
**AMS Advanced
Onboarding Guide
AMS Advanced Account
Onboarding Information
Version October 28, 2020**



AMS Advanced Onboarding Guide: AMS Advanced Account Onboarding Information

Copyright © 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

AWS Managed Services Onboarding Introduction	1
Learning about AMS	1
Key terms	2
AWS Managed Services modes	5
Types of modes and accounts in AMS	6
AMS modes and applications or workloads	12
Real world use cases for AMS modes	16
AMS post-account prescriptive guidance	18
What we do, what we do not do	19
AMS egress traffic management	19
IAM User Role	20
Default AMS multi-account landing zone (MALZ) IAM User Roles	21
Default AMS single-account landing zone (SALZ) IAM User Role	30
Default Access Firewall Rules	38
Linux Stack Instance Ports	39
Windows Stack Instance Ports	39
AMS service management	40
Account governance	40
Service commencement	40
AMS customer relationship management (CRM)	41
CRM Process	41
CRM meetings	42
CRM Meeting Arrangements	43
CRM monthly reports	43
Updates to shared services: Multi-Account Landing Zone	44
AMS planned event management	44
AMS PEM criteria	44
The AMS PEM process	44
Getting help	45
Service hours	45
How do I get offboard assistance from AMS Single-Account Landing Zone accounts?	46
How do I offboard from AMS Multi-Account Landing Zone accounts?	46
How do I offboard a Multi-Account Landing Zone environment?	47
How do I offboard a Multi-Account Landing Zone application account?	47
How do I offboard a Multi-Account Landing Zone application account VPC?	48
AMS Multi-account landing zone onboarding	49
Multi-Account Landing Zone network architecture	49
About Multi-Account Landing Zone network architecture	49
Multi-Account Landing Zone accounts	51
Core account onboarding	77
Create an AWS core account	77
Create an IAM Role for AMS to access your account	78
Secure the New Account with Multi-Factor Authentication (MFA) for the Root User	81
Subscribe to AWS Marketplace for EPS	81
Set up networking	83
Set up access management	85
Application account onboarding	88
Requesting a new application account	89
Setting up Active Directory to federate access to AMS IAM roles	90
Setting Up Networking with the New Application Account	92
Setting Up Additional VPCs in the Application Account	93
Appendix: multi-account landing zone Onboarding Consideration List	93
Account Configuration	93
Multi-Account Landing Zone Monitoring Alerts	94

Network Configuration	94
Active Directory Configuration	95
Trend Micro Endpoint Protection (EPS)	95
Access: Bastions, SSH and RDP	95
Federation	96
AMS Single-account landing zone Onboarding	97
AMS single-account landing zone onboarding process	97
Single-Account Landing Zone network architecture	98
AMS Single-account landing zone shared services	99
Create a New AWS Account for AMS	100
Create an AWS Account	100
Set Up Consolidated Billing -- Link New Account to Payer Account	101
Configure your AWS Account for AMS Access	101
Subscribe to AWS Marketplace for EPS	104
Subscribe to AWS Marketplace for CentOS 7.6	105
Secure the New Account with Multi-Factor Authentication (MFA) for the Root User	105
Set Up Networking	105
Allocate IP Space for your AMS Environment	106
Establish Private Network Connectivity to AWS	106
Set up your Firewall	107
AMS Bastion Options during Application Migrations/Onboarding	107
Set up access management	111
Establish an Active Directory (AD) trust	111
Federate your Active Directory with the AMS IAM roles	115
AMS Single-account landing zone Default Settings	119
Endpoint Security (EPS)	119
Security groups	121
EC2 IAM instance profile	124
Monitored Metrics Defaults	127
Log retention and rotation defaults	135
Continuity Management Defaults	136
Patching Defaults	137
Understand Change Types	138
Change Type Validation	138
Validate the AMS Service	138
Finding your account settings	139
Finding an Instance ID or IP Address	140
DNS Friendly Bastion Names	142
Finding Bastion IP Addresses	143
EC2 Instance, Creating	143
Access, Requesting	152
Other Other RFC, Creating (CLI)	158
Any Stack, Deleting, Rebooting, Starting, Stopping	159
Access Examples	167
Reporting an Incident	175
Creating a Service Request	178
Next Steps	181
Tutorials	182
Appendix: Single-Account Landing Zone Onboarding Questionnaire	200
Deployment Summary	200
Environment/Architecture Considerations	200
Single-Account Landing Zone Monitoring Alerts	201
Maintenance Window	201
Next Steps	201
Appendix: ActiveDirectory Federation Services (ADFS) claim rule and SAML settings	202
ADFS claim rule configurations	202
Web console	202

API and CLI access with SAML	203
Script configuration	203
Windows configuration	203
Linux configuration	204
Document history	206
AWS glossary	208

AWS Managed Services Onboarding Introduction

Topics

- [Learning about AMS \(p. 1\)](#)
- [Key terms \(p. 2\)](#)
- [AWS Managed Services modes \(p. 5\)](#)
- [AMS post-account prescriptive guidance \(p. 18\)](#)
- [What we do, what we do not do \(p. 19\)](#)
- [AMS egress traffic management \(p. 19\)](#)
- [IAM User Role \(p. 20\)](#)
- [Default Access Firewall Rules \(p. 38\)](#)

Welcome to AWS Managed Services (AMS). The purpose of this document is to provide information about, and assistance with, the AMS onboarding process, including details about setting up a new account for AMS, setting up networking and access to AMS, and validating your onboarding setup.

This document is intended for IT administrators tasked with preparing for and carrying out the tasks required to onboard the AMS service to a new AWS account. Onboarding the AMS service requires special privileges to set up Active Directory trusts and complete other networking-level tasks.

Important

This guide is divided into two parts after this introduction: One for multi-account landing zone accounts and one for single-account landing zone accounts. The onboarding is quite different for the two, please go next to the section of the guide that applies to your situation.

Learning about AMS

To understand AMS better, refer to these [AMS User Guide](#) sections:

- [What Is AWS Managed Services](#) introduces the AMS service and describes the key features, operations, and interfaces as well as a typical AMS-managed network architecture. This chapter also provides information on access management including how to access your AMS-managed resources and using bastions.
- [Key Terms](#) provides definitions and explanations for AMS terminology.
- [Understanding AMS Defaults](#) provides the default values AMS uses, including the defaults for basic environment components, IAM and EC2, proxies, monitored metrics, logging, endpoint security (EPS), backups, and patching.
- [Change Management](#) provides details on how requests for change (RFCs) and change types (CTs) work and includes examples of using AMS RFCs.
- Several additional chapters cover accessing the AWS console, the AMS CLI, using the AMS change management system, the AMS SKMS, security, service requests, incidents, monitoring, logs, EPS, backups, and patch management.

To learn more about AMS multi-account landing zone architecture, see [Multi-Account Landing Zone network architecture](#)

To learn more about AMS single-account landing zone architecture, see [Single-Account Landing Zone network architecture](#)

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.
- *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.
- *Billing start date*: AWS Managed Services accounts are activated once you have granted access to AMS to a compatible account and AMS Activation notification occurs as defined in the AWS Managed Services Documentation. If the activation of the AWS Managed Services accounts, Add-on Service Request, or Account tier Service Request is received by AWS on or prior to the 20th day of the month, then the change will be effective as of the first day of the calendar month following the AMS Activation notification or such Service Request. If the activation or Service Request is received by AWS after the 20th day of the month, then the change will be effective as of the first day of the second calendar month following AMS Activation notification or such Service Request.

AMS Activation Notification to the customer 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 Date*: The last day of the calendar month in which the Customer provides the AMS Account Service Termination Request, or the last day of the calendar month following the end of the requisite notice period; provided that, if the Customer provides the AMS Account Service Termination Request after the 20th day of the calendar month, the Service Termination Date will be the last day of the calendar month following the calendar month that such AMS Account Service Termination Request was provided.
 - *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 via a Service Request ("AMS Account Termination Request").

- *Effect of Termination of specified AWS Managed Services Accounts.:* On the Service Termination Date, AWS will (i) hand over the controls of all AMS Accounts or the specified AMS Account, as applicable, to Customer, or (ii) the parties will remove the AWS Identity and Access Management roles that give AWS access from all AMS Accelerate Accounts or the specified AMS Accelerate Account, as applicable.

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 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 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, usually to notify you of upcoming patching.

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 AWS 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.

AWS Managed Services modes

Use this to help you select the appropriate AWS Managed Services (AMS) mode for hosting your applications, based on your desired combination of flexibility and prescriptive governance to achieve your business outcomes.

The intended audience for this information is:

- Customer teams responsible for the strategy and governance of their landing zone. This information will help the team lay out the foundation of an AMS-managed landing zone, with the AMS modes they'd like to offer to their internal and external customers.
- Business and application owners tasked with migrating their application to AMS. This information will help with planning application migration, with the appropriate AMS mode to migrate/host their application. Note, the same application can be hosted in more than one AMS mode during different phases of its Software Development Life Cycle (SDLC) lifecycle.

- AMS partners tasked with guiding customers on the different options to build and migrate to AMS.

This information assumes that you have already made the decision to leverage AMS to accelerate your journey to the cloud. Refer to this paper at two points in your cloud migration journey: First, during the foundation phase of setting up the AMS-managed platform. Second, when you are transitioning from the foundation to the migration phase of your cloud adoption journey, just after onboarding to AMS is complete and you're focusing on application governance and operations.

Types of modes and accounts in AMS

AMS modes can be defined as the ways of interacting with the AMS service under the specific governance framework for each mode. The landing zone differences, multi-account landing zone or MALZ and single-account landing zone or SALZ are noted. The modes are:

- AMS-managed: Standard change management (CM) mode and Operations on demand (OOD)
- AMS-managed: Direct Change mode
- AMS-managed: AWS Service Catalog on AMS
- AMS-managed: Self Service Provisioning (SSP) mode
- AMS-managed: Developer mode
- Customer Managed mode

AMS feature	Standard CM mode / OOD*	Direct Change mode	AWS Service Catalog	Self-service provisioning / Developer mode	Customer Managed
Landing Zone Configuration	MALZ and SALZ	MALZ and SALZ	MALZ and SALZ		
Change Management	Change scheduling, review of manual changes, and change record	Same as Standard CM for high-risk changes like IAM or security groups	None		
Logging, Monitoring, Guardrails, and Event Management	Yes (supported resources)			No	
Continuity management	Yes (supported resources)			Not applicable / No	No
Security management	Instance level security controls and account level controls			Account level controls	AWS Org level controls
Patch management	Yes			Not applicable / No	No
Incident and problem management	Response and resolution SLA for AMS supported resources			Response SLA for resulting resources	No

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Types of modes and accounts in AMS

AMS feature	Standard CM mode / OOD*	Direct Change mode	AWS Service Catalog	Self-service provisioning / Developer mode	Customer Managed
Reporting	Yes			No	
Service request management	Yes			Support requests only	No

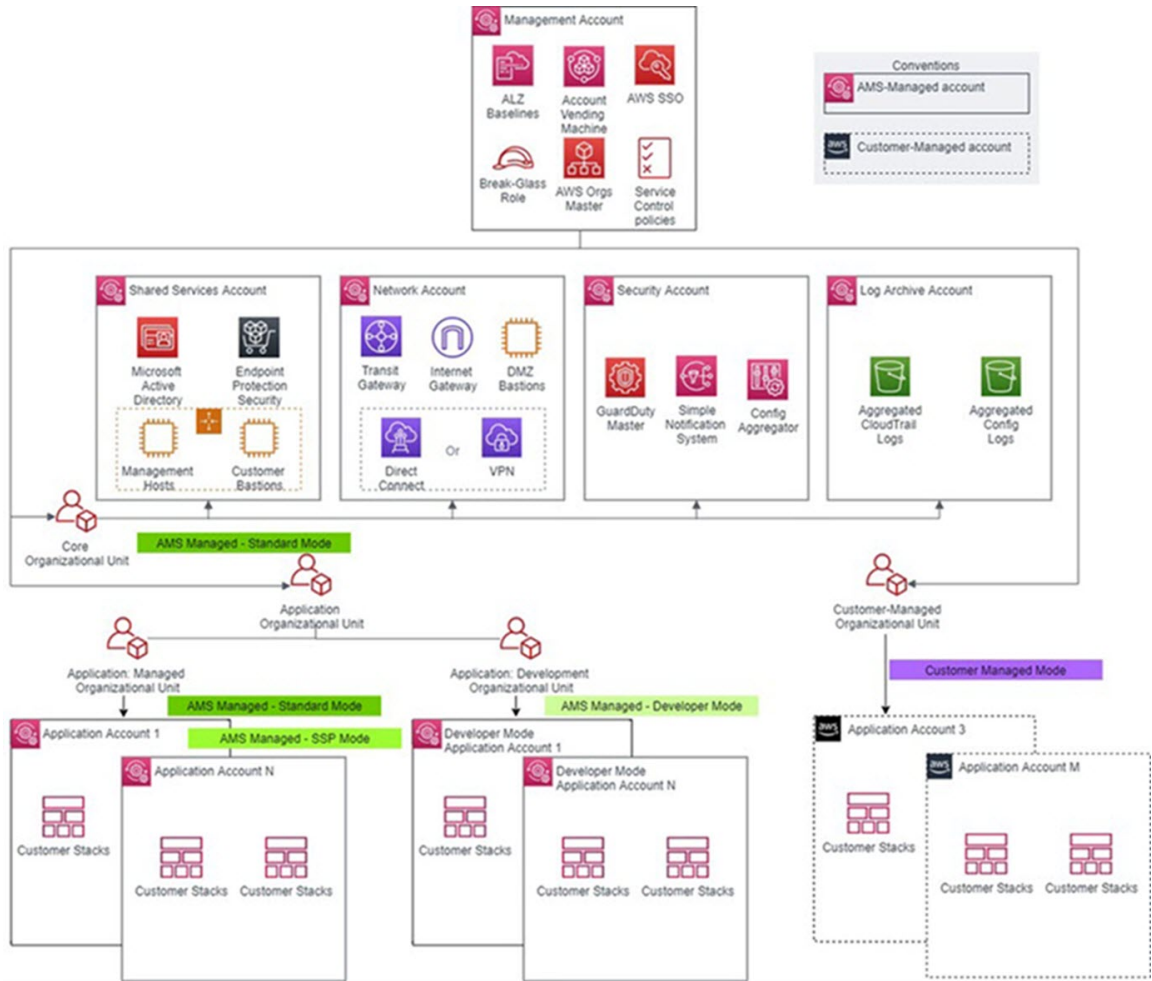
*Operations On Demand (OOD) has an offering for customers using the Standard CM mode to manage their changes through dedicated resourcing. For more details, see the [Operations on Demand catalog of offerings](#) and talk to your cloud service delivery manager (CSDM).

AMS multi-account landing zone (MALZ) gives you the option to automatically provision application accounts (or resource accounts) under the default Organizational Units (OU): Customer Managed OU, Managed OU, or Development OU. The infrastructure provisioned in the application accounts created under each of these OUs is subject to the specific AMS mode offered by those foundational OUs. It is common to find a mix of two or more modes in the same application account. For example: Standard mode and SSP mode can coexist in an AMS managed account that hosts pipeline architecture consisting of API Gateway and Lambda for trigger functions, and EC2, S3, and SQS for ingestion and orchestration. In this case, SSP mode would apply to Lambda and API Gateway.

Figure 1 presents how different modes are offered through the foundational OUs in AMS. When requesting a new application account in AMS, you must select the OU for the account.

MALZ architecture and associated AMS modes

AMS Advanced Onboarding Guide AMS
 Advanced Account Onboarding Information
 Types of modes and accounts in AMS



AMS leverages the foundational OUs based on AWS best practices as a way to logically manage accounts using Service Control Policies (SCPs). This serves as a way to enforce the governance framework with each AMS mode. Any governance and security guardrails (in the form of SCPs) applied to the foundational OUs also get applied to the custom/child OUs automatically. Additional SCPs can be requested for the child OUs. It is important to understand that application accounts are not the same as modes. Modes are applied to the infrastructure provisioned within the accounts and define the operational responsibilities between AMS and customers.

Figure 1: MALZ architecture and associated AMS modes

AMS Modes	Default Governance controls (Guardrails)	
	Preventative Controls	Detective Controls
AMS Managed – Standard CM Mode and OOD	Restrictive	Restrictive
AMS Managed - Direct Change Mode (DCM) AMS Managed – AWS Service Catalog	Restrictive	Restrictive
AMS Managed – Self Service Provisioning (SSP)	Restrictive	Restrictive
AMS Managed – Developer Mode	Permissive	Permissive
Customer Managed	Permissive	Permissive

Note

"Restrictive" implies that you can request custom policies for these OUs, they are approved by AMS on a case-by-case basis to ensure they don't interfere in AMS's capabilities to provide operational excellence. For a detailed list of AMS guardrails see [AMS Guardrails](#) in the user guide.

AMS-managed Standard Change Management (CM) mode

This mode provides standardized governance supported by pre-defined guardrails and a strict set of controls that make accounts in this mode operationally secure.

There can be a learning curve in this mode as application development teams adopt the AMS processes to work within the boundaries of AMS change management, or use AWS Service Catalog to provision resources.

AMS-managed Self Service Provisioning (SSP) mode

This mode provides full access to native AWS service and API Capabilities in AMS managed accounts. You access services through standardized, scoped down, IAM roles. AMS provides service requests

and incident management. Alerting, monitoring, logging, patch, back up, and change management are your responsibility. In many cases, Self-Service Provisioning services (SSPS) are self-managed, or serverless, and don't require management of certain operational tasks like patching. You benefit from using these services within the environment boundary defined by AMS guardrails and any IAM changes (including service linked roles, service roles, cross-account roles, or policy updates) need to be approved by AMS Operations to maintain the baseline security of the platform. You can leverage CloudFormation templates to automate deployment of these services but this is not supported for all SSP services currently. Examples:

- You deploy a data lake using the AMS-managed SSP mode that leverages EC2, S3, Glue and Lambda services. In this case, Glue and Lambda are considered Self Service Provisioning services and you are responsible for monitoring, logging, patch, back up, and change management.
- You deploy containerized applications using this mode; services like ECS, and EKS on Fargate, are SSP services and task monitoring, logging, container level security are the your responsibility. To use additional features like Service Accounts for EKS, you request a change with AMS to enable IAM roles and policies for the specific cluster. You do not need to explicitly request accounts to be provisioned in SSP mode. When you request any of the SSP services be enabled in your application account, SSP mode is automatically activated for those services. Not all AWS services are available as SSP services, a complete list of SSP services can be found in the AMS Service Description document.

For information on all available self-service provisioning services, see [Self-Service Provisioning Services \(SSPS\)](#).

AMS-managed Developer mode

This mode provides two options for provisioning infrastructure, either AMS change management, or access to native AWS service APIs via a highly permissive IAM role ("Developer role"). Resources provisioned through the highly permissive role are managed through a less-restrictive set of guardrails, enforced through detective controls, and have fewer preventative controls.

Depending on which option is selected, you can take on more operational responsibility while also gaining flexibility. It is important to note that Developer mode does not automatically grant access to any AWS service, but only those that have been onboarded to AMS. The two most common use cases for building in this mode are described next. In both use cases, Developer mode may co-exist with Standard mode and/or SSP mode in the same application account.

- Use Case: As a way to expedite deployment or migration of applications in AMS, with the objective that production-ready workloads will operate in AMS Managed – Standard or SSP mode. In this case, you utilize a mix of Developer mode in pre-production phase, and Standard mode or SSP mode in production phase, for the same application. Working through AMS change management may be considered an impediment for application teams due to the learning curve, or the speed of processing manual RFCs. With Developer mode, you can bypass the AMS change management system, while iterating in an account that is protected by the baseline AMS security-hardened network and permissions boundaries. Once the pre-production application design and configuration is finalized, you have the option to re-deploy the production-ready application using CloudFormation templates, or custom AMIs, that are ingested via the AMS change management system. The infrastructure created as an output is consequently managed by AMS.

Example: Setting up a CI/CD pipeline based on open source, or native AWS services, to deploy code to multiple accounts in AMS Managed Landing Zone can save time by building in Developer mode and leveraging the "Developer IAM role" to optimize configurations and permissions. Once finalized, you can re-deploy the infrastructure through AMS change management. Such a pipeline could also be built and iterated upon in Standard mode or SSP mode, however you would need to plan for additional time in processing manual RFCs related to IAM permissions.

- Use Case: As a way to operationalize configurations and tools that are not supported by the AMS change management system. In this case, AMS Managed – Developer mode will be used to host

production workload, with you taking over operational responsibility for the infrastructure provisioned using the "developer IAM role". It is highly recommended that you leverage the AMS change management to provision infrastructure that can be operated by AMS, like EC2, ELB, EBS, S3, etc, so that you can offload operational responsibility for those services to AMS. In this case the application operates in a mixed mode configuration with both Developer mode and Standard mode in the same application account. You can then focus on operational support for services not in scope, this includes monitoring, patching, and continuity management.

- Use Case: Extending a Terraform-based enterprise code repository in AMS can utilize Developer mode to provision infrastructure; however, you are responsible for operating any infrastructure provisioned through Terraform.
- Use Case: You want to deploy EKS on EC2 instead of EKS on Fargate (which is offered in SSP mode). In this case, you can use Developer mode to operate you desired configuration in production while leveraging the security offered by detective controls in AMS Managed Landing Zone.

For usage information, see [Developer mode](#).

Note

Self-Service Provisioning (SSP) mode and Developer mode may both appear to be a suitable fit for an application that has complex architecture rooted in native AWS Services. When architecting workloads, you make trade-offs between operational excellence and agility, based on your business context. This is a good way to think about selecting SSP mode or Developer mode for your application. The selection may also change based on the SDLC phase of the application. For example: When the application is production-ready, then SSP mode maybe a more appropriate option due to stricter AMS guardrails in this mode. The guardrails are enforced in the form of preventative controls like RFC-based change control for IAM updates and SCPs at the application OU level. These business decisions can drive your engineering priorities. You might optimize to increase flexibility for application owners in "pre-prod" phase at the expense of governance and operational support.

AMS-managed Direct Change mode

AMS Direct Change mode (DCM) extends AMS Advanced change management by providing native AWS access to AMS Advanced Plus and Premium accounts to provision and update AWS resources. Use DCM to provision AMS-managed resources using AWS CloudFormation, and to update any AMS-managed resource through the AWS Management Console, AWS APIs or AWS CloudFormation. Use DCM to accelerate migrations by deploying changes via native AWS access.

While DCM unlocks permissions to configure AMS-managed resources using the AWS Management Console, AWS APIs AWS CloudFormation, it also preserves the security boundary of the account. With DCM you can use common tool sets between AWS and AMS migration projects. Depending on use case, you can choose to use AWS CloudFormation during accelerated migrations, or RFCs when you need to leverage AMS curated deployment patterns.

Use DCM to:

- Provision and update fully managed stacks via direct AWS CloudFormation permissions
- Update AMS-managed resources through direct AWS API permissions

To see more use case details for DCM usage, see [Direct Change mode](#).

AMS-managed AWS Service Catalog

AWS Service Catalog provides you with an alternative to the AMS Advanced request for change (RFC) process for provisioning and updating resources in your AMS-managed accounts. AMS Advanced manages all of the infrastructure operations tasks needed to run AWS at scale for all infrastructure

resources provisioned through AWS Service Catalog including security, compliance, provisioning, availability, patch, monitoring, alerting, reporting, incident response, and cost optimization.

Utilizing AWS Service Catalog in your AMS-managed account provides you with a mechanism to centrally manage commonly deployed IT services, and helps you achieve consistent governance, while enabling users to quickly deploy only the approved IT services they need into their managed environments.

AMS Customer managed mode

This mode provides a governance model that is flexible and can be adapted to your requirements. This can be considered a fallback option for services and applications that AMS is unable to operate for you. AMS does not operate infrastructure hosted in accounts created under this mode. However, you can leverage centralized multi-account management in this mode. The following Multi-Account Landing Zone features can be leveraged in this mode:

- Automated Account deployment
- Connectivity through Transit Gateway in networking account
- AMS Config Rules library
- Store copies of logs in logging account
- Aggregation of customer managed Guard Duty alerts to Security account
- Consolidated Billing
- Enablement of custom Service Control Policies.

For example: If you want to run workloads on Ubuntu Pro, which is not an Operating System managed by AMS, you could use a customer managed account for hosting it. You can also consolidate workloads through customer managed accounts, to take advantage of the bulk discount on Reserved Instances/ Sharing Plans available through sharing across an AWS organization.

AMS modes and applications or workloads

Selecting the appropriate AMS mode for your applications or workloads and deciding the Organizational Unit under which you host your application. You could do this by requesting a new application account or hosting in an existing application account.

You should consider operational and governance requirements for your applications when selecting the right fit. The selection of the appropriate AMS mode for each application or workload depends on the following factors:

- The type of SDLC lifecycle function that the environment will provide (e.g., sandbox with unmoderated changes, UAT with some frequent changes, production with minimal changes and highly regulated)
- The governance policies needed (enforced through SCPs at the OU level)
- Operational Model (if you want to own the operational responsibility or want to outsource that to AMS)
- The desired business outcomes, like time to operate in the cloud, and cost of operations.

The following table outlines key considerations for application owners to help decide on the most suitable AMS mode. Application owners should include an assessment phase ahead of application migration to fully understand which mode applies to their specific application. Example: For applications based on cloud-native services or serverless architecture, the best option could be to start building and iterating in Developer mode and deploy the final Infrastructure as Code using AMS Managed – SSP mode. In this case light re-factoring may be required to ensure that any CloudFormation templates created for automated deployment meet the ingest guidelines laid out by AMS. Additionally, any IAM permissions need to be approved by AMS Security to ensure they follow the least privilege model.

The AMS mode selected to host the application, can help enable you to build towards you desired cloud operating model.

Note

More than one cloud operating model can existing in a single AMS Managed Landing Zone based on the different AMS modes selected to host the applications.

Decision issues	Standard CM mode / OOD*	AWS Service Catalog	Direct Change mode	Self-service provisioning	Developer mode	Customer Managed
Operational readiness						
Logging, Monitoring and Event Management	AMS responsible for all managed infrastructure			Customer responsible for Self-Service Provisioned Services (SSP)	Customer responsible for resources provisioned using developer IAM role outside AMS CM system	Customer responsible
Continuity Management	AMS responsibility to execute backup plan selected by customer			Customer responsible for Self-Service Provisioned Services (SSP)	Customer responsible for resources provisioned using developer IAM role outside AMS CM system	Customer responsible
Instance Level Access Management	AMS-managed through one-way AD trust with on-prem domain. Requires managed infrastructure to join AMS domain			Not applicable	Customer responsible for resources provisioned using developer IAM role outside AMS CM system	Customer responsible
Security Management and Account Level Access Management	AMS responsibility for all managed accounts			AMS responsible for all managed accounts	Customer responsible for resources provisioned using developer IAM role outside AMS CM system	Customer responsible
Patch Management	AMS responsibility for all managed accounts			Customer responsible for Self-Service Provisioned	Customer responsible for resources provisioned using	Customer responsible

AMS Advanced Onboarding Guide AMS
 Advanced Account Onboarding Information
 AMS modes and applications or workloads

Decision issues	Standard CM mode / OOD*	AWS Service Catalog	Direct Change mode	Self-service provisioning	Developer mode	Customer Managed
				Services (SSP)	developer IAM role outside AMS CM system	
Change Management	AMS responsibility for all managed accounts			Customer responsible for Self-Service Provisioned Services (SSP)	Customer responsible for resources provisioned using developer IAM role outside AMS CM system	Customer responsible
Provisioning Management	Prescriptive and standardized for the provisioning options offered in AMS	Flexibility to directly use AWS service API for AWS Service Catalog following AMS prescriptive standards	Flexibility to directly use AWS service API following AMS prescriptive standards	Flexibility to directly use AWS service APIs for SSP services	Flexibility to directly use AWS service API for provisioning	Customer responsibility
Incident Management and Audit	AMS responsible for all managed accounts				Customer responsible for resources provisioned using developer IAM role outside AMS Change Management System	Customer responsible
GuardRails and Shared infrastructure (Network) and Security Framework	Prescriptive and standardized leveraging AMS Core Accounts					Flexible and bespoke leveraging AMS Core Accounts
Application readiness						

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
AMS modes and applications or workloads

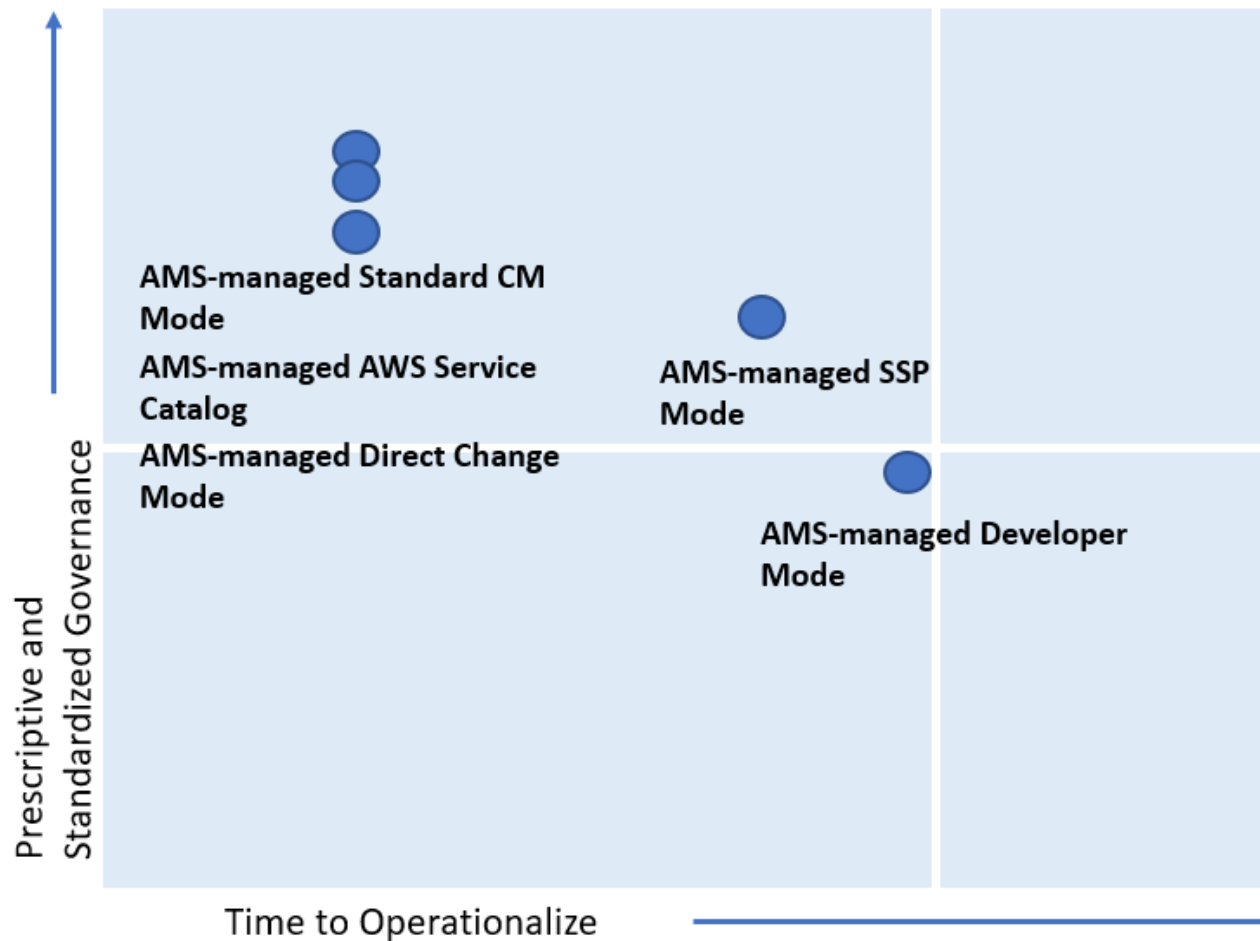
Decision issues	Standard CM mode / OOD*	AWS Service Catalog	Direct Change mode	Self-service provisioning	Developer mode	Customer Managed
Application refactoring	Light refactoring is needed				Light refactoring is needed (if provisioned using AMS Standard CM)	No need for refactoring
Support for AWS services	Limited to what is supported by AMS					Not limited
Business considerations						
Time to operational readiness	Three to six months			6 months + dependent on customer application operations competencies		6-18 months dependent on customer infrastructure and application operations competencies
Costs	\$\$\$\$			\$\$\$	\$\$	\$
Application examples	Webserver with 3 tier stack, apps with compliance and regulatory requirements			Webserver using API Gateway, containerized application leveraging ECS/EKS	Iterating/ optimizing on Data Lake application that uses Lambda, Glue, Athena, etc	De-centralized accounts/ applications like sandbox, third party managed applications

*Operations On Demand (OOD) has an offering for customers using the Standard CM mode to manage their changes through dedicated resourcing. For more details, see the [Operations on Demand catalog of offerings](#) and talk to your cloud service delivery manager (CSDM).

Note

The price comparison between SSP mode and Developer mode assumes that the same AWS services are provisioned.

Comparing AMS Modes against business and IT objectives



As shown, if you are looking for a highly controlled and standardized governance model for your applications, then AMS-managed Standard Change, AWS Service Catalog, or Direct Change modes are the best fit. If you require a bespoke governance model with a focus on application innovation without the need for operational readiness, select Customer Managed mode. With Customer Managed mode, it could take you a longer time to operationalize your applications as you bear the responsibility to establish people, processes, and tools to support operational capabilities such as Incident Management, Configuration Management, Provisioning Management, Security Management, Patch Management, etc.

Real world use cases for AMS modes

Examine these to help determine how to use AMS modes.

- **Use Case 1, business imperative to lower costs with a time-sensitive data center exit:** An enterprise with a compelling business event, like a data center exit, is interested in re-hosting their on-prem applications on the cloud. Most of the on-prem inventory consists of Windows and Linux servers with a mix of operating system versions. In doing so, the customer also wants to take advantage of cost savings that moving to the cloud offers and improving the technical and security posture of their applications. The customer wants to move fast but does not have the in-house cloud operations expertise built out yet. The customer has to find a balance of refactoring, too much refactoring can be risky against a tight timeline. However, with some refactoring, like updating OS versions and optimizing databases, applications can achieve the next level of performance. In this example, the customer can select AMS Managed Standard mode to re-host most of their applications. AMS provides

infrastructure operations, while also guiding the customer operations teams on best practices on securely operating in the cloud.

AMS-managed AWS Service Catalog and AMS-managed Direct Change mode gives the customer an extra flexibility while achieving the same business outcomes and objectives. In addition, the customer can use the AMS Operations On Demand (OOD) offering to have dedicated AMS operations engineers to prioritize the execution of requests for change (RFCs).

While offloading the undifferentiated infrastructure operational tasks (patching, backups, account management, etc) to AMS, the customer can continue to focus on optimizing their application and ramp-up their internal teams on cloud operations. AMS provides monthly reports to the customer on cost savings, and makes recommendations on resource optimizations. In this use case, if there were end-of-life applications hosted on legacy OS versions like Windows 2003 and 2008, that the customer decided not to re-factor, those can also be migrated to AMS and hosted in an account that leverages Customer Managed mode.

- **Use Case 2, building a data lake with Lambda, Glue, Athena within the secure AMS boundary:** An enterprise is looking to set up a Data Lake to meet the reporting needs for multiple applications in AMS. The customer wants to use S3 buckets for the storage of datasets and AWS Athena to query against the dataset for each report. S3 and AWS Athena will be deployed in separate AMS Managed accounts. The account with S3 also has other services like Glue, Lambda, and Step Functions to build a data ingestion pipeline. Glue, Lambda, Athena, and Step Functions are considered Self-Service Provisioning (SSP) services in this case. The customer also deployed an EC2 instance in the account that acts as an ad hoc tooling/scripting server. The customer starts by requesting AMS to enable the SSP services in their AMS Managed account. AMS provisions an IAM role for each service that the customer can assume, once the role is onboarded to the customer's federation solution. For ease of management, the customer can also combine the policies for the separate IAM roles into one custom role, alleviating the need to switch roles when working between the AWS services. Once the role is enabled in the account, the customer is able to configure the services as per their requirements. However, the customer must work with the AMS change management system to request additional permissions, depending on their use case.

For example, for access to Glue Crawlers, additional permissions are needed by Glue. Additional permissions will also be needed to create event sources for Lambda. The customer will work with AMS to update IAM roles to allow cross-account access for Athena to query S3 buckets. Updates to service roles or service-linked roles will also be needed through AMS change management for Lambda to call the Step Functions service, and Glue to read and write to all S3 buckets. AMS works with customers to ensure that the least-privilege access model is followed and the IAM changes requested are not overly permissive and opening up the environment to unnecessary risk. The customer's data lake team spends time planning for all IAM permissions needed for the services specific to the customer's architecture and requests AMS to enable them. This is because all IAM changes are processed manually and undergo review from the AMS Security team. Time to process these requests should be accounted for in the application deployment schedule.

As the SSP services are operational in the account, the customer can request support and report issues through AMS incident management and service requests. However, AMS will not actively monitor performance and concurrency metrics for Lambda, or job metrics for Glue. It is the customer's responsibility to ensure appropriate logging and monitoring is enabled for SSP services. The EC2 instance and S3 bucket in the account are fully managed by AMS.

- **Use Case 3, quick and flexible set up of a CI/CD deployment pipeline in AMS:** A customer is looking to set up a Jenkins-based CI/CD pipeline to deploy code pipeline to all application accounts in AMS. The customer may find it most suitable to host this CI/CD pipeline in the AMS-managed Direct Change mode (DCM) or AMS-managed Developer mode because it gives them flexibility to set up the Jenkins server with required custom configuration on EC2, with the desired IAM permissions to access CloudFormation and S3 buckets that host the artifact repository. While this can also be done in the AMS Managed- Standard mode, the customer team would need to create multiple manual RFCs for IAM roles to iterate on the least permissive set of approved permissions, which are manually reviewed by AMS. DCM allows the customers to achieve their operational goals on AWS while avoiding the need

to create multiple manual RFCs for IAM roles, when using AMS-managed Standard CM mode, to iterate on the least permissive set of approved permissions, which are manually reviewed by AMS. This would take time as well as education on the customer's part to ramp up AMS processes and tools. Working with Developer mode, the customer can start with a "developer role" to provision infrastructure using native AWS APIs. The quickest and most flexible way to set up this pipeline would be to use AMS Managed-Developer mode. Developer mode gives the quickest and easiest way, while compromising on operational integration, while DCM is less flexible but does provide the same level of operational support as Standard CM mode.

- **Use Case 4, bespoke operating model within the AMS foundation:** A customer is looking at a deadline-driven data center exit and one of their enterprise applications is fully managed by a third party MSP, including application operations and infrastructure operations. Assuming that the customer does not have time in the schedule to re-factor this application so that it can be operated by AMS, Customer Managed mode is a suitable option. The customer can take advantage of the automated and quick set up of AMS managed Landing Zone. They can leverage the centralized account management that controls account vending and connectivity through the centralized networking account. It also simplifies their billing by consolidating charges for all customer managed accounts through the AMS Payer account. The customer has flexibility to set up their bespoke access management model with the MSP separate from standard access management used for AMS Managed accounts. This way, using Customer Managed mode, they can set up an AMS managed environment while meeting their business requirement of vacating their on-prem environment. In this case, if the customer also has Windows-based applications that they are migrating to the cloud, and choose to move them to a Customer Managed account, the customer is responsible for creating a cloud operating model. This can be complex, expensive, and time consuming depending on the customer's ability to transform traditional IT processes and train people. The customer can save time and cost by "lift and shift" of such workloads to an AMS Managed account and offload infrastructure operations to AMS.

Note

Customers may sometimes feel the need to move application accounts between the governance framework of Standard or SSP mode and Developer mode. For example, customers may host an application in AMS Managed mode as part of initial lift and shift migration, but overtime want to re-write the application to optimize it for cloud-native AWS services. They could change the mode of the pre-prod account from AMS Managed - Standard to AMS Managed-Developer, giving them the flexibility and agility for provisioning infrastructure. However, once infrastructure provisioning changes have been made using the "developer role", the same infrastructure cannot be moved back to AMS Managed - Standard mode. This is because AMS cannot guarantee operations of infrastructure that was provisioned outside of the AMS change management system. Customers may need to create a new application account that offers AMS Managed - standard mode and then re-deploy the "optimized" infrastructure configuration through CloudFormation templates or custom AMIs ingested into an AMS Managed account. This is a clean way to deploy a production ready configuration. Once deployed, the application will be under prescriptive AMS governance and operations. The same applies to switching modes between Customer Managed and AMS Managed.

AMS post-account prescriptive guidance

As organizations adopt distributed operations and DevOps practices, there are a core set of operational capabilities that should be applied to every account prior to deployment of workloads to meet the pillars of Well Architected.

This link downloads a ZIP file containing a Word document, and a ZIP file with scripts and examples. Automated Account Setup is a set of scripts to automate, or bootstrap, the setup of a new application account.

Once a new account is vended, and before any workloads are deployed, in order to make the account ready from an operational, security and management point of view, you setup default backup plans,

patch windows, and encryption (and more). To help improve the agility, consistency, and responsiveness for application account setup, the following sample "How To" is provided for your reference.

[Automated Account Setup](#).

What we do, what we do not do

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

Note

To request that AMS provide an additional AWS service, file a service request. For more information, see [Making Service Requests](#).

• What we do:

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

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

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

• What we DON'T do:

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

AMS egress traffic management

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

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

- It acts as an implicit (for example, transparent) proxy.
- and
- It allows AMS HTTP and HTTPS dependencies (listed in this section) in order to allow ongoing patching and maintenance of AMS managed infrastructure.

Some examples are:

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

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

Note

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

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

IAM User Role

An IAM role is similar to an IAM user, in that it is an AWS identity with permission policies that determine what the identity can and can't do in AWS. However, instead of being uniquely associated with one person, a role is intended to be assumable by anyone who needs it.

Currently there is one AMS default user role, `Customer_ReadOnly_Role`, for standard AMS accounts and an additional role, `customer_managed_ad_user_role` for AMS accounts with Managed Active Directory.

The role policies set permissions for CloudWatch and S3 log actions, AMS console access, read-only restrictions on most AWS services, restricted access to account S3 console, and AMS change-type access.

Additionally, the `Customer_ReadOnly_Role` has mutative, reserved-instances permissions that allow you to reserve instances. It has some cost-saving values, so, if you know that you're going to need a certain number of EC2 instances for a long period of time, you can call those APIs. To learn more, see [Amazon EC2 Reserved Instances](#).

Note

The AMS service level objective (SLO) for creating custom IAM policies for IAM users is four business days, unless an existing policy is going to be reused. If you want to modify the existing IAM user role, or add a new one, submit an [IAM resource: Update](#) or [IAM resource: Create](#) RFC, respectively.

If you're unfamiliar with Amazon IAM roles, see [IAM Roles](#) for important information.

Multi-Account Landing Zone (MALZ): To see the AMS multi-account landing zone default, un-customized, user role policies, see [Default AMS multi-account landing zone \(MALZ\) IAM User Roles \(p. 21\)](#), next.

Default AMS multi-account landing zone (MALZ) IAM User Roles

JSON policy statements for the default multi-account AMS multi-account landing zone user roles.

Note

The user roles are customizable and may differ on a per-account basis. Instructions on finding your role are provided.

These are examples of the default MALZ user roles. To make sure that you have the policies set that you need, run the AWS command `get-role` or sign in to the AWS Management -> [IAM console](#) and choose **Roles** in the navigation pane.

Core OU account roles

A core account is an MALZ-managed infrastructure account. AMS multi-account landing zone Core accounts include a management account and a networking account.

Core OU account: Common roles and policies

Role	Policy or policies
AWSManagedServicesReadOnlyRole	ReadOnlyAccess (p. 30) (Public AWS Managed Policy).
AWSManagedServicesCaseRole	ReadOnlyAccess (p. 30)
	AWSSupportAccess (p. 28) (Public AWS Managed Policy).
AWSManagedServicesChangeManagementRole (Core account version)	ReadOnlyAccess (p. 30)
	AWSSupportAccess (p. 28)
	AMSChangeManagementReadOnlyPolicy (p. 25)
	AMSChangeManagementInfrastructurePolicy (p. 26)

Core OU account: Master account roles and policies

Role	Policy or policies
AWSManagedServicesBillingRole	AMSBillingPolicy (p. 23) (AMSBillingPolicy).
AWSManagedServicesReadOnlyRole	ReadOnlyAccess (p. 30) (Public AWS Managed Policy).
AWSManagedServicesCaseRole	ReadOnlyAccess (p. 30)
	AWSSupportAccess (p. 28) (Public AWS Managed Policy).
AWSManagedServicesChangeManagementRole (Master account version)	ReadOnlyAccess (p. 30)
	AWSSupportAccess (p. 28)
	AMSChangeManagementReadOnlyPolicy (p. 25)

AMS Advanced Onboarding Guide AMS
 Advanced Account Onboarding Information
 Default AMS multi-account landing
 zone (MALZ) IAM User Roles

Role	Policy or policies
	AMSChangeManagementInfrastructurePolicy (p. 26) AMSMasterAccountSpecificChangeManagementInfrastructurePolicy (p. 26)

Core OU Account: Networking account roles and policies

Role	Policy or policies
AWSManagedServicesReadOnlyRole	ReadOnlyAccess (p. 30) (Public AWS Managed Policy).
AWSManagedServicesCaseRole	ReadOnlyAccess (p. 30)
	AWSSupportAccess (p. 28) (Public AWS Managed Policy).
AWSManagedServicesChangeManagementRole (Networking account version)	ReadOnlyAccess (p. 30)
	AWSSupportAccess (p. 28)
	AMSChangeManagementReadOnlyPolicy (p. 25)
	AMSChangeManagementInfrastructurePolicy (p. 26)
	AMSNetworkingAccountSpecificChangeManagementInfrastructurePolicy (p. 26)

Application Account Roles

Application account roles are applied to your application-specific accounts.

Application account: Roles and policies

Role	Policy or policies
AWSManagedServicesReadOnlyRole	ReadOnlyAccess (p. 30) (Public AWS Managed Policy).
AWSManagedServicesCaseRole	ReadOnlyAccess (p. 30)
	AWSSupportAccess (p. 28) (Public AWS Managed Policy). This policy provides access to all support operations and resources. For information, see Getting Started with AWS Support .
AWSManagedServicesSecurityOpsRole	ReadOnlyAccess (p. 30)
	AWSSupportAccess Example (p. 28)
	This policy provides access to all support operations and resources. AWSCertificateManagerFullAccess information, (Public AWS Managed Policy)

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Default AMS multi-account landing
zone (MALZ) IAM User Roles

Role	Policy or policies
	AWSWAFFullAccess information, (Public AWS Managed policy). This policy grants full access to AWS WAF resources. AMSSecretsManagerSharedPolicy (p. 27)
AWSManagedServicesChangeManagementRole (Application account version)	ReadOnlyAccess (p. 30)
	AWSSupportAccess (p. 28) (Public AWS Managed Policy). This policy provides access to all support operations and resources. For information, see Getting Started with AWS Support .
	AMSSecretsManagerSharedPolicy (p. 27)
	AMSChangeManagementPolicy (p. 27)
	AMSReservedInstancesPolicy (p. 28)
	AMSS3Policy (p. 28)
AWSManagedServicesAdminRole	ReadOnlyAccess (p. 30)
	AWSSupportAccess (p. 28)
	AMSChangeManagementInfrastructurePolicy (p. 26)
	AWSMarketplaceManageSubscriptions (p. 29)
	AMSSecretsManagerSharedPolicy (p. 27)
	AMSChangeManagementPolicy (p. 27)
	AWSCertificateManagerFullAccess (p. 29)
	AWSWAFFullAccess (p. 29)
	AMSS3Policy (p. 28)
AMSReservedInstancesPolicy (p. 28)	

Policy Examples

Examples are provided for most policies used. To view the [ReadOnlyAccess](#) policy (which is pages long as it provides read-only access to all AWS services), you can use this link, if you have an active AWS account: [ReadOnlyAccess](#). Also, a condensed version is included here.

AMSBillingPolicy

AMSBillingPolicy

The new Billing role can be used by your accounting department to view and change billing information or account settings in the Master account. To access information such as Alternate Contacts, view the account resources usage, or keep a tab of your billing or even modify your payment methods, you use this role. This new role comprises of all the permissions listed in the [AWS Billing IAM actions web page](#).

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Default AMS multi-account landing
zone (MALZ) IAM User Roles

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "aws-portal:ViewBilling",
        "aws-portal:ModifyBilling"
      ],
      "Resource": "*",
      "Effect": "Allow",
      "Sid": "AllowAccessToBilling"
    },
    {
      "Action": [
        "aws-portal:ViewAccount",
        "aws-portal:ModifyAccount"
      ],
      "Resource": "*",
      "Effect": "Allow",
      "Sid": "AllowAccessToAccountSettings"
    },
    {
      "Action": [
        "budgets:ViewBudget",
        "budgets:ModifyBudget"
      ],
      "Resource": "*",
      "Effect": "Allow",
      "Sid": "AllowAccessToAccountBudget"
    },
    {
      "Action": [
        "aws-portal:ViewPaymentMethods",
        "aws-portal:ModifyPaymentMethods"
      ],
      "Resource": "*",
      "Effect": "Allow",
      "Sid": "AllowAccessToPaymentMethods"
    },
    {
      "Action": [
        "aws-portal:ViewUsage"
      ],
      "Resource": "*",
      "Effect": "Allow",
      "Sid": "AllowAccessToUsage"
    },
    {
      "Action": [
        "cur:DescribeReportDefinitions",
        "cur:PutReportDefinition",
        "cur>DeleteReportDefinition",
        "cur:ModifyReportDefinition"
      ],
      "Resource": "*",
      "Effect": "Allow",
      "Sid": "AllowAccessToCostAndUsageReport"
    },
    {
      "Action": [
        "pricing:DescribeServices",
        "pricing:GetAttributeValues",
        "pricing:GetProducts"
      ],
      "Resource": "*",

```

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Default AMS multi-account landing
zone (MALZ) IAM User Roles

```
    "Effect": "Allow",
    "Sid": "AllowAccessToPricing"
  },
  {
    "Action": [
      "ce:CreateCostCategoryDefinition",
      "ce>DeleteCostCategoryDefinition",
      "ce:DescribeCostCategoryDefinition",
      "ce>ListCostCategoryDefinitions",
      "ce:UpdateCostCategoryDefinition"
    ],
    "Resource": "*",
    "Effect": "Allow",
    "Sid": "AllowAccessToCostCategories"
  }
]
```

AMSChangeManagementReadOnlyPolicy

AMSChangeManagementReadOnlyPolicy

Permissions to see all AMS change types, and the history of requested change types.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AMSCoreAccountsCMAndSKMSReadOnlyAccess",
    "Effect": "Allow",
    "Action": [
      "amscm:GetChangeTypeVersion",
      "amscm:GetRfc",
      "amscm>ListChangeTypeCategories",
      "amscm>ListChangeTypeClassificationSummaries",
      "amscm>ListChangeTypeItems",
      "amscm>ListChangeTypeOperations",
      "amscm>ListChangeTypeSubcategories",
      "amscm>ListChangeTypeVersionSummaries",
      "amscm>ListRestrictedExecutionTimes",
      "amscm>ListRfcSummaries",
      "amsskms:GetStack",
      "amsskms:GetSubnet",
      "amsskms:GetVpc",
      "amsskms>ListAmis",
      "amsskms>ListStackSummaries",
      "amsskms>ListSubnetSummaries",
      "amsskms>ListVpcSummaries"
    ],
    "Resource": "*"
  }]
}
```

AMSMasterAccountSpecificChangeManagementInfrastructurePolicy

AMSMasterAccountSpecificChangeManagementInfrastructurePolicy

Permissions to request the Deployment | Managed landing zone | Master account | Create application account (with VPC) change type.

