



AWS Whitepaper

Organizing Your AWS Environment Using Multiple Accounts



Organizing Your AWS Environment Using Multiple Accounts: AWS Whitepaper

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

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

Table of Contents

Abstract and introduction	i
Abstract	1
Are you Well-Architected?	1
Introduction	1
Multi-account strategy best practices and recommendations	2
AWS accounts	2
Stages of adoption	3
Best practices	4
Relation to AWS Well-Architected	5
Intended audience	5
Benefits of using multiple AWS accounts	6
Group workloads based on business purpose and ownership	6
Apply distinct security controls by environment	7
Constrain access to sensitive data	7
Promote innovation and agility	7
Limit scope of impact from adverse events	8
Support multiple IT operating models	8
Manage costs	10
Distribute AWS Service Quotas and API request rate limits	10
Core concepts	11
AWS Organizations	11
About organizations	11
Benefits of using OUs	13
Group similar accounts based on function	14
Apply common policies	14
Share common resources	15
Provision and manage common resources	15
Multiple organizations	16
Test changes to your overall AWS environment	16
Support acquisitions and divestments	17
Support large AWS environments	17
Align with your billing requirements	17
Different classification levels for government applications	17
Design principles for your multi-account strategy	19

Organize based on security and operational needs	19
Apply security guardrails to OUs rather than accounts	19
Avoid deep OU hierarchies	20
Start small and expand as needed	20
Avoid deploying workloads to the organization's management account	20
Separate production from non-production workloads	20
Assign a single or small set of related workloads to each production account	21
Use federated access to help simplify managing human access to accounts	21
Use automation to support agility and scale	21
Use multi-factor authentication	22
Break glass access	22
Recommended OUs and accounts	25
Security OU	26
Log archive account	27
Security Tooling (Audit) accounts	29
Infrastructure OU	35
Backup account	36
Identity account	37
Network account	39
Operations Tooling account	41
Monitoring account	45
Sandbox OU	50
Sandbox per builder or team with spend limits	51
Temporary resources and environments	51
Wide-ranging access	51
No access to corporate resources and non-public data	51
Sandbox and development environments	52
Additional Services and Functionalities	52
Example structures	52
Workloads OU	53
Example structure	53
Policy Staging OU	54
Workload-specific policies	54
Recommended testing and promotion workflow	55
Example structure	55
Suspended OU	57

Constraining activity in suspended accounts	57
Tagging suspended accounts	57
Closing suspended accounts	58
Individual Business Users OU	58
Controls	58
Services that do not require direct user access to accounts	58
Exceptions OU	59
Service control policies and scrutiny	59
Consider Workloads OU as an alternative	59
Deployments OU	59
Using CI/CD capabilities residing outside of your AWS environment	59
Separating CI/CD management capabilities from workloads	60
Running CI jobs and CD build stages in deployment accounts	61
Aligning CI/CD accounts with groups of workloads	61
Transitional OU	61
Common reasons for moving accounts into your organization	61
Considerations for moving accounts into your organization	62
After moving accounts	62
Business Continuity OU	62
Controls	63
Organizing workload-oriented OUs	66
Workloads and environments	66
Workloads	66
Workload environments	67
Workload accounts	68
Production and non-production workload environments	68
Workload dependencies across environments	69
Production environments accessing non-production	69
Non-production environments accessing dependencies	69
OU structure for non-production environments	69
Option A: Common guardrails across non-production environments	69
Option B: Different guardrails across non-production environments	70
Worksheets to help decide on workload-oriented OUs	71
Extended workload-oriented OU structure	72
Grouping related workloads	72
Separating business units with significantly different policies	73

Patterns for organizing your AWS accounts	75
Single AWS account	75
Production starter organization	75
Production starter organization with AWS Control Tower	77
Basic organization	79
Basic organization with infrastructure services	81
Basic organization with CI/CD as a separate function	83
Advanced organization	85
Implementation	87
Getting started with your multi-account environment	87
New customers	87
Existing customers	88
Embrace infrastructure as code	89
Examine your operational model	89
Implement identity management and access controls and other security capabilities	89
Separate the production workload environment from non-production environment(s)	90
Separate the workload environments to align with the operational organization units	91
Create the additional organizational units to enable other capabilities	92
Other considerations for implementing these changes	92
Available services	92
AWS Organizations	11
AWS Control Tower	93
AWS Managed Services	93
Conclusion	95
Contributors	96
Additional resources	97
Appendix A: Relation to AWS Well-Architected	98
Operational Excellence Pillar	98
Security Pillar	98
Reliability Pillar	99
Cost Optimization Pillar	99
Appendix B: Worksheet for mapping workload environment purposes to hosting environment types	100
Use the results for internal documentation	100
How to use this worksheet	101
Descriptions of example purposes of workload environments	101

Self-paced learning and experimentation workload environments	101
Development workload environments	101
Static analysis, build, and unit testing workload environments	102
CI jobs and CD pipelines workload environments	102
Smoke testing workload environments	102
Development integration testing workload environments	102
Production-like system testing workload environments	102
Stable shared test workload environments	103
Resiliency testing workload environments	103
Demo workload environments	103
External pre-release workload environments	103
Production workload environments	104
Descriptions of example workload hosting environment types	104
Corporate desktop environments	104
Sandbox environments	104
Development environments	105
Data-oriented development environments	105
Test environments	105
Production environments	105
Example worksheet	106
Empty worksheet	107
Appendix C: Worksheet for identifying attributes of workload hosting environments	110
How to use this worksheet	110
Descriptions of example attributes	111
Owners/tenants	111
Tolerance for extended outages	112
Internet access	112
Internal network access	113
Data	113
Third-party software and cloud services	114
Degree of access	114
Lifespan of resources	115
Direct human write access to workload resources	115
Automated workload provisioning	115
Formal change management for workloads	116
Degree of centrally managed foundation	117

Common enterprise guardrails	117
Example worksheet	117
Empty worksheet	123
Appendix D: Multiple AWS Regions	125
Geographic scopes of data protection	125
Performance considerations	125
Log management	126
Appendix E: How does AWS Control Tower establish your multi-account environment?	127
Establish your multi-account environment with AWS Control Tower	127
Next steps for setting up your multi-account environment	128
Document history	129
Notices	130

Organizing Your AWS Environment Using Multiple Accounts

Publication date: **March 28, 2024** ([Document history](#))

Abstract

Businesses that are starting to adopt Amazon Web Services (AWS), expanding their footprint on AWS, or planning to enhance an established AWS environment need to ensure they have a [foundation on AWS](#) for their cloud environment. One important aspect of their foundation is organizing their AWS environment following a multi-account strategy.

Using multiple AWS accounts to help isolate and manage your business applications and data can help you optimize across most of the [AWS Well-Architected Framework](#) pillars, including operational excellence, security, reliability, and cost optimization. This paper provides best practices and current recommendations for organizing your overall AWS environment. The extent to which you use these best practices depends on your stage of the cloud adoption journey and specific business needs.

Are you Well-Architected?

The [AWS Well-Architected Framework](#) helps you understand the pros and cons of the decisions you make when building systems in the cloud. The six pillars of the Framework allow you to learn architectural best practices for designing and operating reliable, secure, efficient, cost-effective, and sustainable systems. Using the [AWS Well-Architected Tool](#), available at no charge in the [AWS Management Console](#), you can review your workloads against these best practices by answering a set of questions for each pillar.

For more expert guidance and best practices for your cloud architecture—reference architecture deployments, diagrams, and whitepapers—refer to the [AWS Architecture Center](#).

Introduction

Using multiple AWS accounts to help isolate and manage your business applications and data can help you optimize across most of the AWS Well-Architected Framework pillars including operational excellence, security, reliability, and cost optimization.

Topics

- [Multi-account strategy best practices and recommendations](#)
- [AWS accounts](#)
- [Stages of adoption](#)
- [Best practices](#)
- [Relation to AWS Well-Architected](#)
- [Intended audience](#)

Multi-account strategy best practices and recommendations

Businesses can benefit from considering the latest guidance for organizing their AWS environments. A multi-account strategy is key to succeed when customers are starting to adopt AWS, expanding their footprint on AWS, or planning to enhance an established AWS environment.

Customers might have multiple teams with different security and compliance controls that need to be isolated from one another. Some might have different business processes entirely or be part of different business lines that need clarity around costs incurred.

Customers need explicit security boundaries, a mechanism to have direct control and visibility of their limits and any throttling, and a billing separation to directly map costs to underlying projects. The isolation designed into an AWS account can help you meet these needs.

Using multiple AWS accounts to help isolate and manage your business applications and data can help you optimize across most of the [AWS Well-Architected Framework](#) pillars including operational excellence, security, reliability, and cost optimization.

AWS accounts

Your cloud resources and data are contained in an AWS account. An account acts as an identity and access management isolation boundary. When you need to share resources and data between two accounts, you must explicitly allow this access.

By default, no access is allowed between accounts. For example, if you designate different accounts to contain your production and non-production resources and data, no access is allowed between those environments by default.

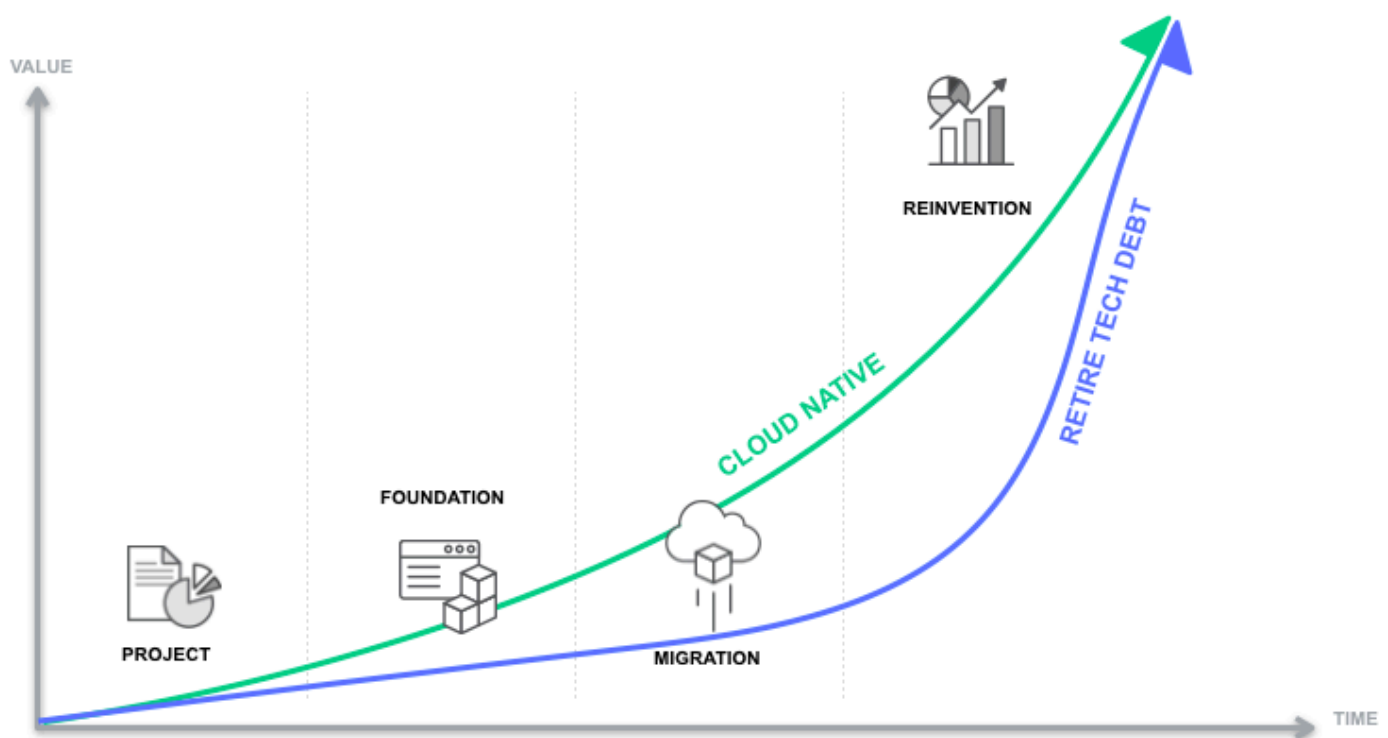
The number of accounts that best meets your needs can range from a few to hundreds or even thousands. Management of many accounts requires use of automation to help minimize your

operational complexity and ensure efficient alignment with your security, governance, and operational requirements. AWS does not charge per account. Rather, you incur charges based on resources used, regardless of account quantity.

Stages of adoption

Through experience working with thousands of customers, AWS has outlined a common set of stages of cloud adoption. These best practices are designed to help you meet your needs throughout your cloud adoption journey. You can start with a small AWS environment and progressively grow and evolve it as you gain experience and expand your adoption.

When your organization is new to AWS, you might start by creating one or more personal or team-managed accounts for initial experimentation. This work is usually done informally and before more concerted efforts are made to evaluate the value of AWS. In this experimental and often informal stage, there's usually little investment made in organizing and rationalizing the number and purpose of accounts.



Stages of cloud adoption

In the **project** stage of adoption, you begin to formally plan for your first few production deployments on AWS. It's common to establish an initial cloud foundation that meets your security, governance, and operational requirements.

A workload identifies a set of components that deliver business value together. A workload is usually the level of detail that business and technology leaders communicate about. Some examples of workloads are:

- Marketing websites
- Ecommerce websites
- Mobile app backends
- Analytic platforms

Workloads vary in levels of architectural complexity, from static websites to complex microservices, each with potentially different requirements on cost or billing identification.

Rather than using a single account, we recommend that you use several accounts to separate your workloads. This approach is designed to make it easier for you to meet your requirements, even in the early project stage of adoption. Based on the success of those first few workloads, you might want to gain further business benefits by expanding your adoption of AWS. This motivation often leads to the *foundation* stage of adoption. In this stage, you invest in evolving your cloud foundational capabilities before greatly expanding adoption.

