLI	Updated CLI parameters (schedule, duration, and cutoff) for Create a maintenance window with the Systems Manager command line interface (CLI)	March 10, 2022
New content: Auto Instance Config	Added definition of AMSInstanceProfile BasePolicy to IAM permissions change details	March 10, 2022
New content: Onboarding	Added a sample Linux command to Outbound internet connectivity in Accelerate	March 10, 2022
New content: Onboarding	Added a least-privilege option to The template to create AMS roles .	March 10, 2022
Updated content: Accelerate escalation instructions	Added guidance, links, and email contacts to Escalation path	March 10, 2022
Updated content: Supported Configurations	AMS expects to end support for RHEL 6 and CentOS on March 14, 2023. See Supported configurations	March 10, 2022
Updated content: Resources table	Added AMS access IAM roles to Resource Inventory resources table	March 10, 2022

Change	Description	Date
Updated content: Onboarding and Backup	Added instructions for opting in to AWS Backup to Onboarding AWS Backup in Accelerate and Continuity management in AMS Accelerate	March 10, 2022
Updated content: Access Management	Removed Advanced-specific instructions from Accelerate guidance How and when to use the root user account in AMS .	March 10, 2022
Updated content: Supported Configurations	AMS now supports Oracle Linux 8.3 and Ubuntu 18.04 and 20.04. See Supported configurations .	February 28, 2022
Updated content: Service Level Agreement	Updated the downloadable Service Level Agreement in Supported services .	February 28, 2022
Updated content: Access Management	Updated How AMS accesses your account with FAQs for AMS operator console roles and a warning not to modify or delete them.	February 28, 2022
Updated content: Alarm Manager	Updated Configuration profile: monitoring . Alarm Manager is no longer limited to single-metric alarms.	February 28, 2022
Updated content: Getting Started	Updated Step 2. Onboarding management resources in Accelerate . Added an IAM role with minimal access for onboarding resources.	February 28, 2022
New content: Scope of Changes in Service Description	Added a new section, Scope of changes performed by AMS Accelerate that emphasize s boundaries and actions that AMS Accelerate does not perform.	February 10, 2022

Change	Description	Date
Updated content: Getting Started	New onboarding process starts with setting up default features and configurations before customizing. Subsections contain feature-specific goals and related links. See Getting Started with AMS Accelerate .	February 10, 2022
Updated content: AMS Backup Management.	Shortened and reorganized the Continuity management in AMS Accelerate chapter for readability.	February 10, 2022
Updated content: Tagging	Added a Tagging Tools section to accommodate code samples for CloudFormation and other tools. See Tagging in AMS Accelerate .	February 10, 2022
Updated content: Baseline Monitoring	Improved trigger condition for RedShift cluster alarm reduces false alarms during maintenance. See Alerts from baseline monitoring in AMS .	February 10, 2022
Updated content: Patching	Updated sample CLI command to register a maintenance window. See Create a maintenance window with the Systems Manager command line interface (CLI) .	February 10, 2022
Updated content: AWS Config Rules Inventory.	Removed deprecated config rule <code>ams-nist-cis-ec2-security-group-attached-to-eni</code> from the AWS Config Rules Inventory table. See Table of Rules .	January 27, 2022
New content: Creating patch maintenance windows.	Added a link to the SSM tutorial and sample commands for creating patch maintenance windows from the command line. See Create a maintenance window with the Systems Manager command line interface (CLI) .	January 27, 2022

Change	Description	Date
New content: Resource Tagger recognizes new Auto Scaling Groups (ASG) resource type.	Added Auto Scaling Groups to the resource types filterable with Resource Tagger configuration profiles. See Syntax and structure .	January 13, 2022
New content: Additional backup plans and vaults.	Added new backup plans and vaults to mitigate high-risk scenarios including ransomware attacks. See View backups in AMS vaults and View backups in AMS vaults .	January 13, 2022

AWS Glossary

For the latest AWS terminology, see the [AWS glossary](#) in the *AWS Glossary Reference*.