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

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Default AMS multi-account landing
zone (MALZ) IAM User Roles

```
"Effect": "Allow",
"Action": [
  "amscm:ApproveRfc",
  "amscm:CancelRfc",
  "amscm:CreateRfc",
  "amscm:RejectRfc",
  "amscm:SubmitRfc",
  "amscm:UpdateRfc",
  "amscm:UpdateRfcActionState",
  "amscm:UpdateRestrictedExecutionTimes"
],
"Resource": [
  "arn:aws:amscm:global:*:changetype/ct-1zdasmc2ewzrs:*"
]
}]
}
```

AMSNetworkingAccountSpecificChangeManagementInfrastructurePolicy

AMSNetworkingAccountSpecificChangeManagementInfrastructurePolicy

Permissions to request the Deployment | Managed landing zone | Networking account | Create application route table change type.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AMSNetworkingAccountAccess",
    "Effect": "Allow",
    "Action": [
      "amscm:ApproveRfc",
      "amscm:CancelRfc",
      "amscm:CreateRfc",
      "amscm:RejectRfc",
      "amscm:SubmitRfc",
      "amscm:UpdateRfc",
      "amscm:UpdateRfcActionState",
      "amscm:UpdateRestrictedExecutionTimes"
    ],
    "Resource": [
      "arn:aws:amscm:global:*:changetype/ct-1urj94c3hdfu5:*"
    ]
  }]
}
```

AMSChangeManagementInfrastructurePolicy

AMSChangeManagementInfrastructurePolicy (for Management | Other | Other CTs)

Permissions to request the Management | Other | Other | Create, and Management | Other | Other | Update change types.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AMSCoreAccountsAccess",
    "Effect": "Allow",
    "Action": [
      "amscm:CancelRfc",
      "amscm:CreateRfc",
      "amscm:SubmitRfc",
      "amscm:UpdateRfc",
    ]
  }]
}
```

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Default AMS multi-account landing
zone (MALZ) IAM User Roles

```
"amscm:UpdateRfcActionState",
"amscm:UpdateRestrictedExecutionTimes",
],
"Resource": [
"arn:aws:amscm:global::changetype/ct-1e1xtak34nx76:*",
"arn:aws:amscm:global::changetype/ct-0xdawir96cy7k:*",
]
}]
}
```

AMSSecretsManagerSharedPolicy

AMSSecretsManagerSharedPolicy

Permissions to view secret passwords/hashes shared by AMS through AWS Secrets manager (e.g. passwords to infrastructure for auditing).

Permissions to create secret password/hashes to share with AMS. (e.g. license keys for products that need to be deployed).

```
{
"Version": "2012-10-17",
"Statement": [{
"Sid": "AllowAccessToSharedNameSpaces",
"Effect": "Allow",
"Action": "secretsmanager:*",
"Resource": [
"arn:aws:secretsmanager:::secret:ams-shared/*",
"arn:aws:secretsmanager:::secret:customer-shared/*"
]
},
{
"Sid": "DenyGetSecretOnCustomerNamespace",
"Effect": "Deny",
"Action": "secretsmanager:GetSecretValue",
"Resource": "arn:aws:secretsmanager:::secret:customer-shared/*"
},
{
"Sid": "AllowReadAccessToAMSNameSpace",
"Effect": "Deny",
"NotAction": [
"secretsmanager:Describe*",
"secretsmanager:Get*",
"secretsmanager:List*"
],
"Resource": "arn:aws:secretsmanager:::secret:ams-shared/*"
}
]
}
```

AMSChangeManagementPolicy

AMSChangeManagementPolicy

Permissions to request and view all AMS change types, and the history of requested change types.

```
{
"Version": "2012-10-17",
"Statement": [{
"Sid": "AMSFullAccess",
"Effect": "Allow",
"Action": [
"amscm:*",
```


AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Default AMS multi-account landing
zone (MALZ) IAM User Roles

```
"amsskms:*"  
],  
"Resource": [  
  "*" ]  
}]  
}
```

AMSReservedInstancesPolicy

AMSReservedInstancesPolicy

Permissions to manage EC2 reserved instances; for pricing information, see [Amazon EC2 Reserved Instances](#).

```
{  
  "Version": "2012-10-17",  
  "Statement": [{  
    "Sid": "AllowReservedInstancesManagement",  
    "Effect": "Allow",  
    "Action": [  
      "ec2:ModifyReservedInstances",  
      "ec2:PurchaseReservedInstancesOffering"  
    ],  
    "Resource": [  
      "*" ]  
    } ]  
}
```

AMSS3Policy

AMSS3Policy

Permissions to create and delete files from existing S3 buckets.

Note

These permissions do not grant the ability to create S3 buckets; that must be done with the Deployment | Advanced stack components | S3 storage | Create change type.

```
{  
  "Version": "2012-10-17",  
  "Statement": [{  
    "Effect": "Allow",  
    "Action": [  
      "s3:AbortMultipartUpload",  
      "s3:DeleteObject",  
      "s3:PutObject",  
    ],  
    "Resource": "*" ]  
}
```

AWSsupportAccess

AWSsupportAccess

Full access to AWS Support. For information, see [Getting Started with AWS Support](#). For Premium Support information, see [AWS Support](#).

```
{
```

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Default AMS multi-account landing
zone (MALZ) IAM User Roles

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

AWSMarketplaceManageSubscriptions

AWSMarketplaceManageSubscriptions (Public AWS Managed Policy)

Permissions to subscribe, unsubscribe, and view AWS Marketplace subscriptions.

```
{  
  "Version": "2012-10-17",  
  "Statement": [{  
    "Action": [  
      "aws-marketplace:ViewSubscriptions",  
      "aws-marketplace:Subscribe",  
      "aws-marketplace:Unsubscribe"  
    ],  
    "Effect": "Allow",  
    "Resource": "*" }  
  ]  
}
```

AWSCertificateManagerFullAccess

AWSCertificateManagerFullAccess

Full access to AWS Certificate Manager. For Certificate Manager information, see [AWS Certificate Manager](#).

[AWSCertificateManagerFullAccess](#) information, (Public AWS Managed Policy).

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

AWSWAFFullAccess

AWSWAFFullAccess

Full access to AWS Web Application Firewall (WAF). For WAF information, see [AWS WAF - Web Application Firewall](#).

[AWSWAFFullAccess](#) information, (Public AWS Managed policy). This policy grants full access to AWS WAF resources.

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

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Default AMS single-account
landing zone (SALZ) IAM User Role

```
"Statement": [{  
  "Action": [  
    "waf:*",  
    "waf-regional:*",  
    "elasticloadbalancing:SetWebACL"  
  ],  
  "Effect": "Allow",  
  "Resource": "*"   
}]  
}
```

ReadOnlyAccess

ReadOnlyAccess (actions a-l only)

Read-only access to all AWS services and resources on the AWS console.

When AWS launches a new service, AMS updates the ReadOnlyAccess policy to add read-only permissions for the new service. The updated permissions are applied to all principal entities that the policy is attached to.

This doesn't grant the ability to log into EC2 hosts or database hosts.

```
{  
  "Version": "2012-10-17",  
  "Statement": [{  
    "Action": [  
      "aws-marketplace:ViewSubscriptions",  
      "aws-marketplace:Subscribe",  
      "aws-marketplace:Unsubscribe"  
    ],  
    "Effect": "Allow",  
    "Resource": "*"   
  }]  
}
```

Single-Account Landing Zone (SALZ): To see the AMS single-account landing zone default, uncustomized, user role policies, see [Default AMS single-account landing zone \(SALZ\) IAM User Role \(p. 30\)](#), next.

Default AMS single-account landing zone (SALZ) IAM User Role

JSON policy statements for the default AMS single-account landing zone user role.

Note

The SALZ default user role is customizable and may differ on a per-account basis. Instructions on finding your role are provided.

This is an example of the default SALZ user role, but to make sure that you have the policies set for you, run the AWS command `get-role` or sign in to the AWS Management -> IAM console at <https://console.aws.amazon.com/iam/>. In the IAM console, in the navigation pane, choose **Roles**.

The customer read-only role is a combination of multiple policies. A breakdown of the role (JSON) follows.

Managed Services Audit Policy:

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

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Default AMS single-account
landing zone (SALZ) IAM User Role

```
"Statement": [
  {
    "Sid": "BasicConsoleAccess",
    "Effect": "Allow",
    "Action": [
      "aws-portal:View*",
      "ec2-reports:View*",
      "support:*"
    ],
    "Resource": [
      "*"
    ]
  },
  {
    "Sid": "AuditAccessToAWSservices",
    "Effect": "Allow",
    "Action": [
      "acm:Describe*",
      "acm:List*",
      "appstream:Get*",
      "autoscaling:Describe*",
      "cloudformation:Describe*",
      "cloudformation:Get*",
      "cloudformation:List*",
      "cloudformation:ValidateTemplate",
      "cloudfront:Get*",
      "cloudfront:List*",
      "cloudsearch:Describe*",
      "cloudsearch:List*",
      "cloudtrail:DescribeTrails",
      "cloudtrail:GetTrailStatus",
      "cloudtrail:LookupEvents",
      "cloudwatch:Describe*",
      "cloudwatch:Get*",
      "cloudwatch:List*",
      "codecommit:Get*",
      "codecommit:List*",
      "codedeploy:BatchGet*",
      "codedeploy:Get*",
      "codedeploy:List*",
      "codepipeline:Get*",
      "codepipeline:List*",
      "config:Describe*",
      "config:Get*",
      "datapipeline:Describe*",
      "datapipeline:EvaluateExpression",
      "datapipeline:GetPipelineDefinition",
      "datapipeline:ListPipelines",
      "datapipeline:ValidatePipelineDefinition",
      "directconnect:Describe*",
      "ds:Describe*",
      "dynamodb:Describe*",
      "dynamodb:List*",
      "ec2:Describe*",
      "ec2:Get*",
      "ecs:Describe*",
      "ecs:List*",
      "elasticache:Describe*",
      "elasticache:List*",
      "elasticbeanstalk:Check*",
      "elasticbeanstalk:Describe*",
      "elasticbeanstalk:List*",
      "elasticbeanstalk:RequestEnvironmentInfo",
      "elasticbeanstalk:RetrieveEnvironmentInfo",
      "elasticfilesystem:Describe*",
      "elasticloadbalancing:Describe*",
```

AMS Advanced Onboarding Guide AMS
 Advanced Account Onboarding Information
 Default AMS single-account
 landing zone (SALZ) IAM User Role

```

    "elasticmapreduce:Describe*",
    "elasticmapreduce:List*",
    "elastictranscoder:List*",
    "events:Describe*",
    "events:Get*",
    "events:List*",
    "guardduty:Get*",
    "guardduty:List*",
    "kinesis:Describe*",
    "kinesis:List*",
    "kms:List*",
    "lambda:Get*",
    "lambda:List*",
    "macie:Describe*",
    "macie:Get*",
    "macie:List*",
    "opsworks:Describe*",
    "opsworks:Get*",
    "rds:Describe*",
    "rds:Download*",
    "rds:List*",
    "redshift:Describe*",
    "redshift:View*",
    "route53:Get*",
    "route53:List*",
    "route53domains:CheckDomainAvailability",
    "route53domains:Get*",
    "route53domains:List*",
    "sdb:Get*",
    "sdb:List*",
    "ses:Get*",
    "ses:List*",
    "sns:Get*",
    "sns:List*",
    "sqs:Get*",
    "sqs:List*",
    "ssm:ListCommands",
    "ssm:ListCommandInvocations",
    "storagegateway:Describe*",
    "storagegateway:List*",
    "swf:Count*",
    "swf:Describe*",
    "swf:Get*",
    "swf:List*",
    "tag:get*",
    "trustedadvisor:Describe*",
    "waf:Get*",
    "waf:List*",
    "waf-regional:Get*",
    "waf-regional:List*"
  ],
  "Resource": [
    "*"
  ]
},
{
  "Sid": "AWSManagedServicesFullAccess",
  "Effect": "Allow",
  "Action": [
    "amscm:*",
    "amsskms:*"
  ],
  "Resource": [
    "*"
  ]
}

```

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Default AMS single-account
landing zone (SALZ) IAM User Role

```
]
}
```

Managed Services IAM ReadOnly Policy

```
{
  "Statement": [
    {
      "Action": [
        "iam:GenerateCredentialReport",
        "iam:GetAccountAuthorizationDetails",
        "iam:GetAccountPasswordPolicy",
        "iam:GetAccountSummary",
        "iam:GetCredentialReport",
        "iam:GetGroup",
        "iam:GetGroupPolicy",
        "iam:GetInstanceProfile",
        "iam:GetPolicy",
        "iam:GetPolicyVersion",
        "iam:GetRole",
        "iam:GetRolePolicy",
        "iam:GetUser",
        "iam:GetUserPolicy",
        "iam:ListAccountAliases",
        "iam:ListAttachedRolePolicies",
        "iam:ListEntitiesForPolicy",
        "iam:ListGroupPolicies",
        "iam:ListGroups",
        "iam:ListGroupsForUser",
        "iam:ListInstanceProfiles",
        "iam:ListInstanceProfilesForRole",
        "iam:ListMFADevices",
        "iam:ListPolicies",
        "iam:ListPolicyVersions",
        "iam:ListRolePolicies",
        "iam:ListRoles",
        "iam:ListSAMLProviders",
        "iam:ListUsers",
        "iam:ListVirtualMFADevices"
      ],
      "Effect": "Allow",
      "Resource": [
        "*"
      ],
      "Sid": "IAMReadOnlyAccess"
    },
    {
      "Action": [
        "iam:*"
      ],
      "Effect": "Deny",
      "Resource": [
        "arn:aws:iam:*:group/mc-*",
        "arn:aws:iam:*:group/mc_*",
        "arn:aws:iam:*:policy/mc-*",
        "arn:aws:iam:*:policy/mc_*",
        "arn:aws:iam:*:role/mc-*",
        "arn:aws:iam:*:role/mc_*",
        "arn:aws:iam:*:role/Sentinel-*",
        "arn:aws:iam:*:role/Sentinel_*",
        "arn:aws:iam:*:user/mc-*",
        "arn:aws:iam:*:user/mc_*"
      ],
      "Sid": "DenyAccessToIamRolesStartingWithMC"
    }
  ]
}
```

],

Managed Services User Policy

```
"Version": "2012-10-17"
}
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowCustomerToListTheLogBucketLogs",
      "Effect": "Allow",
      "Action": [
        "s3:ListBucket"
      ],
      "Resource": [
        "arn:aws:s3:::mc-a*-logs-*"
      ],
      "Condition": {
        "StringLike": {
          "s3:prefix": [
            "aws/*",
            "app/*",
            "encrypted",
            "encrypted/",
            "encrypted/app/*"
          ]
        }
      }
    },
    {
      "Sid": "BasicAccessRequiredByS3Console",
      "Effect": "Allow",
      "Action": [
        "s3:ListAllMyBuckets",
        "s3:GetBucketLocation"
      ],
      "Resource": [
        "arn:aws:s3:::*"
      ]
    },
    {
      "Sid": "AllowCustomerToGetLogs",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject*"
      ],
      "Resource": [
        "arn:aws:s3:::mc-a*-logs-*/aws/*",
        "arn:aws:s3:::mc-a*-logs-*/encrypted/app/*"
      ]
    },
    {
      "Sid": "AllowAccessToOtherObjects",
      "Effect": "Allow",
      "Action": [
        "s3:DeleteObject*",
        "s3:Get*",
        "s3:List*",
        "s3:PutObject*"
      ],
      "Resource": [
        "*"
      ]
    }
  ],
}
```

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Default AMS single-account
landing zone (SALZ) IAM User Role

```
{
  "Sid": "AllowCustomerToListTheLogBucketRoot",
  "Effect": "Allow",
  "Action": [
    "s3:ListBucket"
  ],
  "Resource": [
    "arn:aws:s3:::mc-a*-logs-*"
  ],
  "Condition": {
    "StringEquals": {
      "s3:prefix": [
        "",
        "/"
      ]
    }
  }
},
{
  "Sid": "AllowCustomerCWLConsole",
  "Effect": "Allow",
  "Action": [
    "logs:DescribeLogStreams",
    "logs:DescribeLogGroups"
  ],
  "Resource": [
    "arn:aws:logs:*:*:log-group:*"
  ]
},
{
  "Sid": "AllowCustomerCWLAccessLogs",
  "Effect": "Allow",
  "Action": [
    "logs:FilterLogEvents",
    "logs:GetLogEvents"
  ],
  "Resource": [
    "arn:aws:logs:*:*:log-group:/aws/*",
    "arn:aws:logs:*:*:log-group:/infra/*",
    "arn:aws:logs:*:*:log-group:/app/*",
    "arn:aws:logs:*:*:log-group:RDSOSMetrics:*:*"
  ]
},
{
  "Sid": "AWSManagedServicesFullAccess",
  "Effect": "Allow",
  "Action": [
    "amscm:*",
    "amsskms:*"
  ],
  "Resource": [
    "*"
  ]
},
{
  "Sid": "ModifyAWSBillingPortal",
  "Effect": "Allow",
  "Action": [
    "aws-portal:Modify*"
  ],
  "Resource": [
    "*"
  ]
},
{
  "Sid": "DenyDeleteCWL",
```


AMS Advanced Onboarding Guide AMS
 Advanced Account Onboarding Information
 Default AMS single-account
 landing zone (SALZ) IAM User Role

```

    "Effect": "Deny",
    "Action": [
      "logs:DeleteLogGroup",
      "logs:DeleteLogStream"
    ],
    "Resource": [
      "arn:aws:logs:*:*:log-group:*"
    ]
  },
  {
    "Sid": "DenyMCCWL",
    "Effect": "Deny",
    "Action": [
      "logs:CreateLogGroup",
      "logs:CreateLogStream",
      "logs:DescribeLogStreams",
      "logs:FilterLogEvents",
      "logs:GetLogEvents",
      "logs:PutLogEvents"
    ],
    "Resource": [
      "arn:aws:logs:*:*:log-group:/mc/*"
    ]
  },
  {
    "Sid": "DenyS3MCNamespace",
    "Effect": "Deny",
    "Action": [
      "s3:*"
    ],
    "Resource": [
      "arn:aws:s3:::mc-a*-logs-*/encrypted/mc/*",
      "arn:aws:s3:::mc-a*-logs-*/mc/*",
      "arn:aws:s3:::mc-a*-logs-*audit/*",
      "arn:aws:s3:::mc-a*-internal-*/*",
      "arn:aws:s3:::mc-a*-internal-*"
    ]
  },
  {
    "Sid": "ExplicitDenyS3CfnBucket",
    "Effect": "Deny",
    "Action": [
      "s3:*"
    ],
    "Resource": [
      "arn:aws:s3:::cf-templates-*"
    ]
  },
  {
    "Sid": "DenyListBucketS3LogsMC",
    "Action": [
      "s3:ListBucket"
    ],
    "Effect": "Deny",
    "Resource": [
      "arn:aws:s3:::mc-a*-logs-*"
    ],
    "Condition": {
      "StringLike": {
        "s3:prefix": [
          "auditlog/*",
          "encrypted/mc/*",
          "mc/*"
        ]
      }
    }
  }
}

```

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Default AMS single-account
landing zone (SALZ) IAM User Role

```
    },  
    {  
      "Sid": "DenyS3LogsDelete",  
      "Effect": "Deny",  
      "Action": [  
        "s3:Delete*",  
        "s3:Put*"  
      ],  
      "Resource": [  
        "arn:aws:s3::mc-a*-logs-*/*"  
      ]  
    },  
    {  
      "Sid": "DenyAccessToKmsKeysStartingWithMC",  
      "Effect": "Deny",  
      "Action": [  
        "kms:*"  
      ],  
      "Resource": [  
        "arn:aws:kms::*key/mc-*",  
        "arn:aws:kms::*alias/mc-*"  
      ]  
    },  
    {  
      "Sid": "DenyListingOfStacksStartingWithMC",  
      "Effect": "Deny",  
      "Action": [  
        "cloudformation:*"  
      ],  
      "Resource": [  
        "arn:aws:cloudformation::*:stack/mc-*"  
      ]  
    },  
    {  
      "Sid": "AllowCreateCWMetricsAndManageDashboards",  
      "Effect": "Allow",  
      "Action": [  
        "cloudwatch:PutMetricData"  
      ],  
      "Resource": [  
        "*"   
      ]  
    },  
    {  
      "Sid": "AllowCreateandDeleteCWDashboards",  
      "Effect": "Allow",  
      "Action": [  
        "cloudwatch:DeleteDashboards",  
        "cloudwatch:PutDashboard"  
      ],  
      "Resource": [  
        "*"   
      ]  
    }  
  ]  
}
```

Customer Secrets Manager Shared Policy

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Sid": "AllowSecretsManagerListSecrets",  
      "Effect": "Allow",  
      "Action": [  
        "secretsmanager:ListSecrets"  
      ],  
      "Resource": [  
        "*"   
      ]  
    }  
  ]  
}
```