Evolving your foundation in AWS often includes formalizing and expanding the structure of your accounts to meet the needs of onboarding more teams and workloads. At this stage, it's important for you to design and prepare to manage your AWS environment so that it can scale to meet your needs without the need for a corresponding linear increase in headcount. As you plan for and perform large-scale migrations and deploy net new cloud native workloads, you can continue to adjust and enhance your approach to using multiple accounts.

Best practices

The best practices described in this paper are designed to help you more easily achieve your security, governance, and operational requirements through multiple accounts. The best practices were assembled based on the experiences of thousands of customers who have progressed through their cloud adoption journeys.

The best practices can help you quickly establish the initial right-sized scope of your AWS environment, and adjust and expand your AWS environment as you gain experience both with the AWS services and how you work with the AWS Cloud.

Although your needs will likely be similar to the needs of other organizations, each organization also has some unique requirements. Accordingly, these best practices are intended to offer guidance rather than a one-size-fits-all solution to organizing your AWS environment. As a result, the design of your AWS environment might differ from the examples provided in this document. However, these best practices will help you make informed decisions as you design your environment.

Relation to AWS Well-Architected

[AWS Well-Architected](#) helps cloud architects build secure, high-performing, resilient, and efficient infrastructure for their applications and workloads. Based on six pillars—operational excellence, security, reliability, performance efficiency, cost optimization, and sustainability—AWS Well-Architected provides a consistent approach for customers and partners to evaluate architectures and implement designs that can scale over time.

The best practices for organizing your AWS environment addressed in this guide augment and support the best practices represented in the Well-Architected pillars.

Refer to [Appendix A: Relation to AWS Well-Architected](#) for a detailed list of areas of the Well-Architected framework that are related to how you organize your AWS environment.

Intended audience

These best practices are oriented toward cloud architects and technical leads who are responsible for the overall security and architecture of an AWS environment. Whether you are new to AWS or you have already been using AWS for years, your team will benefit from reviewing these best practices and comparing them to your requirements and current AWS environment.

These best practices are intended to apply to organizations largely independent of their industry, size, expected scale of adopting AWS, and workload portfolio. Depending on your needs, not all of the best practices might apply to your situation.

If you're just starting to experiment and learn about AWS by using a single AWS account, you don't need to consider these best practices until you begin planning for your first few production workloads.

Benefits of using multiple AWS accounts

As you adopt AWS, we recommend that you determine how your business, governance, security, and operational requirements can be met in AWS. Use of multiple AWS accounts plays an important role in how you meet those requirements.

The use of multiple accounts enables you to realize the benefits in the following sections.

Topics

- [Group workloads based on business purpose and ownership](#)
- [Apply distinct security controls by environment](#)
- [Constrain access to sensitive data](#)
- [Promote innovation and agility](#)
- [Limit scope of impact from adverse events](#)
- [Support multiple IT operating models](#)
- [Manage costs](#)
- [Distribute AWS Service Quotas and API request rate limits](#)

Group workloads based on business purpose and ownership

You can group workloads with a common business purpose in distinct accounts. As a result, you can align the ownership and decision making with those accounts and avoid dependencies and conflicts with how workloads in other accounts are secured and managed.

Different business units or product teams might have different processes. Depending on your overall business model, you might choose to isolate distinct business units or subsidiaries in different accounts. Isolation of business units can help them operate with greater decentralized control, but still provides the ability for you to provide overarching guardrails. This approach might also ease divestment of those units over time.

Guardrails are governance rules for security, operations, and compliance that you can define and apply to align with your overall requirements.

If you acquire a business that is already operating in AWS, you can move the associated accounts intact into your existing organization. This movement of accounts can be an initial step toward integrating acquired services into your standard account structure.

Apply distinct security controls by environment

Workloads often have distinct security profiles that require separate control policies and mechanisms to support them. For example, it's common to apply different security and operational policies for the non-production and production environments of a given workload. By using separate accounts for the non-production and production environments, by default, the resources and data that make up a workload environment are separated from other environments and workloads.

Constrain access to sensitive data

When you limit sensitive data stores to an account that is built to manage it, you can more easily constrain the number of people and processes that can access and manage the data store. This approach simplifies the process of achieving least privilege access. Limiting access at the coarse-grained level of an account helps contain exposure to highly sensitive data.

For example, designating a set of accounts to house publicly accessible (Amazon S3) buckets enables you to implement policies for all your other accounts to expressly forbid making Amazon S3 buckets publicly available.

Promote innovation and agility

At AWS, we refer to your technologists as [builders](#) because they are all responsible for building value using AWS products and services. Your builders likely represent diverse roles, such as application developers, data engineers, data scientists, data analysts, security engineers, and infrastructure engineers.

In the early stages of a workload's lifecycle, you can help promote innovation by providing your builders with separate accounts in support of experimentation, development, and early testing. These environments often provide greater freedom than more tightly controlled production-like test and production environments by enabling broader access to AWS services while using guardrails to help prohibit access to and use of sensitive and internal data.

- **Sandbox accounts** are typically disconnected from your enterprise services and do not provide access to your internal data, but offer the greatest freedom for experimentation.
- **Development accounts** typically provide limited access to your enterprise services and development data, but can more readily support day-to-day experimentation with your enterprise approved AWS services, formal development, and early testing work.

In both cases, we recommend security guardrails and cost budgets so that you limit risks and proactively manage costs.

In support of later stages of the workload lifecycle, you can use distinct test and production accounts for workloads or groups of related workloads. Having an environment for each set of workloads can enable owning teams to move faster by reducing dependencies on other teams and workloads and minimizing the impact of changes.


Limit scope of impact from adverse events

An AWS account provides security, access, and billing boundaries for your AWS resources that can help you achieve resource independence and isolation. By design, all resources provisioned within an account are logically isolated from resources provisioned in other accounts, even within your own AWS environment.

This isolation boundary provides you with a way to limit the risks of an application-related issue, misconfiguration, or malicious actions. If an issue occurs within one account, impacts to workloads contained in other accounts can be either reduced or eliminated.

Support multiple IT operating models

Organizations often have multiple IT operating models or ways in which they divide responsibilities among parts of the organization to deliver their application workloads and platform capabilities. The following figure shows three example operating models:



Example operating models

- In the *Traditional Ops* model, teams who own custom and commercial off-the-shelf (COTS) applications are responsible for engineering their applications, but not for their production operations. A cloud platform engineering team is responsible for engineering the underlying platform capabilities. A separate cloud operations team is responsible for the operations of both applications and platform.
- In the *CloudOps* model, application teams are also responsible for production operations of their applications. In this model, a common cloud platform engineering team is responsible for both engineering and operations of the underlying platform capabilities.
- In the *DevOps* model, the application teams take on the additional responsibilities of engineering and operating platform capabilities that are specific to their applications. A cloud platform engineering team is responsible for engineering and operations of shared platform capabilities that are used by multiple applications.

As a practice, IT Service Management (ITSM) is a common element across all of the models. Your overall goals and requirements of ITSM might not change across these models, but the responsible individuals and solutions for meeting those goals and requirements can vary depending on the model.

Given the implications of centralized operations versus more distributed operational responsibilities, you will likely benefit from establishing separate groups of accounts in support of different operating models. Use of separate accounts enables you to apply distinct governance and operational controls that are appropriate for each of your operating models.

To learn more about operating models and their implications on your cloud adoption, refer to the [AWS Well-Architected Operational Excellence Pillar Operating Model](#).

Manage costs

An account is the default means by which AWS costs are allocated. Because of this fact, using different accounts for different business units and groups of workloads can help you more easily report, control, forecast, and budget your cloud expenditures.

In addition to cost reporting at the account level, AWS has built-in support to consolidate and report costs across your entire set of accounts. When you require fine-grained cost allocation, you can apply cost allocation tags to individual resources in each of your accounts.

For more information about cost optimization, see the AWS Well-Architected Cost Optimization Pillar's [Expenditure and Usage Awareness](#) best practices.

Distribute AWS Service Quotas and API request rate limits

[AWS Service Quotas](#), also known as limits, are the maximum number of service resources or operations that apply to an account. For example, the number of Amazon S3 buckets that you can create for each account.

You can use Service Quotas to help protect you from unexpected excessive provisioning of AWS resources and malicious actions that could dramatically impact your AWS costs.

AWS services can also throttle or limit the rate of requests made to their API operations.

Because Service Quotas and request rate limits are allocated for each account, use of separate accounts for workloads can help distribute the potential impact of the quotas and limits.

To learn more about managing service quotas, refer to the AWS Well-Architected Reliability Pillar: [Manage Service Quotas and Constraints](#).

Core concepts

This section covers the following core concepts for defining your multi-account strategy on AWS:

Topics

- [AWS Organizations](#)
- [Benefits of using organizational units \(OUs\)](#)
- [Multiple organizations](#)

AWS Organizations

[AWS Organizations](#) helps you centrally govern your environment as you grow and scale your workloads on AWS. Whether you are a growing startup or a large enterprise, Organizations helps you to centrally provision accounts and resources; secure and audit their environment for compliance; share resources; control access to accounts, regions, and services; as well as optimize costs and simplify billing. Additionally, Organizations supports aggregation of health events, consolidated data on use of access permissions, and centralized management of backups and tagging for multi-account environments.

This section includes best practices for organizing your AWS accounts, including grouping your accounts into organizational units (OUs) so that you can more effectively secure and manage your overall AWS environment.

What is an organization?

An *organization* is an entity that you create to consolidate a collection of accounts so that you can administer them as a single unit. Within each organization, you can organize the accounts in a hierarchical, tree-like structure with a [root](#) at the top and [organizational units](#) (OUs) nested under the root. Each account can be placed directly in the root, or placed in one of the OUs in the hierarchy.

Each organization consists of:

- A management account
- Zero or more member accounts
- Zero or more organizational units (OUs)

- Zero or more policies

Organizations management account

The management account creates the AWS organization's resources, OUs, and policies, to manage the organization's member accounts. Access to the management account must be strictly controlled by a small set of highly-trusted individuals from the organization, following the Principles of Least Privilege based on the activities they need to perform. This account is not used for workloads and should generally not contain customer resources.

Additionally, the organization management account is where automation tooling is installed to automate consistent deployment of guardrails or other standardized infrastructure constructs across accounts in an organization. A trust relationship, which is used by the automation tooling, exists between child AWS accounts in the organization and the organization management account. This relationship is established by default when new AWS accounts are created in the organization, and it enables management account users and roles to assume this cross-account [AWS Identity and Access Management \(IAM\) role](#) in child accounts.

Considerations for setting up the management account:

Most customers start with one AWS account, where they build some Proof of Concepts (PoCs) before deploying their workloads on AWS. In this situation, we recommend creating a new AWS account to be your management account, and [inviting your existing account](#) into your new AWS organization. This allows you to keep any PoCs or workloads that you might already have in that account intact.

When you set up the management account, we recommend using an email address that belongs to a shared mailbox, to avoid losing access to this account if only one individual has access to this email address, and for example, they leave your organization or lose access to the account.

Organizations member accounts

AWS Organizations member accounts belong to the organization and reside in the overall organization's structure. All billing for member accounts is consolidated to the management account of the organization.

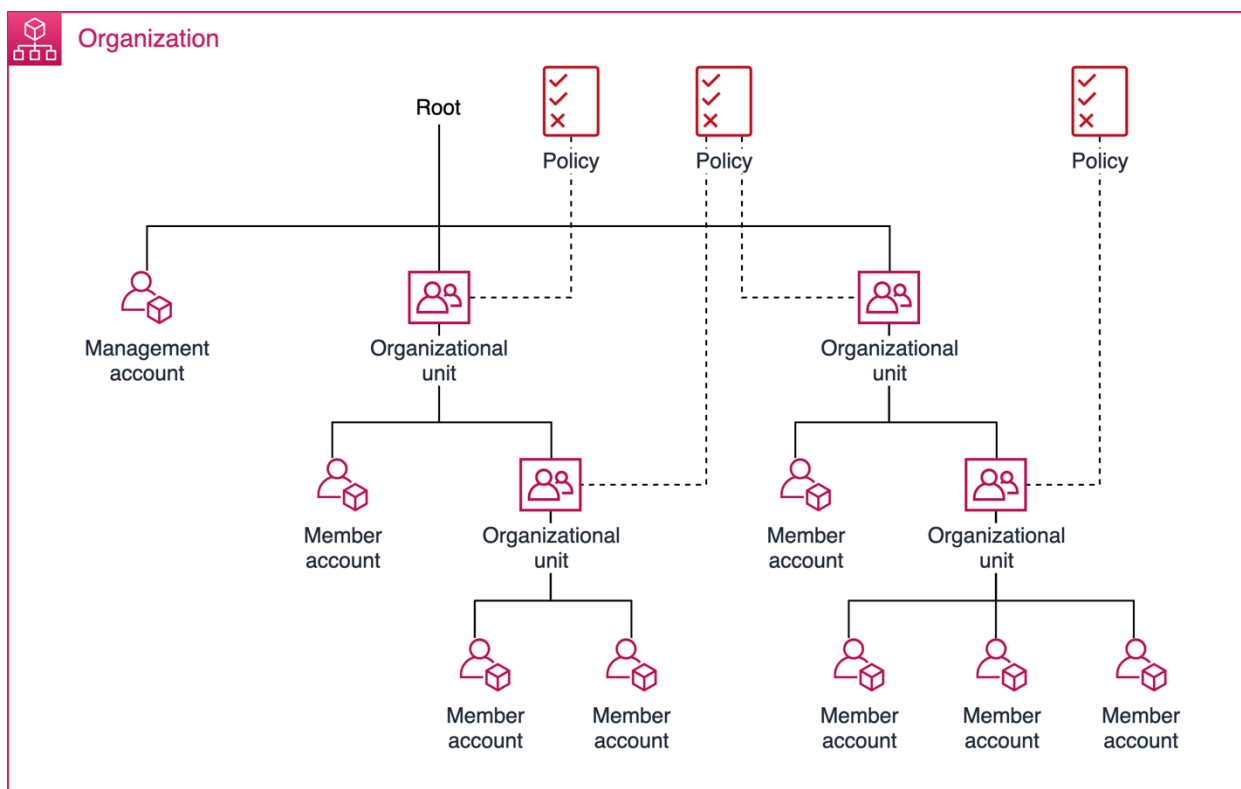
Most of your workloads will reside in member accounts, except for some centrally managed processes that must reside in either the management account or in accounts assigned as designated administrators for specific AWS services.