```
    "Action": "secretsmanager:listSecrets",
    "Resource": "*"
  },
  {
    "Sid": "AllowCustomerAdminAccessToSharedNameSpaces",
    "Effect": "Allow",
    "Action": "secretsmanager:*",
    "Resource": [
      "arn:aws:secretsmanager:*:*:secret:ams-shared/*",
      "arn:aws:secretsmanager:*:*:secret:customer-shared/*"
    ]
  },
  {
    "Sid": "DenyCustomerGetSecretCustomerNamespace",
    "Effect": "Deny",
    "Action": "secretsmanager:GetSecretValue",
    "Resource": "arn:aws:secretsmanager:*:*:secret:customer-shared/*"
  },
  {
    "Sid": "AllowCustomerReadOnlyAccessToAMSNameSpace",
    "Effect": "Deny",
    "NotAction": [
      "secretsmanager:Describe*",
      "secretsmanager:Get*",
      "secretsmanager:List*"
    ],
    "Resource": "arn:aws:secretsmanager:*:*:secret:ams-shared/*"
  }
]
```

Customer Marketplace Subscribe Policy

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowMarketPlaceSubscriptions",
      "Effect": "Allow",
      "Action": [
        "aws-marketplace:ViewSubscriptions",
        "aws-marketplace:Subscribe"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

Default Access Firewall Rules

These are the default firewall rules required to access your instances.

Note

For information on firewall rules and ports required for establishing an AD one-way trust, see the AMS Security Guide by going to the AWS Artifact console -> Reports tab and search for AWS Managed Services.

Linux Stack Instance Ports

These rules are required for your authentication into AMS Linux stacks.

Linux Instance Ports Rules FROM: Linux Stack Instance TO: CORP Domain Controller

Port	Protocol	Service	Direction
389	TCP	LDAP	Ingress
389	UDP	LDAP	Ingress
88	TCP	Kerberos	Ingress
88	UDP	Kerberos	Ingress

Windows Stack Instance Ports

These rules are required for your authentication into AMS Windows stacks.

FROM: Windows Stack Instance TO: CORP Domain Controller

Port	Protocol	Service	Direction
88	TCP UDP	Kerberos	Ingress and Egress
135	TCP UDP	DCE/RPC Locator service	Ingress and Egress
389	TCP UDP	LDAP	Ingress and Egress
3268	TCP UDP	msft-gc, Microsoft Global Catalog (LDAP service which contains data from Active Directory forests)	Ingress and Egress
445	TCP	Microsoft-DS Active Directory, Windows shares	Ingress and Egress
49152 - 65535	TCP	Dynamic or private ports that cannot be registered with IANA. This range is used for private, or customized services or temporary purposes and for automatic allocation of ephemeral ports.	Ingress and Egress

AMS service management

Topics

- [Account governance \(p. 40\)](#)
- [Service commencement \(p. 40\)](#)
- [AMS customer relationship management \(CRM\) \(p. 41\)](#)
- [Updates to shared services: Multi-Account Landing Zone \(p. 44\)](#)
- [AMS planned event management \(p. 44\)](#)
- [Getting help \(p. 45\)](#)
- [Service hours \(p. 45\)](#)
- [How do I get offboard assistance from AMS Single-Account Landing Zone accounts? \(p. 46\)](#)
- [How do I offboard from AMS Multi-Account Landing Zone accounts? \(p. 46\)](#)

How the AMS service works for you.

Account governance

This section covers AMS account governance.

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

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

Service commencement

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

Service Commencement

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

Service commencement

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

AMS customer relationship management (CRM)

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

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

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

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

CRM Process

The CRM process includes these activities:

- Identifying and understanding your business processes and needs. Your agreement with AMS identifies your stakeholders.
- Defining the services to be provided to meet your needs and requirements.
- Meeting with you in the service review meetings to discuss any changes in the AMS service scope, SLA, contract, and your business needs. Interim meetings may be held with you to discuss performance, achievements, issues, and action plans.

- Monitoring your satisfaction by using our customer satisfaction survey and feedback given at meetings.
- Reporting performance on monthly internally-measured performance reports.
- Reviewing the service with you to determine opportunities for improvements. This includes frequent communication with you regarding the level and quality of the AMS service provided.

CRM meetings

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

Meeting	Purpose	Mode	Participants
Weekly status review (optional)	<p>Outstanding issues or incidents, patching, security events, problem records</p> <p>12-week operational trend (+/- 6)</p> <p>Application operator concerns</p> <p>Weekend schedule</p>	On-site customer location/Telecom/Chime	<p>AMS: CSDM and cloud architect (CA)</p> <p>Customer assigned team members (ex: Cloud/Infrastructure, Application Support, Architecture teams, etc.)</p>
Monthly business review	<p>Review service level performance (reports, analysis, and trends)</p> <p>Financial analysis</p> <p>Product roadmap</p> <p>CSAT</p>	On-site customer location/Telecom/Chime	<p>AMS: CSDM, cloud architect (CA), AMS account team, AMS technical product manager (TPM) (optional), AMS OPS manager (optional)</p> <p>You: Application Operator representative</p>
Quarterly business review	<p>Scorecard and service level agreement (SLA) performance and trends (6 months)</p> <p>Upcoming 3/6/9/12 months plans/migrations</p> <p>Risk and risk mitigations</p> <p>Key improvement initiatives</p> <p>Product roadmap items</p> <p>Future direction aligned opportunities</p>	On-site customer location	<p>AMS: CSDM, cloud architect, AMS account team, AMS service director, AMS operation manager</p> <p>You: Application operator representative, service representative, service director</p>

Meeting	Purpose	Mode	Participants
	Financials		
	Cost savings initiatives		
	Business optimization		

CRM Meeting Arrangements

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

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

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

Note

Your CSDM works with you to establish your account governance.

CRM monthly reports

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

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

Note

In addition to the described information, your CSDM also informs you of any material change in scope or terms, including use of subcontractors by AMS for operational activities.

AMS generates reports about patching and backup that your CSDM includes in your monthly report. As part of the report generating system, AMS adds some infrastructure to your account that is not accessible to you:

- An S3 Bucket, with the raw data reported
- An Athena instance, with query definitions to query the data

- A Glue Crawler to read the raw data from the S3 bucket

Updates to shared services: Multi-Account Landing Zone

AMS uses the core OU to provide shared services such as access, networking, EPS, log storage, alert aggregation in your Multi-Account Landing Zone. AMS is responsible for addressing vulnerabilities, patching, and deployments of these shared services. AMS regularly updates the resources used for providing these shared services so that users have access to latest features, and security updates. The updates typically happen on a monthly basis. Resources that are part of these updates are:

- Accounts that are part of the core OU.

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

- New AMS AMIs incorporating the latest updates.

Note

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

AMS planned event management

AWS Managed Services (AMS) planned event management (PEM) is an AMS service offering. PEM is used to engage, plan, and run customer events and projects using AMS Change Management Services and dedicated AMS resources. Change management delivers an individual request for change (RFC). The PEM delivers a set of related RFCs that align with the scope and timeline of the PEM event or project.

AMS PEM criteria

A planned event is defined as a scope-bound and time-bound project. For example, migrations, game days, disaster recovery tests, projects, or events that require dedicated on-site or off-site AMS resources such as Operation Engineers or Cloud Architects.

The AMS PEM process

The PEM process consists of the following phases:

- **Initiation** —: You engage with the Cloud Service Deliver Managers (CSDM), Technical Delivery Managers (TDM), and Cloud Architects (CA) to provide project information and the technical details to AMS. AMS works with you to ensure that the PEM plan information is correct and complete. For PEM acceptance, AMS Operations requires a lead time of 2 weeks to allow the AMS Operations appropriate time to ensure planning, technical review and resource assignment. Additional time may be required for delivery of pre-PEM tasks.

- **Technical Review** —: AMS Cloud Architects review the technical aspects of the PEM plan. They work with AMS Security and Operations to ensure compliance, provide execution optimization and automation, and define pre-PEM execution tasks and deliverables.
- **Planning** —: AMS ensures that the necessary AMS resources are assigned.
- **Readiness and Execution** —: AMS ensures pre-execution tasks are completed, and facilitates internal and customer communications. AMS also ensures execution of the PEM plan and provides execution status and progress reporting.

Getting help

You can reach out to AMS to identify the root cause of your failure. AMS business hours are 24 hours a day, 7 days a week, 365 days a year.

AMS provides several avenues for you to ask for help or make service requests.

- To ask for information or advice, or for access to an AMS-managed IT service, or to request an additional service from AMS, use the AMS console and submit a service request. For details, see [Creating a Service Request](#). For general information about AMS service requests, see [Service Request Management](#).
- To report an AWS or AMS service performance issue that impacts your managed environment, use the AMS console and submit an incident report. For details, see [Reporting an incident](#). For general information about AMS incident management, see [Incident response](#).
- For specific questions about how you or your resources or applications are working with AMS, or to escalate an incident, email one or more of the following:
 1. First, if you are unsatisfied with the service request or incident report response, email your CSDM: ams-csdm@amazon.com
 2. Next, if escalation is required, you can email the AMS Operations Manager (your CSDM will most likely do this): ams-opsmanager@amazon.com
 3. Further escalation would be to the AMS Director: ams-director@amazon.com
 4. Finally, you are always able to reach the AMS VP: ams-vp@amazon.com

Customer contacts with AMS that require escalation will follow the escalation path described next.

Service hours

Feature	AMS Accelerate		AMS Advanced	
Service request	Monday to Friday: 08:00–18:00, local business hours	24/7	Monday to Friday: 08:00–18:00, local business hours	24/7
Incident management (P1)	24/7			
Incident management (P2-P3)	Monday to Friday: 08:00–18:00, local business hours	24/7	Monday to Friday: 08:00–18:00, local business hours	24/7
Backup and recovery	24/7			

Feature	AMS Accelerate	AMS Advanced	
Patch management	24/7		
Monitoring and alerting	24/7		
Automated request for change (RFC)	Not Applicable	24/7	
Non-automated request for change (RFC)	Not Applicable	Monday to Friday: 08:00–18:00, local business hours	24/7
Cloud service delivery manager (CSDM)	Monday to Friday: 08:00–17:00, local business hours		

How do I get offboard assistance from AMS Single-Account Landing Zone accounts?

AMS offers off-boarding assistance within 30 days prior to termination of AMS.

You must request off-boarding assistance at least 7 days before such assistance can be provided. Off-boarding assistance can be offered in two forms:

- Control hand-over: AMS will transfer account control back to the Customer along with access credentials for all AMS Managed Applications, or
- Resource termination and data transfer: AMS backs-up all the data, deletes all the data in customer's Managed Environment, de-provisions any active resources in the account, and hands over the data backup to the Customer. At customer's request AMS can transfer customer data in the existing format using Snowball or any other media with which AWS can interface. In addition to data backups, the following customer data can be provided as part of off-boarding assistance:
 - Data stored in storage services including logs
 - Customer-specific Change type schemas
 - CloudFormation templates for Change type schemas.

If off-boarding activities are not completed upon the termination of AMS, we hand over the controls of the account(s) to enable you to complete any pending activity.

How do I offboard from AMS Multi-Account Landing Zone accounts?

Currently AMS supports 3 types of offboarding for multi-account landing zone accounts:

- Multi-Account Landing Zone environmental offboarding
- Application account offboarding
- Application account VPC offboarding.

How do I offboard a Multi-Account Landing Zone environment?

Scenario:

You want to leave AMS, and close (terminate) your AMS AWS account completely (including the primary account that you provided).

Process:

1. You communicate with your CSDMs or CAs and request offboarding via email or a service request.
2. AMS works on offboarding your accounts.
3. Your CSDM or CA sends you an outbound email containing further instructions; see [How do I close my AWS account?](#)

Prerequisites: Verify that you can access the account email used for account closures.

Offboarding Conclusions:

- All components are disassociated from the AMS services, but are not yet deleted in all accounts. Account closure will eventually delete all resources.
- Billing is *not stopped* until you request the account closures.
- After the account is closed, you can still sign in and file a support case or contact AWS Support for 90 days.
- After 90 days, any content remaining in the account is permanently deleted, and AWS services that aren't already terminated, are terminated.

How do I offboard a Multi-Account Landing Zone application account?

Some offboarding scenarios.

Scenario:

You want to offboard one, or more than one, application accounts from your multi-account landing zone environment, close (terminate) those accounts completely.

Process:

1. You communicate with your CSDMs or CAs to request offboarding via email or a service request.
2. AMS works on offboarding your accounts from AMS.
3. Your CSDM or CA sends you an outbound email containing further instructions; see [How do I close my AWS account?](#)

Prerequisites: Verify that you can access the account email used for account closures.

Offboarding Conclusions:

- All components are disassociated from AMS services, and most of the resources are deleted in the requested accounts. Your resources remain in the account until you request account closure.
- Core accounts and other application accounts function normally after the offboarding request.

- Billing is *not stopped* until you request account closures.
- After the account is closed, you can still sign in and file a support case or contact AWS Support for 90 days.
- After 90 days, any content remaining in the account is permanently deleted, and AWS services that aren't already terminated, are terminated.

How do I offboard a Multi-Account Landing Zone application account VPC?

Scenario:

You want to offboard one of your VPCs from an AMS-managed application account.

Process:

1. You communicate with your CSDMs or CAs via email or a service request, to request VPC offboarding.
2. AMS works on offboarding the VPC from AMS.
3. Your CSDM and CA notify you of the completion of the offboarding.

Prerequisites:

- Verify there are *no running instance stacks* associated with this offboarding VPC.

If there are running instance stacks associated with the offboarding VPC, you are responsible for deleting those instance stacks prior to requesting a VPC offboarding.

- Your application account should always contain at least one VPC. If you request a VPC offboarding and that VPC is the only VPC in the account, follow up with an application VPC create RFC.
- Verification email to confirm the request. Your request should contain the VPC name of the VPC that you want to delete, and the account ID that the VPC is associated with.

Offboarding Conclusions:

Application accounts function normally after VPC offboarding request.

AMS Multi-account landing zone onboarding

Multi-Account Landing Zone network architecture

Topics

- [About Multi-Account Landing Zone network architecture \(p. 49\)](#)
- [Multi-Account Landing Zone accounts \(p. 51\)](#)

About Multi-Account Landing Zone network architecture

Topics

- [Service region \(p. 50\)](#)
- [Organizational units \(p. 50\)](#)
- [Service control policies and AWS Organization \(p. 51\)](#)

Before starting the onboarding process, it is important to understand the baseline architecture, or landing zone, that AMS creates on your behalf, its components, and functions.

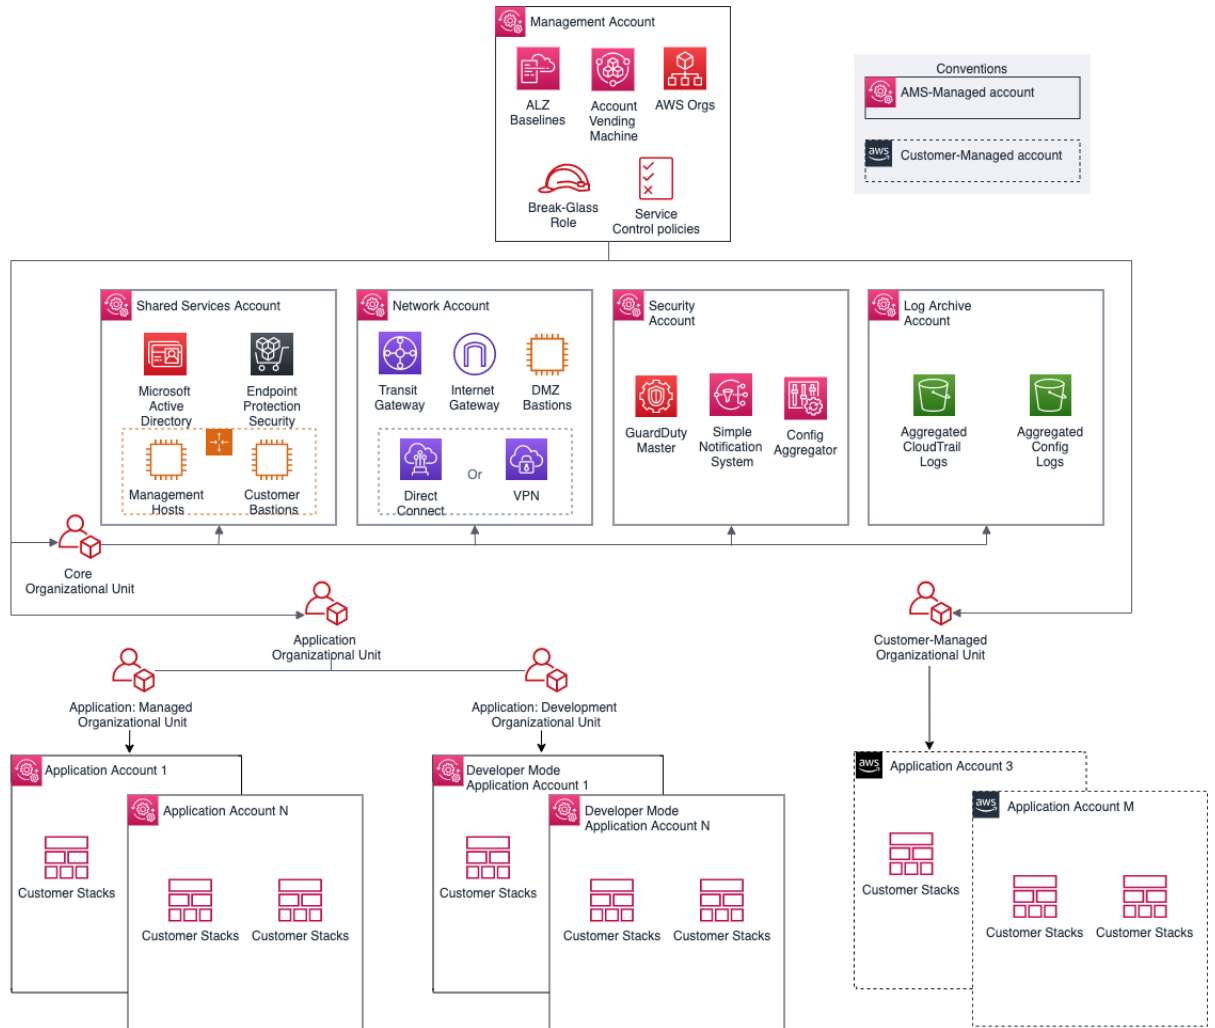
AMS multi-account landing zone is a multi-account architecture, pre-configured with the infrastructure to facilitate authentication, security, networking, and logging.

Note

For estimates of costs, see [AMS Multi-Account Landing Zone environment basic components](#).

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

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
About Multi-Account Landing Zone network architecture



Service region

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

Organizational units

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

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

AMS-managed multi-account landing zone also enables you to create custom OUs for grouping and organizing AWS Accounts and to associate custom SCPs with them; for examples on doing this, see [Management account: Creating a custom OU](#) and [Management account: Creating a custom SCP](#), respectively. AMS provides three existing OUs under which new OUs and accounts can be requested: application > managed, application > development, and customer managed.

- Application > managed OU:

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

- Application > development OU:

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

- Customer Managed OU:

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

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

Service control policies and AWS Organization

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

You can also create custom SCPs and attach them to specific OUs. They can be requested from your Management account using change type ct-33ste5yc7hprs. AMS then reviews the custom SCPs requested before applying them to the target OUs. For examples, see [Management account: Creating a custom OU](#) and [Management account: Creating a custom SCP](#).

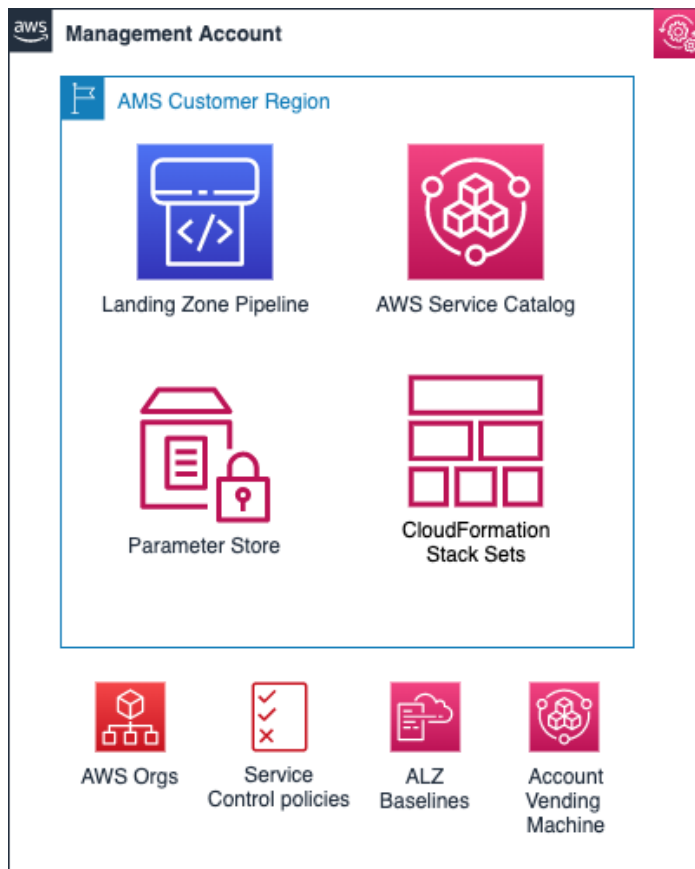
Multi-Account Landing Zone accounts

Topics

- [Management account \(p. 51\)](#)
- [Networking account \(p. 52\)](#)
- [Shared Services account \(p. 65\)](#)
- [Log Archive account \(p. 67\)](#)
- [Security account \(p. 67\)](#)
- [Application accounts: AMS-managed, Developer mode, Customer Managed \(p. 68\)](#)
- [Tools account, Migrating Workloads: CloudEndure Landing Zone \(MALZ\) \(p. 71\)](#)

Management account

The management account is your initial AWS account when you begin onboarding with AMS. It utilizes AWS Organizations as a management account, which gives the account the ability to create and financially manage member accounts. It contains the AWS landing zone (ALZ) framework, account configuration stack sets, AWS Organization service control policies (SCPs), etc. The following diagram provides a high-level overview of the resources contained in the management account.



Resources in the management account

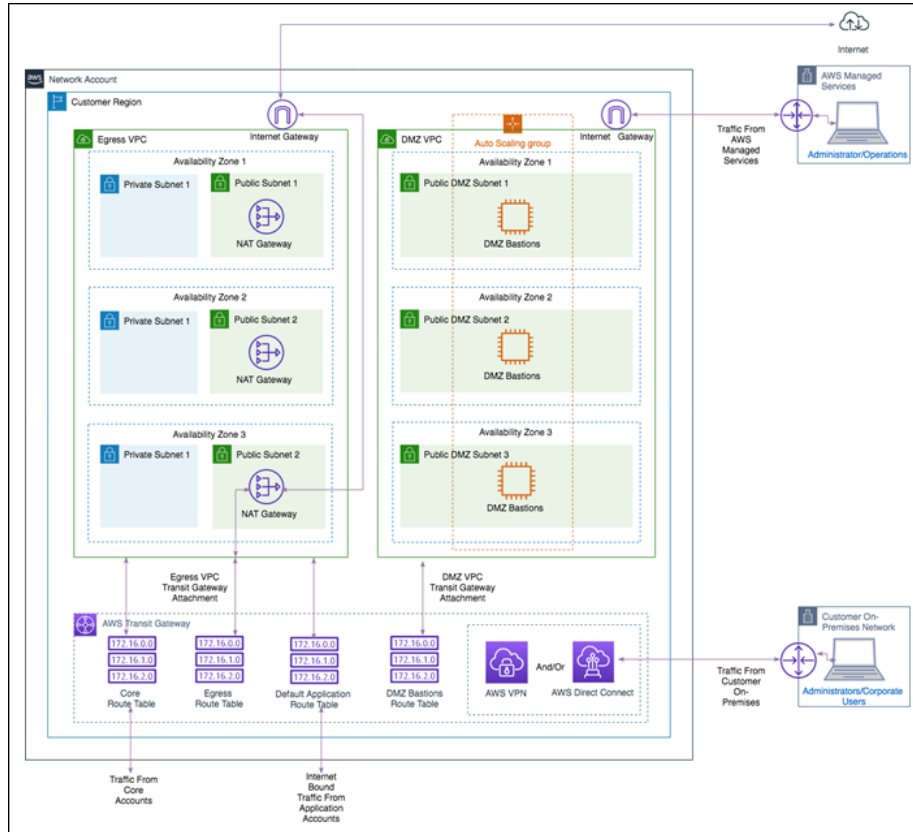
Other than the above standard services, no additional AWS resources are created in the management account during onboarding. The following inputs are required during onboarding to AMS:

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

Networking account

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

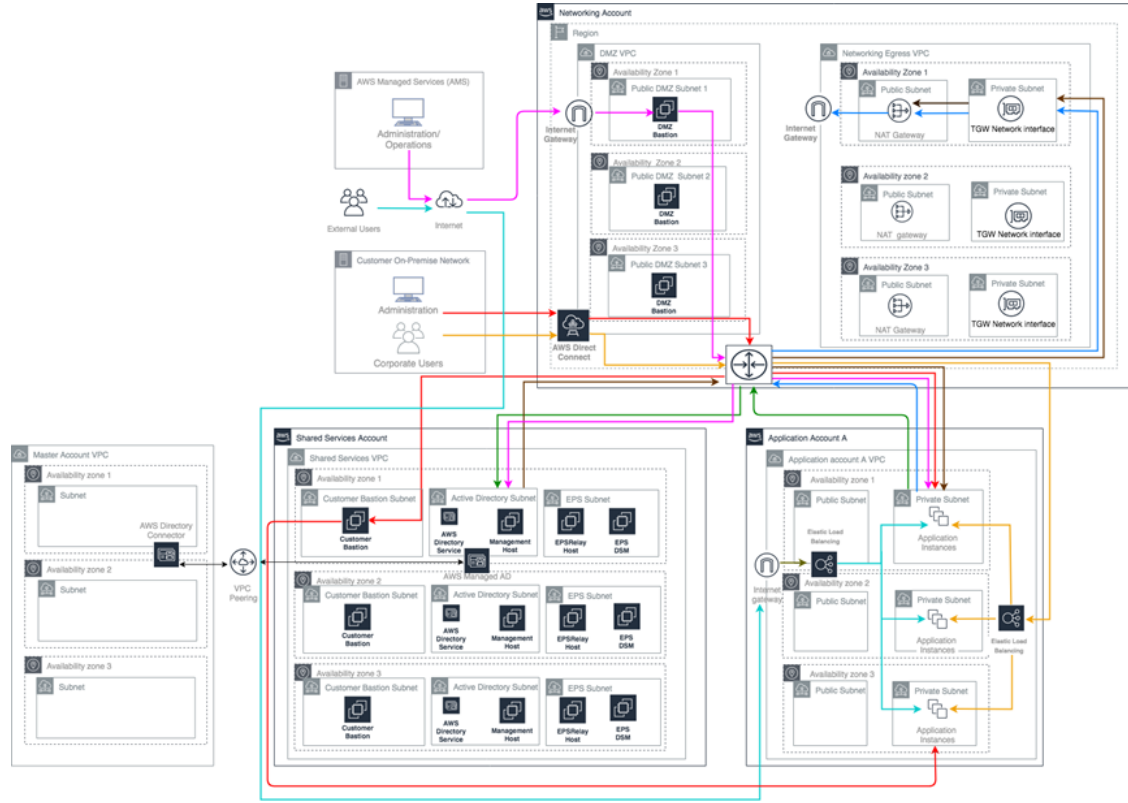
AMS Advanced Onboarding Guide AMS
 Advanced Account Onboarding Information
 Multi-Account Landing Zone accounts



Networking account architecture

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

AMS Advanced Onboarding Guide AMS Advanced Account Onboarding Information Multi-Account Landing Zone accounts



	Egress Internet Traffic from Application Account VPC and Shared Services VPC through Egress VPC (Networking Account) and Transit Gateway
	Egress Traffic from an Application Account VPC to Shared Services VPC via Transit gateway (Networking Account)
	Egress Internet Traffic from Shared Services VPC to Application and Networking Account VPCs via Transit Gateway (Networking Account).
	Ingress through internet with managed internet gateway for AMS administrators and operators through DMZ bastions to Application VPCs and Shared Services Account VPCs via Transit Gateway (Networking Account)
	Ingress through DirectConnect (internal customer network administrators) and Customer Bastions to Application Account's VPC instances via Transit Gateway (Networking Account)
	Ingress through DirectConnect (internal customer network users) for Corporate Users to Application Instances in Application Account VPCs via Transit Gateway (Networking Account).
	Ingress through Internet with managed Internet Gateway (external users), through AWS load balancers in Application (Public) Subnet and then to Application Instances in Application Account VPC.

Customer's AMS environment is categorized into multiple accounts, managed under AWS Organization. The environment is split into AMS Core Infrastructure and Application Infrastructure. Core accounts consist of Master Account, Networking Account, Shared Services Account, Logging account and Security account, whereas Application Infrastructure consists of applications accounts.

Each AMS accounts can have multiple VPCs in one region with resource subnets located in up to three availability zones. Each availability zone can have private and public subnets (depends on configuration selected). Your ("customer") corporate network is connected through a DirectConnect (VPN) tunnel, and AMS operations connects to your Application infrastructure over the internet.

Master account is the central hub to manage and configure member accounts. Landingzone framework and SSO enablement is configured in this account.

The Networking Account serves as the central hub for network routing between AMS Core Accounts, your OnPremise Network, and egress traffic out to the Internet via Transit Gateway. Transit Gateway is an AWS service that enables customers to connect their VPCs and their on-premises networks to a single gateway. Networking account consists of DMZ VPC which contain DMZ bastions hosts that serve as SSG jump boxes for AMS operations team and Egress VPC through which all network traffic is routed.

Shared Services account has a VPC with following subnets: ActiveDirectory Subnet, Customer Bastion Subnet and EPS subnet. AD Subnet consists of AMS Directory service, AD domain controller, and management hosts that automate provisioning and common tasks. And EPS subnets consists of Antivirus (Trend Micro) management servers that include EPS DSM and EPS relay (for scalability). Lastly, customer bastion subnets consists of internal (customer) bastion hosts.

Your "Customer" accounts contain your workloads, EC2 instances, RDS etc

External users connect to your applications for the internet via an AWS load balancer that is located in your application account.

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

Note
For information about default service limits and constraints for all active services, see the [AWS Service Limits](#) documentation.

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

Private network connectivity to AMS Multi-account landing zone environment

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

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

Centralized edge connectivity using transit gateway

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

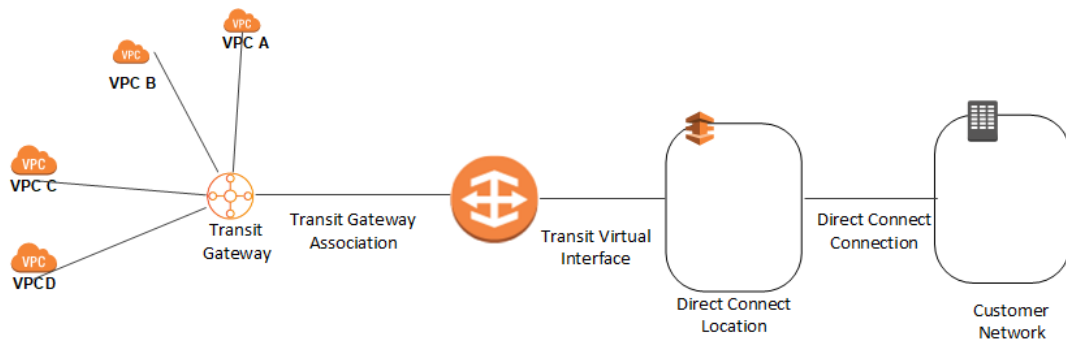
AWS Direct Connect (DX) gateway is used to connect your DX connection over a transit virtual interface to the VPCs or VPNs that are attached to your transit gateway. You associate a Direct Connect gateway with the transit gateway. Then, create a transit virtual interface for your AWS Direct Connect connection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

This configuration offers the following benefits. You can:

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

Note

For information about using a DX with AWS services, see the Resiliency Toolkit section [Classic](#). For more information about Transit Gateway associations, see [Transit Gateway associations](#).



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

Connecting DX or VPN to account VPCs

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

Resources in the networking account

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

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

AWS Network Manager

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

Use one of the following roles to access this resource:

- `AWSManagedServicesCaseRole`
- `AWSManagedServicesReadOnlyRole`
- `AWSManagedServicesChangeManagementRole`

Egress VPC

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

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

Note

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

The AMS multi-account landing zone team recommends the range of 24 (with more IP address) to provide some buffer in case other resources/appliances, are deployed in the future.

Managed Palo Alto egress firewall

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

Traffic control

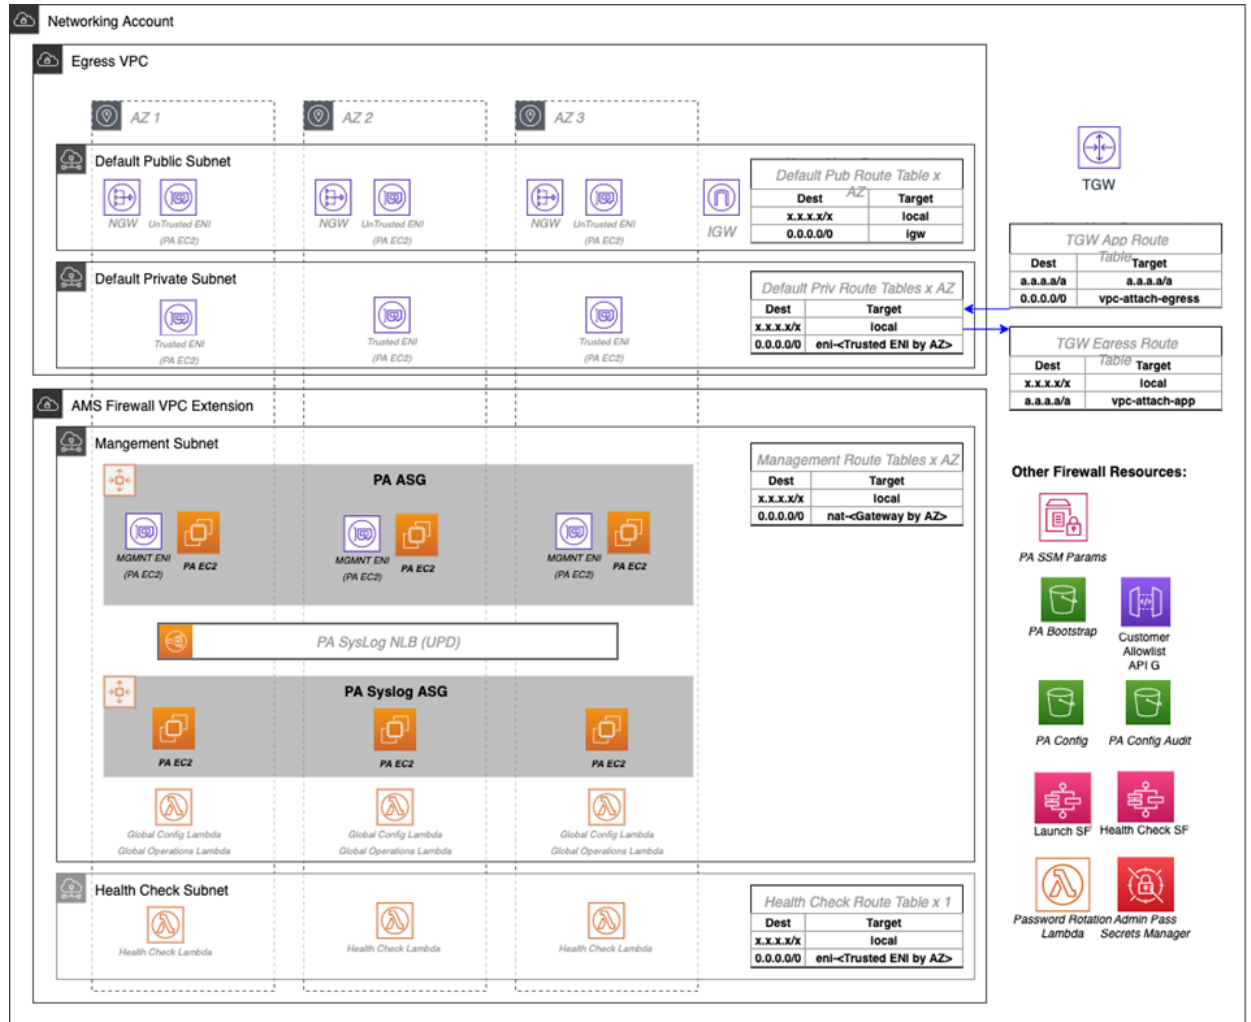
The managed outbound firewall solution manages a domain allow-list composed of AMS-required domains for services such as backup and patch, as well as your defined domains. When outbound

AMS Advanced Onboarding Guide AMS
 Advanced Account Onboarding Information
 Multi-Account Landing Zone accounts

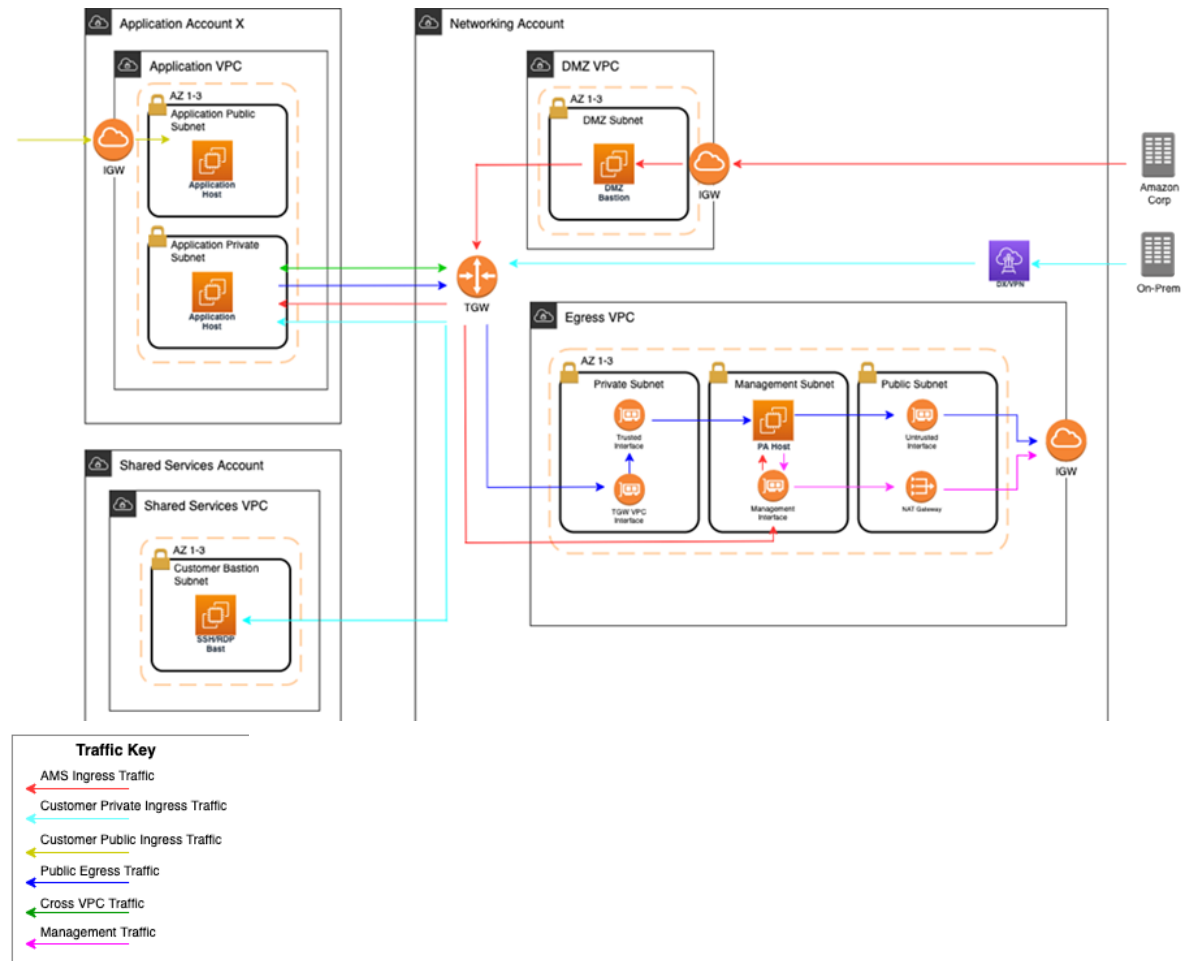
internet traffic is routed to the firewall, a session is opened, traffic is evaluated, and if it matches an allowed domain, the traffic is forwarded to the destination.

Architecture

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



Network flow



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

1. Egress traffic destined for the internet is sent to the Transit Gateway (TGW) through VPC route table
2. TGW routes traffic to the egress VPC via the TGW route table
3. VPC routes traffic to the internet via the private subnet route tables
 - a. In the default Multi-Account Landing Zone environment, internet traffic is sent directly to a network address translation (NAT) gateway. The managed firewall solution reconfigures the private subnet route tables to point the default route (0.0.0.0/0) to a firewall interface instead.

The firewalls themselves contain three interfaces:

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

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

Backup and Restore

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

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

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

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

Updates

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

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

Operator access

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

Default logs

By default, the logs generated by the firewall reside in local storage for each firewall. Overtime, local logs will be deleted based on storage utilization. The AMS solution provides real-time shipment of logs off of the machines to CloudWatch logs; for more information, see [CloudWatch Logs integration \(p. 63\)](#).

AMS engineers still have the ability to query and export logs directly off the machines if required. In addition, logs can be shipped to a customer-owned Panorama; for more information, see [Panorama integration \(p. 63\)](#).

The Logs collected by the solution are the following:

RFC Status Codes

Log Type	Description
Traffic	<p>Displays an entry for the start and end of each session. Each entry includes the date and time, source and destination zones, addresses and ports, application name, security rule name applied to the flow, rule action (allow, deny, or drop), ingress and egress interface, number of bytes, and session end reason.</p> <p>The Type column indicates whether the entry is for the start or end of the session, or whether the session was denied or dropped. A "drop" indicates that the security rule that blocked the traffic specified "any" application, while a "deny" indicates the rule identified a specific application.</p>

Log Type	Description
	If traffic is dropped before the application is identified, such as when a rule drops all traffic for a specific service, the application is shown as "not-applicable".
Threat	<p>Displays an entry for each security alarm generated by the firewall. Each entry includes the date and time, a threat name or URL, the source and destination zones, addresses, and ports, the application name, and the alarm action (allow or block) and severity.</p> <p>The Type column indicates the type of threat, such as "virus" or "spyware;" the Name column is the threat description or URL; and the Category column is the threat category (such as "keylogger") or URL category.</p>
URL Filtering	Displays logs for URL filters, which control access to websites and whether users can submit credentials to websites.
Configuration	Displays an entry for each configuration change. Each entry includes the date and time, the administrator user name, the IP address from where the change was made, the type of client (web interface or CLI), the type of command run, whether the command succeeded or failed, the configuration path, and the values before and after the change.
System	Displays an entry for each system event. Each entry includes the date and time, the event severity, and an event description.
Alarms	The alarms log records detailed information on alarms that are generated by the system. The information in this log is also reported in Alarms. Refer to "Define Alarm Settings".
Authentication	<p>Displays information about authentication events that occur when end users try to access network resources for which access is controlled by Authentication policy rules. Users can use this information to help troubleshoot access issues and to adjust user Authentication policy as needed. In conjunction with correlation objects, users can also use Authentication logs to identify suspicious activity on the users network, such as brute force attacks.</p> <p>Optionally, users can configure Authentication rules to Log Authentication Timeouts. These timeouts relate to the period of time when a user needs authenticate for a resource only once but can access it repeatedly. Seeing information about the timeouts helps users decide if and how to adjust them.</p>
Unified	Displays the latest Traffic, Threat, URL Filtering, WildFire Submissions, and Data Filtering log entries in a single view. The collective log view enables users to investigate and filter these different types of logs together (instead of searching each log set separately). Or, users can choose which log types to display: click the arrow to the left of the filter field and select traffic, threat, url, data, and/or wildfire to display only the selected log types.

Event management

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

Event Alarms:

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

Management Alarms:

- Health Check Monitor Failure Alarm
 - When health check workflow fails unexpectedly
 - This is for the workflow itself, not if a firewall health check fails
- Password Rotation Failure Alarm
 - When password rotation fails
 - API/Service user password is rotated every 90 days

Metrics

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

Custom Metrics:

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

CloudWatch Logs integration

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

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

Note

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

Panorama integration

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

Licensing

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

- Marketplace Licenses: Accept the terms and conditions of the VM-Series Next-Generation Firewall Bundle 1 from the networking account in MALZ.
- BYOL Licenses: Accept the terms and conditions of the VM-Series Next-Generation Firewall (BYOL) from the networking account in MALZ and share the "BYOL auth code" obtained after purchasing the license to AMS.

Limitations

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

Note

The AMS solution runs in Active-Active mode as each PA instance in its AZ handles egress traffic for their respected AZ. So, with two AZs, each PA instance handles egress traffic up to 5 Gbps and effectively provides overall 10 Gbps throughput across two AZs. The same is true for all limits in each AZ. Should the AMS health check fail, we shift traffic from the AZ with the bad PA to another AZ, and during the instance replacement, capacity is reduced to the remaining AZs limits.

AMS does not currently support other Palo Alto bundles available on AWS Marketplace; for example, you cannot ask for the "VM-Series Next-Generation Firewall Bundle 2". Note that the AMS Managed Firewall solution using Palo Alto currently provides only an egress traffic filtering offering, so using advanced VM-Series bundles would not provide any additional features or benefits.

Onboarding requirements

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

Pricing

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

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

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

Perimeter (DMZ) VPC

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

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

Note

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

AWS Transit Gateway

AWS Transit Gateway (TGW) is a service that enables you to connect your Amazon Virtual Private Clouds (VPCs) and your on-premises networks to a single gateway. Transit gateway is the networking backbone

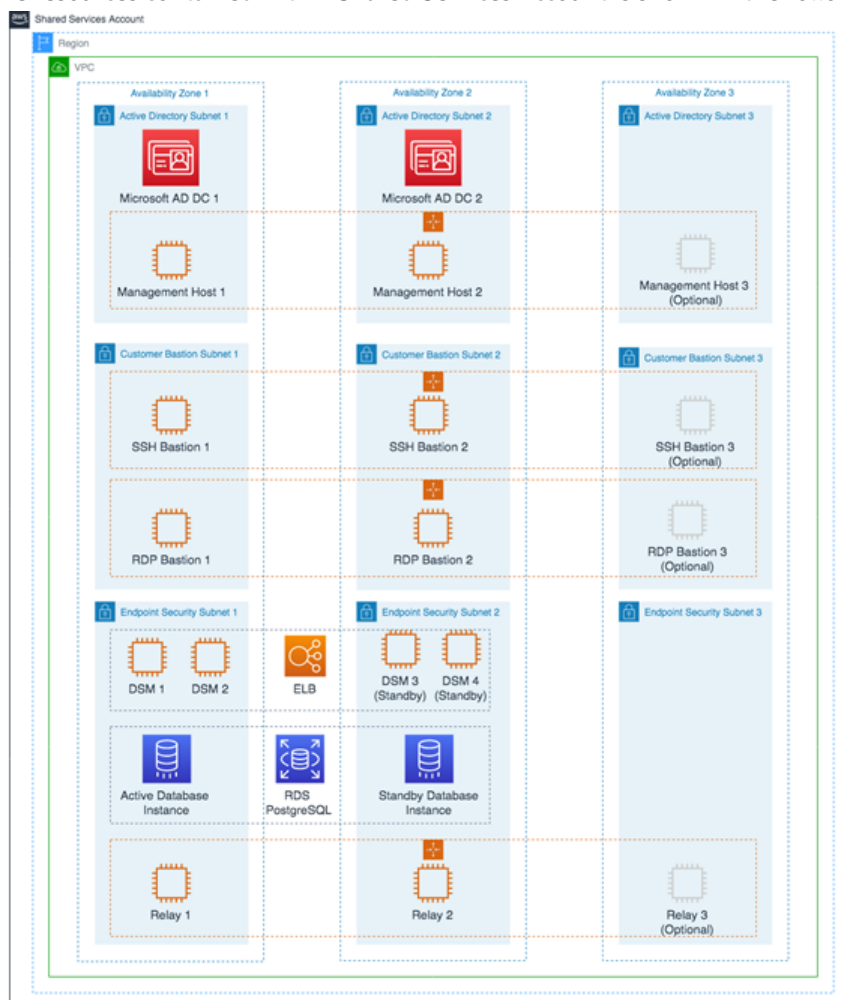
that handles the routing between AMS account networks and external networks. For information about Transit Gateway, see [AWS Transit Gateway](#).

Provide the following input to create this resource:

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

Shared Services account

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



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

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

Note

The AMS team recommends the range of /23.

- *Active Directory Details:* Microsoft Active Directory (AD) is utilized for user/resource management, authentication/authorization, and DNS, across all of your AMS multi-account landing zone accounts. AMS AD is also configured with a one-way trust to your Active Directory for trust-based authentication. The following input is required to create the AD:
 - Domain Fully Qualified Domain Name (FQDN): The fully qualified domain name for the AWS Managed Microsoft AD directory. The domain should not be an existing domain or child domain of an existing domain in your network.
 - Domain NetBIOS Name: If you don't specify a NetBIOS name, AMS defaults the name to the first part of your directory DNS. For example, corp for the directory DNS corp.example.com.
- *Trend Micro – endpoint protection security (EPS):* Trend Micro endpoint protection (EPS) is the primary component within AMS for operating system security. The system is comprised of Deep Security Manager (DSM), EC2 instances, relay EC2 instances, and an agent present within all data plane and customer EC2 instances.

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

The following default inputs are required to create EPS (if you want to change from the defaults):

- Relay Instance Type: Default Value - m5.large
- DSM Instance Type: Default Value - m5.xlarge
- DB Instance Size: Default Value - 200 GB
- RDS Instance Type: Default Value - db.m5.large
- *Customer bastions:* You are provided with SSH or RDP bastions (or both) in the Shared Services Account, to access other hosts in your AMS environment. In order to access the AMS network as a user (SSH/RDP), you must use "customer" Bastions as the entry point. The network path originates from the on-premise network, goes through DX/VPN to the transit gateway (TGW), and then is routed to the Shared Services VPC. Once you are able to access the bastion, you can jump to other hosts in the AMS environment, provided that the access request has been granted.
 - The following inputs are required for SSH bastions.
 - SSH Bastion Desired Instance Capacity: Default Value - 2.
 - SSH Bastion Maximum Instances: Default Value - 4.
 - SSH Bastion Minimum Instances: Default Value -2.
 - SSH Bastion Instance Type: Default Value - m5.large (can be changed to save costs; for example a t3.medium).
 - SSH Bastion Ingress CIDRs: IP address ranges from which users in your network access SSH Bastions.
 - The following inputs are required for Windows RDP bastions.
 - RDP Bastion Instance Type: Default Value - t3.medium.
 - RDP Bastion Desired Minimum Sessions: Default Value - 2.
 - RDP Maximum Sessions: Default Value -10.
 - RDP Bastion Configuration Type: You can choose one of the below configuration
 - SecureStandard = A user receives one bastion and only one user can connect to the bastion.
 - SecureHA = A user receives two bastions in two different AZ's to connect to and only one user can connect to the bastion.

- SharedStandard = A user receives one bastion to connect to and two users can connect to the same bastion at once.
- SharedHA = A user receives two bastions in two different AZ's to connect to and two users can connect to the same bastion at once.
- Customer RDP Ingress CIDRs: IP address ranges from which users in your network will access RDP Bastions.

Log Archive account

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



Security account

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



Application accounts: AMS-managed, Developer mode, Customer Managed

Application accounts are AWS accounts within the AMS-managed landing zone architecture that you use to host your workloads. AMS offers three types of Application Accounts with different operational models, responsibilities and features. Each account type is grouped under an organizational unit (OU) from which you can request additional nested OUs. The three types of application accounts are described in this section.

Application Accounts are provisioned through RFC from the [Management account](#).

AMS-managed application accounts

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

There are some AWS services that you can use in your AMS-managed account without AMS management. The list of services and how to add them into your AMS account are described in the [Self-Provisioned Services](#) section.

Developer mode application accounts

Accounts with Developer mode are a type of AMS-managed account that provide customers with elevated permissions in AMS "Plus" accounts to provision and update AWS resources outside of the AMS change management process. When using an account that has Developer mode enabled, continuity management, patch management, and change management are provided for resources provisioned through the AMS change management process, or by using an AMS Amazon Machine Image (AMI). However, you are responsible for monitoring infrastructure resources that are provisioned outside of the AMS change management process. With elevated permissions, you have an increased responsibility to ensure adherence to internal controls.

For more information, see [Developer mode](#).

Customer Managed application accounts

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

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

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

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

To create your Customer Managed account, see [Management account, customer-managed application account: Creating](#).

Accessing your Customer Managed account

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

To access the CMA:

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

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

Connecting your CMA with Transit Gateway

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

Note

You can only have a VPC attached to the TGW if the CMA is in the same AWS Region. For more information see [Transit gateways](#).

To add your CMA to Transit Gateway, request a new route (use the Management | Other | Other | Create `ct-1e1xtak34nx76`) change type and include this information:

- CMA account number
- Transit Gateway ID
- TGW attachment ID from CMA account (for example, `tgw-attach-04eb40d1e14ec7272`)
- CMA route table ID (for example, `rtb-0ff4d759eb28b2a05`)

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

1. By default this VPC will not be able to communicate with any of the other VPCs in your MALZ network.
2. Decide with your solutions architect what VPCs you want this Customer Managed VPC to communicate with. Submit a Management | Other | Other | Update RFC against the Networking

account to create the TGW routes you need. Include the CMA Account Number, Transit Gateway ID, TWG Attachment ID from the CMA account (e.g. tgw-attach-12345678901234567), and the CMA Route Table ID (e.g. rtb-12345678901234567).

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

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

Associating the TGW attachment to a route table:

Decide which TGW route table you want to associate the VPC with. We recommend creating a new application route table for Customer Managed VPCs. Submit a Management | Other | Other | Update RFC on the Networking account to associate the VPC or TGW attachment to the route table you select.

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

1. By default, this VPC will not be able to communicate with any of the other VPCs in your Multi-Account Landing Zone network.
2. Decide with your solutions architect what VPCs you want this customer-managed VPC to communicate with. Submit a Management | Other | Other | Update RFC against the networking account to create the TGW routes you need.

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

Decide with your solutions architect what traffic you want to send to the AMS Multi-Account Landing Zone transit gateway. Submit a Management | Other | Other | Update RFC against the networking account to create the TGW routes you need.

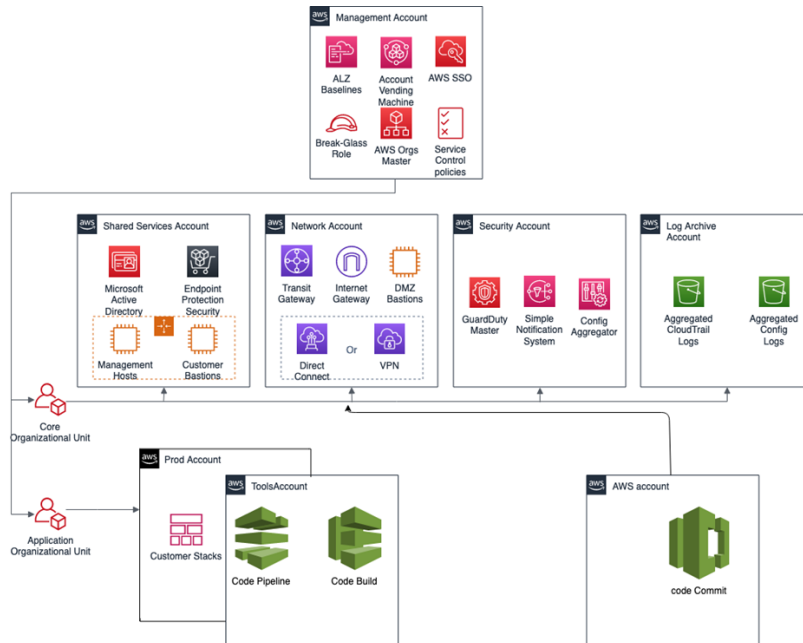
Tools account, Migrating Workloads: CloudEndure Landing Zone (MALZ)

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

A tools account provides the following:

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

High level architecture diagram:



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

Note

We recommend having two availability zones (AZs), since this is a migration hub.

By default, AMS creates the following two security groups (SGs) in every account. Confirm that the two SGs are present, and, if not, open a new Management | Other | Other | Create CT (ct-1e1xtak34nx76) to request them:

- SentinelDefaultSecurityGroupPrivateOnlyEgressAll
- InitialGarden-SentinelDefaultSecurityGroupPrivateOnly

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

Prerequisites:

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

AWS Application Migration Service (AWS MGN)

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

For the most up-to-date information on AWS Region support, see [the AWS Regional Services List](#).

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

Requesting AWS MGN Initialization

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

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

  https://console.aws.amazon.com/mgn/home?region=MALZ_PRIMARY_REGION#/welcome using all
  default values
  to 'Create template' and complete the initialization process.
```

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

Application Migration Service > Set up Application Migration Service

Set up Application Migration Service

In order to use Application Migration Service in this region, the service must first be initialized by creating a Replication Settings template. After the template is created, Application Migration Service will automatically create the IAM roles required for the service to operate. The service can only be initialized by the Admin user of your AWS account.

Create Replication Settings template [Info](#)

Every source server added to this console has Replication Settings that control how data is sent from the source server to AWS. These settings are created automatically based on this template, and can be modified at any time for any source server or group of source servers. The template itself can also be modified at any time (changes made will only affect newly added servers).

Replication Servers [Info](#)

Staging area subnet [Info](#)

Replication Server Instance type [Info](#)

EBS volume type (for replicating disks over 500GiB) [Info](#)

EBS encryption [Info](#)

Security groups [Info](#)

Always use Application Migration Service security group

Additional security groups

Data routing and throttling [Info](#)

Use private IP for data replication (VPN, DirectConnect, VPC peering)

Create public IP

Throttle network bandwidth (per server - in Mbps)

Replication resources tags [Info](#)

Add new tag

You can add up to 50 more tags.

Cancel **Create template**

Enable access to the new Tools account

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

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

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

The tools account also has these additional IAM roles available:

- AMS Migration role
- CloudEndure user role

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

Navigate to the AMS console and file the following RFCs:

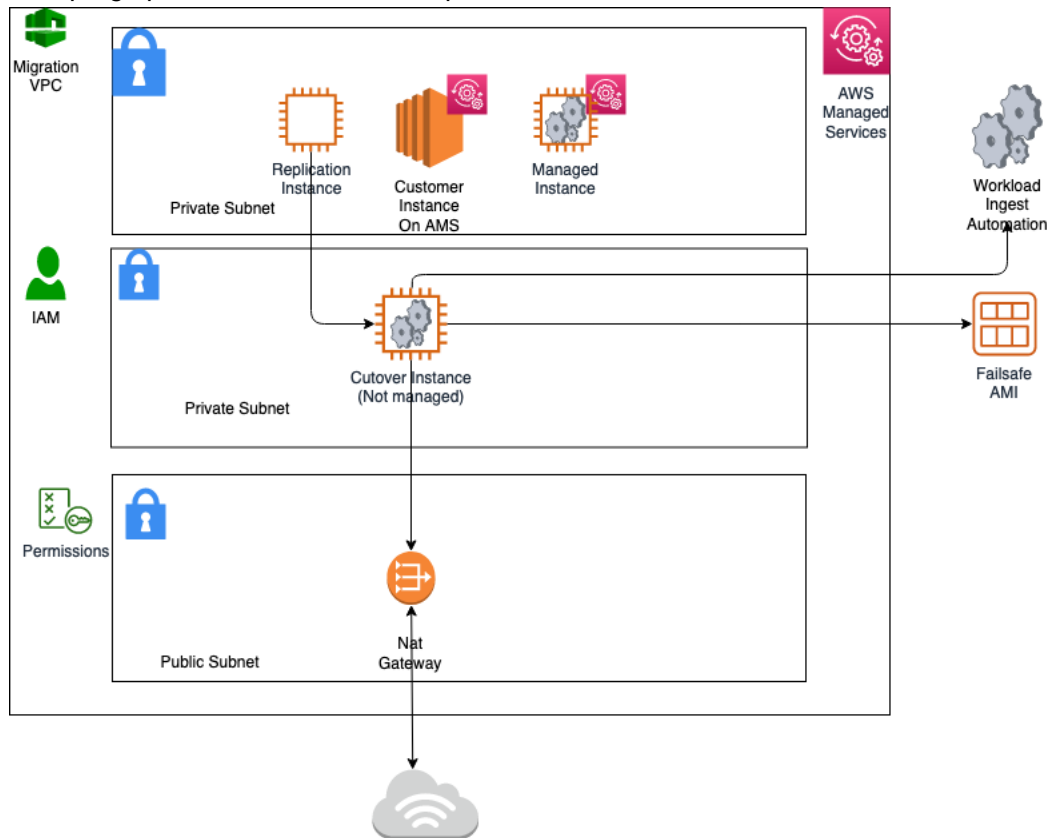
- a. Create KMS key. Use either [Create KMS Key \(auto\)](#) or [Create KMS Key \(review required\)](#).

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

b. Share the KMS key.

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

Example graphic of a final account setup:



Example policy

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

Testing connectivity and end-to-end setup

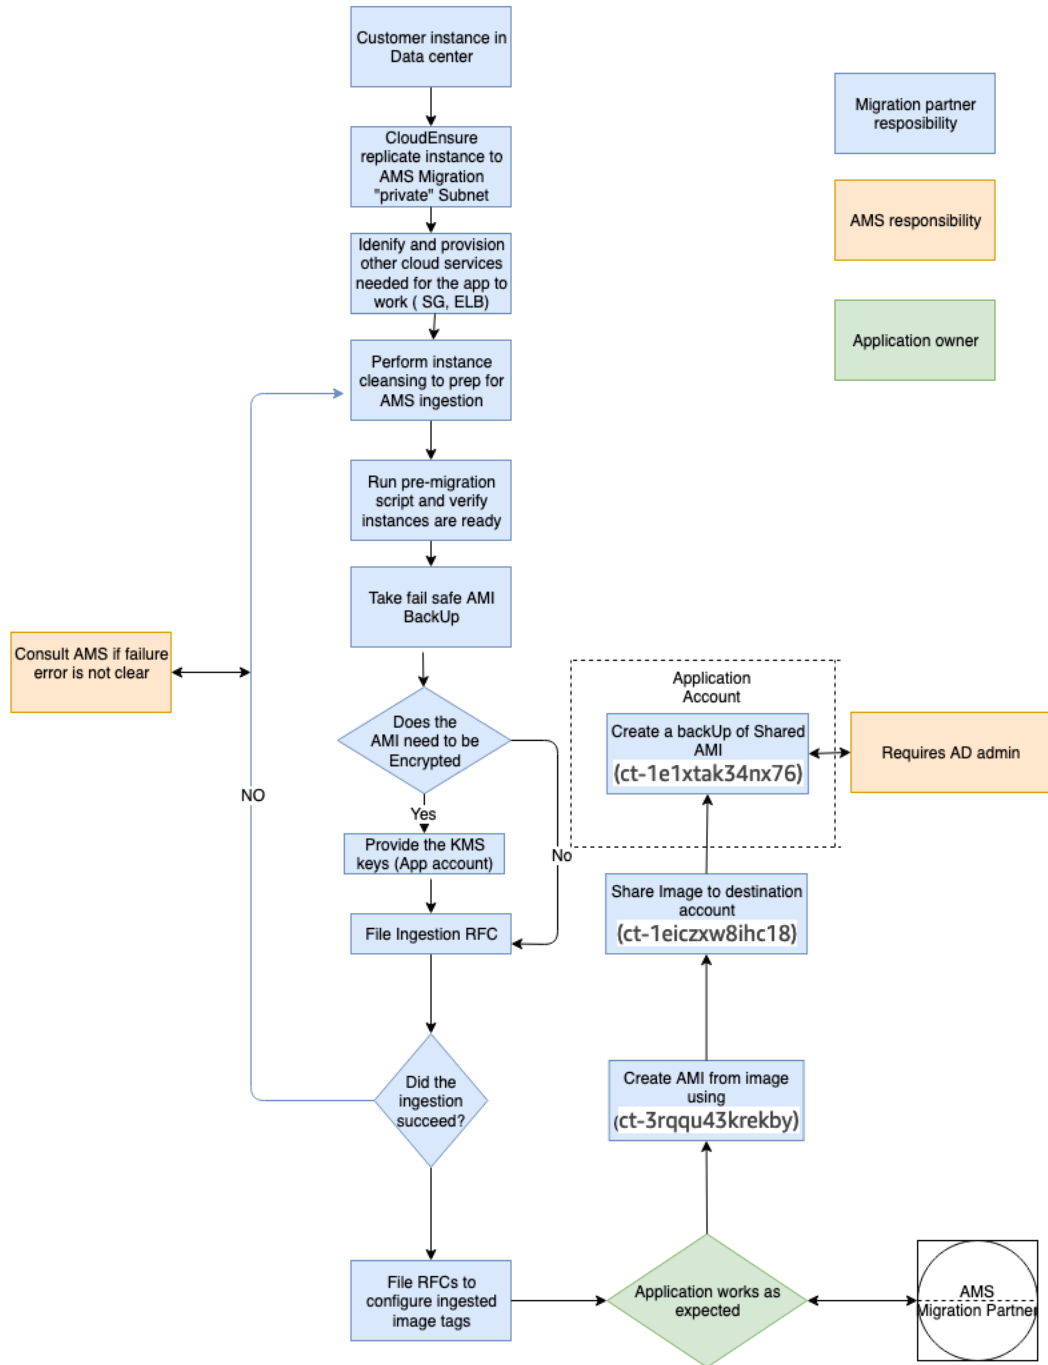
To test the tools account, follow these steps.

1. Start with configuring CloudEndure and installing the CloudEndure agent on a server that will replicate to AMS.
2. Create a project in CloudEndure.
3. Enter the AWS credentials shared when you performed the prerequisites, though secrets manager.
4. In **Replication settings**:
 - a. Select both AMS "Sentinel" security groups (Private Only and EgressAll) for the **Choose the Security Groups to apply to the Replication Servers** option.

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

The AMS ingestion process:

AMS Ingestion Process



Tools account hygiene

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

- Post instance WIGs ingestion:

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

Migration at scale - Migration Factory

See [Introducing AWS CloudEndure Migration Factory Solution](#).

Core account onboarding

Topics

- [Create an AWS core account \(p. 77\)](#)
- [Create an IAM Role for AMS to access your account \(p. 78\)](#)
- [Secure the New Account with Multi-Factor Authentication \(MFA\) for the Root User \(p. 81\)](#)
- [Subscribe to AWS Marketplace for EPS \(p. 81\)](#)
- [Set up networking \(p. 83\)](#)
- [Set up access management \(p. 85\)](#)

For onboarding questions, contact your Cloud Architect.

Create an AWS core account

AMS multi-account landing zone requires the provisioning of a new Amazon Web Services (AWS) account to act as the management account in the AMS multi-account landing zone environment. To create an AWS account, follow these step-by-step instructions: [How do I create and activate a new Amazon Web Services account?](#)

The simple steps are: Go to [Create Account](#), and click **Sign Up Now** and, on the page that opens, click **Create a new AWS account**. Follow the on-screen instructions, which include receiving a phone call and entering a PIN using your phone keypad. You'll also need to enter a credit card. AMS uses this account as the management account, or payer account, for your new multi-account landing zone.

Note

Once you are onboarded, talk to your CSDM about moving billing off of your credit card and onto an invoice system. The following information will be required:

- Billing Company Name
- Billing Contact Name
- Billing Contact Phone Number
- Billing Contact Email
- Billing Address

Your CSDM will help you with this update. Once completed, and to change the payment method, see [Managing your AWS payment methods](#).

Note

Do not link your new account to an existing management account, or payer account. Ensure that your account is not part of an existing AWS Organization; for information, see [What Is AWS Organizations?](#)

Important

It is very important that you ensure that an **email address** (a distribution list, not an individual's email address) and **phone number** are associated with the account so that you receive responses to potential security incidents. The phone number and email address for the account cannot be changed without resetting the account password, which is a significant undertaking for an AMS root account. To ensure that these values are stable, *it is critical to select contact information not associated with individuals*, which can change. Choose an email alias that can point to a group. Follow this same best practice in selecting a phone number: choose a number that can point to a group or to a number owned by the company and not an individual.

For details on the questions you will be asked to onboard your Core account to AMS multi-account landing zone, see [Appendix: multi-account landing zone Onboarding Consideration List \(p. 93\)](#).

Create an IAM Role for AMS to access your account

Now that you've successfully created your new AWS account, the next step in the process is to allow AMS access to the new account to create and configure your AMS environment, and for ongoing change and provisioning requests to be fulfilled. For details, see [Delegate Access Across AWS Accounts Using IAM Roles](#).

AWS Identity and Access Management (IAM) is a web service that helps you securely control access to AWS resources for your users. You use IAM to control who can use your AWS resources (authentication) and what resources they can use and in what ways (authorization).

Activate IAM Access to the AWS Console

1. Sign in to the AWS Management Console with your root account credentials (the email and password that you used to create your AWS account). Do not sign in with other IAM credentials. The AWS Management Console home page opens.
2. In the top navigation bar, open the drop-down menu for your account name, and then choose **My Account**. The Billing home page opens.
3. Scroll down to **IAM User Access -> Billing Information**, and choose **Edit**. An **Activate IAM access** area opens.
4. Select the check box and then choose **Update**. You can now use IAM policies to control which pages a user can access.

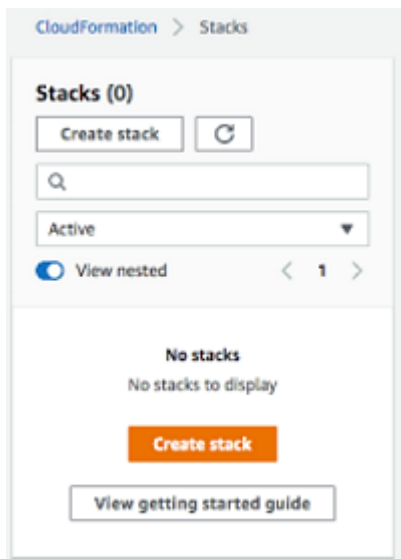
Create an IAM Role for AMS to Use

1. Your AMS Cloud Architect provides you with a JSON or YAML file that contains the IAM role AMS uses for creating infrastructure.

Or you can use this to create the file yourself:

```
{
  "AWSTemplateFormatVersion": "2010-09-09",
  "Description": "AMS Onboarding Role stack (for Prod)",
  "Parameters": {},
  "Conditions": {},
  "Resources": {
    "OnboardingRole": {
      "Type": "AWS::IAM::Role",
      "Properties": {
        "RoleName": "aws_managedservices_onboarding_role",
        "ManagedPolicyArns": ["arn:aws:iam::aws:policy/AdministratorAccess"],
        "AssumeRolePolicyDocument": {
          "Version": "2012-10-17",
          "Statement": [{
            "Action": "sts:AssumeRole",
            "Effect": "Allow",
            "Principal": {
              "AWS": ["328792436863"]
            }
          }]
        }
      }
    }
  }
}
```

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



3. Choose **Create Stack**. You see the following page.

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Create an IAM Role for AMS to access your account

The screenshot shows the 'Create stack' wizard in Step 1, 'Specify template'. The left sidebar lists four steps: Step 1 (Specify template), Step 2 (Specify stack details), Step 3 (Configure stack options), and Step 4 (Review). The main content area is titled 'Create stack' and contains two sections: 'Prerequisite - Prepare template' and 'Specify template'. In the 'Prerequisite' section, three radio buttons are present: 'Template is ready' (selected), 'Use a sample template', and 'Create template in Designer'. The 'Specify template' section explains that a template is a JSON or YAML file and provides options for the 'Template source': 'Amazon S3 URL' and 'Upload a template file' (selected). Below this, there is a file upload area with a 'Choose file' button and the text 'No file chosen'. A note states 'JSON or YAML formatted file'. At the bottom right of the main content area, there is a 'View in Designer' button. At the very bottom of the wizard, there are 'Cancel' and 'Next' buttons.

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

The screenshot shows the 'Specify stack details' wizard in Step 2. The left sidebar lists four steps: Step 1 (Specify template), Step 2 (Specify stack details), Step 3 (Configure stack options), and Step 4 (Review). The main content area is titled 'Specify stack details' and contains two sections: 'Stack name' and 'Parameters'. The 'Stack name' section has a text input field with the placeholder 'Enter a stack name' and a note: 'Stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-)'. The 'Parameters' section has a heading 'Parameters' and a note: 'Parameters are defined in your template and allow you to input custom values when you create or update a stack.' Below this, it says 'No parameters' and 'There are no parameters defined in your template'. At the bottom right of the main content area, there are 'Cancel', 'Previous', and 'Next' buttons.

5. Enter **ams-onboarding-role** into the **Stack name** section and continue scrolling down and selecting next until you reach this page.

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Secure the New Account with Multi-Factor
Authentication (MFA) for the Root User

Rollback on failure
Enabled

Timeout
-

Termination protection
Disabled

► Quick-create link

Capabilities

ⓘ The following resource(s) require capabilities: [AWS::IAM::Role]

This template contains Identity and Access Management (IAM) resources. Check that you want to create each of these resources and that they have the minimum required permissions. In addition, they have custom names. Check that the custom names are unique within your AWS account. [Learn more.](#)

I acknowledge that AWS CloudFormation might create IAM resources with custom names.

Cancel Previous Create change set **Create stack**

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

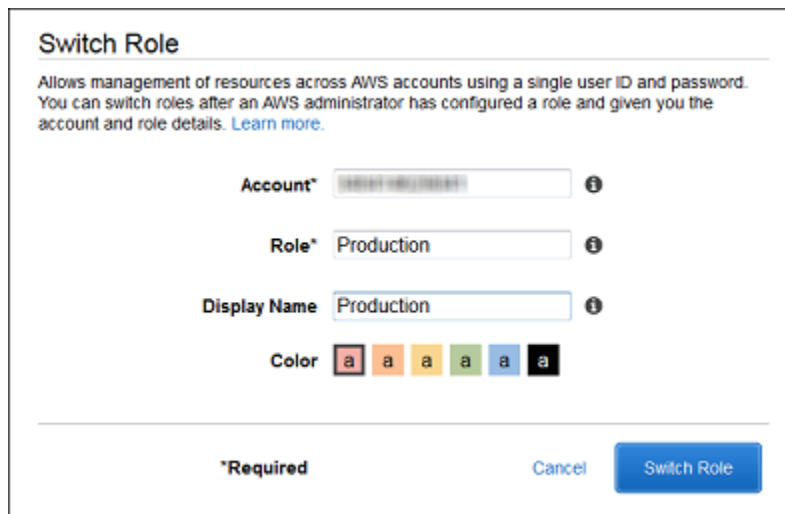
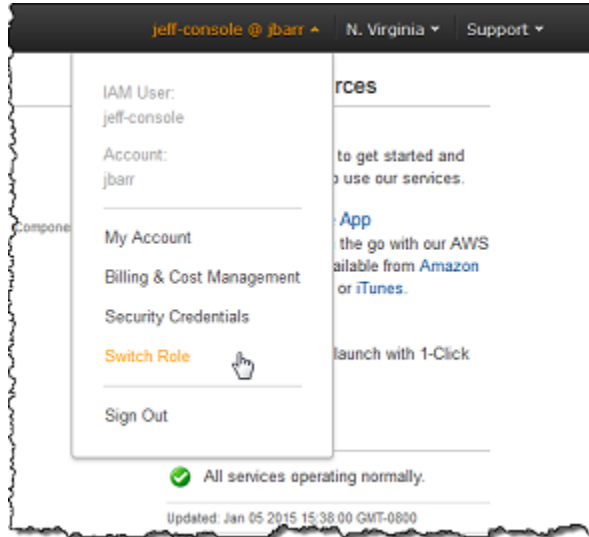
Secure the New Account with Multi-Factor Authentication (MFA) for the Root User

This section has been redacted because it contains sensitive AMS security-related information. This information is available through the AMS console **Documentation**. To access AWS Artifact, you can contact your CSDM for instructions or go to [Getting Started with AWS Artifact](#).

Subscribe to AWS Marketplace for EPS

Trend Micro Endpoint Protection (EPS) is the primary component within AMS for operating system security. In order to set up EPS once AMS landing zone creation is started, you need to log in to the shared services core account and subscribe to the Trend Micro Deep Security AMI on AWS Marketplace. Your CSDM or CA will advise you.

1. Log in to the AWS console using the role or user that you specified in the Onboarding Questionnaire for CustomerEPSSubscriptionIAMRoleOrUser
2. Navigate to the **Switch Role** screen.



- Account: Provided by AMS
- Role: EPSPMarketplaceSubscriptionRole
- Display Name: EPS Subscription Session

To subscribe to Trend Micro Deep Security in the AWS Marketplace, follow these steps after you have switched the role in the console:

1. Navigate to [the AWS Marketplace](#).
2. Under **Find AWS Marketplace products that meet your needs**, select the following options:
 - a. **Vendors:** Trend Micro
 - b. **Pricing Plan:** Bring Your Own License if you have a license or By Hosts Billing
 - c. **Delivery Methods:** Amazon Machine Image
3. Click **Continue to Subscribe** in the right panel.
4. Review the **Terms and Conditions**, and click **Accept Terms** in the upper right corner.
5. Sign out of the account and confirm with your Cloud Architect that the procedure has been completed.

At this point AMS deploys infrastructure into your AMS environment and the environment is ready for you to use once you have connected your network and set up your access.

Set up networking

Networking in the AMS environment is primarily handled in the networking core account.