Organizational units

An organizational unit (OU) provides a means to group accounts within a root. An OU can also contain other OUs. When you attach a policy to one of the nodes in the hierarchy, it flows down and affects all the branches (OUs) and leaves (accounts) beneath it. An OU can have exactly one parent, and each account can be a member of exactly one OU.

OUs are not meant to mirror your own organization's reporting structure. Instead, OUs are intended to group accounts that have common overarching security policies and operational needs. The primary question to ask yourself is: How likely will the group need a set of similar policies?

The following diagram shows a basic organization that consists of seven accounts that are organized into four OUs under the [root](#). The organization also has a few policies that are applied to OUs.



Example of a basic organization

Benefits of using organizational units (OUs)

The following benefits of using OUs helped shape the [Recommended OUs and accounts](#) and [Patterns for organizing your AWS accounts](#).

- [Group similar accounts based on function](#)
- [Apply common policies](#)
- [Share common resources](#)
- [Provision and manage common resources](#)

Group similar accounts based on function

When you have multiple accounts that perform either similar or related functions, you can benefit from grouping these accounts into distinct top-level OUs. Prudent use of top-level OUs can help your teams better understand the overall structure of your AWS accounts.

For example, these best practices recommend [a set of top-level OUs](#) to help you organize different sets of related accounts. At a minimum, the top-level OUs are used to distinguish between overall functions of accounts.

Apply common policies

OUs provide a way for you to organize your accounts so that it's easier to apply common overarching policies to accounts that have similar needs. Policies in AWS Organizations enable you to apply additional types of management to the accounts in your organization.

By attaching policies to OUs rather than to individual accounts, you can simplify management of policies across groups of similar accounts. As the number of accounts in your environment grows, simplifying policy management by attaching policies to OUs becomes more important.

AWS Organizations supports use of authorization and management policies. For a complete list of policy types, refer to [Managing AWS Organizations policies](#).

Authorization policies

AWS Organizations [service control policies](#) (SCPs) are a type of organization policy that you can use to manage permissions in your organization. SCPs offer central control over the maximum available permissions for all accounts in your organization.

SCPs are a means of implementing guardrails in your AWS organization. Your use of SCPs can help ensure that your accounts stay within your access control guidelines. For example, you can use SCPs to constrain the set of AWS services and actions allowed on resources.

Although you can apply SCPs to the root of your organization, you typically associate SCPs with underlying OUs. For example, based on the nature of the workloads deployed in accounts within an OU, you might choose to restrict the set of AWS services and AWS Regions that are allowed to be used by accounts in the OU.

You only apply an SCP at the root when you have an overarching security policy that applies across your entire organization and set of OUs. For example, you might apply an SCP at the root of the organization to deny AWS Organizations member accounts from attempting to leave the organization of their own accord.

Management policies

You can also apply [tag policies](#) to your parts of your organization to help you monitor and ensure compliance with your cloud resource tagging standards.

You can use artificial intelligence (AI) services [opt-out policies](#), which enable you to control data collection for AWS AI services for all of your organization's accounts.

[Backup policies](#) help you centrally manage and apply backup plans to the AWS resources across your organization's accounts.

Share common resources

OUs provide a means for you to organize your accounts so that it's easier for you to share centrally managed resources across similar accounts.

AWS services have been introducing support for sharing their resources through [AWS Resource Access Manager](#) (AWS RAM) and AWS Organizations. For example, with AWS RAM, you can use OUs as the basis for sharing centrally managed network resources such as [Amazon Virtual Private Cloud](#) (Amazon VPC) [subnets](#).

Provision and manage common resources

Sometimes you need to deploy common, centrally managed resource configurations to groups of related accounts. In cases where resource sharing doesn't apply, you can use a variety of AWS services and third-party tools that work with OUs to automatically roll out and update your own custom resources.

For example, you can use OUs as a basis for targeting automation to deploy and update your own sets of IAM roles and [customer managed IAM policies](#) that help establish common baseline and/or workload-specific security controls to groups of related accounts.

Multiple organizations

Most customers are best served using a single production organization, and we recommend that customers manage their accounts within a single organization. This allows you to ensure consistency across accounts in your environment because centrally-applied policies or service-level configurations are programmatically applied across accounts within your organization. Separating your workload accounts across organizations requires additional overhead or customization to ensure central standards are applied within each organization.

There are certain exceptions, outlined in this section, where you might need to work across multiple organizations.

- [Test changes to your overall AWS environment](#)
- [Support acquisitions and divestments](#)
- [Support large AWS environments](#)
- [Align with your billing requirements](#)
- [Different classification levels for government applications](#)

Test changes to your overall AWS environment

You might need to develop code that interacts with APIs and other mechanisms fundamental to the management of your organization. In these cases, to determine whether your changes break something without having to make the changes in your production organization, we recommend that you test your changes in an organization different from the one running your production workloads.

For example, you might need to either modify the automation that creates new accounts to change the configuration baseline of accounts it creates or change the configuration of a workflow management system you're using to modify SCPs. In addition, you might want to test the delegated administration capabilities of various AWS services prior to applying them in your production organization.

In these circumstances, we recommend that you establish an additional organization for testing that resembles your more formally managed production organization. You can perform testing of changes to how you manage your organization in your test organization before applying those changes to be applied to your production organization.

Support acquisitions and divestments

You might acquire an entity that has already established an organization. If you decide to merge the acquired entity's AWS environment with your AWS environment, you can move member accounts from the acquired organization to your mainstream organization. In this case, you can later decommission the acquired entity's organization.

If you plan to potentially divest a portion of your portfolio, you can manage the workloads and supporting AWS accounts for that portion of your portfolio in a separate organization to support simpler divesture and isolated billing.

Support large AWS environments

If you need more accounts than the maximum number supported by an organization, we recommend that you divide your accounts between multiple organizations. For more details about the maximum number of accounts supported in an organization, refer to [Quotas for AWS Organizations](#).

Align with your billing requirements

An organization gathers billing information from all member accounts into a single AWS bill. If you have use cases where different sets of accounts require distinct bills or payments, then multiple organizations might be required.

Different classification levels for government applications

Some customers in government, critical national infrastructure providers, and defense industries need to handle data with defined classification levels. These customers require high-assurance mechanisms to keep data (and, in most cases, metadata) associated with at least some of those markings, separate from each other. Assets within a single account should all handle data at the same protective marking.

As noted in this section, an organization itself contains collective data from multiple accounts. Data such as account names, billing data, organizational unit names, and activity logs can be accessed centrally for those with appropriate permissions, such as a cloud administrator or audit team. This means that commingling of billing and logging data from accounts that are processing data might not meet a customer's requirements for separation by classification levels.

An organization's configuration could be modified to have customized separation for protective marking with proper tags, logging customization (based on protective markings), and defined

permissions for administrative users. However, this might be more easily achieved with separate organizations.

Design principles for your multi-account strategy

The following design principles helped develop the best practices described in this paper. You can also use these principles to help guide your initial account design and evolve it over time.

These design principles complement the [Benefits of using multiple accounts](#) and [Benefits of using OUs](#).

Topics

- [Organize based on security and operational needs](#)
- [Apply security guardrails to OUs rather than accounts](#)
- [Avoid deep OU hierarchies](#)
- [Start small and expand as needed](#)
- [Avoid deploying workloads to the organization's management account](#)
- [Separate production from non-production workloads](#)
- [Assign a single or small set of related workloads to each production account](#)
- [Use federated access to help simplify managing human access to accounts](#)
- [Use automation to support agility and scale](#)
- [Use multi-factor authentication](#)
- [Break glass access](#)

Organize based on security and operational needs

We recommend that you organize accounts using [OUs based on function](#), compliance requirements, or a common set of controls rather than mirroring your organization's reporting structure.

Apply security guardrails to OUs rather than accounts

Where feasible, we recommend that you apply security guardrails (for example, SCPs) to OUs instead of accounts so that you can more efficiently manage the distribution of guardrails across accounts that have the same or similar requirements.

For more information about managing security guardrails, refer to [Permissions management](#) in the AWS Well-Architected Security Pillar.

Avoid deep OU hierarchies

Overly complicated structures can be difficult to understand and maintain. Although AWS Organizations supports a depth of five levels of OUs, the recommended structure strives to use OUs only when there is sufficient benefit.

When you consider the addition of new OU levels, you should review the [Benefits of using OUs](#) and these principles to decide whether the additional complexity adds sufficient value.

Start small and expand as needed

We recommend that you review the example [Patterns for organizing your AWS accounts](#), start with a subset of the [Recommended OUs and accounts](#), and expand the structure of your AWS accounts when your needs call for the creation of new OUs.

You shouldn't need to invest a lot of time at the beginning of your adoption journey designing what you expect your AWS account structure will look like in several years.

We provide examples of successive degrees of building out an AWS account structure in [Patterns for organizing your AWS accounts](#).

Avoid deploying workloads to the organization's management account

Since privileged operations can be performed within an organization's management account and SCPs do not apply to the management account, we recommend that you limit access to an organization's management account. You should also limit the cloud resources and data contained in the management account to only those that must be managed in the management account.

Separate production from non-production workloads

We recommend that you separate production workloads from non-production workloads. For overall recommendations on designing this separation, refer to [Organizing workload-oriented OUs](#).

Assign a single or small set of related workloads to each production account

In support of your production workloads, we recommend that you either assign a single workload to each production account or assign a small set of closely related workloads to each production account.

Consider separating workloads that have different owners into their own production accounts to simplify access management, streamline change approval processes, and limit the scope of impact for misconfigurations.

Use federated access to help simplify managing human access to accounts

We recommend that you use AWS identity federation capabilities by using either IAM Identity Center or IAM integration with a third-party identity provider. These capabilities enable you to use a common identity provider and your existing processes for controlling human user access to your AWS accounts.

By using federated access and a common identity provider, you avoid the need to manage individual users in each account. Instead, your human users can use their existing credentials to access authorized accounts. You also gain the benefit of keeping personally identifiable information (PII) out of IAM.

With federated access, your human users use temporary credentials instead of long-term access keys for programmatic access to their AWS environments.

Use of federated access avoids the creation and management of users in your AWS accounts for humans. Instead, use of users can be limited to those exceptional cases such as third-party applications that do not support the use of roles.

For more information about managing identities, refer to [Identity Management](#) in the AWS Well-Architected Security Pillar and [Identity federation in AWS](#).

Use automation to support agility and scale

It is important to design and manage your accounts so that you can rapidly respond to business needs without the need for a corresponding linear increase in headcount. When you consider

moving beyond managing just a few accounts, you must consider the work to establish processes and automation that will enable you to do so in an efficient manner.

For example, if you implement an account design in which new business initiatives call for the creation of new accounts, then you will benefit from having automation in place so that you can rapidly and reliably provision environments based on your standard configurations. Automation can also help you monitor compliance and apply updates to your baseline configurations over time.

Use multi-factor authentication

Multi-factor authentication (MFA) should be used by your root and all AWS users in your accounts regardless of privilege level or access mechanism. You can follow our current recommendation for MFA best practices for your AWS accounts ([management account](#) or [member accounts](#)) to set up MFA across your AWS environment.

Break glass access

The organization management account is used to provide break glass access to AWS accounts within the organization. Break glass (which draws its name from breaking the glass to pull a fire alarm) refers to a quick means for a person who does not have access privileges to certain AWS accounts to gain access in exceptional circumstances by using an approved process.


The use cases for break glass access include:

- Failure of the organization's IdP.
- A security incident involving the organizations' IdP(s).
- A failure involving IAM Identity Center.
- A disaster involving the loss of an organization's entire cloud or IdP teams. It is important that access to these roles is monitored, and alarms and alerts are triggered when the roles are used to access the environment.

In the case of an incident requiring remediation, we recommend that a user with access to an administrative federated role within the AWS account perform the required remediation. In cases where this user is unavailable to carry out a time sensitive action, we recommend that a highly-restricted group or set of groups be preconfigured within your IdP, each providing appropriate [federated access](#) into the appropriate set of AWS accounts. A user can either be added into one of

these groups using a high-priority and temporary change request, or a select group of privileged and trusted users can be prepopulated into these groups.

Security teams investigating an incident would use this mechanism to access a read-only role in an impacted account, or use the read-only access mechanism provided through the security tooling account. In summary, common high-priority irregular access scenarios need to be incorporated into standard federated access processes and procedures.

 **Note**


AWS Organizations Service Control Policies do not apply to the organization management account, and administrator access to this account would grant privileged status to the entire organization, given the trust relationship to the management account. Therefore, access to break glass IAM users must be tightly controlled, but accessible through a predefined and strict process. This process often involves one trusted individual having access to the password, and a different trusted individual having access to the hardware multi-factor authentication (MFA) key, meaning it typically requires two people to access any one set of break glass credentials.

Human access to AWS accounts within the organization should be provided using federated access. Although the use and creation of AWS IAM users is highly discouraged, break glass users are an exception.

To ensure human break-glass access to your environment, we recommend that you create the following in your AWS organization:

- At least two IAM users with IAM login credentials to prevent lockdown in case one of them is not available, and additional users depending on your operating model. Do not create unnecessary IAM privileged users in your management account. These users will assume roles in the member accounts in your organization through trust policies.
- A break glass role that is deployed to all the accounts in the organization, and that can only be assumed by the break glass users from the management account. These roles are needed to allow access from the management account to apply and update guardrails, to troubleshoot and resolve issues with the automation tooling from the security tooling account, or to remediate security and operational issues in one of the member accounts in the AWS organization. When setting up these roles in your organization, you need to ensure they can be used in emergency

situations, bypassing established controls under the situations described earlier in the paragraph, such as service control policies.

 **Note**


If you are currently using AWS Identity Center and you are not using an external IdP (you are using the IAM Identity Center store or your domain service for your identity source), you can use this break glass access in case of Identity Center failure. Review how to set up [emergency access for your IAM Identity Center](#).

We strongly recommend configuring these users with a hardware-based MFA device, which can be used in exceptional circumstances to gain access to the organization management account or sub-accounts within the organization by assuming a role.

While we recommend the use of the organization management account for break glass access, some organizations might choose to add a dedicated break glass account. This does not eliminate the need for organizational break glass users in the organization management account.

Recommended OUs and accounts

This section provides details on the recommended OUs and, when applicable, a set of recommended AWS accounts.



Recommended OUs

Depending on your requirements, you might not need to establish all the recommended OUs. As you adopt AWS and learn more about your needs, you can expand the overall set of OUs. Refer to the [Patterns for organizing your AWS accounts](#) for examples of how you might begin to organize your AWS accounts.

While the provided OU recommendations are geared towards common use cases, it is your organization's responsibility to define a customized OU structure that aligns with your distinct requirements relevant to isolation and automation.

The recommended OUs consist of:

Topics

- [Security OU](#)
- [Infrastructure OU](#)
- [Sandbox OU](#)
- [Workloads OU](#)

- [Policy Staging OU](#)
- [Suspended OU](#)
- [Individual Business Users OU](#)
- [Exceptions OU](#)
- [Deployments OU](#)
- [Transitional OU](#)
- [Business Continuity OU](#)

The Security OU and the Infrastructure OU are categorized as foundational OUs. Foundational OUs are defined as OUs that contain accounts, workloads, and other AWS resources that provide common security and infrastructure capabilities to secure and support your overall AWS environment.

Accounts, workloads, and data residing in the foundational OUs are typically owned by your centralized Cloud Platform or Cloud Engineering teams made up of cross-functional representatives from your Security, Infrastructure, and Operations teams.

The majority of your accounts are contained in the other OUs. These OUs are intended to contain your business-related workloads. They also contain tools and services that support the entire lifecycle of your business-related services and data.

Security OU

The Security OU is a foundational OU. Your security organization should own and manage this OU along with any child OUs and associated accounts.

We recommend that you create the following accounts in the Security OU:

- Log Archive
- Security Tooling (Audit)

Note

A default deployment of AWS Control Tower will create a Log Archive and Audit (also referred to as Security Tooling) accounts.

Depending on your initial requirements, you might not need to establish all of these accounts. Refer to [Patterns for organizing your AWS accounts](#) for example sets of OUs and accounts that are commonly used in the early stages of adopting AWS.

Log Archive account

The Log Archive is an account that acts as a consolidation point for log data that is gathered from all the accounts in the organization and primarily used by your security, operations, audit, and compliance teams. This account contains a centralized storage location for copies of every account's audit, configuration compliance, and operational logs. It also provides a storage location for any other audit/compliance logs, as well as application/OS logs. For example, in this account, we recommend that you consolidate AWS API access logs recorded in AWS CloudTrail, logs of changes to AWS resources recorded in AWS Config, and other logs that have security implications. If you use VPC peering between accounts, then you might also benefit from consolidating [VPC Flow Logs](#) data in this account. Logs should generally be made directly available for local use by teams working in any account on a shorter-term retention basis. It is common practice to auto-ingest logs from the log archive account into a security information and event management (SIEM) solution.

Note

By utilizing AWS Control Tower for AWS environment management, it automatically enforces best practices, deploying AWS Config and AWS CloudTrail seamlessly across your environment. Their logs are consolidated in an Amazon S3 bucket within the Log Archive account.

Recommended AWS Organization Integrated Service Delegation

AWS service	Implementation Details	Control Tower Enabled
Amazon Security Lake	Amazon Security Lake centralizes security data from cloud, on-premises, and custom sources into a data lake that's stored in your account.	No

Services in the Log Archive account

With [Amazon Security Lake](#), you can automatically centralize security data from AWS and third-party sources into a data lake that's stored in your [Log Archive account](#). Review [Managing access in this account](#) in the following sections to learn how to grant access to the logs from other accounts in your AWS organization.

Logs should be available within the workload account for use by teams on short-term retention basis. It is common practice to auto-ingest logs from the log archive account into a security information and event management (SIEM) solution.

If you are using [AWS Control Tower](#) to manage your overall AWS environment, then AWS Config is automatically enabled in each Control Tower enrolled account, and AWS CloudTrail Org trail is created for all accounts in the Organization. The AWS CloudTrail logs and AWS Config configuration history are consolidated in an Amazon S3 bucket in the log archive account.

Operational log data

Operational log data used by your infrastructure, operations, and workload owning teams often overlaps with the log data used by security, audit, and compliance teams.

We recommend that you consolidate your operational log data into the Log Archive account. Based on your specific security and governance requirements, you might need to filter operational log data saved to this account. You might also need to specify who and what has access to the operational log data in the log archive account.

Immutable log data

Log data housed in the Log Archive account is considered immutable in that it is protected from being changed or deleted. Data retention policies and legislation that apply to your organization might also apply to the data in your log archive account.

Managing access to this account

We strongly recommend that you only house log data in this account. By doing so, access to this account can be greatly limited.

Workloads and tools that need to consume the consolidated log data are typically housed in your other accounts and are granted access through read-only IAM roles to access the log data in a read-only, least privileged manner.

Additionally, to ensure log data is properly protected, we recommend SCPs be applied to the Security OU preventing modification or deletion of files within the centralized logging S3 bucket(s). Additionally, the use of S3 bucket versioning provides visibility into the complete history of all log files.

Security Tooling (Audit) account

In the context of AWS services, this account is used to provide centralized delegated admin access to AWS security tooling and consoles, as well as provide view-only access for investigative purposes into all accounts in the organization. The security tooling account should be restricted to authorized security and compliance personnel and related security. This account is an aggregation point (or points for organizations that split the functionality across multiple accounts) for AWS security services, including [AWS Security Hub](#), [Amazon GuardDuty](#), [Amazon Macie](#), [AWS AppConfig](#), [AWS Firewall Manager](#), [Amazon Detective](#), [Amazon Inspector](#), and [IAM Access Analyzer](#).

Note

ViewOnlyAccess and *ReadOnlyAccess* IAM managed policies provide permissions that do not include mutable actions. The *ReadOnlyAccess* grants read access to all AWS services and resources whereas the *ViewOnlyAccess* access provides read-only access and further restricts read operations to view resources and only metadata.

Recommended AWS Organization Integrated Service Delegation

AWS service	Implementation Details	Control Tower Enabled
AWS Audit Manager	Continuously audit your AWS use across multiple accounts in your organization to simplify how you assess risk and compliance. Recommended to be in same AWS account AWS Security Hub delegated admin exists.	No

AWS service	Implementation Details	Control Tower Enabled
	Delegation needs to be done on home and operational AWS Regions.	
AWS CloudFormation Stacksets	CloudFormation Stacksets can be delegated to multiple accounts within your AWS Organization. Delegation of the service needs to be completed at only one AWS region for the AWS account.	Yes, delegation not configured
AWS CloudTrail	The management of CloudTrail Org Trails can be delegated to one account. It is recommended that the Security team manage the implementation.	Yes, delegation not configured
AWS Config	Organization-wide aggregated view of your AWS resources, your AWS Config rules, and the AWS resources' compliance state. Creating an Organization aggregator can be done across multiple AWS regions into the region the aggregator is being deployed to. Multiple accounts can be delegated the AWS Config aggregator.	Yes, delegation not configured

AWS service	Implementation Details	Control Tower Enabled
AWS Detective	Required to be deployed to same account which is managing Amazon GuardDuty and AWS Security Hub. Requires GuardDuty to be enabled on Security Tooling account prior to delegating AWS Detective. Delegation needs to be done on home and operational AWS Regions.	No
AWS Firewall Manager	Configure full delegated administration support for Security Tooling account. Firewall Manager delegation is a global configuration for all AWS Regions and only needs to be delegated from your home AWS Region.	No
Amazon GuardDuty	Amazon GuardDuty allows for one delegated admin per AWS Organization. It is recommended to delegate Amazon GuardDuty to the same account AWS Security Hub and Amazon Macie are delegated to. Delegation needs to be done on home and operational AWS Regions.	No

AWS service	Implementation Details	Control Tower Enabled
Amazon Inspector	Delegate an administrator to enable or disable scans for member accounts, view aggregated finding data from the entire organization, create and manage suppression rules. Delegation needs to be done on home and operational AWS Regions.	No
Amazon Macie	Amazon Macie allows for one delegated admin per AWS Organization. It is recommended to delegate Amazon Macie to the same account AWS Security Hub and Amazon GuardDuty are delegated to. Delegation needs to be done on home and operational AWS Regions.	No
AWS Security Hub	AWS Security Hub allows for one delegated admin per AWS Organization. It is recommended to delegate AWS Security Hub to the same account Amazon GuardDuty and Amazon Macie are delegated to, for ease of pivoting between these services in the AWS Console. Delegation needs to be done on each operational Region.	No

AWS service	Implementation Details	Control Tower Enabled
Amazon S3 Storage Lens	Allows for multiple delegated admin accounts per AWS Organization. Service is global and only needs to be delegated from the home AWS Region	No
AWS Trusted Advisor	Allows for centralized view of AWS Trusted Advisor information. Requires the management account in your organization must have a Business, Enterprise On-Ramp, or Enterprise Support plan. Service is global and only needs to be delegated from the home AWS Region.	No
IAM Access Analyzer	Configured with the entire AWS organization as the zone of trust so that it's easier for you to quickly look across resource policies and identify resources with public or cross-account access you might not intend. We recommend that you configure this analyzer in one of your security tooling accounts.	No

Additional Services and Functionalities

Common examples of security capabilities that can be centrally accessed and managed using the Security Tooling account include:

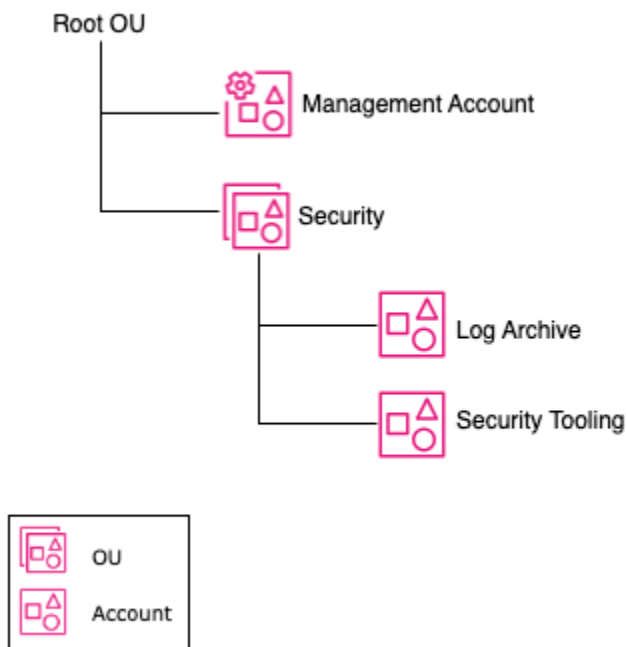
- **Third-party cloud security monitoring tools** — You can also house third-party cloud security monitoring services and tools in your security tooling accounts. For example, these accounts typically contain security information and event management (SIEM) tools and vulnerability scanners
- **Automated detection and response workflows** — Automated detection and response workflows that act on data collected through these types of services are normally contained in your security tooling accounts.
- **Incident response (IR) support** — Tools to support manual incident response (IR) procedures are typically housed in your security tooling accounts. Refer to the [AWS Security Incident Response Guide](#) for more information.

AWS Solutions

AWS Solution	Description
Automated Security Response on AWS	Add-on that works with AWS Security Hub and provides predefined response and remediation actions based on industry compliance standards and best practices for security threats. It helps Security Hub customers to resolve common security findings and to improve their security posture in AWS.
Automations for AWS Firewall Manager	Allows you to centrally configure, manage, and audit firewall rules across all your accounts and resources in AWS Organizations. This solution is a reference implementation to automate the process to set up AWS Firewall Manager security policies.
Security Automations for AWS WAF	Automatically deploys a set of AWS WAF (web application firewall) rules that filter common web-based attacks. Users can select from preconfigured protective features that define the rules included in an AWS WAF web access control list (web ACL).

Example structure

The following example structure represents the recommended Security OU at a basic level. Note that within Control Tower governed environments, the accounts within the Security OU are limited to the Log Archive and Security Tooling (also known as Audit by default for AWS Control Tower deployments).



Example structure of Security OU

Infrastructure OU

The Infrastructure OU is a foundational OU that is intended to contain infrastructure services. The accounts in this OU are also considered administrative and your infrastructure and operations teams should own and manage this OU, any child OUs, and associated accounts.

The Infrastructure OU is used to hold AWS accounts containing AWS infrastructure resources that are shared, utilized by, or used to manage accounts in the organization. This includes centralized operations or monitoring of your organization. No application accounts or application workloads are intended to exist within this OU.

Common use cases for this OU include accounts to centralize management of resources. For example, a Network account might be used to centralize your AWS network, or an Operations Tooling account to centralize your operational tooling.

Note

For guidance on where to contain non-infrastructure shared services, refer to [Workloads OU](#).

In most cases, given the way most AWS Organization integrated services interact with the accounts within the Infrastructure OU, it does not generally make sense to have production and non-production variants of these accounts within the Infrastructure OU. In situations where non-production accounts are required, these workloads should be treated like any other application and placed in an account within the appropriate Workloads OU corresponding with the non-production phase of the SDLC (Dev OU or Test OU).

Backup account

The Backup account serves as a dedicated and centralized hub for backup and disaster recovery management. It provides a unified platform to orchestrate, monitor, and enforce backup policies across AWS accounts within the AWS Organization.

By consolidating backup processes in a central account, organizations can achieve several benefits. It simplifies backup management by eliminating the need to configure and maintain backup settings separately in each member account, streamlining operational efficiency and reducing the potential for errors. It ensures consistent and comprehensive data protection across the entire AWS infrastructure, regardless of the specific AWS services and resources in use. This approach also enhances compliance and governance efforts by enabling centralized auditing and reporting on backup and recovery activities, making it easier to track data protection metrics and maintain necessary records for compliance purposes.