There are several processes that need to be completed to set up networking for AWS Managed Services (AMS):

- Allocating IP space for your AMS environment
- Establishing private network connectivity to AWS
- Setting up your firewall to allow AMS operations

Allocating IP Space for your AMS Environment

You should have already worked with your Cloud Architect in defining the IP space for your AMS environment while filling out the onboarding questionnaire.

Establishing Private Network Connectivity to AWS

AWS offers private connectivity by using VPN connectivity and dedicated lines via AWS Direct Connect. Private Connectivity can be setup in two ways:

- Centralized Edge connectivity using Transit Gateway
- Connecting DX and/or VPN to account VPCs

Centralized Edge Connectivity using Transit Gateway

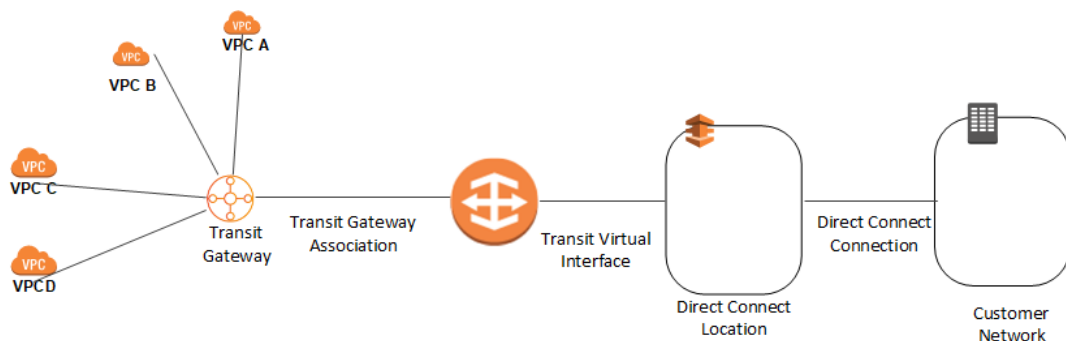
AWS Transit Gateway is a service that enables you to connect your Amazon Virtual Private Clouds (VPCs) and your on-premises networks to a single gateway. Transit Gateway can be used to consolidate your existing edge connectivity and route it through a single ingress/egress point. For more details, see [AWS Transit Gateway](#).

Connecting Direct Connect to Transit Gateway

You can use your existing Direct Connect connection or create a new Direct Connect connection in one of your existing AWS accounts. The Direct Connect connection should be a dedicated or hosted connection running at 1 Gbps or more.

Note

For information about using a DX with AWS services, see [Getting Started at an AWS Direct Connect Location](#).



If you are using your existing Direct Connect dedicated connection, then it should not have any transit virtual interfaces created on the connection. This is because AWS Direct Connect dedicated connection has a limit of one transit virtual interface per connection. If you are using your existing Direct Connect hosted connection, then it should not have any virtual interfaces created on the connection. This is because private, public, or transit virtual interfaces per AWS Direct Connect hosted connection can only be one.

For more details of Direct Connect Limits, please refer to [AWS Direct Connect Limits](#).

Once the Direct Connect connection is available:

1. AMS would create a Direct Connect Gateway in the networking account . You would need to provide an Autonomous System Number (ASN) number for the Direct Connect Gateway and the prefixes that have to be advertised from the Direct Connect Gateway. This ASN will be used as the Amazon side ASN.
2. You will then create a new Transit VIF and set the Virtual interface owner as the networking account.
3. AMS would then login to the networking account and accept the connection proposal.
4. AMS would then associate the transit gateway with the Direct Connect gateway.
5. AMS will the associate the attachment with the on-prem Transit Gateway routing table.

Note

The Autonomous System Number (ASN) provided for the Direct Connect gateway and the Transit gateway must be different.

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

Connecting VPN to Transit Gateway

To attach a VPN connection to your transit gateway, you must specify the customer gateway. For more information about the requirements for a customer gateway, see Requirements for Your Customer Gateway in the Amazon VPC Network Administrator Guide.

You would need to provide the BGP ASN number, static public IP address and routing Option (Static or Dynamic). Once these details are provided, AMS would create the VPN attachment and associate the attachment with the on-prem Transit Gateway routing table.

For more details on Transit Gateway attachments, see [Transit Gateway VPN Attachments](#).

Connecting DX and/or VPN to account VPCs

You can also directly connect your VPCs to Direct Connect or VPN. The traffic would directly flow from the VPCs to Direct Connect or VPN without traversing through the transit gateway.

Note

The shared services VPC and application account VPCs have to be connected to a DX or VPN connection to establish private connectivity.

AWS Direct Connect (DX) Setup

This section describes the basic steps for setting up a AWS Direct Connect (DX) to communicate between your AMS-managed VPC and your internal network.

Note

For information about using a DX with AWS services, see [Getting Started at an AWS Direct Connect Location](#).

To set up a DX connection, you need to complete the following steps:

1. Sign Up for Amazon Web Services
2. Submit AWS Direct Connect Connection Request
3. Complete the Cross Connect
4. (Optional) Configure Redundant Connections with AWS Direct Connect
5. Performed by AMS: Create a Virtual Interface
6. Performed by AMS: Download Router Configuration
7. Verify Your Virtual Interface

VPN Setup

The basic steps that AMS follows for setting up a VPN to communicate between your AMS-managed VPC and your internal network.

Note

To gain overall understanding about using a VPN with AWS services, see [What is AWS Site-to-Site VPN](#) and [Your Customer Gateway](#) (your VPN appliance).

We follow the AWS VPN User Guide [Getting Started](#) and [Testing the Site-to-Site VPN Connection](#) sections to complete the following steps:

1. In your AWS VPC, Create a Customer Gateway.
2. In your AWS VPC, Create a Virtual Private Gateway.
3. In your AWS VPC, Enable Route Propagation in Your Route Table.
4. In your AWS VPC, Update Your Security Group to Enable Inbound SSH, RDP, and ICMP Access.
5. In your internal Network, Create a VPN Connection and Configure the Customer Gateway.
6. Test VPN connectivity between the VPC and your internal network.

Set up access management

Using a network managed by AWS Managed Services (AMS) means giving AMS access to manage your cloud infrastructure. You'll need to configure a means of securely connecting between your private network and AMS. This starts with some decisions:

- *AMS API/CLI and Console access:* You will want to install the AMS CLI (instructions are provided in this document). You use the AMS change management API to make change requests to AMS and the AMS SKMS API to learn about your AMS-managed resources. Using Active Directory Federation Services (AD FS), you can access the AMS Console.
- *User access:* Connectivity needs to be established between AD on the AMS side (via Directory Services) and the directory you use to manage users.
- *Instance access:* Instance-level access is accomplished via a one-way trust configuration. Directory Services trusts credentials in your CORP AD, allowing stacks within the AMS side to allow login with CORP credentials.

Note

Your Active Directory (AD) that AMS sets up the trust to, must be the directory that has the accounts of users authorized by you to gain access to your AWS resources.

Establish an Active Directory Trust

To set up a trust, AMS requires your domain controller **Local Policies -> Security Options -> Network Access: Named Pipes that can be accessed anonymously**, have the **Netlogon** and **lsarpc** pipes listed.

These pipes are listed by default, but are sometimes removed for security concerns. Once the trust is established, they can be removed from the list again.

Configure the Conditional Forwarder

1. In the AD **DNS Manager** -> **Create a New Conditional Forwarder**, under **DNS Domain**: Use the domain name AMS supplied to you; for example, *A523434123.amazonaws.com* (change this to the domain name selected in the onboarding questionnaire).
2. Under **IP addresses of the master servers**: Add the AMS-supplied IP addresses. Make sure there isn't a connection problem by validating both addresses.
3. Select **Store this conditional forwarder in Active Directory and replicate as follows: All DNS servers in this domain** and press **OK**.

Configure the AD trust

Follow this Microsoft AD article [Create a one-way, incoming, forest trust for one side of the trust](#), using the settings and choices described in this section.

1. Open the **Start** -> **Administrative Tools** -> **Active Directory Domains and Trusts** dialog. Right-click the domain node for the domain that you want to establish a trust with, and then click **Properties** -> **Trusts** -> **New Trust** to open the New Trust Wizard. Enter the domain name provided to you by AMS for the **Trust Name** and press **Next**.
2. Under **Trust Type**, select appropriate trust level (e.g. Forest Trust). Press **Next**.
3. Under **Direction of Trust**, select **One-way: incoming**. Press **Next**.
4. Under **Sides of Trust**, select **This domain only**. Press **Next**.
5. Under **Trust Password**, type a password of your choosing. Press **Next**.
6. For **Trust Selections Completed** and **Trust Creation Complete**, just press **Next**.
7. Under **Confirm Incoming Trust**, select **No**, do not confirm the incoming trust. Press **Next**.
8. Under **Completed the New Trust Wizard**, select **Finish**, and then **OK** to close.
9. Provide the trust password (contact us via your CSDM's phone number for security reasons). AMS will complete the trust configuration.

Active Directory sites and services

To reduce login latency, add the VPC CIDR range to your Active Directory sites and services (**Start** -> **Administrative Tools** -> **Active Directory Sites and Services**). Add the VPC CIDR range to an Active Directory Site that contains Domain Controllers that are closest to AWS.

Provide the AD site name of the site that you dedicated for AMS to your CSDM. AMS will rename the default site on the AMS side of AD to match the provided name.

Active Directory name suffix routing

After the one-way forest trust has been established, complete the following steps to validate suffix routing:

1. Under **Start** > **All Programs** > **Administrative Tools**, click **Active Directory Domains and Trusts**.
The Active Directory Domains and Trusts console opens.
2. Right-click your corporate domain and click **Properties**.
The Properties dialog for that domain opens.
3. Click the **Trusts** tab.

The Trusts page opens.

4. Click the Amazon domain name and click **Properties**.

The Properties page for the Amazon domain trust opens.

5. Click **Name Suffix Routing** and click **Refresh**.

Make sure there are no conflicts to ensure that the Service Principal Names (SPNs) can resolve over the trust.

Federate your Active Directory with the AMS IAM roles

The purpose of federating your directory with the AMS IAM roles is to enable corporate users to use their corporate credentials to interact with the AWS Console and the AWS APIs, and, therefore, the AMS console and APIs.

Federation process example

This example uses Active Directory Federation Services (AD FS); however, any technology that supports AWS IAM Federation is supported. For more information on AWS-supported IAM federation, see [IAM Partners](#) and [Identity Providers and Federation](#). Your CSDM will help you through this process, which involves a joint effort with your AD team and AMS.

For detailed information on integrating SAML for API access, refer to this AWS blog, [How to Implement Federated API and CLI Access Using SAML 2.0 and AD FS](#).

For an example that installs the AMS CLI and SAML, see [Appendix: AD FS claim rule and SAML settings](#).

Configuring Federation to the AMS console

The IAM roles and SAML identity provider (Trusted Entity) detailed in the following table have been provisioned as part of the AMS infrastructure. These roles allow you to audit and view the AMS core accounts.

Role	Permissions
AWSManagedServicesReadOnlyRole	Allows you to view the AMS infrastructure in the core accounts.
AWSManagedServicesCaseRole	Allows you to view the resources in your new application account and file AMS incidents and service requests.
AWSManagedServicesChangeManagementRole	Allows you to view the AMS infrastructure in the core accounts, file AWS Support tickets, and request some RFCs.

For more information, see [Core OU account roles](#).

For the full list of the roles available under different accounts see [IAM User Role \(p. 20\)](#).

Verify console access

Once you are set up with ADFS, and have the AMS URL to use for authentication, follow these steps.

With an Active Directory Federated Service (ADFS) configuration, you can follow these steps:

1. Open a browser window and go to the sign in page provided to you for your account. The ADFS **IdpInitiatedSignOn** page for your account opens.
2. Select the radio button next to **Sign in to one of the following sites**. The **Sign in** site picklist becomes active.
3. Choose the **signin.aws.amazon.com** site and click **Sign in**. Options for entering your credentials open.
4. Enter your CORP credentials and click **Sign in**. The AWS Management Console opens.
5. Paste into the location bar the URL of the AMS console and press **Enter**. The AMS console opens.

Verify API access

AMS uses the AWS API, with some AMS-specific operations that you can read about in the [AMS API Reference](#).

AWS provides several SDKs that you can access at [Tools for Amazon Web Services](#). If you don't want to use an SDK, you can make direct API calls. For information on authentication, see [Signing AWS API Requests](#). If you are not using an SDK, or making direct HTTP API requests, you can use the AMS CLIs for Change Management (CM) and SKMS.

Install the AMS CLIs

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

1. To install the AWS CLI, see [Installing the AWS Command Line Interface](#), and follow the appropriate instructions. Note that at the bottom of that page there are instructions for using different installers, [Linux](#), [MS Windows](#), [macOS](#), [Virtual Environment](#), [Bundled Installer](#) (Linux, macOS, or Unix).
2. After the installation, run `aws help` to verify the installation.
3. Once the AWS CLI is installed, to install or upgrade the AMS CLI, download the AMS distributables zip file and unzip. You can access the AMS CLI distributables through the **Documentation** link in the left nav of the AMS console, or ask your cloud service delivery manager (CSDM) to send you the zip file.
4. Open either the **Managed Cloud Distributables -> CLI -> Windows** or the **Managed Cloud Distributables -> CLI -> Linux / MacOS** directory, depending on your operating system, and:
5. For **Windows**, execute the appropriate installer (this method only works on Windows 32 or 64 bits systems):
 - 32 Bits: `ManagedCloudAPI_x86.msi`
 - 64 Bits: `ManagedCloudAPI_x64.msi`
6. For **Mac/Linux**, execute the file named: **MC_CLI.sh** by running this command: `sh MC_CLI.sh`. Note that the **amscm** and **amsskms** directories and their contents must be in the same directory as the **MC_CLI.sh** file.
7. If your corporate credentials are used via federation with AWS (the AMS default configuration) you must install a credential management tool that can access your federation service. For example, you can use this AWS Security Blog [How to Implement Federated API and CLI Access Using SAML 2.0 and AD FS](#) for help configuring your credential management tooling.
8. After the installation, run `aws amscm help` and `aws amsskms help` to see commands and options.

Application account onboarding

Topics

- [Requesting a new application account \(p. 89\)](#)
- [Setting up Active Directory to federate access to AMS IAM roles \(p. 90\)](#)
- [Setting Up Networking with the New Application Account \(p. 92\)](#)
- [Setting Up Additional VPCs in the Application Account \(p. 93\)](#)

For onboarding questions, contact your cloud service delivery manager (CSDM). See also [Application accounts: AMS-managed, Dev-mode, Customer-managed](#). For general information about modes, see [AMS modes AMS service management \(p. 40\)](#).

For information on the different modes of application accounts, see [Application accounts: AMS-managed, Dev-mode, Customer-managed](#). For general information about modes, see [AMS modes](#).

Requesting a new application account

You must have a multi-account AWS Managed Services (AMS) environment set up with core accounts, before requesting a new application account. For information about setting up a multi-account environment with core accounts, see [Core account onboarding \(p. 77\)](#).

You can choose one of the following Amazon VPC types for the initial VPC in the application account:

- **Private:** This VPC has no Internet gateway attached. This is suitable for private applications that require no access to/from the Internet.
- **Public:** This VPC has an Internet gateway attached and has public and private subnets. This is suitable for public applications that require access to/from the Internet.

You can request a new application account by submitting a `Deployment | Managed landing zone | Management account | Create application account (with VPC) (ct-1zdasmc2ewzrs)` RFC and providing the following values in the RFC:

- **Account Name:** A custom name for the account. Note that the Account Name has a maximum length of 50 characters.
- **Account Email:** The distribution list email for the account. This email ID is used for creating the AWS account.
- **Support level:** The AWS Support level, Premium or Plus.
- **VPC Name:** A name for the VPC.
- **Number of Availability Zones (AZs):** 2 or 3.
- **VPC CIDR:** The CIDR block for the VPC.
- **Route Type:** This can be either `routable` or `isolated`. `Routable` means that application VPCs associated with the Transit Gateway (TGW) application route table can connect to this VPC. `Isolated` means that application VPCs associated with the TGW application route table cannot connect to this VPC. The default is `routable`.
- **Transit Gateway Application Route Table:** The Transit Gateway route table to which the application account VPC has to be associated with. If no value is provided, the default `defaultAppRouteDomain` is used, which means that this account will be able to communicate with all other accounts under the same route table.
- **PublicSubnet<1-3>CIDRCIDR for public subnet in AZ 1:** The CIDR for public subnet in Availability Zone 1.
- **PrivateSubnet<1-10>AZ<1-3>CIDRCIDR for public subnet in AZ 1:** The CIDR for public subnet in Availability Zone 1.

At this point, AMS deploys a new application account into your AMS management account, with the specified VPC configuration.

Setting up Active Directory to federate access to AMS IAM roles

Federate your directory with the AMS IAM roles to enable corporate users to use their corporate credentials to interact with the AWS Console and the AWS APIs, and the AMS console and AMS APIs.

Federation process example

This example uses Active Directory Federation Services (ADFS). However, any technology that supports AWS IAM Federation is supported. For more information about AWS-supported IAM federation, see [IAM Partners](#) and [Identity Providers and Federation](#). Your CSDM will help you through this process, which involves a joint effort with your AD team and AMS.

For detailed information about integrating SAML for API access, refer to this AWS blog, [How to Implement Federated API and CLI Access Using SAML 2.0 and AD FS](#).

For an example that installs the AMS CLI and SAML, see [Appendix: AD FS claim rule and SAML settings](#) in the AMS User Guide.

Configuring federation to the AMS console

The IAM roles and SAML identity provider (Trusted Entity) detailed in the following table are provisioned in your new application account. These roles allow you to gain access to the new application account and file RFCs, write to S3 buckets, and perform other actions.

Role	Permissions
AWSManagedServicesReadOnlyRole	Allows you to view the resources in your new application account.
AWSManagedServicesCaseRole	Allows you to view the resources in your new application account and file AWS Support tickets.
AWSManagedServicesChangeManagementRole	Allows you to view the AMS infrastructure in the application accounts, file RFCs, file AWS Support tickets, write to S3 buckets, manage Secrets Manager secrets, and manage Reserved Amazon Elastic Compute Cloud (Amazon EC2) instances.
AWSManagedServicesSecurityOpsRole	Allows you to view the AMS infrastructure in the application accounts, manage Secrets Manager secrets, manage Web Application Firewall rules, manage certificates, and file AWS Support tickets.
AWSManagedServicesAdminRole	Allows you to view the AMS infrastructure in the application accounts, manage Marketplace subscriptions, manage Secrets Manager secrets, manage Web Application Firewall rules, manage certificates, create RFCs, manage Reserved Amazon EC2 instances, write to S3 buckets, file AWS Support tickets, and manage AWS Artifacts agreements.

Submitting the federation request to AMS

If this is your first account, work with your CSDM(s) and/or Cloud Architect(s) to provide the metadata XML file for your identity provider.

If you are onboarding an additional account or Identity Provider and have access to either the management account or the desired application account, follow these steps.

1. Create a service request from the AMS console.

Note

- If creating an identity provider for an application account, submit this request from either the application account itself or the management account.
- If creating an identity provider for an AMS core account, submit this request from the management account.
- If creating an identity provider for the management account, submit this request from the management account, or contact your CSDM for assistance.

In the service request, provide the details necessary to add the identity provider:

- AccountId of the account where the new identity provider will be created.
 - Desired identity provider name, if not provided, the default will be **customer-saml**; typically, this must match the settings configured in your federation provider.
 - For existing accounts, include whether the new identity provider should be propagated to all existing console roles or provide a list of roles that should trust the new identity provider.
 - Attach the metadata XML file exported from your federation agent to the service request as a file attachment.
2. From the same account where you created the service request, create a new RFC using CT-ID ct-1e1xtak34nx76 (Management | Other | Other | Create) with the following information.
 - Title: "Onboard SAML IDP <Name> for Account <AccountId>".
 - AccountId of the account where the identity provider will be created.
 - Identity provider name.
 - For Existing Accounts: Whether the identity provider should be propagated to all existing console roles, or the list of roles which should trust the new identity provider.
 - Case ID of service request created in Step 1, where the metadata XML file is attached.

Verify Console Access

After you are set up with AD FS, and have the AMS URL to use for authentication, you can perform the following procedure.

With an Active Directory Federated Service (AD FS) configuration, you can follow these steps:

1. Open a browser window and go to the sign in page provided to you for your account. The **AD FS IdpInitiatedSignOn** page for your account opens.
2. Select the radio button next to Sign in to one of the following sites. The Sign in site list becomes active.
3. Choose the signin.aws.amazon.com site and choose Sign in. Options for entering your credentials open.
4. Enter your CORP credentials and choose Sign in. The AWS Management Console opens.
5. Paste into the location bar the URL of the AMS console and press **Enter**. The AMS console opens.

Verify API Access

AMS uses the AWS API, with some AMS-specific operations that you can read about in the [AMS API Reference](#).

AWS provides several SDKs that you can access at [Tools for Amazon Web Services](#). If you don't want to use an SDK, you can make direct API calls. For information on authentication, see [Signing AWS API Requests](#). If you are not using an SDK, or making direct HTTP API requests, you can use the AMS CLIs for Change Management (CM) and SKMS.

Setting Up Networking with the New Application Account

Setting up networking for the application account includes configuring firewall rules and potentially setting up additional Transit Gateway (TGW) route tables.

Set Up Your Firewall

To use the applications deployed in your AMS environment, you must create some firewall rules. You do not need these rules to access your instances, you can hop through the bastions into your instances.

Firewall Rules for Application Access

You must open the following ports for traffic through your firewall:

- From your on-premise network to your new application VPC CIDRs in both the ingress and egress directions.
- From your new application VPC CIDRs to your on-premise network in both the ingress and egress directions (if your cloud applications need to reach out to your on-premise applications).

Port	Protocol	Service	From/To	To/From
80	TCP	HTTP Web Access	On Premise Network	AMS Application VPC
443	TCP	HTTPS Web Access	On Premise Network	AMS Application VPC

Setting Up Additional Transit Gateway Application Route Tables

AWS Managed Services (AMS) networking is flexible and supports a variety of networking use cases.

- Communication between application VPCs in the same account.
- Communication between application VPCs in different accounts.
- Isolation between application VPCs in different accounts.
- Isolation between application VPCs in same accounts.

If you have unique/special requirements for networking, contact your AMS Cloud Architect and they will develop a plan for your requirements to be met by AMS network architecture.

Based on the networking decision taken for application account VPCs, you can create multiple Transit Gateway (TGW) application route tables by submitting a Deployment | Managed landing zone | Networking account | Create application route table (ct-1urj94c3hdfu5) RFC.

The change type requires you to specify a `TransitGatewayApplicationRouteTableName`, a meaningful name for the TGW route table.

Note

The route table created is empty. You must file a Management | Other | Other | Update change type, with the route table name, to add routes to it.

Setting Up Additional VPCs in the Application Account

You can request an additional application account VPC by submitting a Deployment | Managed landing zone | Application account | Create VPC (ct-1j3503fres5a5) RFC.

This works in the same way as configuring a VPC for a new application account. For details, see [Requesting a new application account \(p. 89\)](#).

Appendix: multi-account landing zone Onboarding Consideration List

Topics

- [Account Configuration \(p. 93\)](#)
- [Multi-Account Landing Zone Monitoring Alerts \(p. 94\)](#)
- [Network Configuration \(p. 94\)](#)
- [Active Directory Configuration \(p. 95\)](#)
- [Trend Micro Endpoint Protection \(EPS\) \(p. 95\)](#)
- [Access: Bastions, SSH and RDP \(p. 95\)](#)
- [Federation \(p. 96\)](#)

Use this to consider your AMS multi-account landing zone (MALZ) deployment elements and structure so AMS can determine what infrastructure components are needed. Your Cloud Architect will provide you with a similar questionnaire.

Note

For more information on instance types, see [Amazon EC2 Instance Types](#).

For more information on database instance types, see [Amazon RDS Instance Types](#).

If you require Direct connect, see the AMS single-account landing zone Onboarding Guide to create a Direct connect connection.

You will receive an onboarding questionnaire from your CSDM containing questions about your desired configuration settings for your account. Work with your CSDM to complete the questionnaire before proceeding.

Account Configuration

- New Account ID

The AWS account ID that you created for AMS multi-account landing zone. Should not be part of an AWS organization.

- Service Region

The primary region in which the AMS multi-account landing zone environment will be deployed.

- The core account emails for notifications. (these should all be in the same domain). Provide an email address for each:
 - Shared Services Account
 - Networking Account
 - Logging Account
 - Security Account
- Your Service Type, Premium or Plus

This determines the service level agreements (SLAs) for resolving issues in your environment

Multi-Account Landing Zone Monitoring Alerts

AMS provides a way for you to be directly alerted (versus getting AMS service notifications) for certain monitoring alerts. To sign up for this, make sure that your Cloud Architect or Cloud Service Delivery Manager receive this information:

Direct Alerts Email: These are the email addresses that you want AMS to send certain resource-based alerts to. For details of which alerts are sent directly to email, see [Alerts from Baseline Monitoring in AMS](#) in the AMS User Guide for Multi-Account Landing Zone. For more information on AMS monitoring, see [Monitoring Management](#) in the AMS User Guide for Single-Account Landing Zone.

Network Configuration

- Transit Gateway ASN Number

This is the Autonomous System Number (ASN) for the AWS side of a Border Gateway Protocol (BGP) session, it must be unique and cannot be the same one used for your Direct Connect or VPN. The range is 64512 to 65534 (inclusive) for 16-bit ASNs.

- Your AMS multi-account landing zone infrastructure VPC CIDR ranges.

These CIDR ranges cannot overlap with your on-premise network

You can either include a /22 CIDR range, or provide each VPC CIDR individually. Note that only these CIDR ranges are allowed:

- 10.0.0.0 - 10.255.255.255 (10/8 prefix)
- 172.16.0.0 - 172.31.255.255 (172.16/12 prefix)
- 192.168.0.0 - 192.168.255.255 (192.168/16 prefix)

Note that IP range 198.18.0.0/15 may not be used (it is reserved by AWS Directory Service).

- Core Infrastructure VPC CIDR range (/22 range recommended)
- Networking VPC CIDR range (/24 range recommended)
- Shared Services VPC CIDR range (/23 range recommended)
- DMZ VPC CIDR range (/25 range recommended)
- VPN ECMP (enable or disable)

For VPN ECMP support, choose enable if you need Equal Cost Multipath (ECMP) routing support between VPN connections. If connections advertise the same CIDRs, the traffic is distributed equally between them.

Network access control list (ACL)

A network access control list (NACL) is an optional layer of security for your VPC that acts as a firewall for controlling traffic in and out of one or more subnets. You might set up network ACLs with rules similar to your security groups in order to add an additional layer of security to your VPC. For more information about the differences between security groups and network ACLs, see [Comparison of security groups and network ACLs](#).

However, in AMS Managed Multi-Account Landing Zone, in order for AMS to effectively manage and monitor Infrastructure, the use of NACLs is limited to following scope:

- NACLs are not supported in the multi-account landing zone Core accounts i.e. Management, Networking, Shared-Services, Logging and Security.
- NACLs are supported in multi-account landing zone Application accounts as long as they are only used as a "Deny" list and have "Allow All" to ensure AMS monitoring and management operations.

In large scale multi-account environments, you can also leverage features like centralized egress firewalls to control outbound traffic and/or AWS Transit Gateway routing tables in AMS multi-account landing zone to segregate network traffic among VPCs.

Active Directory Configuration

Domain FQDN for AMS managed Active Directory

Trend Micro Endpoint Protection (EPS)

- Instance sizes for your EC2 instances and Auto Scaling groups

Trend Micro Endpoint Protection (EPS) is the primary component within AMS for operating system security. The system is comprised of Deep Security Manager (DSM) EC2 instances, relay EC2 instances, and an agent present within all of AMS data plane and your EC2 instances.

- Relay instance type (minimum supported by AMS is m5.large)
- DB instance size (200 GB recommended)
- RDS instance type (only db.m5.large or db.m5.xlarge allowed)
- DSM License type (Marketplace or BYOL)

If you already have a license, choose BYOL (bring your own license). AMS will contact you to obtain the necessary information about the license.

- AWS IAM Role/User Amazon resource name (ARN) for Trend Micro Deep Security Subscription (Role ARN: `arn:aws:iam::ACCOUNT_ID:role/ROLE_NAME`)

Provide us an IAM role ARN, or an IAM user ARN from one of your existing AWS accounts to which you have access. AMS creates an IAM role in your AMS multi-account landing zone Shared Services account and adds the role/user provided in the trust of an IAM role in Shared Services so that the role can be assumed by you to subscribe to the Trend Micro Deep Security in AWS Marketplace.

Access: Bastions, SSH and RDP

- SSH Bastion settings

AMS provides SSH bastions in your Shared Services account to access hosts in the AMS environment. In order to access the AMS network as an SSH user, you must use SSH Bastions as the entry point. The network path originates from the On-Prem network, goes through DX/VPN to the transit gateway

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

- Desired instance count (2 recommended)
- Maximum instances (4 recommended)
- Minimum instances (2 recommended)
- Instance type (m5.large recommended)
- Ingress CIDRs: IP address ranges from which users in your network will access SSH Bastions (ip range 1, ip range 2, ip range 3, ... etc)
- RDP Bastion settings

AMS optionally provides RDP bastions in your Shared Services account to access hosts in the AMS environment. In order to access the AMS network as an RDP user, you must use RDP Bastions as the entry point. The network path originates from the On-Prem network, goes through DX/VPN to the TGW, and then is routed to Shared Services VPC. Once you are able to access the bastion, you can jump to other hosts in the AMS environment, provided that the proper access request has been granted.

- Instance type (t3.medium recommended)
- Desired minimum sessions (2 recommended)
- Desired maximum sessions (10 recommended)
- RDP Bastion Configuration Type, Shared Standard or Shared HA (default is Shared Standard)

SecureStandard = A user receives one bastion and only one user can connect to the bastion.

SecureHA = A user receives two bastions in two different AZ's to connect to and only one user can connect to the bastion.

SharedStandard = A user receives one bastion to connect to and two users can connect to the same bastion at once.

SharedHA = A user receives two bastions in two different AZ's to connect to and two users can connect to the same bastion at once.

Federation

Identity Provider (IDP) Name

Defaults to `customer-saml`

AMS Single-account landing zone Onboarding

AMS single-account landing zone onboarding process

Onboarding AMS single-account landing zone accounts involves the following tasks:

1. You must create a new AWS account that AMS configures as the networking account to host the firewall. Create the new account within your AWS organization, if you have one. AMS will follow the procedure of creating a normal AMS account, so all the information required must be gathered, for example CIDR, EPS licenses, users, and so forth. Note: A CIDR allocation of /24 is good.
2. Let us know whether or not you want to remove the Internet gateways (IGWs) from the egress traffic accounts.
3. Determine your approved domains. AMS enables destination filtering by maintaining an approved domain list; the list can be modified later.
4. Confirm the instance size you want to use based on your expected throughput. By default, the instance is created in a m4.xlarge instance where we have found that the firewall throughput is 350Mbps. AMS can increase the size to a c4.8xLarge instance where the expected throughput is 1.25 Gbps.
5. Set up networking between AMS and your private network, this involves several tasks:
 - a. Allocate IP space
 - b. Establish private network connectivity to AWS
 - c. Set up your firewall
 - d. Set up access management
 - e. Schedule backups
6. Provide access to the created account to AMS.
7. Validate that the AMS service is operating properly.

AMS will be able to perform the account build-out (onboarding) of your account within 2 weeks (10 business days) from the initial request date. Any follow-up activity can be performed by leveraging [AMS Planned Event Management \(PEM\)](#).

Note

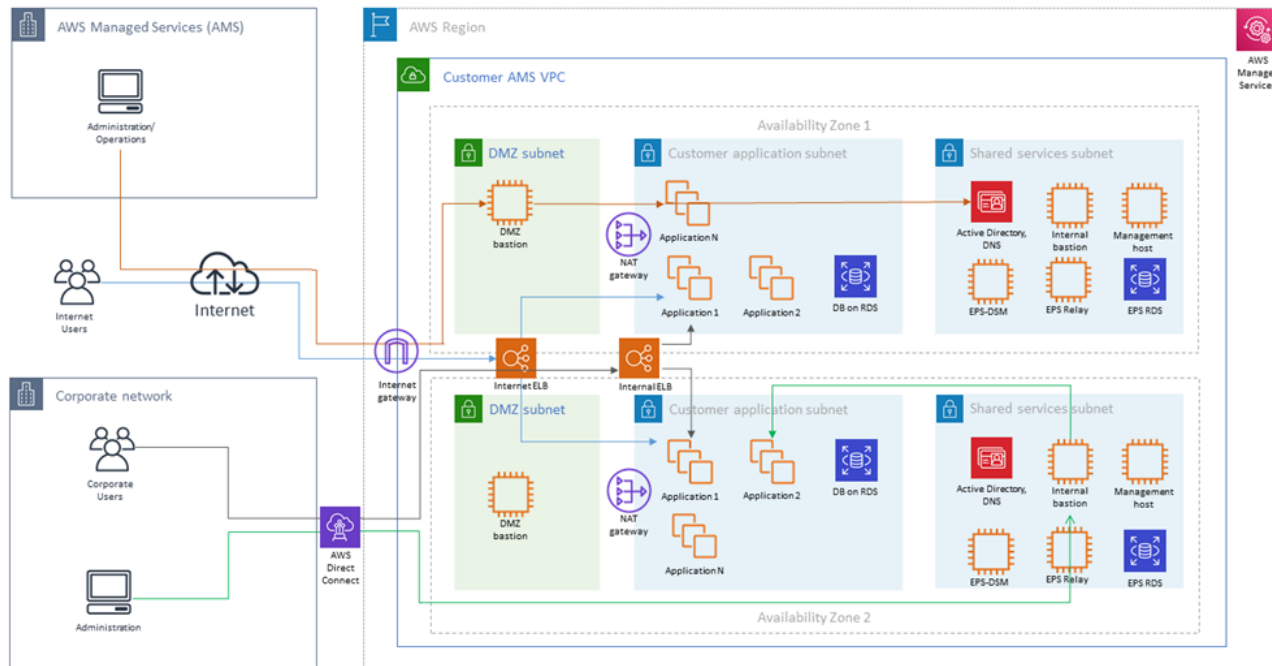
- US East (Virginia)
- US West (N. California)
- US West (Oregon)
- US East (Ohio)
- Canada (Central)
- South America (São Paulo)

- EU (Ireland)
- EU (Frankfurt)
- EU (London)
- EU West (Paris)
- Asia Pacific (Mumbai)
- Asia Pacific (Seoul)
- Asia Pacific (Singapore)
- Asia Pacific (Sydney)
- Asia Pacific (Tokyo)

New regions are added frequently. To learn more, see [AWS regions and availability zones](#).

Single-Account Landing Zone network architecture

The following diagram depicts the AMS single-account landing zone VPC network layout and is an example of the highly available setup.



AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
AMS Single-account landing zone shared services

— 1 →	Ingress through DirectConnect (internal customer network users) and Internet with managed Internet Gateway (external users), through AWS load balancers to customers subnet applications. Note that traffic for external users goes through load balancers in DMZ (Public) Subnet, while traffic for internal users goes through load balancers in Application (Private) Subnet	Each AMS account has a VPC in one region with resource subnets located in two availability zones. Each availability zone has three subnets: DMZ, Customer, and Shared Services. Your (“customer”) corporate network is connected through a DirectConnect (VPN) tunnel, and AMS Operations connects to your managed VPC over the Internet. Shared services subnets contain AMS Directory Services with one AD Domain Controller per shared services subnet, and AMS Management Hosts that automate provisioning and common tasks, Antivirus (TrendMicro) management servers that include EPS DSM and EPS relay (for scalability), and internal (customer) bastion hosts.
— 2 →	Ingress through Internet with managed Internet Gateway for AMS administrators and operators through DMZ bastions to customer and shared services subnets	DMZ subnets contain Internet load balancers, your DMZ instances, and DMZ bastion hosts that serve as SSH jump boxes for the AMS Operations team. DMZ bastions, as well as other AMS infrastructure in the Shared services subnet, have two nodes for high availability. Your “customer” subnets contain your workloads, EC2 instances, RDS, etc.
— 3 →	Ingress through DirectConnect (internal customer network administrators) and internal bastions to customer subnets	External users connect to your applications for the Internet via an AWS Load Balancer that is located in your DMZ.

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

Note

To learn about default service limits and constraints for all active services, see the [AWS Service Limits](#) documentation.

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

AMS Single-account landing zone shared services

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

- AMS Directory Services = AD Domain Controller

Creates an Active Directory in AMS accounts, creates the AMS domain, joins managed stacks to the domain on launch.

- Management hosts = AMS Management Host (automate provisioning and common tasks)

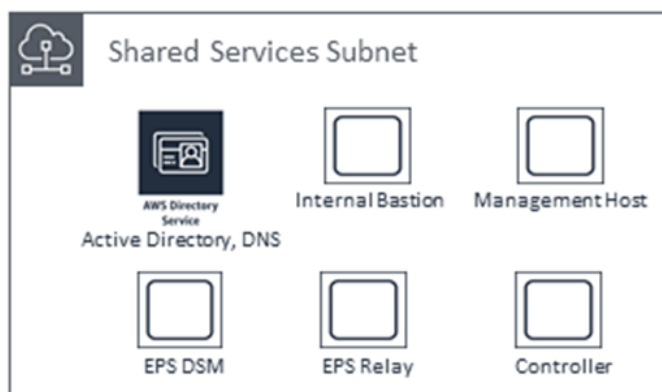
Act as an API endpoint to modify AWS Directory Service, interact with AWS Directory Service domain controllers.

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

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

- Internal bastion hosts = Customer bastions

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



Create a New AWS Account for AMS

The five main components of creating a new AWS account for AWS Managed Services (AMS) are:

1. [Create an AWS Account \(p. 100\)](#)
2. [Set Up Consolidated Billing -- Link New Account to Payer Account \(p. 101\)](#)
3. [Configure your AWS Account for AMS Access \(p. 101\)](#)
4. [Secure the New Account with Multi-Factor Authentication \(MFA\) for the Root User \(p. 105\)](#)
5. [Subscribe to AWS Marketplace for EPS \(p. 104\)](#)

Please contact your customer service delivery manager (CSDM) if you have any questions.

Create an AWS Account

The AMS program requires the provisioning of a new Amazon Web Services (AWS) account. Step by step instructions are available in the following video: [How do I create and activate a new Amazon Web Services account?](#) The simple steps are:

Go to <https://aws.amazon.com/>, and click **Create an AWS Account**. Follow the on-screen instructions, which include receiving a phone call and entering a PIN using your phone keypad. You'll also need to enter a credit card, unless you will set up consolidated billing (described next).

Note

If you already have an account, you can go to the [AWS Pricing](#) page and click **Create a Free Account**. *Be sure to sign up for the EC2 Service*, at least. Signing up for one service allows you access to all services in AWS. You are charged only for the services that you use.

If you plan to link your new account to a payer account for the purposes of consolidated billing, you do not need to enter payment method information when prompted. Instead, once you

reach the screen to enter credit card information, simply navigate away. You will need the email address associated with the payer account to send a consolidated billing/linked account request which is detailed in the next section.

Important

It is critical that you ensure that an email address and phone number are associated with the account so you receive responses to potential security incidents. The phone number and email address for the account cannot be changed without resetting the account password, which is a significant undertaking for an AMS root account. To ensure that these values are stable, it is critical to select contact information not associated with individuals, which can change. Choose an email alias that can point to a group. Follow this same best practice in selecting a phone number: choose a number that can point to a group or to a number owned by the company and not an individual.

Set Up Consolidated Billing -- Link New Account to Payer Account

If you'd like your new AMS-managed AWS account bill to be rolled into a payment for an existing AWS Organizations management account, you need to set up consolidated billing and link the accounts. For details on Consolidated Billing read [Paying Bills for Multiple Accounts Using Consolidated Billing](#) and [AWS Multi Account Billing Strategy](#). These are the basic steps to designate an existing account as a payer account and add existing accounts to it for consolidated billing:

1. Sign up for consolidated billing in the <https://console.aws.amazon.com/billing/>, and designate one of your existing accounts as a *payer account*. If you already have an account designated as a payer account, open the settings for that account and do the following:
2. Add *linked accounts* (up to 20) to your consolidated billing *account family*:
 - a. In the [Billing and Cost Management console](#) -> **Consolidated Billing** (left nav pane), **Manage Requests and Accounts** page, choose **Send a Request**.

The **Send a Consolidated Billing Request** page opens.
 - b. Enter email addresses for the accounts that you would like to link to your payer account. If you choose, you can add notes that will be added to the email body. Click **Send**.

The linked account owner receives a hyperlink in the email, and uses it to log in to the AWS website, and accepts or denies the request to link that account to the payer account.

For details, see [Creating and Editing Consolidated Billing Account Families](#).

Configure your AWS Account for AMS Access

With the above steps completed, you've successfully secured your new AWS account and ensured associated costs are billed appropriately. The final step in the process is to allow AMS access to the new account for initial stack configuration and for ongoing change and provisioning requests to be fulfilled. For details, read [Delegate Access Across AWS Accounts Using IAM Roles](#). The basic steps are described in this section.

Activate Access to the AWS Website

In order to grant your IAM users access to your account's billing information and tools, you must activate the functionality. Follow these steps:

1. Sign in to the AWS Management Console with your *root account* credentials (the email and password that you used to create your AWS account). Don't sign in with your IAM user credentials.

The AWS Management Console home page opens.

2. In the top navigation bar, open the drop-down menu for your account name, and then choose **My Account**.

The Billing home page opens.

3. Scroll down to **IAM User Access to Billing Information**, and click **Edit**.

An **Activate IAM access** area opens.

4. Select the check box and click **Update**.

You can now use IAM policies to control which pages a user can access.

Create an IAM Role with Access to the AWS Website

AWS Identity and Access Management (IAM) is a web service that helps you securely control access to AWS resources for your users. You use IAM to control who can use your AWS resources (authentication) and what resources they can use and in what ways (authorization).

1. Go to the [IAM Management Console](#), click **Roles** in the left nav pane.

The Roles management page opens with information about IAM roles, a **Create role** option, and a list of existing roles.

The screenshot shows the AWS IAM Roles management console. On the left is a navigation sidebar with options: Dashboard, Groups, Users, Roles (highlighted), Policies, Identity providers, Account settings, Credential report, and Encryption keys. The main content area is titled 'Roles' and contains a search bar, a 'What are IAM roles?' section with a list of examples, an 'Additional resources' section with links to FAQs and documentation, and buttons for 'Create role' and 'Delete role'. Below this is a table with columns for 'Role name', 'Description', and 'Trusted'.

2. Click **Create role**.

The Create role **Select type of trusted entity** page opens. Click **Another AWS account** and a settings area opens up below.

Enter the AMS trusted **Account ID** provided to you by AMS. Leave the **Require external ID** and **Require MFA** options de-selected.

Create role



Select type of trusted entity

Four options for trusted entity type:

- AWS service
- Another AWS account (Selected)
- Web identity
- SAML (Saml 2.0 federation)

Allows entities in other accounts to perform actions in this account. [Learn more](#)

Specify accounts that can use this role

Account ID* ⓘ

Options

- Require external ID (Best practice when a third party will assume this role)
- Require MFA ⓘ

* Required

Cancel

Next: Permissions

3. Click **Next: Permissions**.

The Create role **Attach permissions policies** page opens with options for creating a new policy, refreshing the page, and searching existing policies. A list of existing policies is provided.

Filter: Policy type Search


Showing 355 results

	Policy name	Attachments	Description
<input type="checkbox"/>	admin	1	
<input type="checkbox"/>	admin_1	1	
<input checked="" type="checkbox"/>	AdministratorAccess	2	Provides full access to AWS services and resources.
<input type="checkbox"/>	AmazonAPIGatewayAdministrator	0	Provides full access to create/edit/delete APIs in Amazon ...
<input type="checkbox"/>	AmazonAPIGatewayInvokeFullAccess	0	Provides full access to invoke APIs in Amazon API Gateway

4. Select the **AdministratorAccess** policy and then click **Next: Review**.

The Create role **Review** page opens.

Create role



Review

Provide the required information below and review this role before you create it.

Role name*
Maximum 64 characters. Use alphanumeric and '+=, @, _' characters.

Role description
Maximum 1000 characters. Use alphanumeric and '+=, @, _' characters.

Trusted entities The account 123456789999

Policies  AdministratorAccess [↗](#)

* Required

Cancel

Previous

Create role

5. Name the new role **aws_managedservices_onboarding_role** and type "AMS Onboarding Role" for the **Role description**. Review the settings for the new role and, if satisfied, click **Create role**.

The role management page opens with your new role listed.

Subscribe to AWS Marketplace for EPS

Recent changes to AMS endpoint security (EPS) require you to subscribe to TrendMicro Deep Security through the AWS Marketplace and accept the Software Terms.

TrendMicro offers two license models: Per Protected Instance Hour and Bring your own License (BYOL).

- **BYOL:**

1. You use your own license that you have purchased through external channels.
2. You must provide all the license keys to AMS to build the EPS infrastructure. You can provide an activation code that licenses all modules, or individual activation codes that license a certain set of modules. AMS creates only the license files that correspond with the activation codes you provide. Since the license activation occurs during onboarding, in the presence of an AMS lead engineer and CSDM, you can share that information then.
3. Additionally, you must subscribe to BYOL TrendMicro Market Place AMI Subscription. See [Trend Micro Deep Security \(BYOL\)](#).

- **Per Protected Instance Hour:**

1. In this subscription, you are not required to have any previously-procured Trend license.
2. However, you must subscribe to the Marketplace subscription.

3. No license key sharing with AMS is required in this model, as the Trend usage is metered automatically including the software license + EC2 infrastructure usage. See [Trend Micro Deep Security](#).

To subscribe to Trend Micro, follow these steps:

1. Login into your AWS account.
2. Navigate to Trend Micro Deep Security ([BYOL](#) or [Per Protected Instance Hour](#)) product page.
3. Click **Continue to Subscribe** in the right panel.
4. Click **Accept Terms** in the upper right corner.

Enable IDS and IPS in Trend Micro Deep Security

You can request that AMS enable Trend Micro Intrusion Detection System (IDS) and Intrusion Protection Systems (IPS), non-default features, for your account.

To do this, submit an update request (Management | Other | Other | Update) and include a list of email addresses to receive IDS and IPS notifications. These addresses are added to an SNS topic in your account, which AMS creates for you.

Note

AMS cannot add any Trend Micro service that might interfere with our ability to provide other AMS services.

Next step: [Secure the New Account with Multi-Factor Authentication \(MFA\) for the Root User \(p. 105\)](#)

Subscribe to AWS Marketplace for CentOS 7.6

AMS now provides the CentOS 7 (x86_64) - with Updates HVM sold by Centos.org, as an AMS AMI. In order to utilize this AMI, you must opt in to the FREE Cent OS license, and accept the license on all your AMS accounts.

To subscribe, go to [AWS Marketplace](#) and follow the instructions for opting-in.

You will not incur software charges for using this product, but you are still responsible for other AWS charges, including EC2 usage. If this is a "Bring Your Own License" product you must have a valid software license in order to use it.

You can review information for this software at [CentOS 7 \(x86_64\) - with Updates HVM](#).

Secure the New Account with Multi-Factor Authentication (MFA) for the Root User

This section has been redacted because it contains sensitive AMS security-related information. This information is available through the AMS console **Documentation**. To access AWS Artifact, you can contact your CSDM for instructions or go to [Getting Started with AWS Artifact](#).

Set Up Networking

There are several processes that need to be completed to set up networking for AWS Managed Services (AMS):

1. Allocating IP space for your AMS environment
2. Establishing private network connectivity to AWS

3. Setting up your firewall to allow AMS operations

Allocate IP Space for your AMS Environment

AMS was designed and tested using a /16 CIDR block as the recommended network allocation. It is important that the trusted network connected to AMS use a CIDR block that does not overlap with the CIDR block assigned to AMS. These addresses are required to set up your virtual private cloud (VPC) and subnets. For more information about AWS VPCs, see [Amazon VPC Limits](#) and [Amazon VPC FAQs](#).

While a /16 CIDR block may seem like a lot of IP addresses, a VPC, once created, cannot be expanded. So this allocation ensures that your AMS-managed VPC can function for a considerable period. Within the CIDR block, you must allocate IP address ranges for, at least, two private subnets and two public subnets.

AWS accepts connectivity to the AMS environment via native AWS virtual private network (VPN) functionality. On your side, this can be achieved via AWS Direct Connect (DX), hardware VPN, or software VPN. On the AMS side, we use the Virtual Gateway functionality of VPCs.

Basic Environment Components

User Network-to-Amazon VPC Connectivity Options	
Hardware VPN	Establishes a hardware VPN connection from your network equipment on a remote network to AMS-managed network equipment attached to your VPC.
AWS Direct Connect (DX)	Establishes a private, logical (or encrypted if used with a VPN) connection from your remote network to the Amazon VPC, leveraging AWS Direct Connect.
Software VPN	Establishes a VPN connection from your equipment on a remote network to a user-managed software VPN appliance running inside an Amazon VPC.

Note

AMS recommends redundant private VPN to DX connections. For details, see [How can I configure VPN as a backup for my AWS Direct Connect connection?](#) Your customer service delivery manager (CSDM) will assist in setting this up at the time of onboarding your account.

Establish Private Network Connectivity to AWS

Add AMS to your corporate Active Directory to establish connectivity. You may want to perform administrative actions or user access over a private networking connection. AWS offers both VPN connectivity and dedicated lines via AWS Direct Connect. The following steps explain how to work with AMS to establish either (or both) means of connectivity.

VPN Setup

This section describes the basic steps for setting up a VPN to communicate between your AMS-managed VPC and your internal network.

Note

To gain overall understanding about using a VPN with AWS services refer to [What is AWS Site-to-Site VPN](#) and all about [Your Customer Gateway](#) (your VPN appliance).

Follow the AWS VPN User Guide [Getting Started](#) and [Testing the Site-to-Site VPN Connection](#) sections to complete the following steps.

- Step 1: In your AWS VPC, Create a Customer Gateway
- Step 2: In your AWS VPC, Create a Virtual Private Gateway
- Step 3: In your AWS VPC, Enable Route Propagation in Your Route Table
- Step 4: In your AWS VPC, Update Your Security Group to Enable Inbound SSH, RDP, and ICMP Access
- Step 5: In your internal Network, Create a VPN Connection and Configure the Customer Gateway
- Step 6: Test VPN connectivity between the VPC and your internal network

AWS Direct Connect Setup

This section describes the basic steps for setting up a AWS Direct Connect (DX) to communicate between your AMS-managed VPC and your internal network.

Note

For information about using a DX with AWS services, see [Getting Started at an AWS Direct Connect Location](#).

To set up a DX connection, you need to complete the following steps:

1. [Sign Up for Amazon Web Services](#)
2. [Submit AWS Direct Connect Connection Request](#)
3. [Complete the Cross Connect](#)
4. [\(Optional\) Configure Redundant Connections with AWS Direct Connect](#)
5. Performed by AMS: Create a Virtual Interface
6. Performed by AMS: Download Router Configuration
7. [Verify Your Virtual Interface](#)

Set up your Firewall

This section has been redacted because it contains sensitive AMS security-related information. This information is available through the AMS console **Documentation**. To access AWS Artifact, you can contact your CSDM for instructions or go to [Getting Started with AWS Artifact](#).

AMS Bastion Options during Application Migrations/ Onboarding

In order to provide you with the best experience during migration efforts, below are the potential options AMS could currently leverage:

- *Option 1:* Bypass Bastions for migration efforts only (you must sign off on this for security purposes as a temporary measure).

Note: Auditing capabilities will still be in place to ensure AMS has visibility into each request.

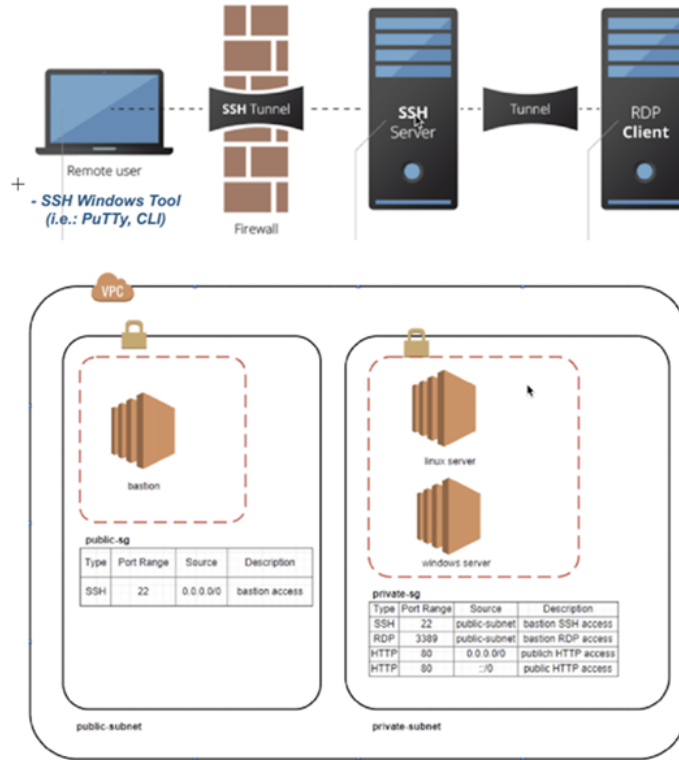
- *Option 2:* SSH Tunneling with a tool of choice; for example, PuTTY, as illustrated.

The environment components described would already need to be in place for this option.

AMS would provide additional notes and instructions.

AMS Advanced Onboarding Guide AMS
 Advanced Account Onboarding Information
 AMS Bastion Options during
 Application Migrations/Onboarding

SSH Tunneling Option
Application Migration Effort

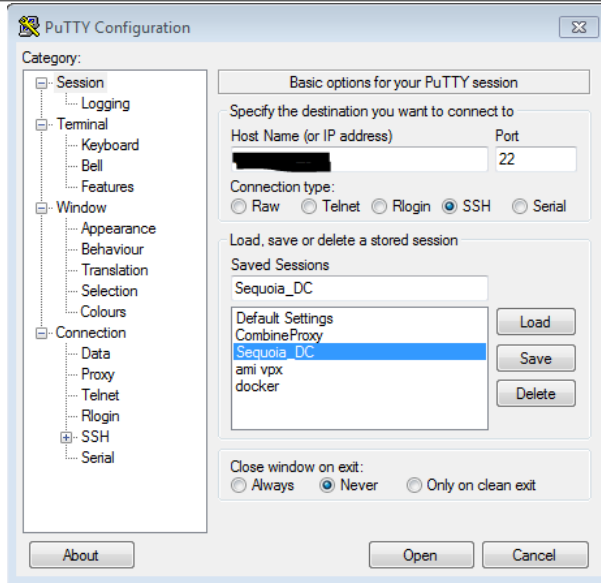


SSH tunneling steps with PuTTY:

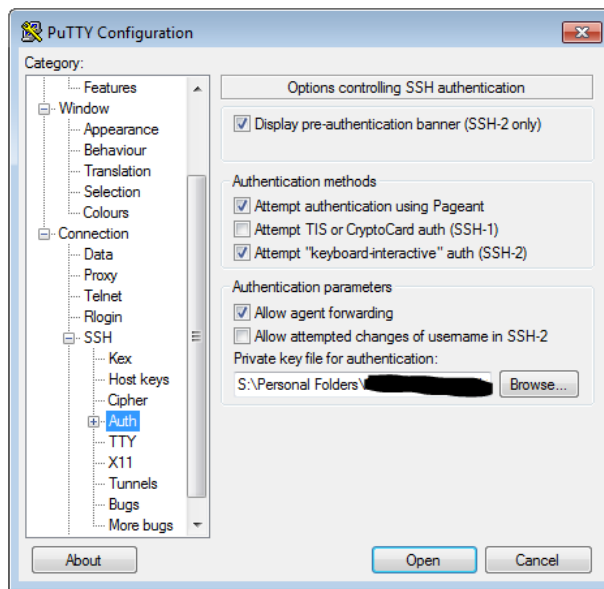
Within PuTTY, you would create an SSH session, with the public IP of the bastion host, provide the PEM key in the AUTH section, and then create a Tunnel. The tunnel's source port should be an unused local port (e.g. 5000) and the IP would be the IP of the destination host (the Windows box you are trying to reach) with the RDP port appended (3389). Be sure to save your configuration, as you don't want to have to do it each time you log into the box. Connect to the bastion host, and log in. Then, start an RDP session for localhost:5000 (or whichever port you choose).

1. Set Host Name or public IP of the bastion host

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
AMS Bastion Options during
Application Migrations/Onboarding

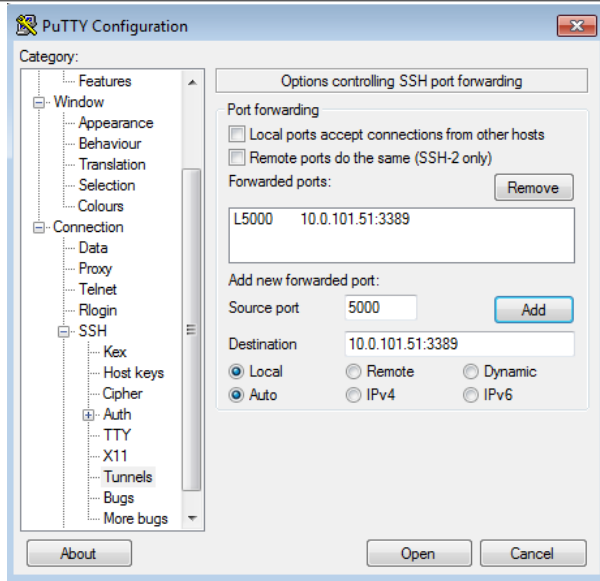


2. In SSH ->Auth, set the private key file in .pk format




3. In SSH ->Tunnels, add the new forwarded port. The Source Port should be the arbitrary unused port, and the Destination should be the IP of the destination server behind the bastion host, with the RDP port appended.


AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
AMS Bastion Options during
Application Migrations/Onboarding



4. Connect to the bastion host via PuTTY and log in.



5. Start an RDP session to localhost:5000 to reach the destination server.



Set up access management

Using a network managed by AWS Managed Services (AMS) means giving AMS access to manage your cloud infrastructure. You'll need to configure a means of securely connecting between your private network and AMS. This starts with some decisions:

- *AMS API/CLI and Console access*: You will want to install the AMS CLI (instructions are provided in this document). You use the AMS change management API to make change requests to AMS and the AMS SKMS API to learn about your AMS-managed resources. Using Active Directory Federation Services (AD FS), you can access the AMS Console.

Note

If you are setting up your own ITSM, you will need to use the AWS Support API (SAPI) for service requests and incident reports. SAPI is documented in the [AWS Support API Reference](#).

- *User access*: Whether you manage users with Windows Active Directory (AD), or a Linux/LDAP solution, connectivity needs to be established between AD on the AMS side (via Directory Services) and your directory.
- *Instance access*: Instance-level access is accomplished via a one-way Forest trust configuration. Directory Services trusts credentials in their CORP AD, allowing stacks within the AMS side to allow login with CORP credentials.

Note that your Active Directory (AD) that AMS sets up the trust to must be the directory that has the accounts of users authorized by you to gain access to your AWS resources.

Important

To set up a Forest trust, AMS requires your domain controller **Local Policies -> Security Options -> Network Access: Named Pipes that can be accessed anonymously**, have the **Netlogon** and **lsarpc** pipes listed. These pipes are listed by default, but are sometimes removed for security concerns. Once the trust is established, they can be removed from the list again.

Establish an Active Directory (AD) trust

Before you begin, ensure that the appropriate firewall ports are open.

The trust from the AMS-managed Active Directory and your corporate directory service allows you to use your corporate-managed credentials to access AMS-managed instances to perform development, test, or administrative functions.

Creating a trust connection is a two-part exercise:

First, configure a conditional forwarder, a DNS configuration so DNS queries know which DNS server to go to.

Second, configure a trust, an Active Directory (AD) construct to allow access from users in one domain to use resources in another domain.

Configure the conditional forwarder

Follow this MicroSoft AD article [Assign a Conditional Forwarder for a Domain Name](#), and use these settings and choices:

1. In the AD **DNS Manager** -> **Create a New Conditional Forwarder**, under **DNS Domain**: Use the domain name AMS supplied to you; for example, `A523434123.amazonaws.com`.
2. Under **IP addresses of the master servers**: Add the AMS-supplied IP addresses. Make sure there isn't a connection problem by validating both addresses.

3. Select **Store this conditional forwarder in Active Directory and replicate as follows: All DNS servers in this domain** and press **OK**.

Configure the trust

Follow this MicroSoft AD article [Create a one-way, incoming, forest trust for one side of the trust](#), using the settings and choices described in this section.

1. Open the **Start -> Administrative Tools -> Active Directory Domains and Trusts** dialog. Right-click the domain node for the domain that you want to establish a trust with, and then click **Properties -> Trusts -> New Trust** to open the New Trust Wizard. Enter the domain name provided to you by AMS for the **Trust Name** and press **Next**.
2. Under **Trust Type**, select **Forest Trust**. Press **Next**.
3. Under **Direction of Trust**, select **One-way: incoming**. Press **Next**.
4. Under **Sides of Trust**, select **This domain only**. Press **Next**.
5. Under **Trust Password**, type a password of your choosing. Press **Next**.
6. For **Trust Selections Completed** and **Trust Creation Complete**, just press **Next**.
7. Under **Confirm Incoming Trust**, select **No**, do not confirm the incoming trust. Press **Next**.
8. Under **Completed the New Trust Wizard**, select **Finish**, and then **OK** to close.
9. Provide the trust password (contact us via your CSDM's phone number for security reasons). AMS will complete the trust configuration.

Active Directory sites and services

To reduce login latency, add the VPC CIDR range to your Active Directory Sites and Services (**Start -> Administrative Tools -> Active Directory Sites and Services**). Add the VPC CIDR range to an Active Directory Site that contains Domain Controllers that are closest to AWS.

Active Directory name suffix routing

After the one-way forest trust has been established, please complete the additional steps.

1. Under **Start > All Programs > Administrative Tools**, click **Active Directory Domains and Trusts**.
The Active Directory Domains and Trusts console opens.
2. Right-click your corporate domain and click **Properties**
The Properties dialog for that domain opens.
3. Click the **Trusts** tab.
The Trusts page opens.
4. Click the Amazon domain name and click **Properties**.
The Properties page for the Amazon domain trust opens.
5. Click **Name Suffix Routing** and click **Refresh**.

These steps ensure that the Service Principal Names (SPNs) can resolve over the trust.

Providing secure access to your applications

Like the operating system access, all application access should be governed using Active Directory (AD) groups. Using Amazon RDS as an example, you must break the mirror to add a new user without AD groups. So the best approach is to create a group in AD and add it at database creation time.

For more information on Active Directory grouping, see [Using Group Nesting Strategy – AD Best Practices for Group Strategy](#)

Troubleshooting

Some things to try if you run into trouble:

- The AMS-managed Active Directory outbound security group needs to be allowed connection through your CIDR block (e.g. 10.27.0.0/16) to your domain controller.
- Trace the route in the AWS Console from domain controller to domain controller checking all security groups along the way.
- Make sure you can ping the AMS-managed Active Directory Domain Controllers if Internet Control Message Protocol (ICMP) is allowed.
- Make sure your Domain Controller can communicate with AWS Directory Services.
- Make sure the conditional forwarders resolve and are validated.
- If you do not see **Forest Trust** in the New Trust wizard, then your conditional forwarders may not be working correctly:
 - Use nslookup to test resolution
 - Try rebooting the Domain Controller

AMS Managed Active Directory

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

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

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

Notes:

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

Access works like this:

- AWS Directory Service team: Has access to domain controllers.
- AMS: Has access to Directory Service APIs to perform certain actions on the domain. These actions include taking AD snapshots, changing AD schema, and others actions.
- You: Have access to the domain (AD) for creating users, groups, and so on.

- We recommend that you perform a proof of concept on Managed AD before migrating your corporate AD, because not all functionality from a traditional AD environment is available in a Managed AD environment.
- AMS will not manage or provide guidance on your AD management. For example, AMS will not provide guidance on Organizational Unit structure, group policy structure, AD user naming conventions, and so forth.

It works like this:

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

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

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

For more information, see [AWS Directory Service FAQs](#).

- Domain FQDN: The FQDN for your AMS Managed AD domain.
- Domain NetBIOS name: The NetBIOS name for your AMS Managed AD domain.
- Account numbers of AMS-standard accounts you would like Managed AD integration to (AMS configures a one way trust from the AMS-standard account's AD to the Managed AD)
- Are Active Directory Schema modifications required and if so, what modifications?
- By default, two domain controllers are provisioned. Do you require more? If so, how many do you require and for what reason?
- Networking for Managed Active Directory information:
 - Managed AD VPC CIDR for domain controllers (a CIDR in your private subnet range for the Managed AD domain controllers):
 - Subnet CIDR 1 for domain controllers: [your CIDR, needs to be part of AMS Managed AD VPC CIDR]
 - Subnet CIDR 2 for domain controllers: [your CIDR, needs to be part of AMS Managed AD VPC CIDR]

For example:

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

To avoid IP address conflicts, be sure that the Managed AD VPC CIDR you specify does not conflict with any other private subnet CIDR you are using in your corporate network.

- VPN Technology (optional): [Direct Connect/Direct Connect and VPN]
 - Your gateway's BGP Autonomous System Number (ASN): [Customer-provided ASN]
 - The Internet-routable IP address for your gateway's outside interface, the address must be static: [Customer Provided IP Address]
 - Whether or not your VPN connection requires static routes: [yes/no]

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

For the most recent RACI table, see the "Roles and Responsibilities" section in the See the AMS FAQs appendix in the User Guide.

Federate your Active Directory with the AMS IAM roles

The purpose of federating your directory with the AMS IAM roles is to enable corporate users to use their corporate credentials to interact with the AWS Console and the AWS APIs, and therefore the AMS console and APIs.

Federation process example

This example uses Active Directory Federation Services (AD FS); however, any technology that supports AWS IAM Federation is supported. For more information on AWS-supported IAM federation, see [IAM Partners](#) and [Identity Providers and Federation](#). Your CSDM will help you through this process, which involves a joint effort with your AD team and AMS.

For detailed information on integrating SAML for API access, refer to this AWS blog, [How to Implement Federated API and CLI Access Using SAML 2.0 and AD FS](#).

Note

For an example that installs the AMS CLI and SAML, see [Appendix: ActiveDirectory Federation Services \(ADFS\) claim rule and SAML settings \(p. 202\)](#).

Configuring federation to the AMS console

The IAM roles and SAML identity provider (Trusted Entity) detailed in the following table have been provisioned as part of your account onboarding. These roles allow you to submit and monitor RFCs, service requests, and incident reports, as well as get information on your VPCs and stacks.

Role	Identity Provider	Permissions
Customer_ReadOnly_Role	SAML	For standard AMS accounts. Allows you to submit RFCs to make changes to AMS-managed infrastructure, as well as create service requests and incidents.

Role	Identity Provider	Permissions
customer_managed_ad_user_role	SAML	For AMS Managed Active Directory accounts. Allows you to login to the AMS Console to create service requests and incidents (no RFCs).

For the full list of the roles available under different accounts see [IAM User Role \(p. 20\)](#).

A member of the onboarding team uploads the metadata file from your federation solution to the pre-configured identity provider. You use a SAML identity provider when you want to establish trust between a SAML-compatible IdP (identity provider) such as Shibboleth or Active Directory Federation Services, so that users in your organization can access AWS resources. SAML identity providers in IAM are used as principals in an IAM trust policy with the above roles.

While other federation solutions provide integration instructions for AWS, AMS has separate instructions. Using the following blog post, [Enabling Federation to AWS Using Windows Active Directory, AD FS, and SAML 2.0](#), along with the amendments given below, will enable your corporate users to access multiple AWS accounts from a single browser.

After creating the relying party trust as per the blog post, configure the claims rules in the following way:

- **Nameld:** Follow the blog post.
- **RoleSessionName:** Use the following values:
 - **Claim rule name:** RoleSessionName
 - **Attribute store:** Active Directory
 - **LDAP Attribute:** SAM-Account-Name
 - **Outgoing Claim Type:** https://aws.amazon.com/SAML/Attributes/RoleSessionName
- Get AD Groups: Follow the blog post.
- Role claim: Follow the blog post, but for the Custom rule, use this:

```
c:[Type == "http://temp/variable", Value =~ "(?i)^AWS-([\d]{12})-"]
=> issue(Type = "https://aws.amazon.com/SAML/Attributes/Role", Value =
RegexReplace(c.Value, "AWS-([\d]{12})-",
"arn:aws:iam::$1:saml-provider/customer-readonly-saml,arn:aws:iam::$1:role/");
```

When using AD FS, you must create Active Directory security groups for each role in the format shown in the following table (customer_managed_ad_user_role is for AMS Managed AD accounts only):

Group	Role
AWS-[AccountNo]-Customer_ReadOnly_Role	Customer_ReadOnly_Role
AWS-[AccountNo]-customer_managed_ad_user_role	customer_managed_ad_user_role

For further information, see [Configuring SAML Assertions for the Authentication Response](#).

Tip

To help with troubleshooting, download the SAML tracer plugin for your browser.

Submitting the federation request to AMS

If this is your first account, work with your CSDM(s) and/or Cloud Architect(s) to provide the metadata XML file for your identity provider.

If you are onboarding an additional account or Identity Provider and have access to either the management account or the desired application account, follow these steps.

1. Create a service request from the AMS console, provide the details necessary to add the identity provider:
 - AccountId of the account where the new identity provider will be created.
 - Desired identity provider name, if not provided, the default will be **customer-saml**; typically, this must match the settings configured in your federation provider.
 - For existing accounts, include whether the new identity provider should be propagated to all existing console roles or provide a list of roles that should trust the new identity provider.
 - Attach the metadata XML file exported from your federation agent to the service request as a file attachment.
2. From the same account where you created the service request, create a new RFC using CT-ID ct-1e1xtak34nx76 (Management | Other | Other | Create) with the following information.
 - Title: "Onboard SAML IDP <Name> for Account <AccountId>".
 - AccountId of the account where the identity provider will be created.
 - Identity provider name.
 - For Existing Accounts: Whether the identity provider should be propagated to all existing console roles, or the list of roles which should trust the new identity provider.
 - Case ID of service request created in Step 1, where the metadata XML file is attached.

Verify console access

Once you are set up with ADFS, and have the AMS URL to use for authentication, follow these steps.

With an Active Directory Federated Service (ADFS) configuration, you can follow these steps:

1. Open a browser window and go to the sign in page provided to you for your account. The ADFS **IdpInitiatedSignOn** page for your account opens.
2. Select the radio button next to **Sign in to one of the following sites**. The **Sign in** site picklist becomes active.
3. Choose the **signin.aws.amazon.com** site and click **Sign in**. Options for entering your credentials open.
4. Enter your CORP credentials and click **Sign in**. The AWS Management Console opens.
5. Paste into the location bar the URL of the AMS console and press **Enter**. The AMS console opens.

Verify API access

AMS uses the AWS API, with some AMS-specific operations that you can read about in the [AMS API Reference](#).

AWS provides several SDKs that you can access at [Tools for Amazon Web Services](#). If you don't want to use an SDK, you can make direct API calls. For information on authentication, see [Signing AWS API Requests](#). If you are not using an SDK, or making direct HTTP API requests, you can use the AMS CLIs for Change Management (CM) and SKMS.

Install the AMS CLIs

See [Appendix: ActiveDirectory Federation Services \(ADFS\) claim rule and SAML settings \(p. 202\)](#) for an example of installing the CLI to use with SAML.

Note

You must have administrator credentials for this procedure.

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

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

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

2. Once the AWS CLI is installed, to install or upgrade the AMS CLI, download either the **AMS CLI** or **AMS SDK** distributables zip file and unzip. You can access the AMS CLI distributables through the **Documentation** link in the left nav of the AMS console, or ask your cloud service delivery manager (CSDM) to send you the zip file.
3. The README file provides instructions for any install.

Open either:

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

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

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

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

4. If your corporate credentials are used via federation with AWS (the AMS default configuration) you must install a credential management tool that can access your federation service. For example, you can use this AWS Security Blog [How to Implement Federated API and CLI Access Using SAML 2.0 and AD FS](#) for help configuring your credential management tooling.
5. After the installation, run `aws amscm help` and `aws amsskms help` to see commands and options.

Scheduling backups at the VPC level

Backup scheduling in the VPC, where the target instances are allocated, is created during account onboarding with a default tag in the VPC creation schema. The backup system schedules the execution of the snapshots depending on that VPC Tag. Modification of the schedule can be made by creating a service request. For more information, see [VPC Tag and Defaults](#).

For backup defaults, see [Understanding AMS Defaults](#)

AMS Single-account landing zone Default Settings

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

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

For an example of infrastructure costs, see [Basic components](#).

Firewall rules are provided in [Set up your Firewall \(p. 107\)](#)

Endpoint Security (EPS)

Resources that you provision in your AMS Advanced environment automatically include the installation of an endpoint security (EPS) monitoring client. This process ensures that the AMS Advanced-managed resources are monitored and supported 24x7. In addition, AMS Advanced monitors all agent activity, and an incident is created if any security event is detected.

Note

Security incidents are handled as incidents; for more information, see [Incident response](#).

Endpoint Security provides anti-malware protection, specifically, the following actions are supported:

- EC2 instances register with EPS
- EC2 instances deregister from EPS
- EC2 instances real-time anti-malware protection
- EPS agent-initiated heartbeat
- EPS restore quarantined file
- EPS event notification
- EPS reporting

AMS Advanced uses Trend Micro for endpoint security (EPS). These are the default EPS settings. To learn more about Trend Micro, see the [Trend Micro Deep Security Help Center](#); note that non-Amazon links may change without notice to us.

AMS Advanced Multi-Account Landing Zone (MALZ) default settings are described in the following sections; for non-default AMS multi-account landing zone EPS settings, see [AMS Advanced Multi-Account Landing Zone EPS non-default settings](#).

Note

You can bring your own EPS, see [AMS bring your own EPS](#).

General EPS settings

Endpoint security general network settings.

EPS defaults

Setting	Default
Firewall Ports (Instances' Security Group)	EPS Deep Security Manager agents (DSMs) must have port 4120 open for the Agent/Relay to

Setting	Default
	Manager communication, and port 4119 for the Manager Console. EPS Relays must have port 4122 open for the Manager/Agent to Relay communication. No specific ports should be open for customer instance inbound communication because agents initiate all requests.
Communication Direction	Agent/Appliance Initiated
Heartbeat Interval	Ten minutes
Number of missed heartbeats before an alert	Two
Maximum allowed drift (difference) between server times	Unlimited
Raise offline errors for inactive (registered, but not online) virtual machines	No
Default policy	Base policy (described next)
Activation of multiple computers with the same host name	Is allowed
Alerts for pending updates are raised	After seven days
Update source	Trend Micro Update Server (https://ipv6-iaus.trendmicro.com/iau_server.dll/)
Event or log data deletion	Events and logs are deleted from the DSM database after seven days.
Agent software versions are held	Up to five
Most recent rule updates are held	Up to ten
Logs storage	By default, log files are stored securely in Amazon S3, but you can also archive them to Amazon Glacier to help meet audit and compliance requirements.

Base policy

Endpoint security base policy default settings.

EPS base policy

Setting	Default
Enabled Modules	Anti-Malware
Disabled Modules	Web Reputation
	Firewall
	Intrusion Protection
	Integrity Monitoring

Setting	Default
	Log Inspection
	Application Control

Anti-malware

Endpoint security anti-malware settings.

EPS anti-malware defaults

Setting	Default	Notes
Real-Time Scan	Scan everything	Quarantine all suspected viruses. Enable IntelliTrap and spyware/grayware protection. Spyware and Grayware trigger Anti-Malware and result in a quarantine of the item.
	Every Day/All Day (24 hours)	
Manual Scan	Scan everything	Must be requested, then follows default real-time scan configuration.
Scheduled Scan	Scan everything	Set for the last Sunday of every month, 6am.
Smart Protection	Disabled	N/A
Quarantined Files	Trend Micro Deep Security Manager (DSM)	Appx 1GB of disk reserved for quarantine.
Scan Limitation	Trend Micro DSM	Scan files of all sizes.
Allowed Spyware or Grayware	None	N/A
Local Event Notification	Yes	N/A

Security groups

In AWS VPCs, AWS Security Groups act as virtual firewalls, controlling the traffic for one or more stacks (an instance or a set of instances). When a stack is launched, it's associated with one or more security groups, which determine what traffic is allowed to reach it:

- For stacks in your public subnets, the default security groups accept traffic from HTTP (80) and HTTPS (443) from all locations (the internet). The stacks also accept internal SSH and RDP traffic from your corporate network, and AWS bastions. Those stacks can then egress through any port to the Internet. They can also egress to your private subnets and other stacks in your public subnet.
- Stacks in your private subnets can egress to any other stack in your private subnet, and instances within a stack can fully communicate over any protocol with each other.

Important

The default security group for stacks on private subnets allows all stacks in your private subnet to communicate with other stacks in that private subnet. If you want to restrict communications between stacks within a private subnet, you must create new security groups that describe the restriction. For example, if you want to restrict communications to a database server so that the stacks in that private subnet can only communicate from a specific application server over a specific port, request a special security group. How to do so is described in this section.

Default Security Groups

MALZ

The following table describes the default inbound security group (SG) settings for your stacks. The SG is named "SentinelDefaultSecurityGroupPrivateOnly-vpc-ID" where **ID** is a VPC ID in your AMS multi-account landing zone account. All traffic is allowed outbound to "mc-initial-garden-SentinelDefaultSecurityGroupPrivateOnly" via this security group (all local traffic within stack subnets is allowed).

All traffic is allowed outbound to 0.0.0.0/0 by a second security group "SentinelDefaultSecurityGroupPrivateOnly".

Tip

If you're choosing a security group for an AMS change type, such as EC2 create, or OpenSearch create domain, you would use one of the default security groups described here, or a security group that you created. You can find the list of security groups, per VPC, in either the AWS EC2 console or VPC console.

There are additional default security groups that are used for internal AMS purposes.

AMS default security groups (inbound traffic)

Type	Protocol	Port range	Source
All traffic	All	All	SentinelDefaultSecurityGroupPrivateOnly (restricts outbound traffic to members of the same security group)
All traffic	All	All	SentinelDefaultSecurityGroupPrivateOnlyEgressAll (does not restrict outbound traffic)
HTTP, HTTPS, SSH, RDP	TCP	80 / 443 (Source 0.0.0.0/0) SSH and RDP access is allowed from bastions	SentinelDefaultSecurityGroupPublic (does not restrict outbound traffic)
MALZ bastions:			
SSH	TCP	22	SharedServices VPC CIDR and DMZ VPC CIDR, plus Customer-provided on-prem CIDRs
SSH	TCP	22	
RDP	TCP	3389	
RDP	TCP	3389	
SALZ bastions:			
SSH	TCP	22	mc-initial-garden-LinuxBastionSG

Type	Protocol	Port range	Source
SSH	TCP	22	mc-initial-garden-LinuxBastionDMZSG
RDP	TCP	3389	mc-initial-garden-WindowsBastionSG
RDP	TCP	3389	mc-initial-garden-WindowsBastionDMZSG

SALZ

The following table describes the default inbound security group (SG) settings for your stacks. The SG is named "mc-initial-garden-SentinelDefaultSecurityGroupPrivateOnly-*ID*" where *ID* is a unique identifier. All traffic is allowed outbound to "mc-initial-garden-SentinelDefaultSecurityGroupPrivateOnly" via this security group (all local traffic within stack subnets is allowed).

All traffic is allowed outbound to 0.0.0.0/0 by a second security group "mc-initial-garden-SentinelDefaultSecurityGroupPrivateOnlyEgressAll-*ID*".

Tip

If you're choosing a security group for an AMS change type, such as EC2 create, or OpenSearch create domain, you would use one of the default security groups described here, or a security group that you created. You can find the list of security groups, per VPC, in either the AWS EC2 console or VPC console.

There are additional default security groups that are used for internal AMS purposes.

AMS default security groups (inbound traffic)

Type	Protocol	Port range	Source
All traffic	All	All	SentinelDefaultSecurityGroupPrivateOnly (restricts outbound traffic to members of the same security group)
All traffic	All	All	SentinelDefaultSecurityGroupPrivateOnlyEgressAll (does not restrict outbound traffic)
HTTP, HTTPS, SSH, RDP	TCP	80 / 443 (Source 0.0.0.0/0) SSH and RDP access is allowed from bastions	SentinelDefaultSecurityGroupPublic (does not restrict outbound traffic)
MALZ bastions:			
SSH	TCP	22	SharedServices VPC CIDR and DMZ VPC CIDR, plus Customer-provided on-prem CIDRs
SSH	TCP	22	
RDP	TCP	3389	
RDP	TCP	3389	
SALZ bastions:			
SSH	TCP	22	mc-initial-garden-LinuxBastionSG
SSH	TCP	22	mc-initial-garden-LinuxBastionDMZSG

Type	Protocol	Port range	Source
RDP	TCP	3389	mc-initial-garden-WindowsBastionSG
RDP	TCP	3389	mc-initial-garden-WindowsBastionDMZSG

Create, Change, or Delete Security Groups

You can request custom security groups. In cases where the default security groups do not meet the needs of your applications or your organization, you can modify or create new security groups. Such a request would be considered approval-required and would be reviewed by the AMS operations team.

To create a security group outside of stacks and VPCs, submit an RFC using the [Management | Other | Other | Create CT \(ct-1e1xtak34nx76\)](#).

To add or remove a user from an Active Directory (AD) security group, submit a request for change (RFC) using the [Management | Other | Other | Update CT \(ct-0xdawir96cy7k\)](#).

Note

When using manual (approval required) CTs, AMS recommends that you use the ASAP option (choose **ASAP** in the console, leave start and end time blank in the API/CLI) as these CTs require an AMS operator to examine the RFC, and possibly communicate with you before it can be approved and run. If you schedule these RFCs, be sure to allow at least 24 hours. If approval does not happen before the scheduled start time, the RFC is rejected automatically.

Find Security Groups

To find the security groups attached to a stack or instance, use the EC2 console. After finding the stack or instance, you can see all security groups attached to it.

For ways to find security groups at the command line and filter the output, see [describe-security-groups](#).

EC2 IAM instance profile

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

MALZ

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

Policy descriptions

Profile	Policies
customer-mc-ec2-instance-profile	AMSInstanceProfileLoggingPolicy: Allows Ec2 instances to push logs to S3 and CloudWatch.
	AMSInstanceProfileManagementPolicy: Allows Ec2 instances to perform booting actions, like joining Active Directory.
	AMSInstanceProfileMonitoringPolicy: Allows Ec2 instances to report findings to AMS monitoring services.

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
EC2 IAM instance profile

Profile	Policies
	AMSInstanceProfilePatchPolicy: Allows Ec2 instances to receive patches.
customer-mc-ec2-instance-profile-s3	AMSInstanceProfileLoggingPolicy: Allows Ec2 instances to push logs to S3 and CloudWatch.
	AMSInstanceProfileManagementPolicy: Allows Ec2 instances to perform booting actions, like joining Active Directory.
	AMSInstanceProfileMonitoringPolicy: Allows Ec2 instances to report findings to AMS monitoring services.
	AMSInstanceProfilePatchPolicy: Allows Ec2 instances to receive patches.
	AMSInstanceProfileS3WritePolicy: Allows Ec2 instances to read/write to customer S3 buckets.

SALZ

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

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

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

EC2 default IAM instance profile permissions

CW = CloudWatch. ARN = Amazon Resource Name. * = wildcard (any).			
Amazon Elastic Compute Cloud (Amazon EC2)			
EC2 Message Actions	Allow	AcknowledgeMessage, DeleteMessage, FailMessage, GetEndpoint, GetMessages, SendReply	Allows EC2 Systems Manager messaging actions in your account.
Ec2 Describe	Allow	* (All)	Allows the console to display configuration details of an EC2 in your account.
Iam Get Role ID	Allow	GetRole	Allows EC2 to get your IAM ID from <code>aws:iam::*:role/customer-*</code> and <code>aws:iam::*:role/customer_*</code> .

AMS Advanced Onboarding Guide AMS
 Advanced Account Onboarding Information
 EC2 IAM instance profile

CW = CloudWatch. ARN = Amazon Resource Name. * = wildcard (any).			
Instance To Upload Log Events	Allow	Create Log Group	Allows logs to be created in: aws:logs:*:*:log-group:i-*
		Create Log Stream	Allows logs to be streamed to: aws:logs:*:*:log-group:i-*
CW For MMS	Allow	DescribeAlarms, PutMetricAlarm, PutMetricData	Allows CloudWatch to retrieve alarms in your account. Allows CW to create or update an alarm and associate it with the specified metric. Allows CW to publish metric data points to your account.
Ec2 Tags	Allow	CreateTags, DescribeTags,	Allows tags to be added, overwritten, and described on the specified instances in your account.
Explicitly Deny CW Logs	Deny	DescribeLogStreams, FilterLogEvents, GetLogEvents	Disallows listing, filtering, or getting the log streams for: aws:logs:*:*:log-group:/mc/*
Amazon EC2 Simple Systems Manager (SSM)			
SSM Actions	Allow	DescribeAssociation, GetDocument, ListAssociations, UpdateAssociationStatus, UpdateInstanceInformation	Allows a variety of SSM functions in your account.
SSM Access In S3	Allow	GetObject, PutObject, AbortMultipartUpload, ListMultipartUploadPorts, ListBucketMultipartUploads	Allows the SSM on the EC2 to get and update objects in, and to abort a multi-part object upload to, and list ports and buckets available for, multi-part uploads in aws:s3::*:mc-*-internal-*/aws/ssm*.
Amazon EC2 Simple Storage Service (S3)			
Get Object In S3	Allow	Get List	Allows EC2 applications to retrieve and list objects in S3 buckets in your account.
Customer Encrypted Log S3 Access	Allow	PutObject	Allows EC2 applications to update objects in aws:s3::*:mc-*-logs-*/encrypted/app/*

CW = CloudWatch. ARN = Amazon Resource Name. * = wildcard (any).			
Patch Data Put Object S3	Allow	PutObject	Allows EC2 applications to upload patching data to your S3 buckets at <code>aws:s3:::awsms-a*-patch-data-*</code>
Uploading Own Logs To S3	Allow	PutObject	Allows EC2 applications to upload custom logs to: <code>aws:s3:::mc-a*-logs-*/aws/instances/*/\${aws:userid}/*</code>
Explicitly Deny MC Namespace S3 Logs	Deny	GetObject* Put*	Disallows EC2 applications getting or putting any objects from or to: <code>aws:s3:::mc-*-logs-*/encrypted/mc*</code> , <code>aws:s3:::mc-*-logs-*/mc/*</code> , <code>aws:s3:::mc-a*-logs-*-audit/*</code>
Explicitly Deny S3 Delete	Deny	* (all)	Disallows EC2 applications taking any action on objects in: <code>aws:s3:::mc-a*-logs-*/*</code> , <code>aws:s3:::mc-a*-internal-*/*</code> ,
Explicitly Deny S3 CFN Bucket	Deny	Delete*	Disallows EC2 applications deleting any objects from: <code>aws:s3:::cf-templates-*</code>
Explicitly Deny List Bucket S3	Deny	ListBucket	Disallows you listing any encrypted, audit log, or reserved (mc) objects from: <code>aws:s3:::mc-*-logs-*</code>

If you're unfamiliar with Amazon IAM policies, see [Overview of IAM Policies](#) for important information.

Note

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

Monitored Metrics Defaults

The following table shows what is monitored and the default alerting thresholds. You can change the defaults with a change management request for change (RFC).

Note

CloudWatch launched extended retention of metrics in November 1, 2016. For more information, see [CloudWatch Limits](#).

Alerts from baseline monitoring

Resource	Security alert	Alert name and trigger condition	Notes
<p>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).</p>			
Application Load Balancer (ALB) instance	No	RejectedConnectionCount sum > 0 for 1 min, 5 consecutive times.	CloudWatch alarm if the number of connections that were rejected because the load balancer reached its maximum.
Application Load Balancer (ALB) target	No	TargetConnectionErrorCount sum > 0 for 1 min, 5 consecutive times.	CloudWatch alarm if number of connections were unsuccessfully established between the load balancer and the registered instances.
		HTTPCode_Target_5XX_Count sum > 0 for 1 min, 5 consecutive times.	CloudWatch alarm on excess number of HTTP 5XX response codes generated by the targets.
Aurora instance	No	CPUUtilization > 85% for 5 mins, 2 consecutive times.	CloudWatch alarm.
EC2 instance - all OSs	No	CPUUtilization* >= 95% for 5 mins, 6 consecutive times.	CloudWatch alarm. High CPU utilization is an indicator of a change in application state such as dead locks, infinite loops, malicious attacks, and other anomalies.
		StatusCheckFailed > 0 for 5 minutes, 3 consecutive times.	CloudWatch alarm.
		Root Volume Usage >= 85% for 5 mins, 6 consecutive times.	
		Memory Free* MemoryFree < 5% for 5 minutes, 6 consecutive times.	
	Yes	EPS Malware Malware found on instance.	CloudWatch event.
Amazon EC2 instance - Linux	No	Root Volume Inode Usage Average >= 95% for 5 mins, 6 consecutive times.	CloudWatch alarm. Applied to Linux instances only.
		Swap Free*	

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Monitored Metrics Defaults

Resource	Security alert	Alert name and trigger condition	Notes
		Memory Swap < 5% for 5 minutes, 6 consecutive times.	
ElastiCache Cluster	No	CurrConnections = 65000	This alarm notifies AMS of the maximum connection limit of an ElastiCache Host. CloudWatch Alarm. If you would like to update this threshold, contact AMS support.
ElastiCache Node	No	CPUUtilization Average > predefined value for 15 mins, 2 consecutive times.	CloudWatch alarm. Default is 90. If Redis, use one the following values based on instance type: <ul style="list-style-type: none"> • cache.t1.micro: 90% • cache.m1.small: 90% • cache.m1.medium: 90% • cache.m1.large: 45% • cache.m1.xlarge: 22.5% • cache.m2.xlarge: 45% • cache.m2.4xlarge: 11.25% • cache.c1.xlarge: 11.25% • cache.t2.micro: 90% • cache.t2.small: 90% • cache.t2.medium: 45% • cache.m3.medium: 90% • cache.m3.large: 45% • cache.m3.xlarge: 22.5% • cache.m3.2xlarge: 11.25% • cache.r3.large: 45% • cache.r3.xlarge: 22.5% • cache.r3.2xlarge: 11.25% • cache.r3.4xlarge: 5.625% • cache.r3.8xlarge: 2.8125%
ElastiCache Node - memcached	No	SwapUsage maximum > 50,000,000 bytes for 5 mins, 5 consecutive times.	CloudWatch alarm. Applied to memcached only.
OpenSearch cluster	No	ClusterStatus.red maximum is >= 1 for 1 minute, 1 consecutive time. <i>AMS takes pro-active actions to reduce operational impact, when this alert is triggered.</i>	CloudWatch alarm. At least one primary shard and its replicas are not allocated to a node. To learn more, see Red Cluster Status .

AMS Advanced Onboarding Guide AMS
 Advanced Account Onboarding Information
 Monitored Metrics Defaults

Resource	Security alert	Alert name and trigger condition	Notes
OpenSearch domain	No	KMSKeyError >= 1 for 1 minute, 1 consecutive time.	CloudWatch alarm. The KMS encryption key that is used to encrypt data at rest in your domain is disabled. Re-enable it to restore normal operations. To learn more, see Encryption of Data at Rest for OpenSearch Service Service .
		ClusterStatus.yellow maximum is >= 1 for 1 minute, 1 consecutive time <i>AMS takes pro-active actions to reduce operational impact, when this alert is triggered.</i>	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 <i>AMS takes pro-active actions to reduce operational impact, when this alert is triggered.</i>	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 <i>AMS takes pro-active actions to reduce operational impact, when this alert is triggered.</i>	The cluster is blocking write requests. To learn more, see ClusterBlockException .
		Nodes minimum is < x for 1 day, 1 consecutive time <i>AMS takes pro-active actions to reduce operational impact, when this alert is triggered.</i>	x is the number of nodes in your cluster. This alarm indicates that at least one node in your cluster has been unreachable for one day. To learn more, see Failed Cluster Nodes .
		CPUUtilization average is >= 80% for 15 minutes, 3 consecutive times <i>AMS takes pro-active actions to reduce operational impact, when this alert is triggered.</i>	100% CPU utilization is common, but sustained high averages are problematic. Consider using larger instance types or adding instances.

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Monitored Metrics Defaults

Resource	Security alert	Alert name and trigger condition	Notes
		<p>JVMMemoryPressure</p> <p>maximum is $\geq 80\%$ for 5 minutes, 3 consecutive times</p> <p><i>AMS takes pro-active actions to reduce operational impact, when this alert is triggered.</i></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 is $\geq 50\%$ for 15 minutes, 3 consecutive times</p> <p><i>AMS takes pro-active actions to reduce operational impact, when this alert is triggered.</i></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 is $\geq 80\%$ for 15 minutes, 1 consecutive time</p> <p><i>AMS takes pro-active actions to reduce operational impact, when this alert is triggered.</i></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>
OpenSearch instance	No	<p>AutomatedSnapshotFailure</p> <p>maximum is ≥ 1 for 1 minute, 1 consecutive time.</p>	<p>CloudWatch alarm. An automated snapshot failed. This failure is often the result of a red cluster health status. See Red Cluster Status.</p>
Elastic Load Balancing instance	No	<p>SurgeQueueLength</p> <p>> 100 for 1 minute, 15 consecutive times.</p>	<p>CloudWatch alarm if an excess number of requests are pending routing.</p>
		<p>SpilloverCount</p> <p>> 1 for 1 minute, 15 consecutive times.</p>	<p>CloudWatch alarm if an excess number of requests that were rejected because the surge queue is full.</p>
GuardDuty service	Yes	<p>Not applicable; all findings (threat purposes) are monitored. Each finding corresponds to an alert.</p>	<p>List of supported GuardDuty finding types are on GuardDuty Active Finding Types.</p>
Health	Varies	<p>Changes in the GuardDuty findings. These changes include newly generated findings or subsequent occurrences of existing findings.</p>	<p>Notifications 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>

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Monitored Metrics Defaults

Resource	Security alert	Alert name and trigger condition	Notes
AWS Managed Microsoft AD	No	Active Directory Status AWS Managed Microsoft AD instance sends an active status event.	Service event. Emitted when the directory is operating normally after an event.
		Impaired Directory Status AWS Managed Microsoft AD instance sends an impaired directory status event.	Service event. Emitted when the directory is running in a degraded state. One or more issues have been detected, and not all directory operations may be working at full operational capacity.
		Inoperable Directory Status AWS Managed Microsoft AD instance sends an inoperable status event.	Service event. Emitted when the directory is not functional. All directory endpoints have reported issues.
		Deleting Directory Status AWS Managed Microsoft AD instance sends a deleting directory status event.	Service event. Emitted when the directory is currently being deleted.
		Failed Directory Status AWS Managed Microsoft AD instance sends a failed status event.	Service event. Emitted when the directory could not be created.
		RestoreFailed Directory Status AWS Managed Microsoft AD instance sends a restore failed directory status event.	Service event. Emitted when restoring the directory from a snapshot failed.
Amazon RDS instance	No	Failover not attempted Amazon RDS is not attempting a requested failover because a failover recently occurred on the DB instance.	Service event. RDS-EVENT-0034, Amazon RDS Event Categories and Event Messages .
		DB instance partial failover recovery complete The instance has recovered from a partial failover.	Service event. RDS-EVENT-0065, Amazon RDS Event Categories and Event Messages .
		DB instance fail The DB instance has failed due to an incompatible configuration or an underlying storage issue. Begin a point-in-time-restore for the DB instance.	Service event. RDS-EVENT-0031, Amazon RDS Event Categories and Event Messages .

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Monitored Metrics Defaults

Resource	Security alert	Alert name and trigger condition	Notes
		<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>	Service event. RDS-EVENT-0036, Amazon RDS Event Categories and Event Messages .
		<p>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 the customer action would be to modify the memory parameter and reboot the DB instance.</p>	Service event. RDS-EVENT-0035, Amazon RDS Event Categories and Event Messages .
		<p>Error create statspack user account</p> <p>Error while creating Statspack user account PERFSTAT. Drop the account before adding the Statspack option.</p>	Service event. RDS-EVENT-0058, Amazon RDS Event Categories and Event Messages .
		<p>DB instance without enhanced monitoring</p> <p>Enhanced Monitoring can't be enabled without the enhanced monitoring IAM role. For information about creating the enhanced monitoring IAM role, see To create an IAM role for Amazon RDS Enhanced Monitoring.</p>	Service event. RDS-EVENT-0079, Amazon RDS Event Categories and Event Messages .
		<p>DB instance enhanced monitoring disabled</p> <p>Enhanced Monitoring was disabled due to an error making the configuration change. It's likely that the enhanced monitoring IAM role is configured incorrectly. For information about creating the enhanced monitoring IAM role, see To create an IAM role for Amazon RDS Enhanced Monitoring.</p>	Service event. RDS-EVENT-0080, Amazon RDS Event Categories and Event Messages .
		<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>	Service event. RDS-EVENT-0081, Amazon RDS Event Categories and Event Messages .

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Monitored Metrics Defaults

Resource	Security alert	Alert name and trigger condition	Notes
		<p>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>Service event. RDS-EVENT-0045, Amazon RDS Event Categories and Event Messages.</p>
		<p>DB instance read replication ended</p> <p>Replication on the Read Replica was ended.</p>	<p>Service event. RDS-EVENT-0057, Amazon RDS Event Categories and Event Messages.</p>
		<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, Amazon RDS Event Categories and Event Messages.</p>
		<p>Low Storage alert triggers when the allocated storage for the DB instance has been exhausted.</p>	<p>RDS-EVENT-0007, see details at Using Amazon RDS event notification.</p>
		<p>Low storage alert when the DB instance has consumed more than 90% of its allocated storage</p>	<p>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>RDS-EVENT-0143, see details at Amazon RDS Event Categories and Event Messages.</p>
		<p>CPUUtilization</p> <p>Average CPU utilization > 75% for 15 mins, 2 consecutive times.</p>	<p>CloudWatch alarm.</p>
		<p>DiskQueueDepth</p> <p>Sum is > 75 for 1 mins, 2 consecutive times.</p>	
		<p>FreeStorageSpace</p> <p>Average < 1,073,741,824 bytes for 5 mins, 2 consecutive times.</p>	

Resource	Security alert	Alert name and trigger condition	Notes
		ReadLatency Average \geq 1.001 seconds for 5 mins, 2 consecutive times.	
		WriteLatency Average \geq 1.005 seconds for 5 mins, 2 consecutive times.	
		SwapUsage Average \geq 104,857,600 bytes for 5 mins, 2 consecutive times.	
Amazon Redshift cluster	No	HealthStatus The health of the cluster \leq 0 for 5 min, 1 consecutive times.	1 represents a healthy cluster.
		MaintenanceMode Cluster maintenance mode \geq 1 for 5 min, 1 consecutive time.	1 represents ON state.
		ReadLatency The average time for disk read \geq 1 for 5 min, 1 consecutive time.	None.
		WriteLatency The average time for disk write \geq 1 for 5 min, 1 consecutive time.	None.
Amazon Macie	Yes	Newly generated alerts and updates to existing alerts. Macie finds any changes in the findings. These changes include newly generated findings or subsequent occurrences of existing findings.	Amazon Macie alert. For a list of supported Macie alert types, see Macie Alerts . Note that Macie is not configured for all accounts.

Log retention and rotation defaults

This section describes AMS log management defaults; for more information, see [Log Management](#).

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

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

Important

By default, EC2 stack backups are disabled (Backup = False). You can enable EC2 instance backups at the time of creation by adding a tag `Key: Backup, Value: True` when requesting an EC2 stack through an RFC (CT ct-14027q0sjyt1h). If you want to add the tag after the instance has been created, submit an RFC with the Management | Other | Other | Update CT (ct-0xdawir96cy7k) and specify the instances that you want to have the backup tag added to.