Recommended AWS Organization Integrated Service Delegation

AWS service	Implementation Details	Control Tower Enabled
AWS Backup	Register the Backup account as the delegated administr	No

AWS service	Implementation Details	Control Tower Enabled
	ator in the AWS Backup console.	
AWS Organizations: AWS Backup policy administration	Delegate AWS Backup Policy administration to the Backup account by enabling delegation of AWS Organizations in the management account and configure a policy that allows the Backup account to create Backup Policies.	No

Additional Services and Functionalities

Common examples of security capabilities that can be centrally accessed and managed using the Backup account includes::

- Leverage centralized AWS KMS customer managed keys for AWS Backup service within the Backup account to centrally manage the encryption for backup operations across accounts.
- 3rd party backup tools that require resources can be created and managed in the Backup account.

Identity account

The Identity account serves as a centralized identity federation account isolated from all other management and workload activities within the AWS Organization. Federated identity management grants you the ability to efficiently manage the access to the accounts in the AWS Organization and authorization to integrated applications. By managing your identities and controlling access to your environment centrally, you can quickly create, update, and delete the permissions and policies you need to meet your business requirements.

Recommended AWS Organization Integrated Service Delegation

AWS service	Implementation Details	Control Tower Enabled
IAM Identity Center	You can delegate administration of IAM Identity Center to this account which will allow you to administer IAM Identity Center outside of the management account.	Enabled - Yes Delegated - No
IAM Access Analyzer	An IAM Access Analyzer can be configured to detect resources that are shared outside of the organization (organization zone of trust). By default, this is managed from the management account. This can be delegated to a member account. This can be delegated to the Identity account or a Security Tooling account depending on who is responsible for auditing external access (Identity Team or Security Team).	No

Additional Services and Functionalities

Common examples of security capabilities that can be centrally accessed and managed using the Identity account includes:

- **AWS Directory Services** - If you are using an AWS-hosted directory or AWS AD Connector, you can create and managed them in your Identity account alongside of AWS IAM Identity Center.

- **SAML 2.0 custom managed applications** - With IAM Identity Center, you can create or connect workforce users and centrally manage their access across all their AWS accounts and applications.

Network account

The Network account serves as the central hub for your network within your AWS Organization. You can manage your networking resources and route traffic between accounts in your environment, your on-premises, and egress/ingress traffic to the internet. Within this account, your network administrators can manage and build security measures to protect network traffic across your cloud environment.

Recommended AWS Organization Integrated Service Delegation

AWS service	Implementation Details	Control Tower Enabled
AWS Network Manger	Centrally manage and monitor your global networks with transit gateways and their attached resources in multiple AWS accounts within your organization.	No
IPAM	Delegated to a single account for your entire AWS Organization. IPAM will inventory and track all active IPs across your AWS Organization.	No
VPC Reachability Analyzer	Trace paths across accounts in your organizations. You can assign multiple delegated admin accounts as needed.	No

Additional Services and Functionalities

Common examples of network capabilities and AWS services that can be centrally accessed and managed via the Network account include:

- **Amazon VPC** - If you plan to implement centralized networking in your AWS environment, we recommend managing your [VPCs](#) within your network account, and sharing resources across your accounts within your AWS organization.
- **Share your AWS Transit Gateway** - Create an [AWS Transit Gateway](#) resource in the networking account and share it across the accounts within your AWS Organization using AWS Resource Access Manager (RAM).
- **Share your Amazon Route 53 Endpoint Resolvers** - If you plan to use a centralized transitive network with [Amazon Route 53 Public Data Plane](#) in your AWS Organization, we recommend managing and sharing your Route 53 Endpoint Resolvers in your network account within your AWS organization.
- **Share your IPAM pools with your organization** - When you delegate an IPAM account, IPAM enables other AWS Organizations member accounts in the organization to allocate CIDRs from IPAM pools that are shared using AWS Resource Access Manager (RAM).
- **Build centralize [AWS Site-to-Site VPN connections](#)** - Using a transitive network architecture centralized in your Network account, a site-to-site VPN can be established and routing enabled across your cloud environment.
- **Centralize [AWS Direct Connect](#)** - Create and attach AWS Direct Connect to your transitive network with [AWS Transit Gateway](#).
- **Centralized network inspection point** - Build inbound and outbound network traffic inspection points routing through the Network account.

AWS Solutions

The following AWS Solutions are commonly deployed or related to the functional operations of the Network account:

AWS Solution	Description
Network Orchestration for AWS Transit Gateway	Automates the process of setting up and managing transit networks in distributed AWS environments. This solution allows customers

AWS Solution	Description
	to visualize and monitor their global network from a single dashboard rather than toggling between Regions from the AWS console. It creates a web interface to help control, audit, and approve transit network changes.
Automations for AWS Firewall Manager	Allows you to centrally configure, manage, and audit firewall rules across all your accounts and resources in AWS Organizations. This solution is a reference implementation to automate the process to set up AWS Firewall Manager security policies.
Security Automations for AWS WAF	Automatically deploys a set of AWS WAF (web application firewall) rules that filter common web-based attacks. Users can select from preconfigured protective features that define the rules included in an AWS WAF web access control list (web ACL).

Operations Tooling account

Operations Tooling accounts can be used for day-to-day operational activities across your organization. The operations tooling account hosts tools, dashboards, and services needed to centralize operations where monitoring and metric tracking are hosted. These tools help the central operations team to interact with their environment from a central location.

Recommended AWS Organization Integrated Service Delegation

AWS service	Implementation Details	Control Tower Enabled
AWS Account Management	Manage alternate contact information for all of the accounts in your organization. Delegation is done on one	No

AWS service	Implementation Details	Control Tower Enabled
	region and for one account within your AWS Organizations.	
AWS Application Migration Service (AMG)	AWS Application Migration Service simplifies, expedites, and reduces the cost of migrating applications to AWS. By integrating with Organizations, you can use the global view feature to manage large-scale migrations across multiple accounts.	No
Amazon DevOps Guru	You can integrate with AWS Organizations to manage insights from all accounts across your entire organization. You delegate an administrator to view, sort, and filter insights from all accounts to obtain organization-wide health of all monitored applications.	No
AWS Health	Get visibility into events that might affect your resource performance or availability issues for AWS services. You can register up to 5 member accounts in your organization as a delegated administrator.	No

AWS service	Implementation Details	Control Tower Enabled
AWS License Manager	If you are planning to use a centralized model to buy and share licenses across your organization, we recommend you specify one of your Shared Services accounts as the delegated administrator for AWS License Manager.	No
AWS Systems Manager Change Manager	You can delegate administration for Systems Manager to the Operations Tooling account to perform administrative tasks for Change Manager, Explorer, and Ops Center.	No
AWS Systems Manager Explorer		No
AWS CloudFormation Stacksets	You can register multiple delegated administrator accounts in your AWS Organizations. CloudFormation Stackset delegation will give the AWS account full administrative access to deploy resources in other AWS accounts in your Organization. Delegation needs to be done only at the home region.	No

AWS service	Implementation Details	Control Tower Enabled
VPC Reachability Analyzer	Trace paths across accounts in your organizations. VPC Reachability Analyzer can have multiple delegated admin accounts.	

AWS Solutions

The following AWS Solutions are commonly deployed or related to the functional operations of the Operations Tooling account:

AWS Solution	Description
Account Assessment for AWS Organizations	Presented in a web UI, this AWS Solution runs configurable scans on all AWS accounts in your AWS Organizations to help you identify dependencies in your underlying resource-based policies.
Instance Scheduler on AWS	Automates the starting and stopping of Amazon Elastic Compute Cloud (Amazon EC2) and Amazon Relational Database Service (Amazon RDS) instances. This solution helps reduce operational costs by stopping resources that are not in use and starting them when they are needed. The cost savings can be significant if you leave all of your instances running at full utilization continuously.
Cost Optimizer for Amazon WorkSpaces	Analyzes all of your Amazon WorkSpaces usage data and automatically converts the WorkSpace to the most cost-effective billing option (hourly or monthly), depending on your individual usage. You can use this solution with a single account, or with AWS Organizat

AWS Solution	Description
	ions across multiple accounts, to help you monitor your WorkSpace usage and optimize costs.
Workload Discovery on AWS	Workload Discovery on AWS (formerly called Amazon Personalize) is a tool to visualize AWS Cloud workloads. Use Workload Discovery on AWS to build, customize, and share detailed architecture diagrams of your workloads based on live data from AWS .

Monitoring account

An AWS monitoring account can be used to monitor resources, applications, log data, and performance in other AWS accounts. AWS offers a number of tools and services that can be used to manage and monitor resources and workloads in an AWS account, including CloudWatch, Amazon Managed Service for Prometheus, Amazon Managed Grafana, and Amazon OpenSearch. These tools can be used to monitor resource and application usage, performance, review log data, and identify potential issues within the infrastructure or application.

Note

Depending on your business requirements and team structures, you may choose to manage your monitoring resources and services in a single account with your other Operational Tooling services or as a dedicated Monitoring account. The core concept of the Monitoring account is to only give read-only functionality. The account in itself is not intended to have the ability to make changes across account your AWS Organization.

Recommended AWS Organization Integrated Service Delegation

AWS service	Implementation Details	Control Tower Enabled
AWS Health	Configure the Monitoring account as the delegated	No

AWS service	Implementation Details	Control Tower Enabled
	admin for AWS health (in the Management account) for ongoing visibility into your resource performance and the availability of your AWS services and accounts within your organization.	
Amazon S3 Storage Lens	Register the Monitoring account as the delegated admin for Amazon S3 storage Lens (in the Management account) for organization-wide visibility into object-storage usage and activity. You can use S3 Storage Lens metrics to generate summary insights, such as finding out how much storage you have across your entire organization or which are the fastest-growing buckets and prefixes.	No

Additional Services and Functionalities

Common examples of monitoring capabilities that can be centrally accessed and managed using the Monitoring account includes:

- **AWS CloudWatch** - Configure AWS [CloudWatch Cross Account observability](#) and configure as the "monitoring account" or hub account.
- **CloudWatch dashboards** that are created at the account level can be shared with the monitoring account which allows for distributed management with centralized monitoring.
- **3rd party monitoring tools** (such as ElasticSearch, Splunk, Prometheus, Grafana) that require resources can be created and managed in the Monitoring account.

- **Customer created automations and reports** can be run from and stored in the Monitoring account.
- **Log Archive log analysis.** In order to analyze Log data stored in the Log Archive account, Amazon Managed Grafana or Amazon QuickSight can be used in the Monitoring account to analyze Log data in an S3 bucket in the Log Archive account by connecting to Amazon Athena in the Log Archive account.
- **Amazon OpenSearch Service** can be deployed and managed in the Monitoring account to analyze logs, monitor applications, and analyze clickstreams.
- **Amazon QuickSight** can be deployed and managed in the Monitoring account and cross account data sources can be used to centrally monitor or report organization data.
- **Amazon Managed Grafana** can be deployed into the monitoring account for centralized monitoring of resources, containers, CloudWatch logs, and applications by connecting to data sources in different accounts or to centralized CloudWatch metrics, logs, and traces.

AWS Solutions

The following AWS solutions are commonly deployed or related to the functional operations of the Monitoring account:

AWS Solution	Description
Centralized Logging on AWS	<p>Helps organizations collect, analyze, and display Amazon CloudWatch Logs in a single dashboard. This solution consolidates, manages, and analyzes log files from various sources, such as such as audit logs for access, configuration changes, and billing events. You can also collect Amazon CloudWatch Logs from multiple accounts and AWS Regions.</p>
Centralized Logging with OpenSearch	<p>Helps organizations collect, ingest, and visualize log data from various sources using Amazon OpenSearch Service. This solution provides a web-based console, which you can use to create log ingestion pipelines with a few clicks.</p>

AWS Solution	Description
DevOps Monitoring Dashboard on AWS	Automates the process of ingesting, analyzing, and visualizing continuous integration/continuous delivery (CI/CD) metrics. These metrics are displayed in Amazon QuickSight dashboards to help DevOps leaders measure the impact of their DevOps initiatives and make data-driven decisions to drive continuous improvement in their development teams.
Application Monitoring with Amazon CloudWatch	Automates the process of setting up Amazon CloudWatch dashboards for your Apache, NGINX, and Puma workloads running on Amazon EC2. This solution uses several features of Amazon CloudWatch and speeds up the getting started experience.

Shared Services accounts

A Shared Services account is an AWS account created and dedicated to hosting and managing centralized IT services and resources that are shared across multiple other AWS accounts within an AWS Organization. The primary purpose of a Shared Services account is to consolidate similar shared services to give a single access point to manage, interface and consume. You may create multiple Shared Service accounts depending on your need to securely isolate the functionality of the grouped services in the account.

Note

AWS account workload isolation is a best practice for enhancing security and operational efficiency in cloud environments. It involves grouping AWS resources and workloads into separate AWS accounts based on their functionality and security requirements. A Shared Service account should contain resources and workloads that can be grouped together in order to ensure security, compliance, and operational separation of duties.

Recommended AWS Organization Integrated Service Delegation

AWS service	Implementation Details	Control Tower Enabled
AWS Service Catalog	Create and manage catalogs of IT services that are approved for use on AWS.	No
AWS Compute Optimizer	AWS Compute Optimizer can be delegated to one AWS account in your AWS Organization. It is recommended to deploy to a Shared Services account or the Monitoring account.	No


Additional Services and Functionalities

Common examples of security capabilities that can be centrally accessed and managed using the Shared Services account includes:

- **EC2 Image Builder** - EC2 Image Builder integrates with AWS Resource Access Manager (AWS RAM) to allow you to share certain resources with any AWS account or through AWS Organizations.

Example structure

The following example structure represents the recommended Infrastructure OU at a basic level. For general guidance on separating production and non-production workloads, refer to [Organizing workload-oriented OUs](#).



Example structure of Infrastructure OU

Sandbox OU

The Sandbox OU contains accounts in which your builders are generally free to explore and experiment with AWS services and other tools and services subject to your acceptable use policies. These environments are typically disconnected from your internal networks and internal services. Sandbox accounts should not be promoted to any other type of account or environment within the Workloads OU.