EC2 Instance Tag and Defaults

The **EC2 stack backup tag** specifies whether the stack requires a snapshot of the attached EBS volumes or not.

Tag Key: Backup

Tag Value: True, False

By default, the value is `False` the backup tag is not present, and the stack does not have scheduled backups.

Change the tag `Key: Backup` to `Value: True` to enable backups, which are then done on the schedule set with the VPC backup tag.

Note

The casing for the tag value (Value only) is insensitive, so `True/true` or `False/false` are all acceptable.

RDS Instance Backup and Defaults

The Amazon Relational Database Service (RDS) default values are defined in the stack templates:

Backup: Yes

Backup Window: 22:00-23:00 (RDS local time zone)

Retention Period: 7 (7 snapshots stored)

Patching Defaults

This section describes AMS patching defaults; for more information on AMS patching, see the AMS User Guide Patch Management chapter.

AMS releases patched AMIs on a monthly basis; all new stack requests should be configured with the latest AMS AMI.

Important

AMS Patch Orchestrator, tag-based patching, uses AWS Systems Manager (SSM) functionality to allow you to tag, or have AMS tag for you, instances and have those instances patched using a baseline and a window that you configure. To learn more, see [Patch Orchestrator: a tag-based patching model](#).

AMS-standard, account-based, patching: For each account with stacks that receive in-place patching, a notification of upcoming applicable patches is sent out shortly after “patch Tuesday”. The notification contains a list of all stacks and the applicable patches as well as the suggested patch window. For critical patches, the window is set no longer than 10 days in advance, and for standard patching no more than 14 days in advance. If you do not reply to the notification, patching does not occur. If you would like to exclude certain patches, reply to the notification, or submit a service request. If you reply with consent to patching, but don’t specifically request a different schedule, patches are applied as described in the notification that you receive.

Note

The patch service notification is an email sent to the account contacts and contains a link to the AWS Support console. You can reply through the AWS Support console or through the AMS service request page, where the notification appears as a service notification.

At the time of the AMS-standard patching process, AMS performs the following:

1. You are sent a patching service notification fourteen days before the proposed patch window. The patching service notification is sent via email to the contact email address that you have on file for your account.
2. Identifies all reachable EC2 instances in the stack based on the list of stacks provided in the patching notification. In this case, "Reachable" means instances that are in the "Running" EC2 state, and have the EC2 Run Command agent fully operational.
3. AMS performs patching in a manner that ensures that a sufficient number of EC2 instances are running concurrently (configured through the `healthy-host-threshold` setting) so that the stack remains healthy.
4. After the patching operation is complete for all EC2 instances, AMS updates the RFC with the patching status: Success, Partial Success or Failure. In the case of any status other than Success, a ticket is created for an operator to follow up on the patching results and take any corrective actions.

Understand Change Types

You use change types to request specific resources and resource updates, and access to resources.

Change Type Validation

Validation exercises require submitting RFCs with change types (CTs), some of which use stack templates. This section provides an overview of AMS change types and stack templates with links to more information about the resources that they create. Also included are some simple examples of creating and EC2 stack, obtaining access to it,

To learn more about change types (CTs), see [Understanding Change Types](#) and the [AMS Change Type Reference](#).

All CTs have a version number and an identifier (ID) and are associated with at least one category-subcategory-item-operation (CSIO). Because change types are now referred to by their change type name and ID, they can belong in multiple CSIOs and you can incorporate CTs into your own IT service management (ITSM) system.

Note

For the most up-to-date change type list run this change management command:

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

For details on change types, including which are manual, see the [AMS Change Type Reference](#) for syntax and metadata, and see [Understanding Change Types](#) for additional details and walkthrough examples.

Validate the AMS Service

To validate that the AWS Managed Services (AMS) service is working as expected, some exercise that you can do are described in this chapter.

Finding your account settings

Specific to every account are certain settings that are used to create RFCs, set schedules, and determine who receives notifications. This section describes how to find out what those settings are.

Some settings are created during onboarding and require a service request to change. You should make a note of these account details because you will use them when communicating with AMS:

- **Credentials:** If you need to retrieve your AMS user name or password, contact your local IT administrator--AMS uses your corporate Active Directory.
- **Cloud Service Delivery Manager (CSDM):** This person is your liaison with AMS and is available to answer service questions. You are given this person's contact information at onboarding and should keep it available to all in your organization who interact with AMS. You can expect to receive monthly reports on your AMS service from this person.
- **Console access:** You access the AMS console at a URL set up specifically for your account. You can get the URL from your CSDM.
- **AMS CLI:** You can obtain the AMS CLI through the AMS console, or the distributables package that you get from your CSDM. After you have the distributables package, follow the steps outlined in [Installing or upgrading the AMS CLI](#).
- **Maintenance window:** Your maintenance window determines when patching happens for your EC2 instances. The AWS Managed Services Maintenance Window (or Maintenance Window) performs maintenance activities for AWS Managed Services (AMS) and recurs the second Thursday of every month from 3 PM to 4 PM Pacific Time. AMS may change the maintenance window with 48 hours notice. You may have chosen a different window at onboarding--keep a record of your chosen maintenance window.
- **Monitoring:** AMS provides a set of CloudWatch metrics by default, but you can also request additional metrics. If you do, keep record of those.
- **Logs:** By default, your logs are stored at `ams-a-ACCOUNT_ID-log-management-REGION` where `REGION` is the region where the log was generated.
- **Mitigation:** At onboarding, AMS records the mitigation action of your choice in case a malware attack against your resources is identified. For example, contact certain people. Keep this information available to all in your organization who interact with AMS.
- **Region:** You can look at the VPC details page in the AMS console. You can also run this command (uses a SAML profile, remove if your authentication method is different):

```
aws --profile saml amsskms get-vpc --vpc-id VPC_ID
```

Finding your FQDN

Access CTs require the fully qualified domain name, or FQDN, of your AMS-trusted domain, in the form of `C844273800838.amazonaws.com`. To discover your FQDN:

- **Console:** Look in the AWS Directory Service console (under Security, Identity, and Compliance category) in the **Directory name** column
- **API/CLI:** Use these commands while logged into your domain:

Windows (returns user and FQDN):

```
whoami /upn
```

or (DC+DC+DC=FQDN)


```
whoami /fqdn
```

Linux:

```
hostname --fqdn
```

Finding your availability zones

Availability Zone: All accounts have at least two availability zones. To accurately find your availability zone names, you must first know the associated subnet ID.

- Console: In the navigation pane click **Network**, and then click the relevant VPC, if necessary. The **Network** page includes a table of subnets. Select the relevant subnet to open the subnet details page with the name of the associated availability zone.
- API/CLI:

```
aws amsskms list-subnet-summaries --output table
```

```
aws amsskms get-subnet --subnet-id SUBNET_ID
```

Finding your SNS settings

To discover your SNS topics:

- Console: Use the SNS console to view all topics, applications, and subscriptions, and a graph of messages. Also create, delete, subscribe to, and publish to topics.
- API/CLI (when logged into your AMS account):

List your SNS topics:

```
aws sns list-topics
```

List your SNS subscriptions:

```
aws sns list-subscriptions
```

Finding your backup settings

Backups and Snapshots are managed by AMS through the native [AWS Backup](#) service.

The configuration is managed through AWS Backup plans. You can have multiple AWS Backup plans that associate tagged resources with backup schedules and retention policies. For more information, see [Continuity Management](#).

Finding an Instance ID or IP Address

You need an instance IP address to log into the instance.

- To request access to an instance, to log in to an instance, or to create an AMI, you must have the instance ID. For an EC2 instance (either a standalone instance or a part of a stack), or a database instance, you can find the ID in a few different ways:
 - The AMS Console for an instance in an ASG stack: Look on the RFC detail page for the RFC that created the stack. In the Execution Output section, you will find the stack ID for the ASG stack and you can then go to the EC2 Console **Auto Scaling Groups** page and search for that stack ID and find instances for it. When you find the instance, select it and an area opens at the bottom of the page with details, including the IP address.
 - The AMS Console for a standalone EC2 or database (DB) instance: Look on the RFC detail page for the RFC that created the EC2 stack or DB instance. In the Execution Output section, you will find the Instance ID and IP address.
- AWS EC2 Console:
 1. In the navigation pane, select **Instances**. The **Instances** page opens.
 2. Click the instance that you want the ID for. The instance details page opens and displays the ID and IP address.
- AWS Database Console:
 1. On the Home page, select **DB Instances**. The **Instances** page opens.
 2. Filter for the DB instance that you want the ID for. The instance details page opens and displays the ID.
- AMS CLI/API.

Note

The AMS CLI must be installed for these commands to work. To install the AMS API or CLI, go to the AMS console **Developers Resources** page. For reference material on the AMS CM API or AMS SKMS API, see the AMS Information Resources section in the User Guide.

Run the following command to get stack execution output details:

```
aws amsskms get-stack --stack-id STACK_ID
```

The output looks similar to this with the InstanceId appearing near the bottom, under Outputs (values shown are examples):

```
{
  "Stack": {
    "StackId": "stack-7fa52bd5eb8240123",
    "Status": {
      "Id": "CreateCompleted",
      "Name": "CreateCompleted"
    },
    "VpcId": "vpc-01234567890abcdef",
    "Description": "Amazon",
    "Parameters": [
      {
        "Value": "sg-01234567890abcdef,sg-01234567890abcdef",
        "Key": "SecurityGroups"
      },
      {
        "Value": "subnet-01234567890abcdef",
        "Key": "InstanceSubnetId"
      },
      {
        "Value": "t2.large",
        "Key": "InstanceType"
      },
      {
        "Value": "ami-01234567890abcdef",

```

```
        "Key": "InstanceAmiId"
      }
    ],
    "Tags": [],
    "Outputs": [
      {
        "Value": "i-0b22a22eec53b9321",
        "Key": "InstanceId"
      },
      {
        "Value": "10.0.5.000",
        "Key": "InstancePrivateIP"
      }
    ],
    "StackTemplateId": "stm-s6xvs000000000000",
    "CreatedTime": "1486584508416",
    "Name": "Amazon"
  }
}
```

DNS Friendly Bastion Names

MALZ

For Multi-account landing zone (MALZ), DNS records are created for the bastions in the FQDN of the AMS-managed Active Directory. AMS replaces Linux and Windows bastions as required. For example, if there is a new bastion AMI that must be deployed, the bastion DNS records dynamically update to point to new, valid bastions.

1. To access SSH (Linux) bastions, use DNS records like this:
sshbastion(1-4).Your_Domain.com

For example, where the domain is Your_Domain:

- sshbastion1.Your_Domain.com
- sshbastion2.Your_Domain.com
- sshbastion3.Your_Domain.com
- sshbastion4.Your_Domain.com

2. To access RDP (Windows) bastions, use DNS records like this:
rdp-Username.Your_Domain.com.

For example, where the user name is alex, test, demo, or bob, and the domain is Your_Domain.com:

- rdp-alex.Your_Domain.com
- rdp-test.Your_Domain.com
- rdp-demo.Your_Domain.com
- rdp-bob.Your_Domain.com

SALZ

Single-account landing zone (SALZ) replaces Linux and Windows bastions as required. For example, if there is a new bastion AMI that must be deployed, the bastion DNS records dynamically update to point to new, valid bastions.

1. To access SSH (Linux) bastions, use DNS records like this:
`sshbastion(1-4).AACCOUNTNUMBER.amazonaws.com`.

For example, where 123456789012 is the account number:

- `sshbastion1.A123456789012.amazonaws.com`
- `sshbastion2.A123456789012.amazonaws.com`
- `sshbastion3.A123456789012.amazonaws.com`
- `sshbastion4.A123456789012.amazonaws.com`

2. To access RDP (Windows) bastions, use DNS records like this:
`rdpbastion(1-4).AACCOUNT_NUMBER.amazonaws.com`.

For example, where 123456789012 is the account number:

- `rdpbastion1.A123456789012.amazonaws.com`
- `rdpbastion2.A123456789012.amazonaws.com`
- `rdpbastion3.A123456789012.amazonaws.com`
- `rdpbastion4.A123456789012.amazonaws.com`

Finding Bastion IP Addresses

AMS customers can use SSH and RDP bastions, either the [DNS Friendly Bastion Names \(p. 142\)](#) described previously, or bastion IP addresses.

To find bastion IP addresses, SSH and RDP, for your account:

1. For multi-account landing zone only: Log in to the Shared Services account.
2. Open the EC2 Console and choose **Running Instances**.

The **Instances** page opens.

3. In the filter box at the top, enter either **ssh-bastion** or **rdp-bastion**.

In the filter box at the top, enter either **customer-ssh** or **customer-rdp**.

The SSH and/or RDP bastions for your account display.

Note that in addition to your SSH bastions, you may see AMS perimeter network bastions in the list, which are unavailable for this.

4. Select an SSH or RDP bastion. If you're using a Windows computer and want to log in to a Linux instance, you use an SSH bastion. If you want to log in to a Windows instance, you use an RDP bastion. If you're on a Linux OS and want to log in to a Windows instance, you use an SSH bastion through an RDP tunnel (this is so you can access the Windows desktop). To access a Linux instance from a Linux OS, you use an SSH bastion.

EC2 Instance, Creating

You can use the AMS console or API/CLI to create an Amazon EC2 and an Amazon EC2 with additional volumes.

Create EC2 stack instance

Create an Amazon Elastic Compute Cloud, EC2 instance, using the AMS console or the AMS API/CLI.

Classification and CT ID: Deployment | Advanced stack components | EC2 stack | Create

Change type ID: ct-14027q0sjyt1h

Version: 3.0

Note

Starting with version 3.0 of this change type, AMS does not attach the default AMS security groups if you specify your own security groups. If you do not specify your own security groups in the request, AMS attaches the AMS default security groups. In previous versions, AMS attached the default security groups whether or not you provided your own security groups.

The EC2 instance you create comes with default alarms and security groups:

- Alarms (for details see [Alerts from Baseline Monitoring in AMS](#)):
 - CPU Too High
 - EPS DSM Communication
 - Greatest iowait
 - Log Agent HardFailure
 - Memory Free
 - Root Volume Inode Usage
 - Root Volume Usage
 - Swap Free
 - System Status
- **Default Security Groups:**
 - InitialGarden-SentinelDefaultSecurityGroupPrivateOnly-*ID*
 - InitialGarden-SentinelDefaultSecurityGroupPrivateOnlyEgressAll-*ID*

For more information about Amazon EC2, including size recommendations, see [Amazon Elastic Compute Cloud Documentation](#).

To update your EC2 stack after it's created, see [EC2 Stack: Updating](#) .

Note

To create an EC2 stack with additional volumes, see [Create EC2 stack instance with additional volumes \(p. 147\)](#).

Required data:

- **Subject:** A title for the request.
- **Description:** A reason for the request.
- **Name:** A name for the stack or stack component; this becomes the Stack Name.
- **VpcId:** The VPC to use. For information about finding VPC IDs, see [Find VPC](#).
- **TimeoutInMinutes:** The number of minutes allowed for the creation of the stack before the RFC fails. This setting will not delay the RFC execution, but you must give enough time (for example, don't specify "5"). Valid values are "60" up to "360," for long-running UserData.
- **Parameters:**
 - **InstanceAmiId:** The AMI to use to create the EC2 instance. We recommend using the most recent AMI that begins with "customer-". For information about finding AMIs, see [Finding an AMI](#).

To use the CentOS AMIs, you must opt in to the Cent OS license from the AWS Marketplace. Either submit a Service Request to AMS to subscribe. Or, go to AWS Marketplace and follow the instructions for re-opting-in. You do not incur software charges for using this product, but you're responsible for other AWS charges, including EC2 usage.

- **InstanceSubnetId:** The subnet that you want to launch the instance into. For information about finding subnet IDs, see [Find Subnet](#).

Optional data (available with the **Additional configuration** view):

Note

You can add up to 50 tags, but to do so you must enable the **Additional configuration** view.

- **InstanceDetailedMonitoring:** **True** to enable detailed monitoring on the instance, **false** to use only basic monitoring. Default is false.
- **InstanceEBSOptimized:** **True** for the instance to be optimized for Amazon Elastic Block Store I/O, **false** for it to not be. If you set this to **true**, choose an **InstanceType** that supports EBS optimization. Default is **false**, which means that you get basic EBS storage.
- **InstanceProfile:** An IAM instance profile defined in your account for the EC2 instance. The default is the AWS-provided role, **customer-mc-ec2-instanceprofile**.
- **InstanceRootVolumeIops:** The IOPS (Input/Output Operations Per Second) to use for the root volume, if the volume type is io1, io2 or gp3. The default is 100 for io1 or io2 volume type, whereas it is 3,000 if the volume type is gp3.
- **InstanceRootVolumeName:** The name of the root volume to use. The default is **/dev/xvda** for Linux, and **/dev/sda** for Windows.
- **InstanceRootVolumeSize:** The size of the root volume for the instance. The default is **8 GiB** for Linux, and **30 GiB** for Windows.
- **InstanceRootVolumeType:** Choose io1, io2, gp2, or gp3 for SSD-backed volumes optimized for transactional workloads. Choose standard for HDD-backed volumes suitable for workloads where data is infrequently accessed. The default is gp3.
- **InstanceType:** The type of EC2 instance to deploy. If **InstanceEBSOptimized** = true, specify an **InstanceType** that supports EBS optimization. The default is **t2.large**. NOTE: EC2 instances need enough capacity to support AMS tools such as EPS, SSM, and Cloudwatch in addition to the application workload. AMS does not recommend the **t2.micro/t3.micro** and **t2.nano/t3.nano** types. These are smaller instance types, and can degrade the performance of your application and AMS tools. For more information, see [Choosing the Right EC2 Instance Type for Your Application](#).
- **InstanceUserData:** A newline-delimited list where each element is a line of script to be run on boot. For a new line, press ENTER.
- **SecurityGroupIds:** IDs of existing custom security groups to associate with the instance, in the form sg-0123abcd or sg-01234567890abcdef. Up to three custom security groups may be specified. If nothing is specified, the default AMS security groups are applied.

Note

Currently, if you specify custom security groups, you must also specify the IDs of the default AMS security groups for your account, **mc-initial-garden-SG-name** and **mc-initial-garden-SG-name**.

Creating an EC2 instance with the console

The following shows this change type in the AMS console.

▼ **Change type: Create EC2 stack**

Description
Use to create an Amazon Elastic Compute Cloud (EC2) instance.

ID	Version
ct-14027q0sjyt1h	3.0

Execution mode
Automated

How it works:

1. Navigate to the **Choose change type** page: **RFCs -> Create RFC**.
2. Choose a change type from the drop-down lists. Optionally, open the **Additional configuration** area to select a change type version. After your selections are complete, a **Change type: details** area opens. Choose **Next**.
3. Configure the request for change. A **Subject** is required. Optionally, open the **Additional configuration** area to add information about the RFC. Choose **Next**.
4. Choose the execution parameters. At the top, in the **RFC configuration** area, enter values for the change type required parameters. These vary by change type. Open the **Additional configuration** area to add Tags or additional settings. Some change types also provide a **Parameters** area where only the required settings are visible. In that case, open the **Additional configuration** area to view optional parameters.
5. When finished, choose **Create**. If there are no errors, the **RFC successfully created** page displays with the submitted RFC details, and the initial **Execution output**.
6. Open the **Execution parameters** area to see the configurations you submitted. Refresh the page to update the RFC execution status. Optionally, cancel the RFC or create a copy of it with the options at the top of the page.

Creating an EC2 instance with the CLI

How it works:

1. Use either the Inline Create (you issue a `create-rfc` command with all RFC and execution parameters included), or Template Create (you create two JSON files, one for the RFC parameters and one for the execution parameters) and issue the `create-rfc` command with the two files as input. Both methods are described here.
2. Submit the RFC: `aws amscm submit-rfc --rfc-id ID` command with the returned RFC ID.

Monitor the RFC: `aws amscm get-rfc --rfc-id ID` command.

To check the change type version, use this command:

```
aws amscm list-change-type-version-summaries --filter Attribute=ChangeTypeId,Value=CT_ID
```

Note

You can use any `CreateRfc` parameters with any RFC whether or not they are part of the schema for the change type. For example, to get notifications when the RFC status changes, add this line, `--notification "{\"Email\": {\"EmailRecipients\": [\"email@example.com\"]}}\"` to the RFC parameters part of the request (not the execution parameters). For a list of all `CreateRfc` parameters, see the [AMS Change Management API Reference](#).

INLINE CREATE:

Issue the create RFC command with execution parameters provided inline (escape quotation marks when providing execution parameters inline), and then submit the returned RFC ID. For example, you can replace the contents with something like this:

```
aws amscm create-rfc --change-type-id "ct-14027q0sjyt1h" --change-type-version "3.0" --title "EC2-Create-RFC" --execution-parameters "{\"Description\": \"Create a new EC2 Instance stack\", \"VpcId\": \"vpc-0a60eb65b4EXAMPLE\", \"Name\": \"My-EC2\", \"TimeoutInMinutes\": 60, \"Parameters\": {\"InstanceAmiId\": \"ami-1234567890EXAMPLE\", \"InstanceDetailedMonitoring\": false, \"InstanceEBSOptimized\": false, \"InstanceProfile\": \"customer-mc-ec2-instance-profile\", \"InstanceRootVolumeIops\": 3000,
```

```
\"InstanceRootVolumeType\": \"gp3\", \"InstanceType\": \"t2.large\", \"InstanceUserData\":  
\"\", \"InstanceSubnetId\": \"subnet-0bb1c79de3EXAMPLE\"}}"
```

TEMPLATE CREATE:

1. Output the execution parameters for this change type to a JSON file; this example names it CreateEC2Params.json:

```
aws amscm get-change-type-version --change-type-id "ct-14027q0sjyt1h" --query  
"ChangeTypeVersion.ExecutionInputSchema" --output text > CreateEC2Params.json
```

2. Modify and save the CreateEC2Params file. For example, you can replace the contents with something like this:

```
{  
  "Description": "Create a new EC2 Instance stack",  
  "VpcId": "vpc-0a60eb65b4EXAMPLE",  
  "Name": "My-EC2",  
  "TimeoutInMinutes": 60,  
  "Parameters": {  
    "InstanceAmiId": "ami-1234567890EXAMPLE",  
    "InstanceDetailedMonitoring": false,  
    "InstanceEBSOptimized": false,  
    "InstanceProfile": "customer-mc-ec2-instance-profile",  
    "InstanceRootVolumeIops": 3000,  
    "InstanceRootVolumeType": "gp3",  
    "InstanceType": "t2.large",  
    "InstanceUserData": "",  
    "InstanceSubnetId": "subnet-0bb1c79de3EXAMPLE"  
  }  
}
```

3. Output the RFC template to a file in your current folder; this example names it CreateEC2Rfc.json:

```
aws amscm create-rfc --generate-cli-skeleton > CreateEC2Rfc.json
```

4. Modify and save the CreateEC2Rfc.json file. For example, you can replace the contents with something like this:

```
{  
  "ChangeTypeVersion": "3.0",  
  "ChangeTypeId": "ct-14027q0sjyt1h",  
  "Title": "EC2-Create-RFC"  
}
```

5. Create the RFC, specifying the CreateEC2Rfc file and the CreateEC2Params file:

```
aws amscm create-rfc --cli-input-json file://CreateEC2Rfc.json --execution-parameters  
file://CreateEC2Params.json
```

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

If needed, see [EC2 instance stack create fail](#).

Create EC2 stack instance with additional volumes

Create an Amazon Elastic Compute Cloud, EC2 instance, with up to five additional volumes, using the AMS console or the AMS API/CLI.

Classification: Deployment | Advanced stack components | EC2 stack | Create (with additional volumes)

Change type ID: ct-1aqsjf86w6vxg

Version: 4.0

For more information about Amazon EC2, including size recommendations, see [Amazon Elastic Compute Cloud Documentation](#).

To update your EC2 stack with additional volumes after they're created, see [EC2 stack: updating with additional volumes](#)

Important

There is a new version of this change type, v 4.0, that uses a different StackTemplateId (stm-nn8v8ffhcal611bmo). This is important if you're submitting the RFC with this change type at the command line. The new version introduces two new parameters (**RootVolumeKmsKeyId** and **CreditSpecification**) and changes the default for one existing parameter (**InstanceType**).

Required data:

- **Description:** A reason for the request.
- **VpcId:** The VPC to use. For information about finding VPC IDs, see [Find VPC](#).
- **Name:** A name for the stack or stack component. You must use a "Name" tag in order for the instance to have a name in the EC2 console **Instances** details page.
- **TimeoutInMinutes:** The number of minutes allowed for the creation of the stack before the RFC fails. This setting won't delay the RFC execution, but you must give enough time (for example, don't specify "5"). Valid values are "60" up to "360," for long-running UserData.
- **Parameters:**
 - **InstanceAmiId:** The AMI to use to create the EC2 instance. We recommend using the most recent AMS AMI that begins with "customer-". For help finding AMIs, see [Finding an AMI](#).
 - **InstanceSubnetId:** The subnet that you want to launch the instance into. For information about finding subnet IDs, see [Find Subnet](#).

Optional data:

When using the AMS console, to see the optional parameters, you must enable the **Additional configuration** view.

- **Tags:** Up to 40 tags (key/value pairs) to categorize the EC2.
- **RootVolumeKmsKeyId:** The ID, or ARN, of the KMS master key to be used to encrypt the root volume. Specify default to use the default EBS KMS Key. Leave blank to not encrypt the root volume. Note that, if a value is set, the InstanceRootVolumeName must also be specified for KMS encryption settings on the root volume to take effect.
- **CreditSpecification:** The credit option for CPU usage. This is only supported with t2, t3, and t3a, instance types. If your instance is unlikely to require CPU bursting, choose standard, but note that, once all the CPU credits for that instance are used up, it will be throttled. For better burst handling, and to not allow throttling, choose unlimited, but note that additional charges may apply when additional credits are used.
- **InstanceTerminationProtection:** True to prevent the instance from being terminated through the API, false to allow it. Default is false. Termination protection must be disabled with an update (ct-1o1x2itfd6rk8) before deleting the stack or performing an update where instance replacement is required, otherwise failures occur.
- **InstanceType:** The type of EC2 instance to deploy. If **InstanceEBSOptimized** = true, specify an **InstanceType** that supports EBS optimization. Default is **t3.large**. NOTE: EC2 instances need enough capacity to support AMS tools such as EPS, SSM, and Cloudwatch in addition to the application

workload. AMS does not recommend the **t2.micro/t3.micro** and **t2.nano/t3.nano** types. These are smaller instance types, and can degrade the performance of your application and AMS tools. For more information, see [Choosing the Right EC2 Instance Type for Your Application](#). In version 4.0, the default type was raised from **t2.large** to **t3.large**. T3 instances launch with 'unlimited credits' by default. You won't experience CPU throttling even if the instance consumes all CPU credits. You can, instead, choose to launch T2 instances and use the **CreditSpecification** unlimited option.

- **InstanceSecurityGroupIds**: IDs for the security groups you want to attach to the instance. These security groups control access to the EC2 instance. Default AMS security groups are attached to the instance in addition to the security groups specified here.
- **InstanceDetailedMonitoring**: **True** to enable detailed monitoring on the instance, **false** to use only basic monitoring. The default is false.
- **InstanceEBSOptimized**: **True** for the instance to be optimized for Amazon Elastic Block Store I/O, **false** for it to not be. If you set this to **true**, choose an **InstanceType** that supports EBS optimization. Default is **false**, which means that you get basic EBS storage.
- **InstanceProfile**: An IAM instance profile defined in your account for the EC2 instance. The default is **customer-mc-ec2-instanceprofile**.
- **InstanceRootVolumeName**: The name of the root volume to use. Default is **/dev/xvda** for Linux, and **/dev/sda** for Windows.
- **InstanceRootVolumeSize**: The size of the root volume for the instance. The default is **8 GiB** for Linux, and **30 GiB** for Windows.
- **InstanceRootVolumeIops**: The IOPS to use for the root volume, if io1 volume type is specified. The default is **0**.
- **InstanceRootVolumeType**: Choose SSD-backed volumes optimized for transactional workloads (io1, gp2 and gp3), or HDD-backed volumes optimized for large streaming workloads (standard). The default is standard. The additional volumes support all the supported EBS types (gp2, gp3, io1, sc1, st1 and standard), encryption using KMS, and creation using snapshots.
- **InstancePrivateStaticIp**: The static IP address that the instance can support.
- **InstanceSecondaryPrivateIpAddressCount**: The number of secondary private IP addresses that EC2 automatically assigns to the primary network interface. The number of secondary IP addresses that can be assigned is dependent on the type of instance used.
- **InstanceType**: See previous **Important** note.
- **InstanceUserData**: A newline-delimited list where each element is a line of script to be run on boot. To separate lines, use the literal: `"\n"`.
- **Volume1-5Encrypted**: True if the volume is encrypted. False if it is not.
- **Volume1-5Iops**: The IOPS to use for the volume if the type = io1.
- **Volume1-5Throughput**: The throughput to use for the volume volume, if the volume type = gp3. If the volume type is not gp3, any value provided here is ignored. Default is 125.
- **Volume1-5KmsKeyId**: Amazon Resource Name (ARN) of the KMS master key to be used to encrypt the volume.
- **Volume1-5Name**: The device name for the volume (for example, **/dev/sdh** or **xvdh**).
- **Volume1-5Size**: The size of the volume in GiB. The default is 1 GiB.
- **Volume1-5Snapshot**: Snapshot ID for the volume.
- **Volume1-5Type**: The type for the volume. Choose io1, gp2 or gp3 for SSD-backed volumes optimized for transactional workloads. Choose sc1 or st1 for HDD-backed volumes optimized for large streaming workloads. Choose standard for HDD-backed volumes suitable for workloads where data is infrequently accessed.

Creating an EC2 instance and additional volumes with the console

The following shows this change type in the AMS console.

Create EC2 Stack With Additional Volumes Modify version

Description
Create an Amazon Elastic Compute Cloud (EC2) instance with up to five additional volumes.

ID	Version
ct-1aqsjf86w6vxg	4.0 (most recent version)

How it works:

1. Navigate to the **Choose change type** page: **RFCs -> Create RFC**.
2. Choose a change type from the drop-down lists. Optionally, open the **Additional configuration** area to select a change type version. After your selections are complete, a **Change type:** details area opens. Choose **Next**.
3. Configure the request for change. A **Subject** is required. Optionally, open the **Additional configuration** area to add information about the RFC. Choose **Next**.
4. Choose the execution parameters. At the top, in the **RFC configuration** area, enter values for the change type required parameters. These vary by change type. Open the **Additional configuration** area to add Tags or additional settings. Some change types also provide a **Parameters** area where only the required settings are visible. In that case, open the **Additional configuration** area to view optional parameters.
5. When finished, choose **Create**. If there are no errors, the **RFC successfully created** page displays with the submitted RFC details, and the initial **Execution output**.
6. Open the **Execution parameters** area to see the configurations you submitted. Refresh the page to update the RFC execution status. Optionally, cancel the RFC or create a copy of it with the options at the top of the page.

Creating an EC2 instance and additional volumes with the CLI

How it works:

1. Use either the Inline Create (you issue a `create-rfc` command with all RFC and execution parameters included), or Template Create (you create two JSON files, one for the RFC parameters and one for the execution parameters) and issue the `create-rfc` command with the two files as input. Both methods are described here.
2. Submit the RFC: `aws amscm submit-rfc --rfc-id ID` command with the returned RFC ID.

Monitor the RFC: `aws amscm get-rfc --rfc-id ID` command.

To check the change type version, use this command:

```
aws amscm list-change-type-version-summaries --filter Attribute=ChangeTypeId,Value=CT_ID
```

Note

You can use any `CreateRfc` parameters with any RFC whether or not they are part of the schema for the change type. For example, to get notifications when the RFC status changes, add this line, `--notification "{\"Email\": {\"EmailRecipients\": [\"email@example.com\"]}}\"` to the RFC parameters part of the request (not the execution parameters). For a list of all `CreateRfc` parameters, see the [AMS Change Management API Reference](#).

INLINE CREATE:

Issue the create RFC command with execution parameters provided inline (escape quotation marks when providing execution parameters inline), and then submit the returned RFC ID (example shows required parameters only). For example, you can replace the contents with something like this:

```
aws amscm create-rfc --change-type-id "ct-1aqsjf86w6vxg" --change-type-version "4.0"
--title "EC2-Create-A-V-QC" --execution-parameters "{\"Description\": \"My EC2 stack
with addl vol\", \"VpcId\": \"VPC_ID\", \"Name\": \"My Stack\", \"StackTemplateId\": \"stm-
nn8v8ffhcal611bmo\", \"TimeoutInMinutes\": 60, \"Parameters\": {\"InstanceAmiId\": \"AMI_ID\",
\"InstanceSubnetId\": \"SUBNET_ID\"}}
```

TEMPLATE CREATE:

1. Output the execution parameters for this change type to a JSON file named CreateEC2AVParams.json.

```
aws amscm get-change-type-version --change-type-id "ct-1aqsjf86w6vxg" --query
"ChangeTypeVersion.ExecutionInputSchema" --output text > CreateEC2AVParams.json
```

2. Modify and save the CreateEC2AVParams file (example shows most parameters). For example, you can replace the contents with something like this:

```
{
  "Description": "EC2-Create-1-Addl-Volumes",
  "VpcId": "VPC_ID",
  "StackTemplateId": "stm-nn8v8ffhcal611bmo",
  "Name": "My-EC2-1-Addl-Volume",
  "TimeoutInMinutes": 60,
  "Parameters": {
    "InstanceAmiId": "AMI_ID",
    "InstanceSecurityGroupIds": "SECURITY_GROUP_ID",
    "InstanceDetailedMonitoring": "true",
    "InstanceEBSOptimized": "false",
    "InstanceProfile": "customer-mc-ec2-instance-profile",
    "InstanceRootVolumeIops": 100,
    "InstanceRootVolumeName": "/dev/xvda",
    "InstanceRootVolumeSize": 50,
    "InstanceRootVolumeType": "io1",
    "RootVolumeKmsKeyId": "default",
    "InstancePrivateStaticIp": "10.27.0.100",
    "InstanceSecondaryPrivateIpAddressCount": 0,
    "InstanceTerminationProtection": "false",
    "InstanceType": "t3.large",
    "CreditSpecification": "unlimited",
    "InstanceUserData": "echo $",
    "Volume1Encrypted": "true",
    "Volume1Iops": "IOPS",
    "Volume1KmsKeyId": "KMS_MASTER_KEY_ID",
    "Volume1Name": "xvdh",
    "Volume1Size": "2 GiB",
    "Volume1Snapshot": "SNAPSHOT_ID",
    "Volume1Type": "io1",
    "InstanceSubnetId": "SUBNET_ID"
  }
}
```

3. Output the RFC template to a file in your current folder; this example names it CreateEC2AVRfc.json:

```
aws amscm create-rfc --generate-cli-skeleton > CreateEC2AVRfc.json
```

4. Modify and save the CreateEC2AVRfc.json file. For example, you can replace the contents with something like this:

```
{
  "ChangeTypeVersion": "4.0",
  "ChangeTypeId": "ct-1aqsjf86w6vxg",
  "Title": "EC2-Create-1-Addl-Volume-RFC"
}
```

5. Create the RFC, specifying the CreateEC2AVRfc file and the CreateEC2AVParams file:

```
aws amscm create-rfc --cli-input-json file://CreateEC2AVRfc.json --execution-parameters file://CreateEC2AVParams.json
```

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

Access, Requesting

Request administrative access

How to request administrative access to a stand-alone EC2 stack instance or to instances that are part of an EC2 Auto Scaling group (ASG) stack using the AMS console or the AMS API/CLI.

For an example about requesting ReadOnly access, see [ReadOnly access: requesting](#).

Classification: Management | Access | Stack admin access | Grant

Change type ID: ct-1dmlg9g1l91h6

Version: 2.0

The only AMS-managed resources that you can request access to are EC2 standalone instances and instances that are part of ASG stacks. You log in to RDS resources with your credentials. You can add and remove objects from S3 buckets by using the S3 console or API. All other resource changes you accomplish through an RFC to AMS.

Note

You can submit an update to your access request before it expires. For information, see [Accessing, Administrative, Update](#).

To log in to an instance that is part of an ASG, you request access to the ASG stack, which gives you access to all associated instances.

Required data:

- **Subject:** This value displays on the RFC dashboard.
- **DomainFQDN:** Provide the Fully Qualified Domain Name (FQDN) of your AMS-trusted domain. If you're unsure what it is, you can find it in the AWS Management Console for Directory Services (under Security and Identity) **Directory Name** tab.
- **StackIds:** An array of stack identifiers that you want to access. You can find stack IDs by looking on the **Stacks** dashboard of the AMS console or by running the [ListStackSummaries](#) operation of the SKMS API (`list-stack-summaries` in the CLI).
- **Username:** The Active Directory user name of the person who wants access.
- **VpcId:** The ID of the VPC where the stacks are that you want access to. You can find this by looking on the **VPCs** dashboard of the AMS console or by running the [ListVpcSummaries](#) operation of the SKMS API (`list-vpc-summaries` in the CLI).

Optional data:

- **TimeRequestedInHours:** The amount of time, in hours, requested for access to the instance. Access is ended after this time.

Requesting administrator access with the console

The following shows this change type in the AMS console.

▼ Change type: Grant stack admin access	
Description	
Use to request admin access to a single or multiple stacks. Access can be requested for a maximum of 8 hours.	
ID	Version
ct-1dmlg9g1l91h6	2.0
Execution mode	
Automated	

How it works:

1. Navigate to the **Choose change type** page: **RFCs -> Create RFC**.
2. Choose a change type from the drop-down lists. Optionally, open the **Additional configuration** area to select a change type version. After your selections are complete, a **Change type:** details area opens. Choose **Next**.
3. Configure the request for change. A **Subject** is required. Optionally, open the **Additional configuration** area to add information about the RFC. Choose **Next**.
4. Choose the execution parameters. At the top, in the **RFC configuration** area, enter values for the change type required parameters. These vary by change type. Open the **Additional configuration** area to add Tags or additional settings. Some change types also provide a **Parameters** area where only the required settings are visible. In that case, open the **Additional configuration** area to view optional parameters.
5. When finished, choose **Create**. If there are no errors, the **RFC successfully created** page displays with the submitted RFC details, and the initial **Execution output**.
6. Open the **Execution parameters** area to see the configurations you submitted. Refresh the page to update the RFC execution status. Optionally, cancel the RFC or create a copy of it with the options at the top of the page.

Requesting administrator access with the CLI

How it works:

1. Use either the Inline Create (you issue a `create-rfc` command with all RFC and execution parameters included), or Template Create (you create two JSON files, one for the RFC parameters and one for the execution parameters) and issue the `create-rfc` command with the two files as input. Both methods are described here.
2. Submit the RFC: `aws amscm submit-rfc --rfc-id ID` command with the returned RFC ID.

Monitor the RFC: `aws amscm get-rfc --rfc-id ID` command.

To check the change type version, use this command:

```
aws amscm list-change-type-version-summaries --filter Attribute=ChangeTypeId,Value=CT_ID
```

Note

You can use any `CreateRfc` parameters with any RFC whether or not they are part of the schema for the change type. For example, to get notifications when the RFC status changes, add this line, `--notification "{\"Email\": {\"EmailRecipients\" : [\"email@example.com\"]}}\"` to the RFC parameters part of the request (not the execution parameters). For a list of all `CreateRfc` parameters, see the [AMS Change Management API Reference](#).

INLINE CREATE:

Issue the create RFC command with execution parameters provided inline (escape quotation marks when providing execution parameters inline), and then submit the returned RFC ID. For example, you can replace the contents with something like this:

```
aws --profile saml amscm create-rtc --change-type-id "ct-1dmlg9g1191h6" --change-type-version "2.0" --title "Stack-Admin-Access-QC" --execution-parameters "{\"DomainFQDN\": \"TEST.com\", \"StackIds\": [\"stack-01234567890abcdef\"], \"TimeRequestedInHours\": 1, \"Username\": \"TEST\", \"VpcId\": \"VPC_ID\"}"
```

TEMPLATE CREATE:

1. Output the execution parameters JSON schema for this change type to a file; this example names it `GrantAdminAccessParams.json`:

```
aws amscm get-change-type-version --change-type-id "ct-1dmlg9g1191h6" --query "ChangeTypeVersion.ExecutionInputSchema" --output text > GrantAdminAccessParams.json
```

Modify and save the `GrantAdminAccessParams` file. For example, you can replace the contents with something like this:

```
{
  "DomainFQDN":      "mycorpdomain.acme.com",
  "StackIds":        [STACK_ID, STACK_ID],
  "TimeRequestedInHours": 8,
  "Username":        "USERNAME",
  "VpcId":           "VPC_ID"
}
```

Note that the `TimeRequestedInHours` option defaults to one hour. You can request up to eight hours.

2. Output the RFC template to a file in your current folder; this example names it `GrantAdminAccessRfc.json`:

```
aws amscm create-rtc --generate-cli-skeleton > GrantAdminAccessRfc.json
```

3. Modify and save the `GrantAdminAccessRfc.json` file. For example, you can replace the contents with something like this:

```
{
  "ChangeTypeId":      "ct-1dmlg9g1191h6",
  "ChangeTypeVersion": "2.0",
  "Title":              "Request-Admin-Access-to-EC2-RFC"
}
```

4. Create the RFC, specifying the `GrantAdminAccessRfc` file and the `GrantAdminAccessParams` file:

```
aws amscm create-rfc --cli-input-json file://GrantAdminAccessRfc.json --execution-parameters file://GrantAdminAccessParams.json
```

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

To log in to the instance through a bastion, follow the next procedure, [Instance access examples](#).

Request ReadOnly access

How to request ReadOnly access to a stand-alone EC2 stack instance or to instances that are part of an EC2 Auto Scaling group (ASG) stack using the AMS console or the AMS API/CLI.

For a walkthrough on requesting Admin access, see [Admin Access: requesting](#).

Classification: Management | Access | Stack read-only access | Grant

Change type ID: ct-199h35t7uz6jl

Version: 2.0

The only AMS-managed resources that you can request access to are EC2 standalone instances and instances that are part of ASG stacks. S3 buckets you can add and remove objects from using the S3 console or API. All other resource changes you accomplish through an RFC to AMS.

Note

You can submit an update to your access request before it expires. For details, see [Accessing, Read Only, Update](#).

To log into an instance that is part of an EC2 Auto Scaling group (ASG), you request access to the ASG stack, which gives you access to all associated instances.

Required data:

- **Subject:** This value displays on the RFC dashboard.
- **DomainFQDN:** Provide the Fully Qualified Domain Name (FQDN) of your AMS-trusted domain. If you're unsure what it is, you can use the AWS Management Console for Directory Services (under Security and Identity) **Directory Name** tab.
- **StackIds:** An array of stack identifiers that you want to access. You can find stack IDs by looking on the **Stacks** dashboard of the AMS console or by running the [ListStackSummaries](#) operation of the SKMS API (`list-stack-summaries` in the CLI).
- **Username:** The Active Directory user name of the person who wants access.
- **VpcId:** The ID of the VPC where the stacks are that you want access to. You can find this by looking on the **VPCs** dashboard of the AMS console or by running the [ListVpcSummaries](#) operation of the SKMS API (`list-vpc-summaries` in the CLI).

Optional data:

- **TimeRequestedInHours:** The amount of time, in hours, requested for access to the instance. Access is ended after this time.

Requesting ReadOnly access with the console

The following shows this change type in the AMS console.

▼ **Change type: Grant stack read-only access**

Description

Use to request read-only access as an operator to a single or multiple stacks. Access can be requested for a maximum of 8 hours.

ID	Version
ct-199h35t7uz6jl	2.0

Execution mode

Automated

How it works:

1. Navigate to the **Choose change type** page: **RFCs -> Create RFC**.
2. Choose a change type from the drop-down lists. Optionally, open the **Additional configuration** area to select a change type version. After your selections are complete, a **Change type: details** area opens. Choose **Next**.
3. Configure the request for change. A **Subject** is required. Optionally, open the **Additional configuration** area to add information about the RFC. Choose **Next**.
4. Choose the execution parameters. At the top, in the **RFC configuration** area, enter values for the change type required parameters. These vary by change type. Open the **Additional configuration** area to add Tags or additional settings. Some change types also provide a **Parameters** area where only the required settings are visible. In that case, open the **Additional configuration** area to view optional parameters.
5. When finished, choose **Create**. If there are no errors, the **RFC successfully created** page displays with the submitted RFC details, and the initial **Execution output**.
6. Open the **Execution parameters** area to see the configurations you submitted. Refresh the page to update the RFC execution status. Optionally, cancel the RFC or create a copy of it with the options at the top of the page.

Requesting ReadOnly access with the CLI

How it works:

1. Use either the Inline Create (you issue a `create-rfc` command with all RFC and execution parameters included), or Template Create (you create two JSON files, one for the RFC parameters and one for the execution parameters) and issue the `create-rfc` command with the two files as input. Both methods are described here.
2. Submit the RFC: `aws amscm submit-rfc --rfc-id ID` command with the returned RFC ID.

Monitor the RFC: `aws amscm get-rfc --rfc-id ID` command.

To check the change type version, use this command:

```
aws amscm list-change-type-version-summaries --filter Attribute=ChangeTypeId,Value=CT_ID
```

Note

You can use any `CreateRfc` parameters with any RFC whether or not they are part of the schema for the change type. For example, to get notifications when the RFC status changes, add this line, `--notification "{\"Email\": {\"EmailRecipients\" : [\"email@example.com\"]}}\"` to the RFC parameters part of the request (not the execution

parameters). For a list of all CreateRfc parameters, see the [AMS Change Management API Reference](#).

INLINE CREATE:

Issue the create RFC command with execution parameters provided inline (escape quotes when providing execution parameters inline), and then submit the returned RFC ID. For example, you can replace the contents with something like this:

```
aws --profile saml amscm create-rtc --change-type-id "ct-199h35t7uz6jl" --change-type-version "2.0" --title "Stack-RO-Access-QC" --execution-parameters "{\"DomainFQDN\": \"TEST.com\", \"StackIds\": [\"stack-01234567890abcdef\"], \"TimeRequestedInHours\": 1, \"Username\": \"TEST\", \"VpcId\": \"VPC_ID\"}"
```

TEMPLATE CREATE:

1. Output the execution parameters JSON schema for this change type to a file; this example names it GrantReadOnlyAccessParams.json:

```
aws amscm get-change-type-version --change-type-id "ct-199h35t7uz6jl" --query "ChangeTypeVersion.ExecutionInputSchema" --output text > GrantReadOnlyAccessParams.json
```

Modify and save the GrantReadOnlyAccessParams file. For example, you can replace the contents with something like this:

```
{
  "DomainFQDN":      "mycorpdomain.acme.com",
  "StackIds":        [STACK_ID, STACK_ID],
  "TimeRequestedInHours": 8,
  "Username":        "USERNAME",
  "VpcId":           "VPC_ID"
}
```

Note that the TimeRequestedInHours option defaults to one hour. You can request up to eight hours.

2. Output the RFC template to a file in your current folder; this example names it GrantReadOnlyAccessRfc.json:

```
aws amscm create-rtc --generate-cli-skeleton > GrantReadOnlyAccessRfc.json
```

3. Modify and save the GrantReadOnlyAccessRfc.json file. For example, you can replace the contents with something like this:

```
{
  "ChangeTypeId":      "ct-199h35t7uz6jl",
  "ChangeTypeVersion": "2.0",
  "Title":              "Request-ReadOnly-Access-to-EC2-RFC"
}
```

4. Create the RFC, specifying the GrantReadOnlyAccessRfc file and the GrantReadOnlyAccessParams file:

```
aws amscm create-rtc --cli-input-json file://GrantReadOnlyAccessRfc.json --execution-parameters file://GrantReadOnlyAccessParams.json
```

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

To log in to the instance through a bastion, follow the next procedure, [Instance access examples](#).

Other | Other RFC, Creating (CLI)

This example shows how to request a change that none of the available CTs address, by using the Management | Other | Other | Create CT (ct-1e1xtak34nx76).

Use this CT when you can't find a change type for what you want; however, if you are unsure about specifying parameters in an existing CT, it is better to submit a service request for help. For information on submitting service requests, see [Service Request Examples](#).

This type of RFC is Approval-required, meaning that it requires AMS approval before it can be implemented. After submitting the RFC, an AMS operator will contact you to discuss the stack that you want to deploy.

Note

When using manual (approval required) CTs, AMS recommends that you use the ASAP option (choose **ASAP** in the console, leave start and end time blank in the API/CLI) as these CTs require an AMS operator to examine the RFC, and possibly communicate with you before it can be approved and run. If you schedule these RFCs, be sure to allow at least 24 hours. If approval does not happen before the scheduled start time, the RFC is rejected automatically.

REQUIRED DATA:

- **Comment:** What the RFC is for.
- **ChangeTypeId** and **ChangeTypeVersion:** Use Other | Create (ct-1e1xtak34nx76) to request new resources, use Other | Update (ct-0xdawir96cy7k) to change existing resources; both are v1.

OPTIONAL DATA: **Priority:** Acceptable values are High, Medium, or Low.

INLINE CREATE:

- Issue the create RFC command with execution parameters provided inline (escape quotes when providing execution parameters inline). Example uses Other | Create.

```
aws amscm create-rfc --change-type-id "ct-1e1xtak34nx76" --change-type-version "1.0" --title "TITLE" --execution-parameters "{\"Comment\": \"What you want created\"}"
```

- Submit the RFC using the RFC ID returned in the create RFC operation. Until submitted, the RFC remains in the Editing state and is not acted on.

```
aws amscm submit-rfc --rfc-id RFC_ID
```

- Monitor the RFC status and view execution output:

```
aws amscm get-rfc --rfc-id RFC_ID
```

TEMPLATE CREATE:

1. Create and save a JSON file for the execution parameters; example names it OtherParams.json and includes the optional **Priority** parameter:

```
{  
  "Comment": "What you want created",  
  "Priority": "Medium"  
}
```

```
}
```

2. Create and save a JSON file for the RFC parameters; example names it OtherRfc.json.

```
{  
  "ChangeTypeId":      "ct-1e1xtak34nx76",  
  "ChangeTypeVersion": "1.0",  
  "Title":             "TITLE"  
}
```

3. Create the RFC, specifying the OtherRfc file and the OtherParams file:

```
aws amscm create-rfc --cli-input-json file://OtherRfc.json --execution-parameters  
file://OtherParams.json
```

You receive the RfcId of the new RFC in the response. For example:

```
{  
  "RfcId": "RFC-ID"  
}
```

4. Submit the RFC:

```
aws amscm submit-rfc --rfc-id RFC-ID
```

If no errors are reported, the operation was successful.