Sandbox per builder or team with spend limits

A common practice is to provide a sandbox account to either each builder or each small team of builders, along with cloud spend budgets to ensure that their AWS spending aligns within your policies. In more advanced scenarios, you might provide your builders and teams with the option to have multiple sandbox accounts so that they can experiment more freely with configurations that entail use of multiple accounts (for example, experimenting with cross-account IAM roles).

There is a maximum number of accounts in an organization. If you have thousands of builders and expect to allocate a sandbox environment for each builder, you might encounter the maximum quota for accounts. Refer to [Quotas for AWS Organizations](#) for more details on the maximum number of accounts in an organization.

In cases where you need more than several thousand sandbox accounts, you might consider either creating one or more separate organizations to contain the sandbox accounts or establishing a process to recycle sandboxes when they are no longer in use.

Temporary resources and environments

Unlike more persistent development environments, it's common to set expectations with your builders that the resources they create in sandbox environments are temporary in nature. As a cost control measure and to reinforce the temporary nature of sandbox resources, you can put automated procedures in place to periodically purge the resources created in these environments. As a further measure to reduce costs, you can use the [Instance Scheduler on AWS](#) solution to automate the starting and stopping of Amazon Elastic Compute Cloud (Amazon EC2) and Amazon Relational Database Service (Amazon RDS) instances based on provided schedules

Wide-ranging access

Wide-ranging access is typically provided in sandbox-oriented accounts including administrative-like access within each account, full access to most AWS services, and possibly outbound and inbound access to the internet. Access to the internet might be required to connect to AWS service APIs, download externally accessible software packages, and integrate with publicly available services.

No access to corporate resources and non-public data

Given the extent of access provided in sandbox environments, businesses typically employ a combination of guardrails and internal usage agreements to limit builders from accessing

corporate resources and data from their sandbox accounts. Use of non-public data and intellectual property, including proprietary source code and binaries, is typically not allowed in sandbox environments.

Sandbox and development environments

Due to use of non-public data and the more formal nature of the work being performed in development environments, we recommend that you make a high-level distinction between sandbox environments and development environments. For example, in development environments your teams are performing more formal experiments, day-to-day development, and early testing work that requires access to your intellectual property and to enterprise services, such as source code and artifact management.

For more information about potential distinctions between your sandbox, development, and other environments, refer to the following appendices:

- [Appendix B – Worksheet for mapping workload environment purposes to hosting environment types](#)
- [Appendix C – Worksheet for identifying attributes of workload hosting environments](#)

Additional Services and Functionalities

Common examples of monitoring capabilities that can be centrally accessed and managed using a Shared Services account include:

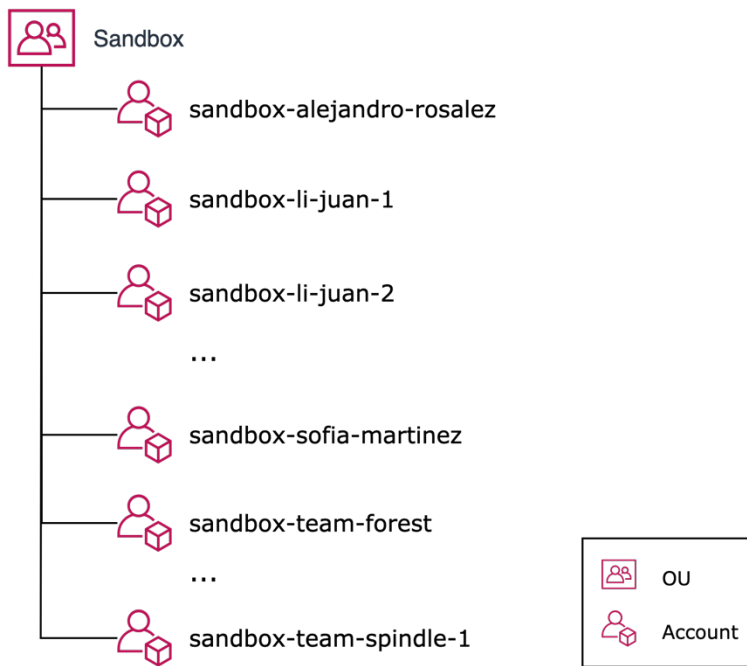
- [Instance Scheduler on AWS](#) solution to automate the starting and stopping of Amazon Elastic Compute Cloud (Amazon EC2) and Amazon Relational Database Service (Amazon RDS) instances based on provided schedules.

Example structures

Sandbox per builder or team

In the following example, sandbox accounts are represented for individual builders and teams. One user has two sandbox accounts so that they can perform experiments that require multiple accounts.

In support of hackathons and other events, you might also find value in creating transient sandbox accounts for temporary teams of people.



Example structure of Sandbox OU

Workloads OU


The Workloads OU is intended to house most of your business-specific workloads including both production and non-production environments. These workloads can be a mix of commercial off-the-shelf (COTS) applications and your own internally developed custom applications and data services.

Workloads in this OU often include shared application and data services that are used by other workloads.

Example structure

The following example represents a basic structure in which sets of workloads owned by diverse business units or teams reside in two child OUs: *Prod* and *Test*. In this example, a common governance and operating model applies across those areas. The *data-lake-prod* account shown in this example contains data services that are shared with other production workloads and accounts.

For general guidance on separating production and non-production workloads and resources, refer to [Organizing workload-oriented OUs](#).



Example structure of Workloads OU

Policy Staging OU

The Policy Staging OU is intended to help teams that manage overall policies for your AWS environment to safely test potentially broadly impacting policy changes before applying them to the intended OUs or accounts. For example, SCPs and tag policies should be tested prior to applying them to the intended OUs or accounts. Similarly, broadly applicable account baseline IAM roles and policies should also be tested using the Policy Staging OU

Workload-specific policies

Development and testing of workload-specific IAM roles and policies do not need to use the Policy Staging OU. Rather, workload owning teams typically develop and test these resources alongside other workload-specific resources in development and test accounts within your Security, Infrastructure, and Workloads OUs.

Recommended testing and promotion workflow

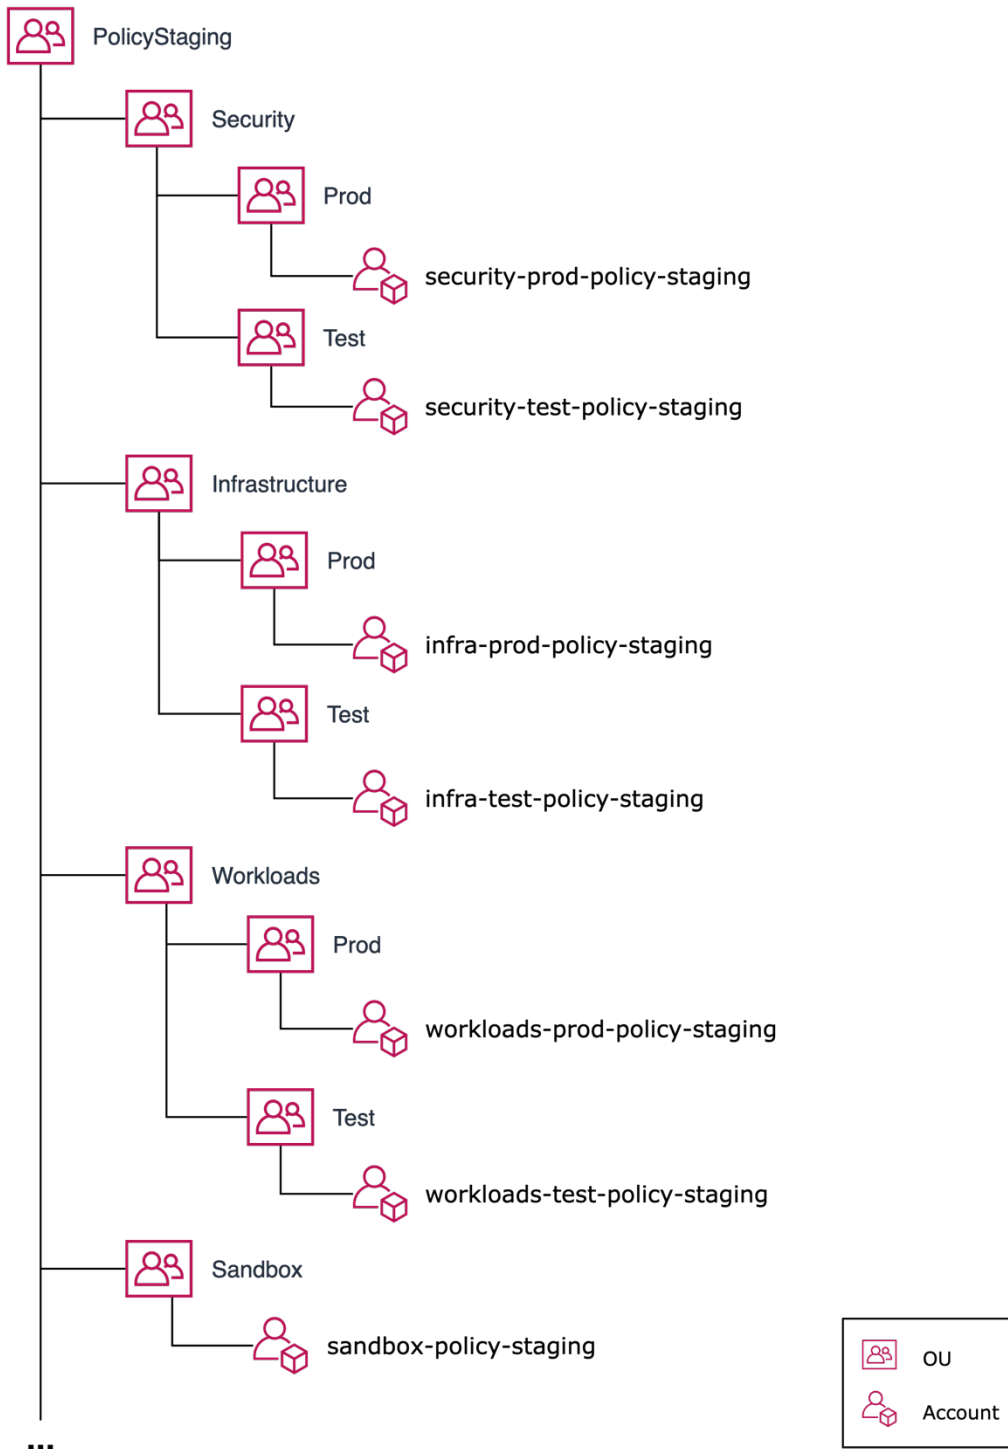
Once you have tested changes in the Policy Staging OU, we recommend that you temporarily associate the policy changes to a single account in the intended OU. If the changes are ultimately targeting an OU, apply the changes to the intended OU and remove the changes from the account only after you have validated that the changes are working as intended.

This approach enables you to validate the changes in production before more broadly applying them.

Example structure

In this example, a set of child OUs mirrors an overall OU structure. At least one test account is included under each child OU.

In support of testing SCPs and tag policies that are intended to be applied at the OU level, your teams should first apply them to one of the test child OUs. SCPs and tag policies that are applied to a specific account require creation of a test account under the appropriate test child OU.



Example structure of Policy Staging OU

Suspended OU

The Suspended OU is used as a temporary holding area for accounts that are required to have their use suspended.

Moving an account to this OU does not automatically change the overall status of the account. For example, in cases where you intend to permanently stop using an account, you would follow the [Closing an account](#) process to permanently close the account.

Examples of using the Suspended OU include:

- A person's sandbox account is no longer needed due to the departure of the person from the company.
- A workload account is no longer needed due to the resources having been either retired or migrated to another account.

Constraining activity in suspended accounts

You can use service control policies (SCPs) to inhibit users other than your security and cloud platform teams from using AWS APIs in each account. Additionally, you can remove application-level access so that users can no longer access and manage application resources for each suspended account.

To reduce risk and potentially minimize costs, you can also stop any running resources and applications in each suspended account.

Resources should not be deleted from a suspended account unless the account is intended to be closed.

Tagging suspended accounts

Because you might use the Suspended OU for a variety of use cases, we recommend that you apply tags to each account to record the reason for moving the account and the OU from where the account originated. Each process that you establish to support your suspension use cases can use the tag to automatically process the suspended accounts. This tag can also aid in your internal tracing and auditing of an account's lifecycle.

Closing suspended accounts

If an account is moved to this OU prior to the start of the closure process, you can implement a policy and process to automatically start the account closing process a certain number of days after an account has been moved to this OU.

Once the account closure process has been completed, the account is no longer visible in your organization.

Individual Business Users OU

The Individual Business Users OU houses accounts for individual business users and teams who need access to directly manage AWS resources outside the context of resources managed within your Workloads OU.

In some cases, you can consider a small number of AWS resources as something other than a workload. For example, a business team might require write access to Amazon S3 buckets to share marketing videos and data with a business partner. In these cases, you might choose to manage these resources in accounts within the individual business users OU rather than in accounts in the Workloads OU.

Controls

We recommend that you apply a combination of SCPs and IAM permissions to this OU and authorized users. This ensures that only those AWS services, resources, and actions needed are granted. Depending on the nature of the use cases, you can apply guardrails to individual accounts in this OU.

Services that do not require direct user access to accounts

The individual business users OU does not apply when users can authenticate and be authorized to interact with applications and services without requiring direct access to an account. For example, business users often need access to Amazon QuickSight for business intelligence (BI) purposes. Assuming that you consider your QuickSight-based BI capability a workload, you can position the QuickSight resources and data in a workloads account in the Workloads OU. In this case, BI users are authorized to access the QuickSight service directly without needing access at the account level.

Exceptions OU

The Exceptions OU houses accounts that require an exception to the security policies that are applied to your Workloads OU. Normally, there should be a minimal number of accounts, if any, in this OU.

Service control policies and scrutiny

Given the unique nature of the exceptions, SCPs are typically applied at the account level in this OU. Due to the customized security controls that apply to these accounts, owners of these accounts can expect to experience greater scrutiny from security monitoring systems.

Consider Workloads OU as an alternative

If you observe a pattern in which multiple accounts require the same set of exceptions, we recommend that you examine either your existing workloads policies or an extended form of the Workloads OU structure and house the accounts under the Workloads OU. You can introduce another level of OU under the Workloads OU to represent a common set of security policies and/or operational processes that can be applied to multiple workload environments. For more information, refer to [Organizing workload-oriented OUs](#).

Deployments OU

The Deployments OU contains resources and workloads that support how you build, validate, promote, and release changes to your workloads. You might already be using continuous integration/continuous delivery (CI/CD) capabilities to help manage and automate how changes to various types of source code are processed.

Using CI/CD capabilities residing outside of your AWS environment

If you already use on-premises and/or managed CI/CD and related capabilities that reside outside of your AWS environment and you do not expect to use and/or manage CI/CD services within your AWS environment in the near term, you might not immediately need to establish the Deployments OU and an associated set of CI/CD oriented accounts.

In this scenario, you must work through any access and potential network connectivity dependencies between your CI/CD capabilities residing outside of your AWS environment and your workloads environments in AWS.

Separating CI/CD management capabilities from workloads

If you intend to deploy and/or manage your own CI/CD capabilities in AWS or use AWS managed CI/CD services, we recommend that you use a set of production deployment accounts within the Deployments OU to house the CI/CD management capabilities.

Reasons for separating your CI/CD management capabilities from your workload environments include:

- **Critical roles played by CI/CD capabilities** — Your CI/CD capabilities are responsible for orchestrating quality validation, security compliance checks, building and publishing production candidate artifacts, promoting artifacts, and ultimately triggering release of artifacts to production environment.

Given the critical nature of these roles, it's important that you can apply appropriate policies and operational practices to your CI/CD capabilities that are different than those applied to your workload environments.

For example, your CI jobs and CD pipelines typically need write access to publish and promote candidate artifacts to an artifact management service. However, your production workload environments should only require read access to artifact management services in order to obtain the already built and promoted artifacts.

- **CD pipelines affect non-production and production workload environments** – When CD pipelines orchestrate the validation of changes and ultimately trigger the release of changes to production, the pipelines often need to access workloads residing in both non-production test and production workload environments. For example, if you manage your CI/CD capabilities in your production workload environments, then you must allow the production workload environments to access your non-production environments. By centralizing your CI/CD capabilities in your CI/CD accounts, you can avoid enabling your production workload environments access to non-production environments.
- **CI/CD capabilities depend on unique tooling** – Your CI/CD management capabilities, CI jobs, and CD pipelines often depends on tools that are different from those required to run and operate your workloads. Limiting the use of these tools to your CI/CD accounts can help you reduce the complexity and attack surface of your workload environments.

Running CI jobs and CD build stages in deployment accounts

Because CI jobs and CD pipeline build stages are responsible for generating formal candidate artifacts, we recommend that you perform these activities in a production environment. Rather than perform these activities in your production workload environments, we recommend that you run them in your production CI/CD accounts.

Aligning CI/CD accounts with groups of workloads

We recommend that you define CI/CD accounts in the Deployments OU that are aligned with how you group related workloads in your workload-oriented OUs. By doing so, you can more easily align the security policies and operational requirements of each group of workloads with their companion CI/CD account.

This approach helps limit the scope of impact of an issue in a CI/CD account to a single workload or group of workloads. Any access issues and resource conflicts that arise in one CI/CD account will most likely impact only the associated group of workloads and the accounts in which the workloads reside.

Transitional OU

The Transitional OU is intended as a temporary holding area for existing accounts and workloads that you move to your organization before formally integrating them into the more standardized areas of your AWS environment structure.

Common scenarios for moving accounts into your organization

Common scenarios for moving accounts into your organization include:

- Acquisition of a company that is already using AWS and has a set of accounts
- Existence of your own accounts that were created before you established your newer AWS environment structure
- Movement of accounts that have previously been managed by a third party
- Divestment of specific workload to be migrated out of your AWS Organization

Considerations for moving accounts into your organization

If you plan to move an account from an existing organization, you must first remove the account from the organization. For more information, refer to [Removing a member account from your organization](#). Once an account is removed from an organization, it is referred to as a standalone account.

Moving a standalone account that does not have dependencies on other accounts is a straightforward process. In this case, there's generally no need to migrate or modify the existing workloads in the account to be moved. For more information, refer to [Inviting an account to join your organization](#).

If the standalone account to be moved has dependencies on other accounts, then you should evaluate those dependencies to determine if they should be addressed before moving the account.

In your target organization, we recommend that you review SCPs in the organization's root to ensure that those SCPs won't adversely impact the accounts to be moved.

If you're moving a set of related accounts to your organization, you can create a child OU under the Transitional OU for the related set of accounts.

After moving accounts

Over time, as you better understand the direction for these accounts and the workloads contained in them, you can either move the accounts to your Workloads OU as is, invest in migrating the workloads to other accounts, or decommission either the workloads or accounts.

Business Continuity OU

Note

The Business Continuity OU is an advanced use-case topic where your AWS Organization requires data isolation and data residency controls based on unique workloads requirements. In general, most cross account disaster recovery strategies can be implemented through using the Backup account within the Infrastructure OU.

The Business Continuity OU is intended to help teams implement a cross-account disaster recovery strategy. The data is as close to air-gapped as possible and the OU has no workload resources. This

creates a secure data bunker to help protect your organization and allow for recovery from severe disasters like ransomware. The secure data bunker should only be accessed when the disaster recovery data for a workload is unavailable, untrustworthy, or destroyed.

The Business Continuity OU does not replace normal disaster recovery plans of your workloads. It's an additional layer of protection that is meant to enhance the resiliency of your organization. General recommendations for disaster recovery for workloads can be found in [Disaster Recovery of Workloads on AWS: Recovery in the Cloud](#).

For organizations that have data residency requirements and are in a geographical area that has only a single AWS Region available, using [AWS Outposts](#) can assist in maintaining compliance. The blog [Ensure Workload Resiliency and Comply with Data Residency Requirements with AWS Outposts](#)

Controls

For the Business Continuity OU to be a secure data bunker, access should be heavily restricted to prevent the data from being compromised. Ideally, users with access to the data within the Business Continuity OU should not have access to the regular environment and users with access to the regular environment should not have access to the Business Continuity OU. Apply a combination of SCPs and IAM permissions to this OU and authorized users to ensure that only those AWS services, resources, and actions needed are granted.

Additional consideration


- Place restrictions on the Backup Administrator role so that backup policies for the Business Continuity OU are not altered.
- Implement monitoring notifications to confirm that backups have not been interrupted. Refer to the documentation on [AWS Backup monitoring](#).
- Require all Backup Vaults use [AWS Backup Vault Lock](#) in Compliance mode with a minimum retention of 14 days or more.
- Audit backups regularly to ensure compliance of your backup policies. Refer to the documentation on [AWS Backup Audit Manager](#).

Example structures

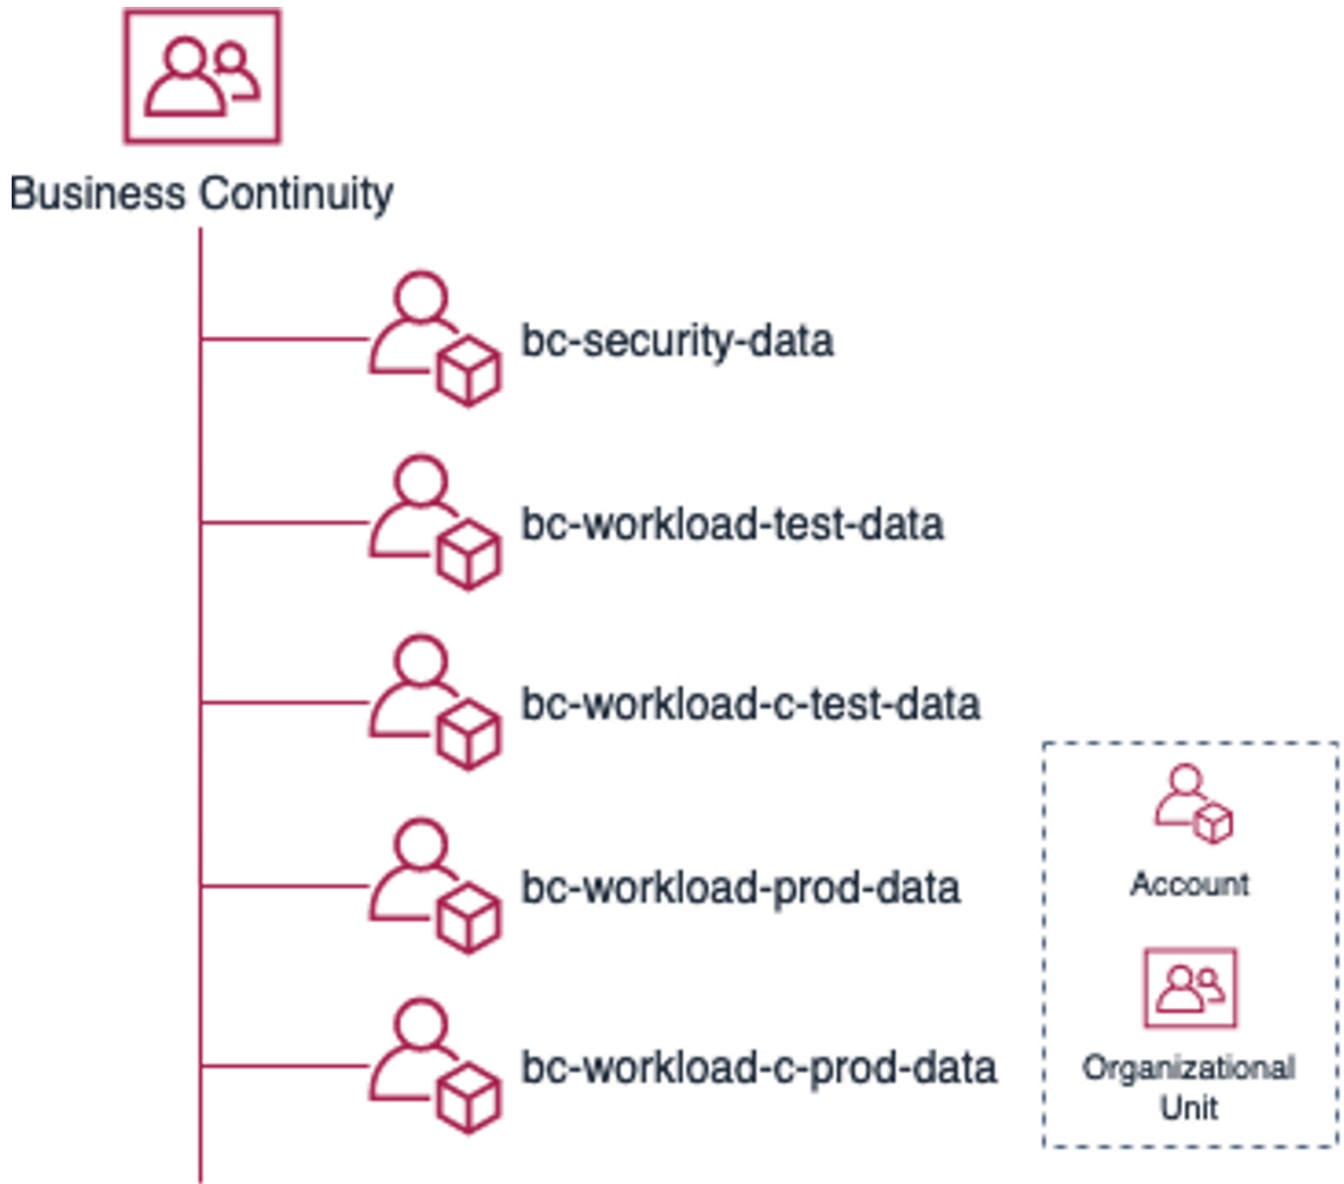
In this example, a company, Rainbows, has a starter AWS environment that follows the [production starter organization](#) guidance. Rainbows has three workloads called A, B, and C. They have separate

testing and production accounts for each. The data for workloads A and B are non-sensitive and unregulated. However, the data for workload C is highly sensitive and regulated and must be isolated from other workload data.

Following is an example of the Business Continuity OU for the Rainbows company. Because the data for workloads A and B is non-sensitive and unregulated, you can keep it in the same business continuity (bc) account, `bc-workload-test-data` and `bc-workload-prod-data`. However, for workload C, the business continuity data is isolated in separate accounts, `bc-workload-c-test-data` and `bc-workload-c-prod-data`, because the workload is highly sensitive and regulated.

 **Note**

Some companies choose to keep each workload's data separated in individual accounts for an enhanced security posture, regardless of a regulatory or compliance need.



Business Continuity OU example structure

Organizing workload-oriented OUs

The recommended [Security OU and accounts](#), [Infrastructure OU and accounts](#), [Workloads OU](#), and [Deployments OU](#) are top-level OUs that contain workloads. This section outlines considerations for organizing these workload-oriented OUs.

Topics


- [Workloads and environments](#)
- [Production and non-production workload environments](#)
- [Workload dependencies across environments](#)
- [OU structure for non-production environments](#)
- [Extended workload-oriented OU structure](#)

Workloads and environments

This section defines basic terms and concepts related to workloads. Becoming aware of these concepts helps you understand our recommendations for organizing your workload-oriented OUs.

Workloads

Many of your top-level OUs will house collections of applications, cloud resources, and data in the form of workloads. A *workload* is a discrete collection of components and data that you manage. A workload can be a commercial off-the-shelf (COTS) application or your own custom application and data service.