5. To monitor the status of the request and to view Execution Output:

```
aws amscm get-rfc --rfc-id RFC-ID
```

Any Stack, Deleting, Rebooting, Starting, Stopping

You can use the AMS console or API/CLI to delete, reboot, start, or stop, an AMS stack.

Delete stack

Delete a stack using the AMS console or the AMS API/CLI.

Classification: Management | Standard stacks | Stack | Delete and Management | Advanced stack components | Stack | Delete

Change type ID: ct-0q0bic0ywqk6c

Note

If deleting an S3 bucket, it must be emptied of objects first.

Important

Deleting stacks can have unwanted and unanticipated consequences. For important caveats, see RFC Troubleshooting section [RFCs for Delete Stack](#).

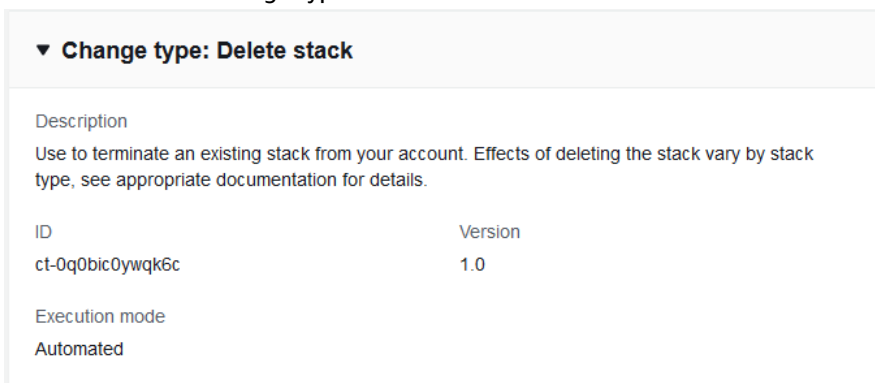
Required Data:

- **Subject:** A title for the request.
- **StackId:** The stack that you want deleted. For help finding stack IDs, see [Find Stack Id](#).
- **TimeoutInMinutes:** The maximum amount of time, in minutes, to allow for execution of deleting the stack. The minimum value is 0 and the maximum value is 720; if you do not provide a value, the default is 60. The value you specify does not prolong the runtime. If the delete is not completed in

the specified time, the RFC is considered failed and you are notified that the delete is over time, but continuing. It is important to note that the delete operation continues even if the RFC fails, because delete operations cannot be rolled back. This parameter is to notify you of delete stack problems in a timely manner.

Deleting a Stack with the Console

Screenshot of this change type in the AMS console:



How it works:

1. Navigate to the **Choose change type** page: **RFCs -> Create RFC**.
2. Choose a change type from the drop-down lists. Optionally, open the **Additional configuration** area to select a change type version. After your selections are complete, a **Change type: details** area opens. Choose **Next**.
3. Configure the request for change. A **Subject** is required. Optionally, open the **Additional configuration** area to add information about the RFC. Choose **Next**.
4. Choose the execution parameters. At the top, in the **RFC configuration** area, enter values for the change type required parameters. These vary by change type. Open the **Additional configuration** area to add Tags or additional settings. Some change types also provide a **Parameters** area where only the required settings are visible. In that case, open the **Additional configuration** area to view optional parameters.
5. When finished, choose **Create**. If there are no errors, the **RFC successfully created** page displays with the submitted RFC details, and the initial **Execution output**.
6. Open the **Execution parameters** area to see the configurations you submitted. Refresh the page to update the RFC execution status. Optionally, cancel the RFC or create a copy of it with the options at the top of the page.

Deleting a Stack with the CLI

How it works:

1. Use either the Inline Create (you issue a `create-rfc` command with all RFC and execution parameters included), or Template Create (you create two JSON files, one for the RFC parameters and one for the execution parameters) and issue the `create-rfc` command with the two files as input. Both methods are described here.
2. Submit the RFC: `aws amscm submit-rfc --rfc-id ID` command with the returned RFC ID.

Monitor the RFC: `aws amscm get-rfc --rfc-id ID` command.

To check the change type version, use this command:

```
aws amscm list-change-type-version-summaries --filter Attribute=ChangeTypeId,Value=CT_ID
```

Note

You can use any `CreateRfc` parameters with any RFC whether or not they are part of the schema for the change type. For example, to get notifications when the RFC status changes, add this line, `--notification "{\"Email\": {\"EmailRecipients\" : [\"email@example.com\"]}}\"` to the RFC parameters part of the request (not the execution parameters). For a list of all `CreateRfc` parameters, see the [AMS Change Management API Reference](#).

INLINE CREATE:

Issue the create RFC command with execution parameters provided inline (escape quotation marks when providing execution parameters inline), and then submit the returned RFC ID. For example, you can replace the contents with something like this:

```
aws amscm create-rtc --change-type-id "ct-0q0bic0ywqk6c" --change-type-version "1.0" --title "Delete My Stack" --execution-parameters "{\"StackId\": \"STACK_ID\"}"
```

TEMPLATE CREATE:

1. Output the RFC template to a file in your current folder; this example names it `DeleteStackRfc.json`:

```
aws amscm create-rtc --generate-cli-skeleton > DeleteStackRfc.json
```

2. Modify and save the `DeleteStackRfc.json` file.

The internal quotation marks in the `ExecutionParameters` JSON extension must be escaped with a backslash (`\`). Example without start and end time:

```
{
  "ChangeTypeVersion": "1.0",
  "ChangeTypeId": "ct-0q0bic0ywqk6c",
  "Title": "Delete-My-Stack-RFC",
  "ExecutionParameters": "{
    \"StackId\": \"STACK_ID\"}"
}
```

3. Create the RFC:

```
aws amscm create-rtc --cli-input-json file://DeleteStackRfc.json
```

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

Reboot stack

Reboot a stack using the AMS console or the AMS API/CLI.

Classification: Management | Standard stacks | Stack | Reboot

Change type ID: `ct-02u0hoaa9grat`

Required Data:

- **Subject:** A title for the request.
- **StackId:** The stack that you want rebooted. For help finding stack IDs, see [Find Stack Id](#).

Rebooting a Stack with the Console

Screenshot of this change type in the AMS console:



How it works:

1. Navigate to the **Choose change type** page: **RFCs -> Create RFC**.
2. Choose a change type from the drop-down lists. Optionally, open the **Additional configuration** area to select a change type version. After your selections are complete, a **Change type: details** area opens. Choose **Next**.
3. Configure the request for change. A **Subject** is required. Optionally, open the **Additional configuration** area to add information about the RFC. Choose **Next**.
4. Choose the execution parameters. At the top, in the **RFC configuration** area, enter values for the change type required parameters. These vary by change type. Open the **Additional configuration** area to add Tags or additional settings. Some change types also provide a **Parameters** area where only the required settings are visible. In that case, open the **Additional configuration** area to view optional parameters.
5. When finished, choose **Create**. If there are no errors, the **RFC successfully created** page displays with the submitted RFC details, and the initial **Execution output**.
6. Open the **Execution parameters** area to see the configurations you submitted. Refresh the page to update the RFC execution status. Optionally, cancel the RFC or create a copy of it with the options at the top of the page.

Rebooting a Stack with the CLI

How it works:

1. Use either the Inline Create (you issue a `create-rfc` command with all RFC and execution parameters included), or Template Create (you create two JSON files, one for the RFC parameters and one for the execution parameters) and issue the `create-rfc` command with the two files as input. Both methods are described here.
2. Submit the RFC: `aws amscm submit-rfc --rfc-id ID` command with the returned RFC ID.

Monitor the RFC: `aws amscm get-rfc --rfc-id ID` command.

To check the change type version, use this command:

```
aws amscm list-change-type-version-summaries --filter Attribute=ChangeTypeId,Value=CT_ID
```

Note

You can use any `CreateRfc` parameters with any RFC whether or not they are part of the schema for the change type. For example, to get notifications when the RFC status

changes, add this line, `--notification "{\"Email\": {\"EmailRecipients\" : [\"email@example.com\"]}}\"` to the RFC parameters part of the request (not the execution parameters). For a list of all CreateRfc parameters, see the [AMS Change Management API Reference](#).

INLINE CREATE:

Issue the create RFC command with execution parameters provided inline (escape quotation marks when providing execution parameters inline), and then submit the returned RFC ID. For example, you can replace the contents with something like this:

```
aws amscm create-rtc --change-type-id "ct-02u0hoaa9grat" --change-type-version "1.0" --title "Reboot My Stack" --execution-parameters "{\"StackId\": \"STACK_ID\"}"
```

TEMPLATE CREATE:

1. Output the RFC template to a file in your current folder. This example names it RebootStackRfc.json. Note that since there is only one execution parameter for stopping (rebooting, or starting) an instance, the execution parameter can be in the schema JSON file itself and there is no need to create a separate execution parameters JSON file.

```
aws amscm create-rtc --generate-cli-skeleton > StopInstanceRfc.json
```

2. Modify and save the RebootStackRfc.json file.

The internal quotation marks in the ExecutionParameters JSON extension must be escaped with a backslash (\). Example:

```
{
  "ChangeTypeId": "ct-02u0hoaa9grat",
  "Title": "Reboot-My-EC2-RFC",
  "TimeoutInMinutes": 60,
  "ExecutionParameters": "{
    \"StackId\": \"STACK_ID\"
  }"
}
```

3. Create the RFC:

```
aws amscm create-rtc --cli-input-json file://RebootStackRfc.json
```

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

Start stack

Start a stack using the AMS console or the AMS API/CLI.

Classification: Management | Standard stacks | Stack | Start

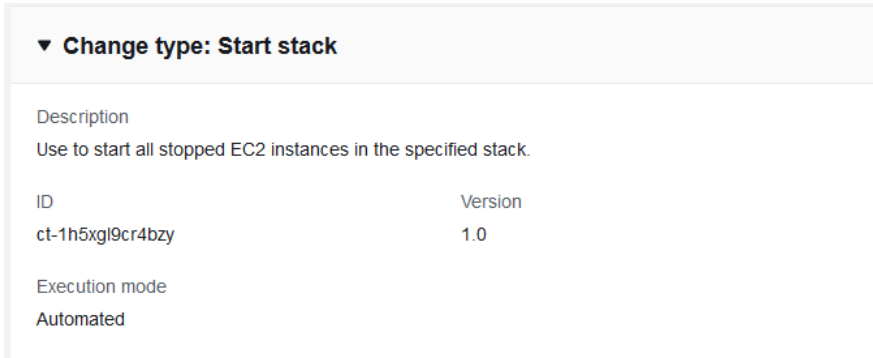
Change type ID: ct-1h5xgl9cr4bzy

Required Data:

- **Subject:** A title for the request.
- **StackId:** The stack that you want started. For help finding stack IDs, see [Find Stack Id](#).

Starting a Stack with the Console

Screenshot of this change type in the AMS console:



How it works:

1. Navigate to the **Choose change type** page: **RFCs -> Create RFC**.
2. Choose a change type from the drop-down lists. Optionally, open the **Additional configuration** area to select a change type version. After your selections are complete, a **Change type:** details area opens. Choose **Next**.
3. Configure the request for change. A **Subject** is required. Optionally, open the **Additional configuration** area to add information about the RFC. Choose **Next**.
4. Choose the execution parameters. At the top, in the **RFC configuration** area, enter values for the change type required parameters. These vary by change type. Open the **Additional configuration** area to add Tags or additional settings. Some change types also provide a **Parameters** area where only the required settings are visible. In that case, open the **Additional configuration** area to view optional parameters.
5. When finished, choose **Create**. If there are no errors, the **RFC successfully created** page displays with the submitted RFC details, and the initial **Execution output**.
6. Open the **Execution parameters** area to see the configurations you submitted. Refresh the page to update the RFC execution status. Optionally, cancel the RFC or create a copy of it with the options at the top of the page.

Starting a Stack with the CLI

How it works:

1. Use either the Inline Create (you issue a `create-rfc` command with all RFC and execution parameters included), or Template Create (you create two JSON files, one for the RFC parameters and one for the execution parameters) and issue the `create-rfc` command with the two files as input. Both methods are described here.
2. Submit the RFC: `aws amscm submit-rfc --rfc-id ID` command with the returned RFC ID.

Monitor the RFC: `aws amscm get-rfc --rfc-id ID` command.

To check the change type version, use this command:

```
aws amscm list-change-type-version-summaries --filter Attribute=ChangeTypeId,Value=CT_ID
```

Note

You can use any `CreateRfc` parameters with any RFC whether or not they are part of the schema for the change type. For example, to get notifications when the RFC status

changes, add this line, `--notification "{\"Email\": {\"EmailRecipients\" : [\"email@example.com\"]}}\"` to the RFC parameters part of the request (not the execution parameters). For a list of all CreateRfc parameters, see the [AMS Change Management API Reference](#).

INLINE CREATE:

Issue the create RFC command with execution parameters provided inline (escape quotation marks when providing execution parameters inline), and then submit the returned RFC ID. For example, you can replace the contents with something like this:

```
aws amscm create-rtc --change-type-id "ct-1h5xgl9cr4bzy" --change-type-version "1.0" --title "Start My Stack" --execution-parameters "{\"StackId\": \"STACK_ID\"}"
```

TEMPLATE CREATE:

1. Output the RFC template to a file in your current folder. This example names it StartInstanceRfc.json. Note that since there is only one execution parameter for starting a stack, the execution parameter can be in the schema JSON file itself and there is no need to create a separate execution parameters JSON file.

```
aws amscm create-rtc --generate-cli-skeleton > StartStackRfc.json
```

2. Modify and save the StartStackRfc.json file. For example, you can replace the contents with something like this:

```
{
  "ChangeTypeId":      "ct-1h5xgl9cr4bzy",
  "Title":             "Start-My-EC2-RFC",
  "TimeoutInMinutes": 60,
  "ExecutionParameters": "{
    \"StackId\": \"STACK_ID\"
  }"
}
```

3. Create the RFC:

```
aws amscm create-rtc --cli-input-json file://StartStackRfc.json
```

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

Stop stack

Stop a stack using the AMS console or the AMS API/CLI.

Classification: Management | Standard stacks | Stack | Stop

Change type ID: ct-3dgbnh6gpst4d

Tip

This CT stops all EC2 instances in the stack, but it doesn't stop instances launched through an Auto Scaling group or stop load balancers.

Stopped instances retain their instance IDs; deleting instances releases their IDs and removes the instance from the account. Stopped instances do not incur charges, but the disk storage on any EBS volumes are still charged.

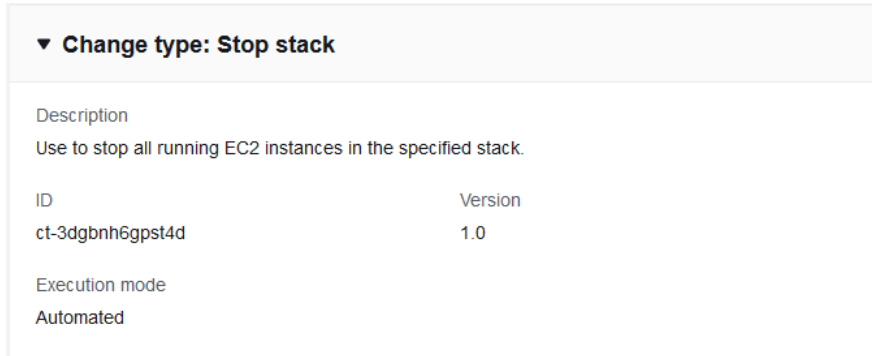
Stopped instances remain stopped unless you have scheduled restarts using the [AMS Resource Scheduler](#).

Required Data:

- **Subject:** A title for the request.
- **StackId:** The stack that you want stopped. For help finding stack IDs, see [Find Stack Id](#).

Stopping a Stack with the Console

Screenshot of this change type in the AMS console:



How it works:

1. Navigate to the **Choose change type** page: **RFCs -> Create RFC**.
2. Choose a change type from the drop-down lists. Optionally, open the **Additional configuration** area to select a change type version. After your selections are complete, a **Change type:** details area opens. Choose **Next**.
3. Configure the request for change. A **Subject** is required. Optionally, open the **Additional configuration** area to add information about the RFC. Choose **Next**.
4. Choose the execution parameters. At the top, in the **RFC configuration** area, enter values for the change type required parameters. These vary by change type. Open the **Additional configuration** area to add Tags or additional settings. Some change types also provide a **Parameters** area where only the required settings are visible. In that case, open the **Additional configuration** area to view optional parameters.
5. When finished, choose **Create**. If there are no errors, the **RFC successfully created** page displays with the submitted RFC details, and the initial **Execution output**.
6. Open the **Execution parameters** area to see the configurations you submitted. Refresh the page to update the RFC execution status. Optionally, cancel the RFC or create a copy of it with the options at the top of the page.

Stopping a Stack with the CLI

How it works:

1. Use either the Inline Create (you issue a `create-rfc` command with all RFC and execution parameters included), or Template Create (you create two JSON files, one for the RFC parameters and one for the execution parameters) and issue the `create-rfc` command with the two files as input. Both methods are described here.
2. Submit the RFC: `aws amscm submit-rfc --rfc-id ID` command with the returned RFC ID.

Monitor the RFC: `aws amscm get-rfc --rfc-id ID` command.

To check the change type version, use this command:

```
aws amscm list-change-type-version-summaries --filter Attribute=ChangeTypeId,Value=CT_ID
```

Note

You can use any `CreateRfc` parameters with any RFC whether or not they are part of the schema for the change type. For example, to get notifications when the RFC status changes, add this line, `--notification "{\"Email\": {\"EmailRecipients\" : [\"email@example.com\"]}}\"` to the RFC parameters part of the request (not the execution parameters). For a list of all `CreateRfc` parameters, see the [AMS Change Management API Reference](#).

INLINE CREATE:

Issue the create RFC command with execution parameters provided inline (escape quotation marks when providing execution parameters inline), and then submit the returned RFC ID. For example, you can replace the contents with something like this:

```
aws amscm create-rtc --change-type-id "ct-3dgbnh6gpst4d" --change-type-version "1.0" --title "Stop My Stack" --execution-parameters "{\"StackId\": \"STACK_ID\"}"
```

TEMPLATE CREATE:

1. Output the RFC template to a file in your current folder. This example names it `StopStackRfc.json`. Note that since there is only one execution parameter for stopping (rebooting, or starting) an instance, the execution parameter can be in the schema JSON file itself and there is no need to create a separate execution parameters JSON file.

```
aws amscm create-rtc --generate-cli-skeleton > StopStackRfc.json
```

2. Modify and save the `StopStackRfc.json` file. For example, you can replace the contents with something like this:

```
{
  "ChangeTypeId":      "ct-3dgbnh6gpst4d",
  "Title":              "Stop-My-EC2-RFC",
  "TimeoutInMinutes":  60,
  "ExecutionParameters": "{
                        \"StackId\": \"STACK_ID\"
                      }"
```

3. Create the RFC:

```
aws amscm create-rtc --cli-input-json file://StopInstanceRfc.json
```

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

If needed, see [EC2 instance stack stop fail](#).

Access Examples

These examples show how to log in to an instance via a bastion once you have been granted access through an RFC. For details on getting access granted, see

Note

An EC2 instance created through an Auto Scaling group will have an IP address that cycles in and out and you will have to use your EC2 console to find that IP address.

REQUIRED DATA:

- **Bastion DNS friendly name or IP address:** Use a DNS friendly name as described in [DNS Friendly Bastion Names \(p. 142\)](#) or find bastion IP addresses as described in [Finding Bastion IP Addresses \(p. 143\)](#).
- **Username** (for example `username@customerdomain.com`) and **Password:** Credentials for the account.
- **Stack IP address:** Get this by looking at the AMS console **Stacks** page for the stack you want to log into and then filtering on that stack ID in the EC2 console for your account. For a single EC2 instance, you can also use the AMS SKMS command [ListStackSummaries](#) to find the stack ID and then [GetStack](#) to find the stack IP address.

Access the bastion IP address, either SSH or RDP, as appropriate, and log in using one of the following procedures.

Linux Computer to Linux Instance

Use SSH to connect to the SSH bastion and then to the Linux instance.

MALZ

For more information about the friendly bastion names, see [DNS bastions](#).

In order to connect to the Linux instance, you must first connect to an SSH bastion.

1. Open a shell window and enter:

```
ssh Domain_FQDN\\Username@SSH_bastion_name  
or SSH_bastion_IP
```

Which would look like this if your `Domain_FQDN` is "corp.domain.com", your account number is "123456789123", Your_Domain is "amazonaws.com", you choose bastion "4", and your user name is "JoeSmith":

```
ssh corp.domain.com\\JoeSmith sshbastion4.A123456789123.amazonaws.com
```

2. Log in with your corporate Active Directory credentials.
3. When presented with a Bash prompt, SSH in to the instance, and then enter:

```
ssh Domain_FQDN\\Username@Instance_IP
```

Or, you can use the Login flag (-l):

```
ssh -l Domain_FQDN\\Username@Instance_IP
```

SALZ

For more information about the friendly bastion names, see [DNS bastions](#).

In order to connect to the Linux instance, you must first connect to an SSH bastion.

1. Open a shell window and enter:

```
ssh DOMAIN_FQDN\\USERNAME@SSH_BASTION_name  
or SSH_BASTION_IP
```

Which would look like this if your account number is 123456789123, you choose bastion 4, and your user name is JoeSmith:

```
ssh corp.domain.com\\JoeSmith sshbastion1.A123456789123.amazonaws.com
```

2. Log in with your corporate Active Directory credentials.
3. When presented with a Bash prompt, SSH in to the instance, and then enter:

```
ssh DOMAIN_FQDN\\USERNAME@INSTANCE_IP
```

Or, you can use the Login flag (-l):

```
ssh -l DOMAIN_FQDN\\USERNAME@INSTANCE_IP
```

Linux Computer to Windows Instance

Use an SSH tunnel and an RDP client to connect to a Windows instance from your Linux computer.

MALZ

This procedure requires a Remote Desktop Connection client for Linux; the example uses Microsoft Remote Desktop (an open source UNIX client for connecting to Windows Remote Desktop Services). Rdesktop is an alternative.

Note

How you log in to Windows instances might change based on the remote desktop client being used.

First you establish an SSH tunnel, and then log in.

For more information about the friendly bastion names, see [DNS Friendly Bastion Names \(p. 142\)](#).

Before you begin:

- Request access to the instance that you want to connect to; for information, see [Access requests](#).
- Choose a friendly DNS SSH bastion name to connect to; for example:

```
sshbastion(1-4).Your_Domain
```

Which would look like this if your Domain_FQDN is "corp.domain.com", your AMS-managed Your_Domain is "amazonaws.com", you choose bastion "4", and your user name is "JoeSmith":

```
ssh corp.domain.com\\JoeSmith sshbastion4.amazonaws.com
```

- Find the IP address of the instance that you want to connect to; for information, see [Finding an instance ID or IP address](#).

1. Set up RDP over an SSH tunnel from a Linux desktop to a Windows instance. In order to issue the ssh command with the right values, there are a couple of ways to proceed:

- In the Linux shell, set the variables, and then enter the SSH connection command:

```
BASTION="sshbastion(1-4).Your_Domain"  
WINDOWS="Windows_Instance_Private_IP"
```

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Access Examples

```
AD="AD_Account_Number"  
USER="AD_Username"  
ssh -L 3389:$WINDOWS:3389 A$AD\\\\"$USER@$BASTION
```

Example, if the following values are used:

```
BASTION="sshbastion4.A123456789123.amazonaws.com"
```

```
WINDOWS="172.16.3.254"
```

```
AD="ACORP_example"
```

```
USER="john.doe"
```

- Add the variable values directly to the ssh command.

In either case, this is what the rendered request would be (assuming the same set of variable values):

```
ssh -L 3389:172.16.3.254:3389 ACORP_example\\\\"john.doe@myamsadomain.com
```

2. Either: Open your Remote Desktop Client, enter the loopback address and port, 127.0.0.1:3389, and then open the connection.

Or, log in to the Windows instance from a new Linux desktop shell. If you use RDesktop, the command looks like this:

```
rdesktop 127.0.0.1:3389
```

A remote desktop window for the Windows instance appears on your Linux desktop.

Tip

If the remote desktop session fails to start, verify that network connectivity to the Windows instance from the SSH bastion is allowed on port 3389 from the shell in step 1 (replace *private_ip_address_of_windows_instance* appropriately):

```
nc private_ip_address_of_windows_instance 3389 -v -z
```

Success:

```
nc 172.16.0.83 3389 -v -z  
Connection to 172.16.0.83 3389 port [tcp/ms-wbt-server] succeeded  
netstat -anvp | grep 3389  
tcp 0 0 172.16.0.253:48079 172.16.3.254:3389 ESTABLISHED
```

SALZ

This procedure for a single-account landing zone requires a Remote Desktop Connection client for Linux; the example uses Microsoft Remote Desktop (an open source UNIX client for connecting to Windows Remote Desktop Services). Rdesktop is an alternative.

Note

How you log in to Windows instances might change based on the remote desktop client being used.

First you establish an SSH tunnel, and then log in.

For more information about the friendly bastion names, see [DNS Friendly Bastion Names \(p. 142\)](#).

Before you begin:

- Request access to the instance that you want to connect to; for information, see [Access requests](#).
- Choose a friendly DNS SSH bastion name to connect to; for example:

```
sshbastion(1-4).AMSAccountNumber.amazonaws.com
```

Which would look like this if your account number is 123456789123 and you choose bastion 4:

```
sshbastion4.A123456789123.amazonaws.com
```

- Find the IP address of the instance that you want to connect to; for information, see [Finding an instance ID or IP address](#).
1. Set up RDP over an SSH tunnel from a Linux desktop to a Windows instance. In order to issue the `ssh` command with the right values, there are a couple of ways to proceed:
 - In the Linux shell, set the variables, and then enter the SSH connection command:

```
BASTION="sshbastion(1-4).AMSAccountNumber.amazonaws.com"  
WINDOWS="WINDOWS_INSTANCE_PRIVATE_IP"  
AD="AD_ACCOUNT_NUMBER"  
USER="AD_USERNAME"  
ssh -L 3389:$WINDOWS:3389 A$AD\\\$USER@$BASTION
```

Example, if the following values are used:

```
BASTION="sshbastion4.A123456789123.amazonaws.com"  
  
WINDOWS="172.16.3.254"  
  
AD="ACORP_example"  
  
USER="john.doe"
```

- Add the variable values directly to the `ssh` command.

In either case, this is what the rendered request would be (assuming the same set of variable values):

```
ssh -L 3389:172.16.3.254:3389 ACORP_example\\\  
\john.doe@sshbastion4.A123456789123.amazonaws.com
```

2. Either: Open your Remote Desktop Client, enter the loopback address and port, 127.0.0.1:3389, and then open the connection.

Or, log in to the Windows instance from a new Linux desktop shell. If you use RDesktop, the command looks like this:

```
rdesktop 127.0.0.1:3389
```

A remote desktop window for the Windows instance appears on your Linux desktop.

Tip

If the remote desktop session fails to start, verify that network connectivity to the Windows instance from the SSH bastion is allowed on port 3389 from the shell in step 1 (replace `private_ip_address_of_windows_instance` appropriately):


```
nc private_ip_address_of_windows_instance 3389 -v -z
```

Success:

```
nc 172.16.0.83 3389 -v -z
Connection to 172.16.0.83 3389 port [tcp/ms-wbt-server] succeeded
netstat -anvp | grep 3389
tcp    0    0 172.16.0.253:48079 172.16.3.254:3389 ESTABLISHED
```


Windows Computer to Windows Instance

Use Windows Remote Desktop Connection client to connect to a Windows instance from your Windows computer.

MALZ

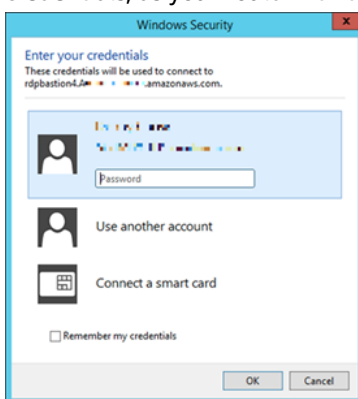
For more information about the friendly bastion names, see [DNS Friendly Bastion Names \(p. 142\)](#).

1. Open the Remote Desktop Connection program, a standard Windows program, and enter the friendly DNS name of the Windows bastion in the hostname field.




2. Choose **Connect**. The Remote Desktop Connection attempts an RDP connection to the bastion.

If successful, a credentials dialog box opens. To gain access, use your corporate Active Directory credentials, as you would with the Windows instance.




3. Open the Remote Desktop Connection program on the bastion and enter the IP address of the Windows instance you would like to connect to (for example, 10.0.0.100), and then choose **Connect**. Your corporate Active Directory credentials are again required before you connect to the Windows instance.



SALZ

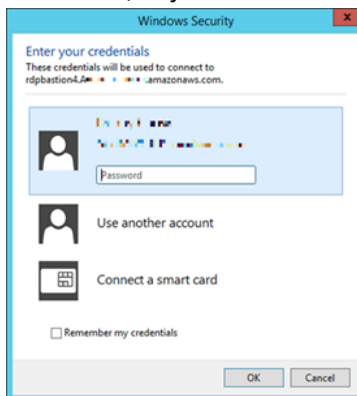
For more information about the friendly bastion names, see [DNS Friendly Bastion Names \(p. 142\)](#).

1. Open the Remote Desktop Connection program, a standard Windows program, and enter the friendly DNS name of the Windows bastion in the hostname field; for example, `rdpbastion(1-4).AAMSAccountNumber.amazonaws.com`, which would look like this if your account number is 123456789123 and you choose bastion 4, `rdpbastion4.A123456789123.amazonaws.com`.




2. Choose **Connect**. The Remote Desktop Connection attempts an RDP connection to the bastion.

If successful, a credentials dialog box opens. To gain access, use your corporate Active Directory credentials, as you would with the Windows instance.



3. Open the Remote Desktop Connection program on the bastion and enter the IP address of the Windows instance you would like to connect to (for example, 10.0.0.100), and then choose **Connect**. Your corporate Active Directory credentials are again required before you connect to the Windows instance.



Windows Computer to Linux Instance

To RDP to an SSH bastion from a Windows environment, follow these steps.

MALZ

Before you begin:

- Request access to the instance that you want to connect to; for information, see [Access requests](#).
- Choose a friendly DNS SSH bastion name to connect to; for example:

```
sshbastion(1-4).YOUR_DOMAIN
```

Which would look like this if YOUR_DOMAIN is myamsaddomain.com" and you choose bastion 4:

```
sshbastion4.myamsaddomain.com
```

- Find the IP address of the instance that you want to connect to; for information, see [Finding an instance ID or IP address](#).

In order to connect to the Linux instance from your Windows machine, you must first connect to an SSH bastion.

Use the native Windows [OpenSSH client](#) or install [PuTTY](#) on your local machine. To learn more about OpenSSH, see [OpenSSH in Windows](#).

1. Use the native Windows or open PuTTY and enter the SSH bastion hostname or the IP address of the SSH bastion. For example, 10.65.2.214 (22 is the port used for SSH; it will be set by default).
2. OpenSSH or PuTTY attempts an SSH connection to the bastion and open a shell window.
3. Use your corporate Active Directory credentials as you would with the RDP hosts to gain access.
4. When presented with a Bash prompt, SSH into the instance. Enter:

```
ssh DOMAIN_FQDN\USERNAME@INSTANCE_IP
```

SALZ

Before you begin:

- Request access to the instance that you want to connect to; for information, see [Access requests](#).
- Choose a friendly DNS SSH bastion name to connect to; for example:

```
sshbastion(1-4).AMSAccountNumber.amazonaws.com
```

Which would look like this if your account number is 123456789123 and you choose bastion 4:

```
sshbastion4.A123456789123.amazonaws.com
```

- Find the IP address of the instance that you want to connect to; for information, see [Finding an instance ID or IP address](#).

In order to connect to the Linux instance from your Windows machine, you must first connect to an SSH bastion.

Use the native Windows [OpenSSH client](#) or install [PuTTY](#) on your local machine. To learn more about OpenSSH, see [OpenSSH in Windows](#).

1. Use the native Windows or open PuTTY and enter the SSH bastion hostname or the IP address of the SSH bastion. For example, 10.65.2.214 (22 is the port used for SSH; it will be set by default).
2. OpenSSH or PuTTY attempts an SSH connection to the bastion and open a shell window.
3. Use your corporate Active Directory credentials as you would with the RDP hosts to gain access.
4. When presented with a Bash prompt, SSH into the instance. Enter:

```
ssh DOMAIN_FQDN\USERNAME@INSTANCE_IP
```

Reporting an Incident

Use the AMS console to report an incident. It's important to create a new incident for each new issue or question. When opening cases related to old inquiries, it's helpful to include the related case number so we can refer to previous correspondence.

Note

If case correspondence strays from the original issue, an AMS operator might ask you to report a new incident.

To report an incident using the AMS console:

1. From the left navigation, choose **Incidents**

The **Incidents** list opens:

Managed Services > Incidents

AWS Managed Services phone and chat operational support
Connect with AMS engineers through phone or chat, in addition to using case correspondence, using Support Center. Click the button below to directly create an incident in Support Center. When going to Support Center on your own, choose incidents or service requests using the Service dropdown menu under Technical Support.

Create incident in Support Center

Incidents Create incident

All open < 1 ... >

Created	Subject	ID	Status
---------	---------	----	--------

If your incident list is empty, the **Clear filter** option resets the filter to **Any status**.

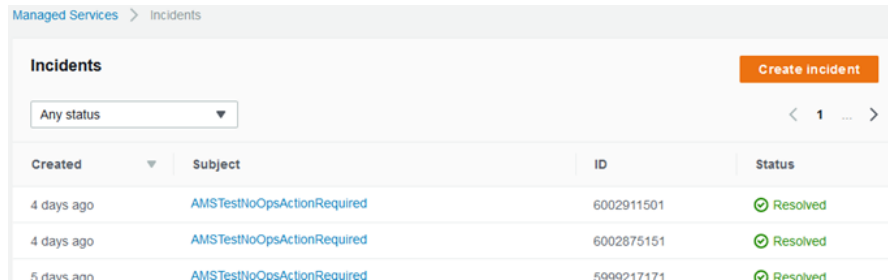
If you know you want to use phone or chat, click **Create incident in Support Center** to open the incident **Create** page in the AWS Support Center, auto-populated with the AMS service type.

Note

Phone calls initiated with AWS Support center are recorded, to better improve response. If the call drops, you must call back through the Support Center case, AWS has no mechanism for calling you back.

Important

Phone and chat support is designed to help with support cases, incidents and service requests. For RFC issues, use the correspondence option on the relevant RFC details page, to reach an AMS engineer.



2. If you want to find an existing incident, select an incident status filter in the drop-down list.

<p>The dropdown menu lists the following status options: All open, Unassigned, Open, Reopened, Work in progress, Pending customer action, Customer action completed, Resolved, and Any status.</p>	<ul style="list-style-type: none">• All incidents that are not yet resolved.• A new incident that is not yet assigned.• An incident that has been assigned.• An incident that you reopened.• An assigned, complicated incident.• Incidents that require your feedback before the next step.• Incidents to which you have recently submitted information.• An incident that has concluded.• All incidents in the account.
--	--

3. Choose **Create**.

The **Create an incident** page opens:

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Reporting an Incident

Incident details

Priority

Low
Non-critical functions of your business service or application related to AWS/AMS resources are impacted.

Medium
A business service or application related to AWS/AMS resources is moderately impacted and functioning in a degraded state.

High
Your business is significantly impacted. Critical functions of your application related to AWS/AMS resources are unavailable. Reserved for the most critical outages affecting production systems.

Access Issues ▼

Subject

Can't Access Instance

CC Emails - *optional*
Email addresses added here will receive notifications when this case is updated

johndoe@example.com X

Details
Use the template below to help describe your issue

What is not functioning properly? Not sure.

When did you notice the disruption? Just now.

What is the impact of the disruption? Can't deploy.

Any additional details to help solve the incident:

Attach files - *optional*

Add Attachment

4. Select a **Priority**:

- **Low**: Non-critical functions of your business service or application related to AWS/AMS resources are impacted.
- **Medium**: A business service or application related to AWS/AMS resources is moderately impacted and functioning in a degraded state.
- **High**: Your business is significantly impacted. Critical functions of your application related to AWS/AMS resources are unavailable. Reserved for the most critical outages affecting production systems.

5. Select a **Category**:


- **Access Issues**: You have a question about accessing your AMS-managed resources.
- **Availability**: A resource appears to be unavailable.
- **Performance Issue**: A resource seems to be under-performing.
- **Security Related**: You have a security concern about your AMS-managed resources.
- **Other**: None of the other categories apply.

Note

If you are going to test incident functionality, AMS asks that you add the no-action flag (AMSTestNoOpsActionRequired) to your incident title.

6. Enter information for:

- **Subject:** A descriptive title for the incident report.
- **CC emails:** A list of email addresses for people you want informed about the incident report and resolution.
- **Details:** A comprehensive description of the incident, the systems impacted, and the expected outcome of the resolution. Answer the pre-set questions, or delete them and enter any relevant information.

To add an attachment, choose **Add Attachment**, browse to the attachment you want, and click **Open**. To delete the attachment, click the Delete icon: .

7. Choose **Submit**.


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

Click **Reply** to open a correspondence area and provide additional details or updates in status.

Click **Close Case** when the incident has been resolved.

Click **Load More** if there is more correspondence than will fit on one page.

Don't forget to rate the communication!



Your incident displays on the **Incidents** list page.


Creating a Service Request

To create a service request using the AMS console:

1. From the left navigation, choose **Service requests**.

The **Service requests** list opens.

Managed Services > Service requests

 **AWS Managed Services phone and chat operational support**
Connect with AMS engineers through phone or chat, in addition to using case correspondence, using Support Center. Click the button below to directly create a service request in Support Center. When going to Support Center on your own, choose incidents or service requests using the Service dropdown menu under Technical Support.

[Create service request in Support Center](#)

Service requests [Create service request](#)

All open < 1 ... >

Created	Subject	ID	Status
---------	---------	----	--------

If your service request list is empty, the **Clear filter** option resets the filter to **Any status**.

Managed Services > Service requests

Service requests [Create service request](#)

Any status < 1 ... >

Created	Subject	ID	Status
4 days ago	AMSTestNoOpsActionRequired	6002895311	Resolved
4 days ago	AMSTestNoOpsActionRequired	6002955301	Resolved
4 days ago	AMSTestNoOpsActionRequired	6002955301	Resolved

If you know you want to use phone or chat, click **Create service request in Support Center** to open the service request **Create** page in the AWS Support Center, auto-populated with the AMS service type.

Note

Phone calls initiated with AWS Support center are recorded, to better improve response. If the call drops, you must call back through the Support Center case, AWS has no mechanism for calling you back.

Important

Phone and chat support is designed to help with support cases, incidents and service requests. For RFC issues, use the correspondence option on the relevant RFC details page, to reach an AMS engineer.

2. If you want to find an existing service request, select a service request status filter in the drop-down list.

<ul style="list-style-type: none">All openUnassignedOpenReopenedWork in progressPending customer actionCustomer action completedResolvedAny status	<ul style="list-style-type: none">All service requests that are not yet resolved.A new service request that is not yet assigned.A service request that has been assigned.A service request that you reopened.An assigned, complicated, service request.Service requests that require your feedback before the next step.Service requests to which you have recently submitted information.A service request that has concluded.All service requests in the account.
--	---

3. Choose **Create**.

The **Create a service request** page opens.

Managed Services > Service requests > Create a service request

Create a service request

Category
Click the dropdown menu to select the category of this service request

Feature Request

Subject
Type the subject of this service request

AMSTestNoOpsActionRequired

CC Emails - optional
Email addresses added here will receive notifications when this case is updated

Details
Type the details of the service request

Testing service request functionality

Add Attachment

Submit

4. Select a **Category**:


- Access:** Use this when you have a question about accessing your AMS-managed resources. To request access to an AMS-managed resource, submit an RFC with the AccessManagement category.
- Alert notification:** Use this when you have an alert and have not heard from AMS.
- Feature Request:** Use this to request that AMS add a feature.
- General Guidance:** Use this for non-resource specific questions.
- Security Related:** Use this when you have a security concern about your AMS-managed resources. Note that while we use encryption, you should exercise caution with the information you submit here.
- Service Reporting Query:** Use this to request a specific report.
- Other:** Use this when none of the other categories apply.

Note

If you are going to test service request functionality, AMS asks that you add the no-action flag (AMSTestNoOpsActionRequired) to your service request title.

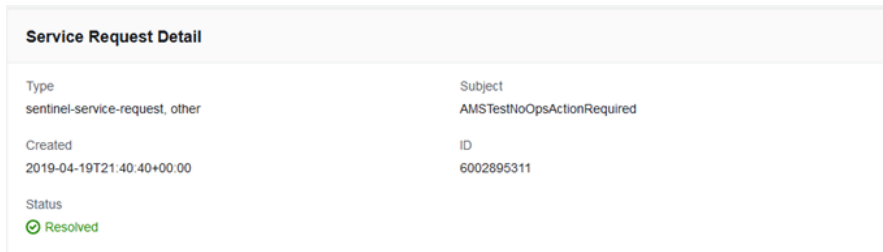
5. Enter information for:


- **Subject:** This creates a link to the service request details on the list page.
- **CC emails:** These emails receive correspondence in addition to your default email contacts.
- **Details:** Provide as much information here as possible.

To add an attachment, choose **Add Attachment**, browse to the attachment you want, and click **Open**. To delete the attachment, click the Delete icon: .

6. Choose **Submit**.

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



Service Request Detail	
Type	Subject
sentinel-service-request, other	AMSTestNoOpsActionRequired
Created	ID
2019-04-19T21:40:40+00:00	6002895311
Status	
 Resolved	

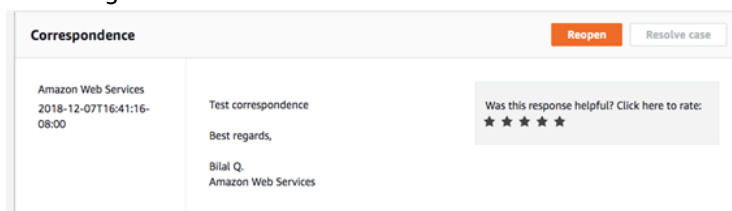
Additionally, your service request displays on the **Service Request** list page. Use this when you have an alert but have not yet heard from AMS.

Click **Reply** to open a correspondence area and provide additional details or status updates.

Click **Resolve Case** when the service request has been resolved.

Click **Load More** to view additional correspondences that do not fit on the initial page.

Don't forget to rate the communication!



Correspondence		Reopen	Resolve case
Amazon Web Services 2018-12-07T16:41:16-08:00	Test correspondence Best regards, Bilal Q. Amazon Web Services	Was this response helpful? Click here to rate: ★ ★ ★ ★ ★	

For billing-related queries, use the **Other** Category in the AMS console; the `ChangeTypeId ct-1e1xtak34nx76` in the AMS CM API, or the `IssueType=AMS` in the AWS Support API.

Next Steps

Now that you've onboarded an AMS account, you'll want to read more AMS documentation. See these documents:

- The tutorials for using the HA Two-Tier stack CT to create a fully-functioning WordPress stack, next, provides a full AMS experience.

- [AMS User Guide](#): The AMS User Guide describes AMS functionality, lists key terms, operations, interfaces and provides an overview of a typical AMS managed-infrastructure architecture. Additionally, access management details and AMS defaults are given. Also provided are detailed descriptions of how to use the AMS change management system and several walkthroughs are provided. Additional management concepts are described as well.
- [AMS API Reference](#): This API reference provides descriptions of all API calls, including request, response, and examples.
- [AMS Application Guide](#): The AMS Application Guide describes different options and methods for deploying and maintaining your applications in AMS.

Tutorials

The following tutorials detail the steps to creating a two-tier stack with the High Availability (advanced) CT (ct-06mjngx5flwto), using the CLI and using the Console. A tutorial is given for deploying a Linux Auto Scaling group (ASG) and for deploying a Windows ASG.

Descriptions for all CT options, including ChangeTypeId can be found in the [AMS Change Type Reference](#).

CLI Tutorial: High Availability Two-Tier Stack (Linux/RHEL)

This section describes how to deploy a high availability (HA) two-tier stack into an AMS environment using the AMS CLI.

Note

This deployment walkthrough has been tested in AMZN Linux and RHEL environments.

Summary of tasks and required RFCs:

1. Create infrastructure (HA two-tier stack)
2. Create an S3 bucket for CodeDeploy applications
3. Create the WordPress application bundle and upload it to the S3 bucket
4. Deploy the application with CodeDeploy
5. Access the WordPress site and log in to validate the deployment

Before You Begin

The Deployment | Advanced Stack Components | High Availability Two Tier Stack Advanced | Create CT creates an Auto Scaling group, a load balancer, a database, and a CodeDeploy application name and deployment group (with the same name that you give the application). For information on CodeDeploy see [What is CodeDeploy?](#)

This walkthrough uses a High Availability Two-Tier Stack (Advanced) RFC that includes UserData and also describes how to create a WordPress bundle that CodeDeploy can deploy.

The `UserData` shown in the example gets instance metadata such as instance ID, region, etc, from within a running instance by querying the EC2 instance metadata service available at `http://169.254.169.254/latest/meta-data/`. This line in the user data script: `REGION=$(curl 169.254.169.254/latest/meta-data/placement/availability-zone/ | sed 's/[a-z]$/')`, retrieves the availability zone name from the meta-data service into the `$REGION` variable for our supported regions, and uses it to complete the URL for the S3 bucket where the CodeDeploy agent is downloaded. The 169.254.169.254 IP is routable only within the VPC (all VPCs can query the service). For information about the service, see [Instance Metadata and User Data](#). Note also that scripts entered as `UserData` are executed as the "root" user and do not need to use the "sudo" command.

This walkthrough leaves the following parameters at the default value (shown):

- Auto Scaling group: `Cooldown=300, DesiredCapacity=2, EBSOptimized=false, HealthCheckGracePeriod=600, IAMInstanceProfile=customer-mc-ec2-instance-profile, InstanceDetailedMonitoring=true, InstanceRootVolumeIops=0, InstanceRootVolumeType=standard, InstanceType=m3.medium, MaxInstances=2, MinInstances=2, ScaleDownPolicyCooldown=300, ScaleDownPolicyEvaluationPeriods=4, ScaleDownPolicyPeriod=60, ScaleDownPolicyScalingAdjustment=-1, ScaleDownPolicyStatistic=Average, ScaleDownPolicyThreshold=35, ScaleMetricName=CPUUtilization, ScaleUpPolicyCooldown=60, ScaleUpPolicyEvaluationPeriods=2, ScaleUpPolicyPeriod=60, ScaleUpPolicyScalingAdjustment=2, ScaleUpPolicyStatistic=Average, ScaleUpPolicyThreshold=75.`
- Load Balancer: `HealthCheckInterval=30, HealthCheckTimeout=5.`
- Database: `BackupRetentionPeriod=7, Backups=true, InstanceType=db.m3.medium, IOPS=0, MultiAZ=true, PreferredBackupWindow=22:00-23:00, PreferredMaintenanceWindow=wed:03:32-wed:04:02, StorageEncrypted=false, StorageEncryptionKey="", StorageType=gp2.`
- Application: `DeploymentConfigName=CodeDeployDefault.OneAtATime.`
- S3 bucket: `AccessControl=Private.`

ADDITIONAL SETTINGS:

`RequestedStartTime` and `RequestedEndTime` if you want to schedule your RFC: You can use [Time.is](#) to determine the correct UTC time. The examples provided must be adjusted appropriately. An RFC cannot proceed if the start time has passed. Alternatively, you can leave those values off to create an ASAP RFC that executes as soon as approvals are passed.

Note

There are many parameters that you might choose to set differently than as shown. The values for those parameters shown in the example have been tested but may not be right for you.

Create the Infrastructure

Gathering the following data before you begin will make the deployment go more quickly.

REQUIRED DATA HA STACK:

- AutoScalingGroup:
 - `UserData`: This value is provided in this tutorial. It includes commands to set up the resource for CodeDeploy and start the CodeDeploy agent.
 - `AMI-ID`: This value determines what kind of EC2 instances your Auto Scaling group (ASG) will spin up. Be sure to select an AMI in your account that starts with "customer-" and is of the operating system that you want. Find AMI IDs with the [ListAmis](#) operation (CLI: `list-amis`) or in the AMS Console VPCs -> VPCs details page. This walkthrough is for ASGs configured to use a Linux AMI.
- Database:
 - These parameters, `DBEngine`, `EngineVersion`, and `LicenseModel` should be set according to your situation though the values shown in the example have been tested.
 - These parameters, `RDSSubnetIds`, `DBName`, `MasterUsername`, and `MasterUserPassword` are required when deploying the application bundle. For `RDSSubnetIds`, use two Private subnets.
- LoadBalancer:
 - These parameters, `DBEngine`, `EngineVersion`, and `LicenseModel` should be set according to your situation though the values shown in the example have been tested.
 - `ELBSubnetIds`: Use two Public subnets.
- Application: The `ApplicationName` value sets the CodeDeploy application name and CodeDeploy deployment group name. You use it to deploy your application. It must be unique in the account.

To check your account for CodeDeploy names, see the CodeDeploy Console. The example uses "WordPress" but, if you will use that value, make sure that it is not already in use.

This procedure utilizes the High availability two-tier stack (advanced) CT (ct-06mjngx5flwto) and the Create S3 storage CT (ct-1a68ck03fn98r). From your authenticated account, follow these steps at the command line.

1. Launch the infrastructure stack.
 - a. Output the execution parameters JSON schema for the HA two tier stack CT to a file in your current folder named CreateStackParams.json.

```
aws amscm get-change-type-version --change-type-id "ct-06mjngx5flwto" --query  
"ChangeTypeVersion.ExecutionInputSchema" --output text > CreateStackParams.json
```

- b. Modify the schema. Replace the *variables* as appropriate. For example, use the OS that you want for the EC2 instances the ASG will create. Record the ApplicationName as you will use it later to deploy the application. Note that you can add up to 50 tags.

```
{  
  "Description":      "HA two tier stack for WordPress",  
  "Name":             "WordPressStack",  
  "TimeoutInMinutes": 360,  
  "Tags": [            
    {                  
      "Key": "ApplicationName",  
      "Value": "WordPress"  
    }                  
  ],  
  "AutoScalingGroup": {  
    "AmiId": "AMI-ID",  
    "UserData": "#!/bin/bash \\  
REGION=$(curl 169.254.169.254/latest/meta-data/placement/availability-  
zone/ | sed 's/[a-z]$/') \\  
yum -y install ruby httpd \\  
chkconfig httpd on \\  
service httpd start \\  
touch /var/www/html/status \\  
cd /tmp \\  
curl -O https://aws-codedeploy-$REGION.s3.amazonaws.com/latest/install  
\  
  \n  
    chmod +x ./install \\  
    ./install auto \\  
    chkconfig codedeploy-agent on \\  
    service codedeploy-agent start"  
  },  
  "LoadBalancer": {  
    "Public": true,  
    "HealthCheckTarget": "HTTP:80/status"  
  },  
  "Database": {  
    "DBEngine": "MySQL",  
    "DBName": "wordpress",  
    "EngineVersion": "8.0.16",  
    "LicenseModel": "general-public-license",  
    "MasterUsername": "admin",  
    "MasterUserPassword": "p4ssw0rd"  
  },  
  "Application": {  
    "ApplicationName": "WordPress"  
  }  
}
```

```
}
```

- c. Output the CreateRfc JSON template to a file in your current folder named CreateStackRfc.json:

```
aws amscm create-rtc --generate-cli-skeleton > CreateStackRfc.json
```

- d. Modify the RFC template as follows and save it, you can delete and replace the contents. Note that RequestedStartTime and RequestedEndTime are now optional; excluding them creates an ASAP RFC that executes as soon as it is approved (which usually happens automatically). To submit a scheduled RFC, add those values.

```
{  
  "ChangeTypeVersion": "3.0",  
  "ChangeTypeId": "ct-06mjngx5flwto",  
  "Title": "HA-Stack-For-WP-RFC"  
}
```

- e. Create the RFC, specifying the CreateStackRfc.json file and the CreateStackParams.json execution parameters file:

```
aws amscm create-rtc --cli-input-json file://CreateStackRfc.json --execution-  
parameters file://CreateStackParams.json
```

You receive the RFC ID in the response. Save the ID for subsequent steps.

- f. Submit the RFC:

```
aws amscm submit-rtc --rtc-id RFC_ID
```

If the RFC succeeds, you receive no output.

- g. To check RFC status, run

```
aws amscm get-rtc --rtc-id RFC_ID
```

Keep note of the RFC ID.

2. Launch an S3 bucket

Gathering the following data before you begin will make the deployment go more quickly.

REQUIRED DATA S3 BUCKET:

- **VPC-ID:** This value determines where your S3 Bucket will be. Use the same VPC ID that you used previously.
- **BucketName:** This value sets the S3 Bucket name, you use it to upload your application bundle. It must be unique across the region of the account and cannot include upper-case letters. Including your account ID as part of the BucketName is not a requirement but makes it easier to identify the bucket later. To see what S3 bucket names exist in the account, go to the Amazon S3 Console for your account.

- a. Output the execution parameters JSON schema for the S3 storage create CT to a JSON file named CreateS3StoreParams.json.

```
aws amscm get-change-type-version --change-type-id "ct-1a68ck03fn98r" --query  
"ChangeTypeVersion.ExecutionInputSchema" --output text > CreateS3StoreParams.json
```

- b. Modify the schema as follows, you can delete and replace the contents. Replace `VPC_ID` appropriately. The values in the example have been tested, but may not be right for you.

Tip

The `BucketName` must be unique across the region of the account and cannot include upper-case letters. Including your account ID as part of the `BucketName` is not a requirement but makes it easier to identify the bucket later. To see what S3 bucket names exist in the account, go to the Amazon S3 Console for your account.

```
{
  "Description":      "S3BucketForWordPressBundle",
  "VpcId":            "VPC_ID",
  "StackTemplateId": "stm-s2b72beb000000000",
  "Name":             "S3BucketForWP",
  "TimeoutInMinutes": 60,
  "Parameters":      {
    "AccessControl": "Private",
    "BucketName":    "ACCOUNT_ID-BUCKET_NAME"
  }
}
```

- c. Output the JSON template for `CreateRfc` to a file, in your current folder, named `CreateS3StoreRfc.json`:

```
aws amscm create-rtc --generate-cli-skeleton > CreateS3StoreRfc.json
```

- d. Modify and save the `CreateS3StoreRfc.json` file, you can delete and replace the contents. Note that `RequestedStartTime` and `RequestedEndTime` are now optional; excluding them creates an ASAP RFC that executes as soon as it is approved (which usually happens automatically). To submit a scheduled RFC, add those values.

```
{
  "ChangeTypeVersion": "1.0",
  "ChangeTypeId":      "ct-1a68ck03fn98r",
  "Title":              "S3-Stack-For-WP-RFC"
}
```

- e. Create the RFC, specifying the `CreateS3StoreRfc.json` file and the `CreateS3StoreParams.json` execution parameters file:

```
aws amscm create-rtc --cli-input-json file://CreateS3StoreRfc.json --execution-parameters file://CreateS3StoreParams.json
```

You receive the `RfcId` of the new RFC in the response. Save the ID for subsequent steps.

- f. Submit the RFC:

```
aws amscm submit-rtc --rtc-id RFC_ID
```

If the RFC succeeds, you receive no output.

- g. To check RFC status, run

```
aws amscm get-rtc --rtc-id RFC_ID
```

Create, Upload, and Deploy the Application

First, create a WordPress application bundle, and then use the CodeDeploy CTs to create and deploy the application.

1. Download WordPress, extract the files and create a `./scripts` directory.

Linux command:

```
wget https://github.com/WordPress/WordPress/archive/master.zip
```

Windows: Paste `https://github.com/WordPress/WordPress/archive/master.zip` into a browser window and download the zip file.

Create a temporary directory in which to assemble the package.

Linux:

```
mkdir /tmp/WordPress
```

Windows: Create a "WordPress" directory, you will use the directory path later.

2. Extract the WordPress source to the "WordPress" directory and create a `./scripts` directory.

Linux:

```
unzip master.zip -d /tmp/WordPress_Temp  
cp -paf /tmp/WordPress_Temp/WordPress-master/* /tmp/WordPress  
rm -rf /tmp/WordPress_Temp  
rm -f master  
cd /tmp/WordPress  
mkdir scripts
```

Windows: Go to the "WordPress" directory that you created and create a "scripts" directory there.

If you are in a Windows environment, be sure to set the break type for the script files to Unix (LF). In Notepad ++, this is an option at the bottom right of the window.

3. Create the CodeDeploy `appspec.yml` file, in the WordPress directory (if copying the example, check the indentation, each space counts). IMPORTANT: Ensure that the "source" path is correct for copying the WordPress files (in this case, in your WordPress directory) to the expected destination (`/var/www/html/WordPress`). In the example, the `appspec.yml` file is in the directory with the WordPress files, so only `/` is needed. Also, even if you used a RHEL AMI for your Auto Scaling group, leave the "os: linux" line as-is. Example `appspec.yml` file:

```
version: 0.0  
os: linux  
files:  
  - source: /  
    destination: /var/www/html/WordPress  
hooks:  
  BeforeInstall:  
    - location: scripts/install_dependencies.sh  
      timeout: 300  
      runas: root  
  AfterInstall:  
    - location: scripts/config_wordpress.sh  
      timeout: 300  
      runas: root  
ApplicationStart:
```



```
- location: scripts/start_server.sh
  timeout: 300
  runas: root
ApplicationStop:
- location: scripts/stop_server.sh
  timeout: 300
  runas: root
```

4. Create bash file scripts in the WordPress `./scripts` directory.

First, create `config_wordpress.sh` with the following content (if you prefer, you can edit the `wp-config.php` file directly).

Note

Replace `DBName` with the value given in the HA Stack RFC (for example, `wordpress`).

Replace `DB_MasterUsername` with the `MasterUsername` value given in the HA Stack RFC (for example, `admin`).

Replace `DB_MasterUserPassword` with the `MasterUserPassword` value given in the HA Stack RFC (for example, `p4ssw0rd`).

Replace `DB_ENDPOINT` with the endpoint DNS name in the execution outputs of the HA Stack RFC (for example, `srt1cz23n45sfg.clgvd67uvydk.us-east-1.rds.amazonaws.com`). You can find this with the `GetRfc` operation (CLI: `get-rtc --rtc-id RFC_ID`) or in the AMS Console RFC details page for the HA Stack RFC that you previously submitted.

```
#!/bin/bash
chmod -R 755 /var/www/html/WordPress
cp /var/www/html/WordPress/wp-config-sample.php /var/www/html/WordPress/wp-config.php
cd /var/www/html/WordPress
sed -i "s/database_name_here/DBName/g" wp-config.php
sed -i "s/username_here/DB_MasterUsername/g" wp-config.php
sed -i "s/password_here/DB_MasterUserPassword/g" wp-config.php
sed -i "s/localhost/DB_ENDPOINT/g" wp-config.php
```

5. In the same directory create `install_dependencies.sh` with the following content:

```
#!/bin/bash
yum install -y php
yum install -y php-mysql
yum install -y mysql
service httpd restart
```

Note

HTTPS is installed as part of the user data at launch in order to allow health checks to work from the start.

6. In the same directory create `start_server.sh` with the following content:

- For Amazon Linux instances, use this:

```
#!/bin/bash
service httpd start
```

- For RHEL instances, use this (the extra commands are policies that allow SELINUX to accept WordPress):

```
#!/bin/bash
setsebool -P httpd_can_network_connect_db 1
setsebool -P httpd_can_network_connect 1
chcon -t httpd_sys_rw_content_t /var/www/html/WordPress/wp-content -R
restorecon -Rv /var/www/html
```

```
service httpd start
```

7. In the same directory create `stop_server.sh` with the following content:

```
#!/bin/bash  
service httpd stop
```

8. Create the zip bundle.

Linux:

```
$ cd /tmp/WordPress  
$ zip -r wordpress.zip .
```

Windows: Go to your "WordPress" directory and select all of the files and create a zip file, be sure to name it `wordpress.zip`.

1. Upload the application bundle to the S3 bucket.

The bundle needs to be in place in order to continue deploying the stack.

You automatically have access to any S3 bucket instance that you create. You can access it through your bastions, or through the S3 console, and upload the WordPress bundle with drag-and-drop or browsing to and selecting the zip file.

You can also use the following command in a shell window; be sure that you have the correct path to the zip file:

```
aws s3 cp wordpress.zip s3://BUCKET_NAME/
```

2. Deploy the WordPress application bundle.

Gathering the following data before you begin will make the deployment go more quickly.

REQUIRED DATA:

- **VPC-ID:** This value determines where your S3 Bucket will be. Use the same VPC ID that you used previously.
 - **CodeDeployApplicationName and CodeDeployApplicationName:** The ApplicationName value you used in the HA 2-Tier Stack RFC set the CodeDeployApplicationName and the CodeDeployDeploymentGroupName. The example uses "WordPress" but you may have used a different value.
 - **S3Location:** For S3Bucket, use the BucketName that you previously created. The S3BundleType and S3Key are from the bundle that you put on your S3 store.
- a. Output the execution parameters JSON schema for the CodeDeploy application deploy CT to a JSON file named `DeployCDAppParams.json`.

```
aws amscm get-change-type-version --change-type-id "ct-2edc3sd1sqmrb" --query  
"ChangeTypeVersion.ExecutionInputSchema" --output text > DeployCDAppParams.json
```

- b. Modify the schema as follows and save it as, you can delete and replace the contents.

```
{  
  "Description": "DeployWPCDApp",  
  "VpcId": "vpc-2892020",
```

```
"Name": "WordPressCDAppDeploy",
"TimeoutInMinutes": 60,
"Parameters": {
  "CodeDeployApplicationName": "WordPress",
  "CodeDeployDeploymentGroupName": "WordPress",
  "CodeDeployIgnoreApplicationStopFailures": false,
  "CodeDeployRevision": {
    "RevisionType": "S3",
    "S3Location": {
      "S3Bucket": "BUCKET_NAME",
      "S3BundleType": "zip",
      "S3Key": "wordpress.zip" }
    }
  }
}
```

- c. Output the JSON template for CreateRfc to a file, in your current folder, named DeployCDAppRfc.json:

```
aws amscm create-rtc --generate-cli-skeleton > DeployCDAppRfc.json
```

- d. Modify and save the DeployCDAppRfc.json file, you can delete and replace the contents. Note that RequestedStartTime and RequestedEndTime are now optional; excluding them creates an ASAP RFC that executes as soon as it is approved (which usually happens automatically). To submit a scheduled RFC, add those values.

```
{
  "ChangeTypeVersion": "1.0",
  "ChangeTypeId": "ct-2edc3sd1sqmrb",
  "Title": "CD-Deploy-For-WP-RFC"
}
```

- e. Create the RFC, specifying the DeployCDAppRfc file and the DeployCDAppParams execution parameters file:

```
aws amscm create-rtc --cli-input-json file://DeployCDAppRfc.json --execution-parameters file://DeployCDAppParams.json
```

You receive the RfcId of the new RFC in the response. Save the ID for subsequent steps.

- f. Submit the RFC:

```
aws amscm submit-rtc --rtc-id RFC_ID
```

If the RFC succeeds, you receive no output.

- g. To check RFC status, run

```
aws amscm get-rtc --rtc-id RFC_ID
```

Validate the Application Deployment

Navigate to the endpoint (ELB CName) of the previously-created load balancer, with the WordPress deployed path: /WordPress. For example:

```
http://stack-ID-FOR-ELB.us-east-1.elb.amazonaws.com/WordPress
```

Tear Down the Application Deployment

Once you are finished with the tutorial, you will want to tear down the deployment so you are not charged for the resources.

The following is a generic stack delete operation. You'll want to submit it twice, once for the HA 2-Tier stack and once for the S3 bucket stack. As a final follow-through, submit a service request that all snapshots for the S3 bucket (include the S3 bucket stack ID in the service request) be deleted. They are automatically deleted after 10 days, but deleting them early saves a little bit of cost.

This walkthrough provides an example of using the AMS console to delete an S3 stack; this procedure applies to deleting any stack using the AMS console.

Note

If deleting an S3 bucket, it must be emptied of objects first.

REQUIRED DATA:

- **StackId**: The stack to use. You can find this by looking at the AMS Console **Stacks** page, available through a link in the left nav. Using the AMS SKMS API/CLI, run the [ListStackSummaries](#) operation (`list-stack-summaries` in the CLI).
- The change type ID for this walkthrough is `ct-0q0bic0ywqk6c`, the version is "1.0", to find out the latest version, run this command:

```
aws amscm list-change-type-version-summaries --filter  
Attribute=ChangeTypeId,Value=ct-0q0bic0ywqk6c
```

INLINE CREATE:

- Issue the create RFC command with execution parameters provided inline (escape quotes when providing execution parameters inline). E

```
aws amscm create-rfc --change-type-id "ct-0q0bic0ywqk6c" --change-type-version "1.0" --  
title "Delete My Stack" --execution-parameters "{\"StackId\":\"STACK_ID\"}"
```

- Submit the RFC using the RFC ID returned in the create RFC operation. Until submitted, the RFC remains in the `Editing` state and is not acted on.

```
aws amscm submit-rfc --rfc-id RFC_ID
```

- Monitor the RFC status and view execution output:

```
aws amscm get-rfc --rfc-id RFC_ID
```

TEMPLATE CREATE:

1. Output the RFC template to a file in your current folder; example names it `DeleteStackRfc.json`:

```
aws amscm create-rfc --generate-cli-skeleton > DeleteStackRfc.json
```

2. Modify and save the `DeleteStackRfc.json` file. Since deleting a stack has only one execution parameter, the execution parameters can be in the `DeleteStackRfc.json` file itself (there is no need to create a separate JSON file with execution parameters).

The internal quotation marks in the ExecutionParameters JSON extension must be escaped with a backslash (`\`). Example without start and end time:

```
{
  "ChangeTypeVersion": "1.0",
  "ChangeTypeId": "ct-0q0bic0ywqk6c",
  "Title": "Delete-My-Stack-RFC"
  "ExecutionParameters": "{
    \"StackId\": \"STACK_ID\"
  }
}
```

3. Create the RFC:

```
aws amscm create-rfc --cli-input-json file://DeleteStackRfc.json
```

You receive the RfcId of the new RFC in the response. For example:

```
{
  "RfcId": "daaa1867-ffc5-1473-192a-842f6b326102"
}
```

Save the ID for subsequent steps.

4. Submit the RFC:

```
aws amscm submit-rfc --rfc-id RFC_ID
```

If the RFC succeeds, you receive no confirmation at the command line.

5. To monitor the status of the request and to view Execution Output:

```
aws amscm get-rfc --rfc-id RFC_ID --query "Rfc.
{Status:Status.Name,Exec:ExecutionOutput}" --output table
```

Console Tutorial: High Availability Two Tier Stack (Linux/RHEL)

This section describes how to deploy a high availability (HA) WordPress site into an AMS environment using the AMS console.

Note

This deployment walkthrough has been tested in AMZN Linux and RHEL environments.

Summary of tasks and required RFCs:

1. Create infrastructure (HA two-tier stack)
2. Create an S3 bucket for CodeDeploy applications
3. Create the WordPress application bundle and upload it to the S3 bucket
4. Deploy the application with CodeDeploy
5. Access the WordPress site and log in to validate the deployment
6. Tear down the deployment

Descriptions for all CT options, including ChangeTypeId, can be found in [AMS Change Type Reference](#).

Before You Begin

The Deployment | Advanced Stack Components | High Availability Two Tier Stack | Create CT creates an Auto Scaling group, a load balancer, a database, and a CodeDeploy application name and deployment

group (with the same name that you give the application). For information on CodeDeploy see [What is CodeDeploy?](#)

This walkthrough uses a High Availability Two-Tier Stack RFC that includes UserData and also describes how to create a WordPress bundle that CodeDeploy can deploy.

The UserData shown in the example gets instance metadata such as instance ID, region, etc, from within a running instance by querying the EC2 instance metadata service available at <http://169.254.169.254/latest/meta-data/>. This line in the user data script: `REGION=$(curl 169.254.169.254/latest/meta-data/placement/availability-zone/ | sed 's/[a-z]$/')`, retrieves the availability zone name from the meta-data service into the \$REGION variable for our supported regions, and uses it to complete the URL for the S3 bucket where the CodeDeploy agent is downloaded. The 169.254.169.254 IP is routable only within the VPC (all VPCs can query the service). For information about the service, see [Instance Metadata and User Data](#). Note also that scripts entered as UserData are executed as the "root" user and do not need to use the "sudo" command.

This walkthrough leaves the following parameters at the default value (shown):

- Auto Scaling group: `Cooldown=300, DesiredCapacity=2, EBSoptimized=false, HealthCheckGracePeriod=600, IAMInstanceProfile=customer-mc-ec2-instance-profile, InstanceDetailedMonitoring=true, InstanceRootVolumeIops=0, InstanceRootVolumeType=standard, InstanceType=m3.medium, MaxInstances=2, MinInstances=2, ScaleDownPolicyCooldown=300, ScaleDownPolicyEvaluationPeriods=4, ScaleDownPolicyPeriod=60, ScaleDownPolicyScalingAdjustment=-1, ScaleDownPolicyStatistic=Average, ScaleDownPolicyThreshold=35, ScaleMetricName=CPUUtilization, ScaleUpPolicyCooldown=60, ScaleUpPolicyEvaluationPeriods=2, ScaleUpPolicyPeriod=60, ScaleUpPolicyScalingAdjustment=2, ScaleUpPolicyStatistic=Average, ScaleUpPolicyThreshold=75.`
- Load Balancer: `HealthCheckInterval=30, HealthCheckTimeout=5.`
- Database: `BackupRetentionPeriod=7, Backups=true, InstanceType=db.m3.medium, IOPS=0, MultiAZ=true, PreferredBackupWindow=22:00-23:00, PreferredMaintenanceWindow=wed:03:32-wed:04:02, StorageEncrypted=false, StorageEncryptionKey="", StorageType=gp2.`
- Application: `DeploymentConfigName=CodeDeployDefault.OneAtATime.`

Variable Parameters:

The Console provides an **ASAP** option for the start time and this walkthrough recommends using it. **ASAP** causes the RFC to be executed as soon as approvals are passed.

Note

There are many parameters that you might choose to set differently than as shown. The values for those parameters shown in the example have been tested but may not be right for you. Only required values are shown in the examples. Values in *replaceable* font should be changed as they are particular to your account.

Create the Infrastructure

This procedure utilizes the High availability two-tier stack CT followed by the Create S3 storage CT.

Gathering the following data before you begin will make the deployment go more quickly.

REQUIRED DATA HA STACK:

- **AutoScalingGroup:**
 - **UserData:** This value is provided in this tutorial. It includes commands to set up the resource for CodeDeploy and start the CodeDeploy agent.

- **AMI-ID:** This value determines the operating system of EC2 instances your Auto Scaling group (ASG) will spin up. Select an AMI in your account that starts with "customer-" and is of the operating system that you want. Find AMI IDs in the AMS Console VPCs -> VPCs details page. This walkthrough is for ASGs configured to use an Amazon Linux or RHEL AMI.
- **Database:**
 - These parameters, **DBEngine**, **EngineVersion**, and **LicenseModel** should be set according to your situation though the values shown in the example have been tested. The tutorial uses these values, respectively: *MySQL, 8.0.16, general-public-license*.
 - These parameters, **DBName**, **MasterUserPassword**, and **MasterUsername** are required when deploying the application bundle. The tutorial uses these values, respectively: *wordpressDB, p4ssw0rd, admin*. Note that DBName can only contain alphanumeric characters.
 - When you enter the **MasterUsername** for the RDS DB, it will appear in cleartext, so log in to the database as soon as possible and change the password to ensure your security.
 - For **RDSSubnetIds**, use two Private subnets. Enter them one at a time pressing "Enter" after each. Find Subnet IDs with the [ListSubnetSummaries](#) operation (CLI: list-subnet-summaries) or in the AMS Console VPCs -> VPC details page.
- **LoadBalancer:**
 - Set this parameter, **Public** to **true** because the tutorial uses Public ELB subnets.
 - **ELBSubnetIds:** Use two Public subnets. Enter them one at a time pressing "Enter" after each. Find Subnet IDs with the [ListSubnetSummaries](#) operation (CLI: list-subnet-summaries) or in the AMS Console VPCs -> VPC details page.
- **Application:** The **ApplicationName** value sets the CodeDeploy application name and CodeDeploy deployment group name. You use it to deploy your application. It must be unique in the account. To check your account for CodeDeploy names, see the CodeDeploy Console. The example uses *WordPress* but, if you will use that value, make sure that it is not already in use.

1. Launch the high availability stack.

- a. On the **Create RFC** page, select the category **Deployment**, subcategory **Standard Stacks**, item **High availability two-tier stack** and operation **Create**, from the list.
- b. IMPORTANT: Choose **Advanced** and set the values as shown.

You only need to enter values for starred (*) options, tested values are shown in the example; you can leave not-required empty options blank.

- c. For the **RFC Description** section:

```
Subject: WP-HA-2-Tier-RFC
```

- d. For the **Resource information** section, set parameters for **AutoScalingGroup**, **Database**, **LoadBalancer**, **Application**, and **Tags**.

Also, the purpose of the "AppName" tag key is so you can easily search for the ASG instances in the EC2 console; you can call this tag key "Name" or any other key name that you want. Note that you can add up to 50 tags.

```
UserData:
#!/bin/bash
REGION=$(curl 169.254.169.254/latest/meta-data/placement/availability-zone/ |
sed 's/[a-z]$/')
yum -y install ruby httpd
chkconfig httpd on
service httpd start
touch /var/www/html/status
cd /tmp
curl -O https://aws-codedeploy-$REGION.s3.amazonaws.com/latest/install
```

```
chmod +x ./install
./install auto
chkconfig coddeploy-agent on
service coddeploy-agent start
AmiId: AMI-ID
Description: WP-HA-2-Tier-Stack

Database:
LicenseModel: general-public-license (USE RADIO BUTTON)
EngineVersion: 8.0.16
DBEngine: MySQL
RDSSubnetIds: PRIVATE_AZ1 PRIVATE_AZ2 (ENTER ONE AT A TIME PRESSING "ENTER"
AFTER EACH)
MasterUserPassword: p4ssw0rd
MasterUsername: admin
DBName: wordpressDB

LoadBalancer:
Public: true (USE RADIO BUTTON)
ELBSubnetIds: PUBLIC_AZ1 PUBLIC_AZ2

Application:
ApplicationName: WordPress

Tags:
Name: WP-Rhel-Stack
```

- e. Click **Submit** when finished.
2. Log in to the database that you created and change the password.
3. Launch an S3 bucket Stack.

Gathering the following data before you begin will make the deployment go more quickly.

REQUIRED DATA S3 BUCKET:

- **VPC-ID:** This value determines where your S3 Bucket will be. Find VPC IDs with the [ListVpcSummaries](#) operation (CLI: list-vpc-summaries) or in the AMS Console VPCs page.
- **BucketName:** This value sets the S3 Bucket name, you use it to upload your application bundle. It must be unique across the region of the account and cannot include upper-case letters. Including your account ID as part of the BucketName is not a requirement but makes it easier to identify the bucket later. To see what S3 bucket names exist in the account, go to the Amazon S3 Console for your account.

- a. On the **Create RFC** page, select the category **Deployment**, subcategory **Advanced Stack Components**, item **S3 storage**, and operation **Create** from the RFC CT pick list.
- b. Keep the default **Basic** option and set the values as shown.

```
Subject: S3-Bucket-WP-HA-RFC
Description: S3BucketForWordPressBundles
BucketName: ACCOUNT_ID-BUCKET_NAME
AccessControl: Private
VpcId: VPC_ID
Name: S3-Bucket-WP-HA-Stack
TimeoutInMinutes: 60
```

- c. Click **Submit** when finished. The bucket deployed with this change type allows full read/write access to the whole account.

Create, Upload, and Deploy the Application

First, create a WordPress application bundle, and then use the CodeDeploy CTs to create and deploy the application.

1. Download WordPress, extract the files and create a `./scripts` directory.

Linux command:

```
wget https://github.com/WordPress/WordPress/archive/master.zip
```

Windows: Paste `https://github.com/WordPress/WordPress/archive/master.zip` into a browser window and download the zip file.

Create a temporary directory in which to assemble the package.

Linux:

```
mkdir /tmp/WordPress
```

Windows: Create a "WordPress" directory, you will use the directory path later.

2. Extract the WordPress source to the "WordPress" directory and create a `./scripts` directory.

Linux:

```
unzip master.zip -d /tmp/WordPress_Temp  
cp -paf /tmp/WordPress_Temp/WordPress-master/* /tmp/WordPress  
rm -rf /tmp/WordPress_Temp  
rm -f master  
cd /tmp/WordPress  
mkdir scripts
```

Windows: Go to the "WordPress" directory that you created and create a "scripts" directory there.

If you are in a Windows environment, be sure to set the break type for the script files to Unix (LF). In Notepad ++, this is an option at the bottom right of the window.

3. Create the CodeDeploy `appspec.yml` file, in the WordPress directory (if copying the example, check the indentation, each space counts). IMPORTANT: Ensure that the "source" path is correct for copying the WordPress files (in this case, in your WordPress directory) to the expected destination (`/var/www/html/WordPress`). In the example, the `appspec.yml` file is in the directory with the WordPress files, so only `/` is needed. Also, even if you used a RHEL AMI for your Auto Scaling group, leave the `"os: linux"` line as-is. Example `appspec.yml` file:

```
version: 0.0  
os: linux  
files:  
  - source: /  
    destination: /var/www/html/WordPress  
hooks:  
  BeforeInstall:  
    - location: scripts/install_dependencies.sh  
      timeout: 300  
      runas: root  
  AfterInstall:  
    - location: scripts/config_wordpress.sh  
      timeout: 300  
      runas: root  
ApplicationStart:
```

```
- location: scripts/start_server.sh
  timeout: 300
  runas: root
ApplicationStop:
- location: scripts/stop_server.sh
  timeout: 300
  runas: root
```

4. Create bash file scripts in the WordPress `./scripts` directory.

First, create `config_wordpress.sh` with the following content (if you prefer, you can edit the `wp-config.php` file directly).

Note

Replace `DBName` with the value given in the HA Stack RFC (for example, `wordpress`).

Replace `DB_MasterUsername` with the `MasterUsername` value given in the HA Stack RFC (for example, `admin`).

Replace `DB_MasterUserPassword` with the `MasterUserPassword` value given in the HA Stack RFC (for example, `p4ssw0rd`).

Replace `DB_ENDPOINT` with the endpoint DNS name in the execution outputs of the HA Stack RFC (for example, `srt1cz23n45sfg.clgvd67uvydk.us-east-1.rds.amazonaws.com`). You can find this with the `GetRfc` operation (CLI: `get-rtc --rtc-id RFC_ID`) or in the AMS Console RFC details page for the HA Stack RFC that you previously submitted.

```
#!/bin/bash
chmod -R 755 /var/www/html/WordPress
cp /var/www/html/WordPress/wp-config-sample.php /var/www/html/WordPress/wp-config.php
cd /var/www/html/WordPress
sed -i "s/database_name_here/DBName/g" wp-config.php
sed -i "s/username_here/DB_MasterUsername/g" wp-config.php
sed -i "s/password_here/DB_MasterUserPassword/g" wp-config.php
sed -i "s/localhost/DB_ENDPOINT/g" wp-config.php
```

5. In the same directory create `install_dependencies.sh` with the following content:

```
#!/bin/bash
yum install -y php
yum install -y php-mysql
yum install -y mysql
service httpd restart
```

Note

HTTPS is installed as part of the user data at launch in order to allow health checks to work from the start.

6. In the same directory create `start_server.sh` with the following content:

- For Amazon Linux instances, use this:

```
#!/bin/bash
service httpd start
```

- For RHEL instances, use this (the extra commands are policies that allow SELINUX to accept WordPress):

```
#!/bin/bash
setsebool -P httpd_can_network_connect_db 1
setsebool -P httpd_can_network_connect 1
chcon -t httpd_sys_rw_content_t /var/www/html/WordPress/wp-content -R
restorecon -Rv /var/www/html
```

```
service httpd start
```

7. In the same directory create `stop_server.sh` with the following content:

```
#!/bin/bash  
service httpd stop
```

8. Create the zip bundle.

Linux:

```
$ cd /tmp/WordPress  
$ zip -r wordpress.zip .
```

Windows: Go to your "WordPress" directory and select all of the files and create a zip file, be sure to name it `wordpress.zip`.

1. Upload the application bundle to the S3 bucket

The package needs to be in place in order to continue deploying the stack.

You automatically have access to any S3 bucket instance that you create. You can access it through your Bastions (see [Accessing Instances](#)), or through the S3 console, and upload the CodeDeploy package with drag-and-drop or browsing to and selecting the file.

You can also use the following command in a shell window; be sure that you have the correct path to the zip file:

```
aws s3 cp wordpress/wordpress.zip s3://BUCKET_NAME/
```

2. Deploy the WordPress CodeDeploy Application Bundle

REQUIRED DATA CODEDEPLOY APPLICATION DEPLOYMENT:

- **CodeDeployApplicationName:** The name you gave the CodeDeploy application.
 - **CodeDeployGroupName:** Since the CodeDeploy application and group were both created from the name you gave the CodeDeploy application in the HA stack RFC, this is the same name as the **CodeDeployApplicationName**.
 - **S3Bucket:** The name you gave the S3 bucket.
 - **S3BundleType** and **S3Key:** These are part of the WordPress application bundle you deployed.
 - **VpcId:** The relevant VPC.
- a. On the **Create RFC** page, select the category **Deployment**, subcategory **Applications**, item **CodeDeploy application**, and operation **Deploy** from the RFC CT pick list.
 - b. Keep the default **Basic** option, and set the values as shown.

Note

Reference the CodeDeploy application, CodeDeploy deployment group, S3 bucket and bundle previously created.

```
Subject: WP-CD-Deploy-RFC  
Description: DeployWordPress  
S3Bucket: BUCKET_NAME  
S3Key: wordpress.zip  
S3BundleType: zip  
CodeDeployApplicationName: WordPress
```

AMS Advanced Onboarding Guide AMS
Advanced Account Onboarding Information
Tutorials

```
CodeDeployDeploymentGroupName: WordPress
CodeDeployIgnoreApplicationStopFailures: false
RevisionType: S3

VpcId: VPC_ID
Name: WP-CD-Deploy-Op
TimeoutInMinutes: 60
```

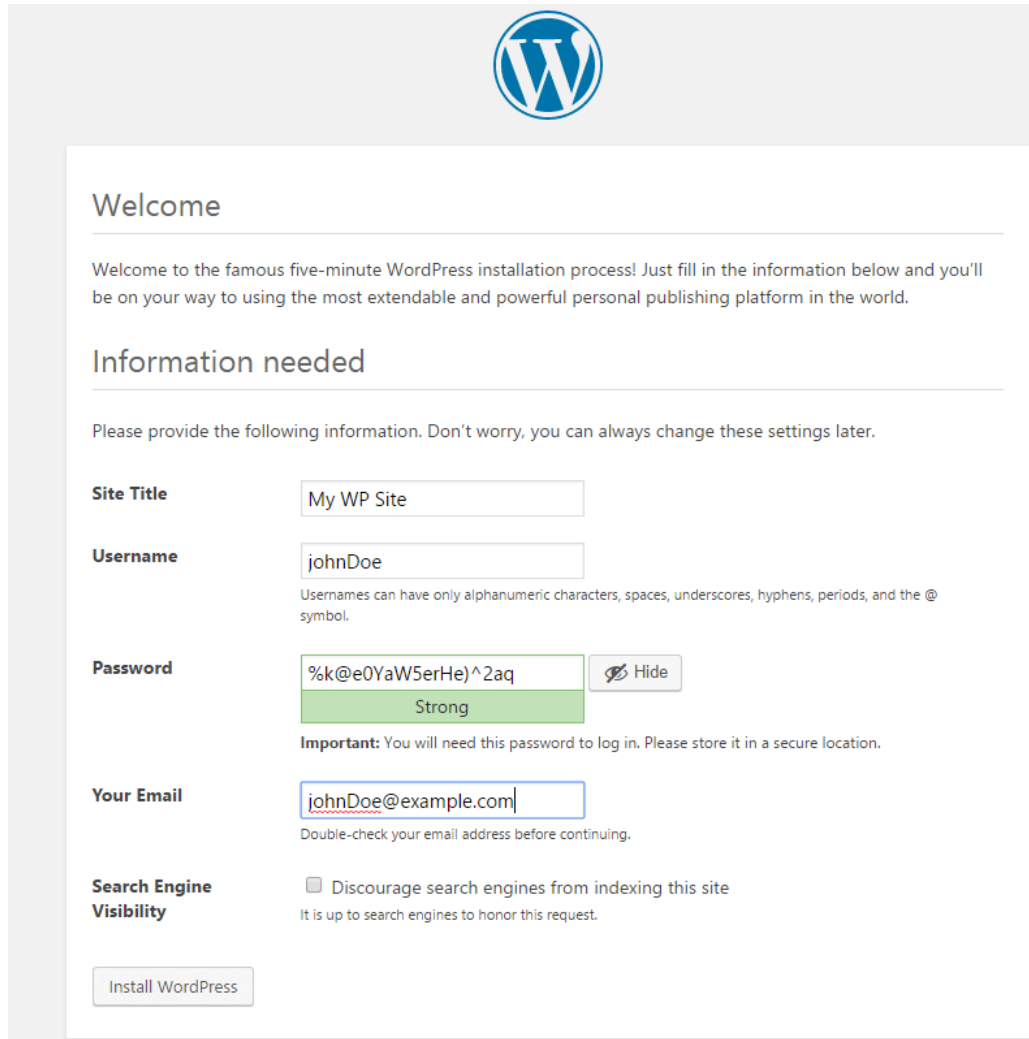
- c. Click **Submit** when finished.

Validate the Application Deployment

Navigate to the endpoint (LoadBalancerCName) of the previously-created load balancer, with the WordPress deployed path: /WordPress. For example:

```
http://stack-ID-FOR-ELB.us-east-1.elb.amazonaws.com/WordPress
```

You should see a page like this:



The screenshot shows the WordPress installation welcome page. At the top center is the WordPress logo. Below it is the heading "Welcome" followed by a horizontal line. The text reads: "Welcome to the famous five-minute WordPress installation process! Just fill in the information below and you'll be on your way to using the most extendable and powerful personal publishing platform in the world." Below this is the heading "Information needed" followed by another horizontal line. The text says: "Please provide the following information. Don't worry, you can always change these settings later." There are four input fields: "Site Title" with the value "My WP Site", "Username" with the value "johnDoe", "Password" with a masked value "%k@e0YaW5erHe)^2aq" and a "Hide" button, and "Your Email" with the value "johnDoe@example.com". Below the password field is a strength indicator showing "Strong" and an "Important" note: "You will need this password to log in. Please store it in a secure location." Below the email field is a note: "Double-check your email address before continuing." At the bottom, there is a "Search Engine Visibility" section with a checkbox labeled "Discourage search engines from indexing this site" which is currently unchecked. Below this is the text: "It is up to search engines to honor this request." At the very bottom left is a button labeled "Install WordPress".

Tear Down the High Availability Deployment

To tear down the deployment, you submit the Delete Stack CT against the HA Two-Tier stack, and the S3 bucket, and you can request that RDS snapshots be deleted (they are deleted automatically after ten days, but they do cost a small amount while there). Gather the stack IDs for the HA stack and the S3 bucket and then follow these steps. See [Any Stack, Deleting, Rebooting, Starting, Stopping](#) (p. 159).

Appendix: Single-Account Landing Zone Onboarding Questionnaire

Topics

- [Deployment Summary](#) (p. 200)
- [Environment/Architecture Considerations](#) (p. 200)
- [Single-Account Landing Zone Monitoring Alerts](#) (p. 201)
- [Maintenance Window](#) (p. 201)
- [Next Steps](#) (p. 201)

This section describes some of the information that you will need to think about before onboarding an account.

Deployment Summary

A description of the deployment. For example:

- This account is for a Line-of-Business application deployment (as opposed to a Product application deployment).
- The deployment involves an auto-scaled ARP (authenticated reverse proxy) within the account's public/DMZ subnet.
- Web and application servers will be deployed within the account's private subnet.
- An RDS (AWS Relational Database Service) instance will also be deployed within the account's private Subnet.
- The servers (ARP, web, application, database, load balancer, etc.) are separated into distinct security groups.
- The account requires an HA (high availability) design spread across availability zones (AZs) i.e. "Multi-AZ".

Environment/Architecture Considerations

- Will your virtual data center connect back to your corporate network?
 - Do you have an existing AWS DirectConnect service or do you require a new DirectConnect service?
 - Do you have an existing VPN connection or do you require a new VPN service?
- What is the available CIDR block range of internal addresses that you could allocate? (/16 recommended, must not overlap corporate network ranges)
- Will your virtual data center require Internet access?
- Which region(s) do you intend to use? (Sydney/N. Virginia/Dublin)
- Will you require a Shared Services subnet to host applications that have connectivity to all other subnets?

- What are your organizational divisions that you would like to be hosted as separate subnets. For each:
 - What connectivity to other subnets do you need?
 - Does the subnet require Internet access?
 - Are there any application deployment restrictions to that subnet?
 - Are there any particular network requirements for that subnet?
- Would you like separate development and/or test environments? (Will include shared services duplicate for anytime access)
- What are your snapshot backup requirements?
- Do you have an existing maintenance process or patch window(s) that you would like to keep?
- What are your domain registration requirements?
- Do you have any single sign-on requirements? (e.g., AD, LDAP)
- What are your overall expected operating system and anticipated capacity requirements?

Single-Account Landing Zone Monitoring Alerts

AMS provides a way for you to be directly alerted (versus getting AMS service notifications) for certain monitoring alerts. To sign up for this, make sure that your Cloud Architect or Cloud Service Delivery Manager receive this information:

Direct Alerts Email: These are the email addresses that you want AMS to send certain resource-based alerts to. For details of which alerts are sent directly to email, see [Alerts from Baseline Monitoring in AMS](#) in the AMS User Guide for Single-Account Landing Zone. For more information on AMS monitoring, see [Monitoring Management](#) in the AMS User Guide for Single-Account Landing Zone.

Maintenance Window

You will want to create a maintenance window that considers different application needs, different AWS Regions, and different stress periods. Your maintenance window is when AMS will apply patching. Here are some guidelines:

- To limit the impact on users, plan your maintenance window according to the AWS Region where your environments are deployed.
- Schedule a window outside of regular business hours and when the least traffic is expected on production servers.
- Typically, infrastructure stacks require monthly updates.
- Schedule a maintenance window for at least 300 minutes. Operating system patching takes 60-90 minutes, infrastructure stack patching takes 180-300 minutes.

Next Steps

The AMS onboarding team will assist you in every step of onboarding your account to AMS. These are onboarding requirements:

- Provision a new AWS account to use for Sentinel and provide AWS Account ID.
- Sign up for the desired level of Support.
- Create a cross-account IAM role to grant the AMS provisioning account access and provide the role name to AMS.
- Add the account 753102745277 as a Trusted Entity.

Appendix: ActiveDirectory Federation Services (ADFS) claim rule and SAML settings

For detailed step-by-step instructions on how to install and configure AD FS see [Enabling Federation to AWS Using Windows Active Directory, ADFS, and SAML 2.0](#).

ADFS claim rule configurations

If you already have an ADFS implementation, configure following:

- Relying party trust
- Claims rules

The relying party trust and claims rules steps are taken from [Enabling Federation to AWS Using Windows Active Directory, AD FS, and SAML 2.0](#)blog

- Claims rules:
 - **Nameid**: Configuration per blog post
 - **RoleSessionName**: Configure as follows
 - **Claim rule name**: **RoleSessionName**
 - **Attribute store**: **Active Directory**
 - **LDAP Attribute**: **SAM-Account-Name**
 - **Outgoing Claim Type**: **https://aws.amazon.com/SAML/Attributes/RoleSessionName**
 - **Get AD Groups**: Configuration per [blog post](#)
 - **Role claim**: Configure as follows

```
c:[Type == "http://temp/variable", Value =~ "(?i)^AWS-([\d]{12})-"]
```

```
=> issue(Type = "https://aws.amazon.com/SAML/Attributes/Role", Value =  
  RegexReplace(c.Value, "AWS-([\d]{12})-", "arn:aws:iam::$1:saml-provider/customer-  
  readonly-saml,arn:aws:iam::$1:role/");
```

Web console

You can access the AWS Web console by using the link below replacing [\[ADFS-FQDN\]](#) with the FQDN of your ADFS implementation.

[https://\[ADFS-FQDN\]/adfs/ls/IdpInitiatedSignOn.aspx](https://[ADFS-FQDN]/adfs/ls/IdpInitiatedSignOn.aspx)

Your IT department can deploy the above link to the user population via a Group Policy.

API and CLI access with SAML

How to configure API and CLI access with SAML.

The python packages are sourced from the blog posts below:

- NTLM: [How to Implement Federated API and CLI Access Using SAML 2.0 and AD FS](#)
- Forms: [How to Implement a General Solution for Federated API/CLI Access Using SAML 2.0](#)
- PowerShell: [How to Set Up Federated API Access to AWS by Using Windows PowerShell](#)

Script configuration

1. Using Notepad++, change the default region to the correct region
2. Using Notepad++, disable SSL verification for test and dev environments
3. Using Notepad++, configure idpentryurl

```
https://[ADFS-FDQN]/adfs/ls/IdpInitiatedSignOn.aspx?  
loginToRp=urn:amazon:webservices
```

Windows configuration

The instructions below are for the python packages. The credentials generated will be valid for 1 hour.

1. [Download and install python \(2.7.11\)](#)
2. [Download and install AWS CLI tools](#)
3. Install the AMS CLI:
 - a. Download the AMS distributables zip file provided by your cloud service delivery manager (CSDM) and unzip.

Several directories and files are made available.
 - b. Open either the **Managed Cloud Distributables -> CLI -> Windows** or the **Managed Cloud Distributables -> CLI -> Linux / MacOS** directory, depending on your operating system, and:

For **Windows**, execute the appropriate installer (this method only works on Windows 32 or 64 bits systems):
 - 32 Bits: ManagedCloudAPI_x86.msi
 - 64 Bits: ManagedCloudAPI_x64.msi
For **Mac/Linux**, execute the file named: **MC_CLI.sh**. You can do this by running this command:
`sh MC_CLI . sh`. Note that the **amscm** and **amsskms** directories and their contents must be in the same directory as the **MC_CLI.sh** file.
 - c. If your corporate credentials are used via federation with AWS (the AMS default configuration) you must install a credential management tool that can access your federation service. For example, you can use this AWS Security Blog [How to Implement Federated API and CLI Access Using SAML 2.0 and AD FS](#) for help configuring your credential management tooling.
 - d. After the installation, run `aws amscm help` and `aws amsskms help` to see commands and options.

4. Download the required SAML script

Download to c:\aws\scripts

5. [Download PIP](#)

Download to c:\aws\downloads

6. Using PowerShell, install PIP

```
<pythondir>.\python.exe c:\aws\downloads\get-pip.py
```

7. Using PowerShell, install boto module

```
<pythondir\scripts>pip install boto
```

8. Using PowerShell, install requests module

```
<pythondir\scripts>pip install requests
```

9. Using PowerShell, install requests security module

```
<pythondir\scripts>pip install requests[security]
```

10. Using PowerShell, install beautifulsoup module

```
<pythondir\scripts>pip install beautifulsoup4
```

11. Using PowerShell, create a folder called .aws in the users profile (%userprofile%\aws)

```
mkdir .aws
```

12. Using PowerShell, create a credential file in the .aws folder

```
New-Item credentials -type file -force
```

The credentials file mustn't have a file extension

The filename must be all lowercase and have the name credentials

13. Open the credentials file with notepad and paste in the following data, specifying the correct region

```
[default]
output = json
region = us-east-1
aws_access_key_id =
aws_secret_access_key =
```

14. Using PowerShell, the SAML script and logon

```
<pythondir>.\python.exe c:\aws\scripts\samlapi.py
```

Username: [USERNAME]@upn

Choose the role you would like to assume

Linux configuration

The credentials generated will be valid for 1 hour.

1. Using WinSCP, transfer the SAML script
2. Using WinSCP, transfer the Root CA certificate (ignore for test and dev)
3. Add the ROOT CA to the trusted root certificates (ignore for test and dev)

```
$ openssl x509 -inform der -in [certname].cer -out certificate.pem (ignore for test and dev)
```

Add contents of certificate.pem to end of /etc/ssl/certs/ca-bundle.crt file ((ignore for test dev)

4. Create .aws folder in home/ec2-user 5

```
[default]
output = json
region = us-east-1
aws_access_key_id =
aws_secret_access_key =
```

5. Using WinSCP, transfer the credentials file to .aws folder

6. Install boto module

```
$ sudo pip install boto
```

7. Install requests module

```
$ sudo pip install requests
```

8. Install beautifulsoup module

```
$ sudo pip install beautifulsoup4
```

9. Copy the script to home/ec2-user

Set the required permissions

Execute the script: samlapi.py

Document history

The following table describes the important changes to the documentation since the last release of AMS.

- **API version: 2019-05-21**
- **Latest documentation update:** October 28, 2021

Change	Description	Date
AMS Modes	AMS has a new mode, Direct Change mode. See AMS modes and applications or workloads (p. 12)	October 28, 2021
AD Trust, Default Access Firewall Rules	Added a note about finding information on firewall rules and ports in the AMS Security Guide, accessed through the AWS Artifact console > Reports tab, Managed Services filter. See Default Access Firewall Rules (p. 38)	October 14, 2021
IAM default roles	A section on default roles was available only in the SALZ onboarding section but it applied to both. That section was moved to the main introduction section. See IAM User Role (p. 20)	September 30, 2021
SALZ account onboarding	Updated information on single-account landing zone (SALZ) onboarding timing and process.	August 26, 2021
Account setup automation	A ZIP with scripts and examples has been added to help you automate many of the required RFCs for account onboarding. See AMS post-account prescriptive guidance (p. 18) .	July 15, 2021
Multi-Account Landing Zone Single-Account Landing Zone Content Merge	The multi-account landing zone onboarding content has been merged into the single-account landing zone Onboarding Guide.	June 17, 2021
Redacted the 'How do I offboard a Multi-Account Landing Zone environment' and 'How do I offboard a Multi-Account Landing Zone application account?' sections and put the redacted content in the new private security guide, which is available on AWS Artifact.	To access AWS Artifact, you can contact your CSDM for instructions or go to Getting Started with AWS Artifact .	June 17, 2021
Book title update.	The user guide title now reflects that this content is for the AMS Advanced operations plan.	April 15, 2021
Updated sections: The FAQs appendix was added to the Service management chapter. Additionally, the AMS Service Description document was added to that same chapter	See AMS service management (p. 40) .	December 17, 2020

Change	Description	Date
and will no longer be a separate document.		
New section: Additional instructions for submitting AD federation configuration have been added.	See Submitting the federation request to AMS (p. 117) .	November 12, 2020
A new section on configuring Direct Alert Emails was added.	See Single-Account Landing Zone Monitoring Alerts (p. 201) .	April 23, 2020
The Priority parameter in the Other Other Create and Other Other Update change types is now optional. The walkthrough was updated.	See Other Other RFC, Creating (CLI) (p. 158) .	January 23, 2019
Replaced the Change Types table with links to the Change Type Reference document.	Understanding Change Types and the AMS Change Type Reference .	November 21, 2019
Removed Information about AMS Managed Firewall	AMS managed firewall is not yet available, so the information was removed.	August 27, 2019
Updated Egress Traffic Management Information	AMS egress traffic management (p. 19) .	July 25, 2019
Updated Default Monitoring Metrics (Alerts from Baseline Monitoring)	Monitored Metrics Defaults (p. 127) .	July 25, 2019

AWS glossary

For the latest AWS terminology, see the [AWS glossary](#) in the *AWS General Reference*.