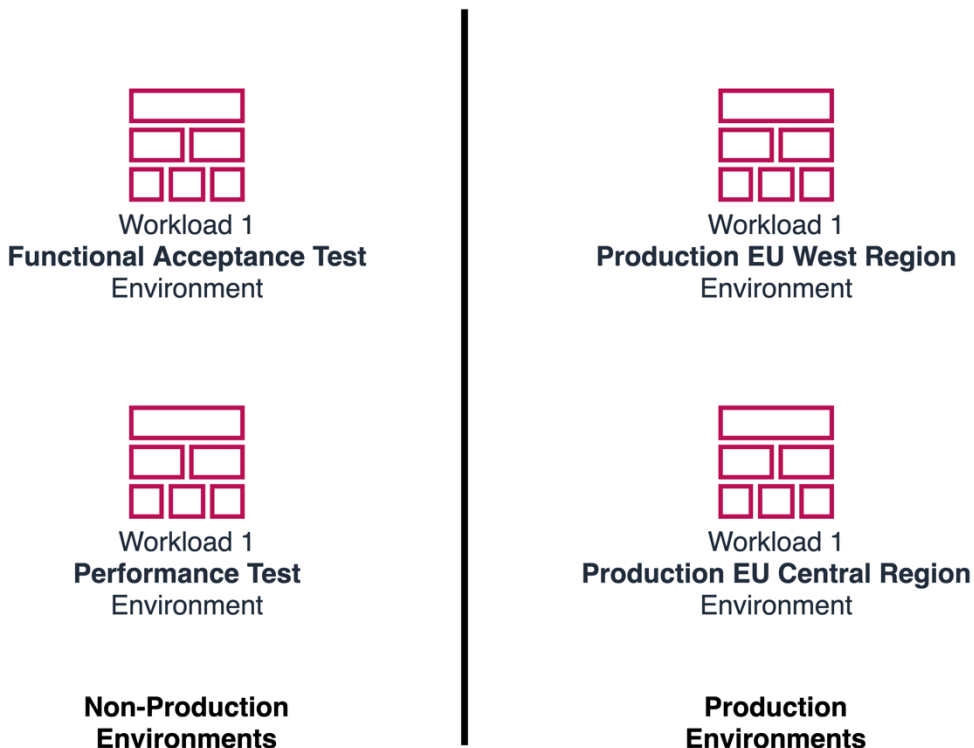
Composition of a workload

Workload environments

For a given type of workload, you typically have multiple instances. This setup means that you can experiment, develop, and test changes to the workload before you promote and deploy those changes to the production instances of the workload. A given instance of a workload is a *workload environment*.

Whether your workload is a COTS application, a custom application, a custom data service, or a foundational security or infrastructure capability, you often need separate non-production workload environments to support your software development lifecycle (SDLC) processes. You can have multiple SDLC processes depending on the diversity of your workload portfolio and your company organization.

The following example shows multiple environments of a workload across non-production test and production workload environments.



Example of multiple environments of a workload

With COTS applications, you might not perform custom development, apart from implementing custom integrations with your own systems. However, you can experiment with and formally test new versions of the COTS applications in non-production environments before deploying them to production.

For detailed examples of common purposes of workload environments in relation to an SDLC, refer to [Appendix B – Worksheet: Mapping workload environment purposes to hosting environment types](#).

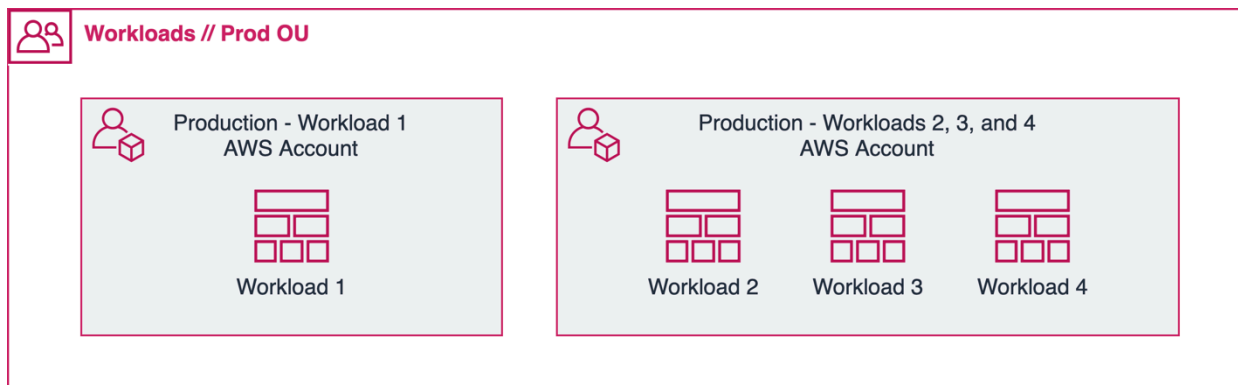
Workload accounts

In your on-premises context, you can refer to the places where your workloads reside as *hosting environments*. For example, you might have a dedicated production hosting environment in which your production workload environments reside.

In AWS, each of your workload environments is typically contained in an account, where each account is similar to a distinct hosting environment.

Depending on how you choose to scope your workload accounts, you might have a single workload environment per account. Or, you might have multiple workload environments and perhaps multiple workload types in the same workloads account.

The following diagram shows two degrees of scoping production workload accounts. In one example, a workload account is dedicated to a single workload environment. In the other example, multiple workload types reside in a single production workload account.



Example workload accounts with different degrees of scoping

Production and non-production workload environments

We recommend that you isolate production workload environments and data in production accounts housed within production OUs, under your top-level workload-oriented OUs. Apart from production OUs, we recommend that you define one or more non-production OUs that contain accounts and workload environments that are used to develop and test workloads.

Workload dependencies across environments

When you consider the structure of your workload-oriented OUs, you should decide on the extent to which you expect access between production and non-production environments.

Production environments accessing non-production

Generally, workloads deployed to your production environments should not depend on workloads contained in your non-production environments.

Non-production environments accessing dependencies

In non-production environments, it is common for workloads to depend on stable shared application, data, and infrastructure services. Where feasible, we recommend that these shared services be non-production test instances. These non-production test instances should use test data so that your non-production workloads do not depend on access to your production environments and data.

For example, you can configure workloads in a non-production test environment that depend on integrating with a data service to use a stable, shared test instance of the service that is populated with test data.

However, in some cases non-production environments might need access to production shared services. For example, it's typical for non-production development and test environments to require read-only access to shared source code and artifact management services. Providing access to these shared services enables you to deploy candidate and promoted changes and artifacts to your non-production environments in support of development and testing activities.


OU structure for non-production environments

You can use OUs to organize your non-production environments in a couple ways.

Option A: Common guardrails across non-production environments

When non-production workloads require the same set of overall access policies or benefit from being operationally managed together, you can define a single *NonProd* OU to contain all the accounts that support non-production forms of your workloads.

The following example shows the *Workloads* OU where a *Prod* child OU contains production accounts and workloads, and a *NonProd* child OU combines both development and test accounts and workloads.




Example Workloads OU with common policies across a NonProd child OU

Option B: Different guardrails across non-production environments

Sometimes your process for developing and testing changes involves workload environments that have fundamentally different access policies or ways in which you manage and apply foundational resources. In these cases, it makes sense to create distinct OUs to support these diverse requirements.

For example, you want to support development environments that provide teams with more freedom to experiment, iterate, and develop largely on their own (rather than more formally managed and controlled production-like test environments). In this case, overall access policies and management of baseline resources for the development environments is significantly different than those used to support test environments. It makes sense for you to create a distinct OU for development work and another OU for your test workloads.

The following example represents a simple form of this structure where *Test* and *Dev* OUs reside adjacent to the recommended *Prod* OU.



Example Workloads OU with different policies for Test and Dev child OUs

The preceding example shows two different approaches to scoping development environment accounts. One approach is where development environments are aligned with the same groupings of workloads as used in test and production OUs. The other approach is one in which development environments are aligned based on teams.

Worksheets to help decide on workload-oriented OUs

The following appendices include a set of worksheets and example considerations for identifying your overall types of workload environments and supporting OUs:

- [Appendix B – Worksheet for mapping workload environment purposes to hosting environment types](#)
- [Appendix C – Worksheet for identifying attributes of workload hosting environments](#)

Appendix B helps you identify the overall types of work you perform from design through production and helps you identify the corresponding workload environments in which you expect to perform work and house workloads.

Appendix C helps you further refine the overall types of workload environments by identifying key distinguishing access and management attributes of each overall type of workload environment.

By understanding commonalities of, and contrasts between, your overall types of workload environments, you can make an informed decision about the set of child OUs that can best support your workload-oriented OUs.

Extended workload-oriented OU structure

An extended form of the workload-oriented OU structure can be used to support cases in which you need to either organize workloads for visibility and management purposes or apply different security and operational policies to either a workload or group of related workloads.


When workloads have diverse security and operational policy requirements, you cannot effectively manage guardrails and other controls at the level of the workload-oriented OU. By adding child OUs to a workload-oriented OU, you can group related workloads in the same child OU. You can then apply distinct security and operational policies to the child OUs.

For example, a workload or a group of related workloads might benefit from having a distinct allow list of AWS services that is implemented via a service control policy (SCP). This policy might be different than the requirements associated with other workloads. Rather than applying the SCP to each of the related workload accounts, it is recommended that you apply the SCP to an OU that groups the related accounts.

Grouping related workloads


When you have groups of related workloads that require the same overall set of security and operational policies, you can create a child OU for each group of workloads.

For example, if you manage a series of database services that are shared across your organizations and have common security and operational policy requirements, you might find value in grouping those data services under a common child OU.



Group of workloads with distinct policy requirements

The following example represents a shared backup capability you can provide across your AWS environment. If this capability requires a set of security and operational policies that are distinct from other infrastructure workloads, then you can allocate a distinct OU for this workload.

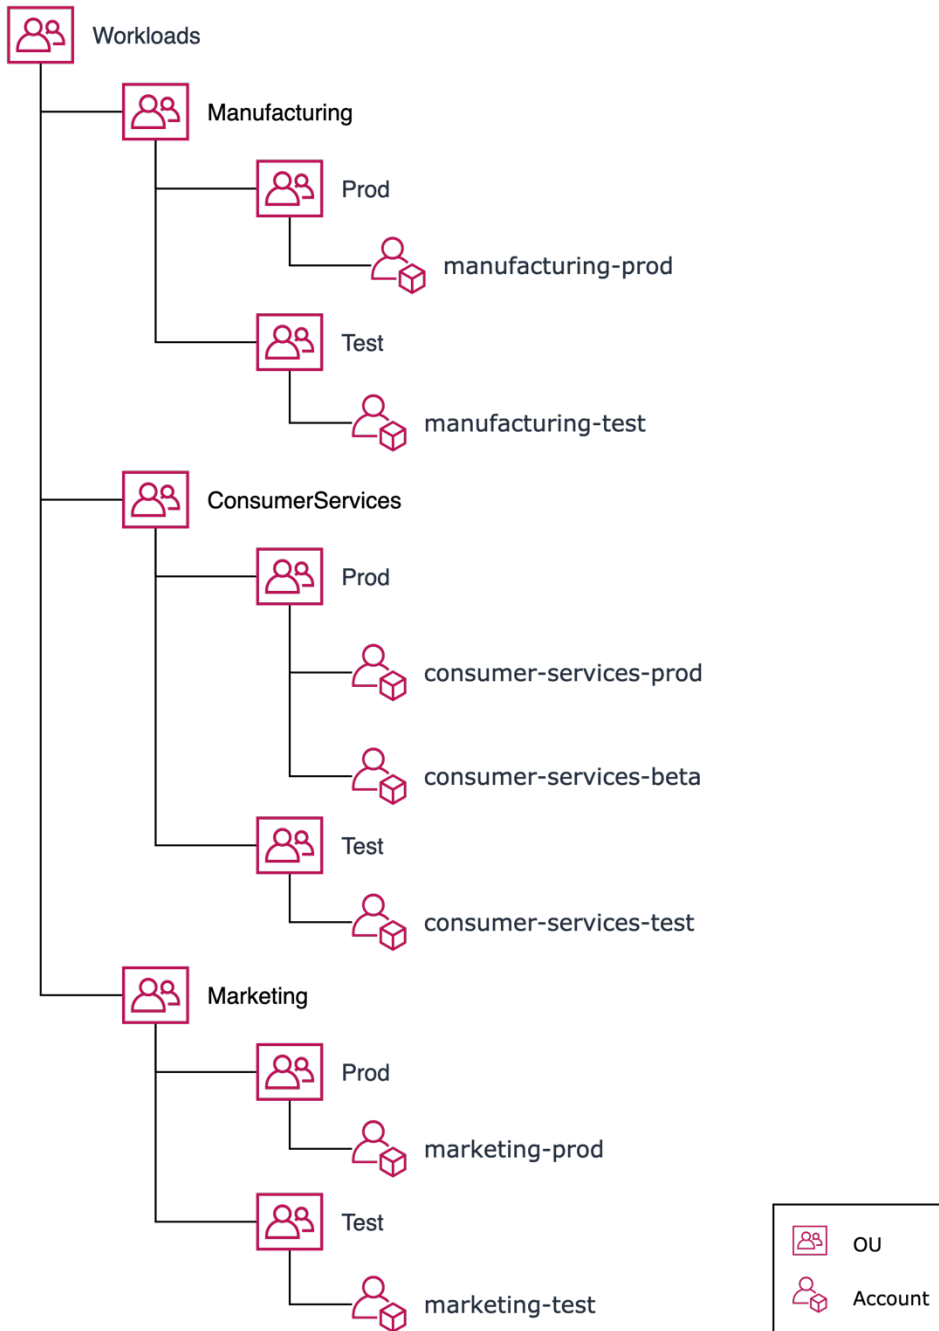


Single workload with distinct policy requirements

Separating business units with significantly different policies

If you have largely autonomous business units (BUs) that manage workloads in your common AWS organization and the BUs have significantly different security and operational policies, you can create a child OU under your Workloads OU for each BU.

In the following example, each BU is provided with its own OU so that different SCPs and/or operational policies can be applied independently from the other OUs.



Example business unit separation

Patterns for organizing your AWS accounts

This section introduces a series of example AWS account structures based on the principles and best practices contained in this document.

It's common for customers to start with a basic structure and incrementally expand it as their needs evolve and their experience with AWS grows. Accordingly, the examples start simple and progressively grow to represent this typical evolution.

More detailed descriptions of the OUs represented in these examples are addressed in [Recommended OUs and accounts](#) and [Organizing workload-oriented OUs](#).

Single AWS account

If you have experimented with AWS, you might have used a single AWS account in which to perform some of your initial non-production work. You might have even started to manage some of your early workloads as production resources in the single AWS account.

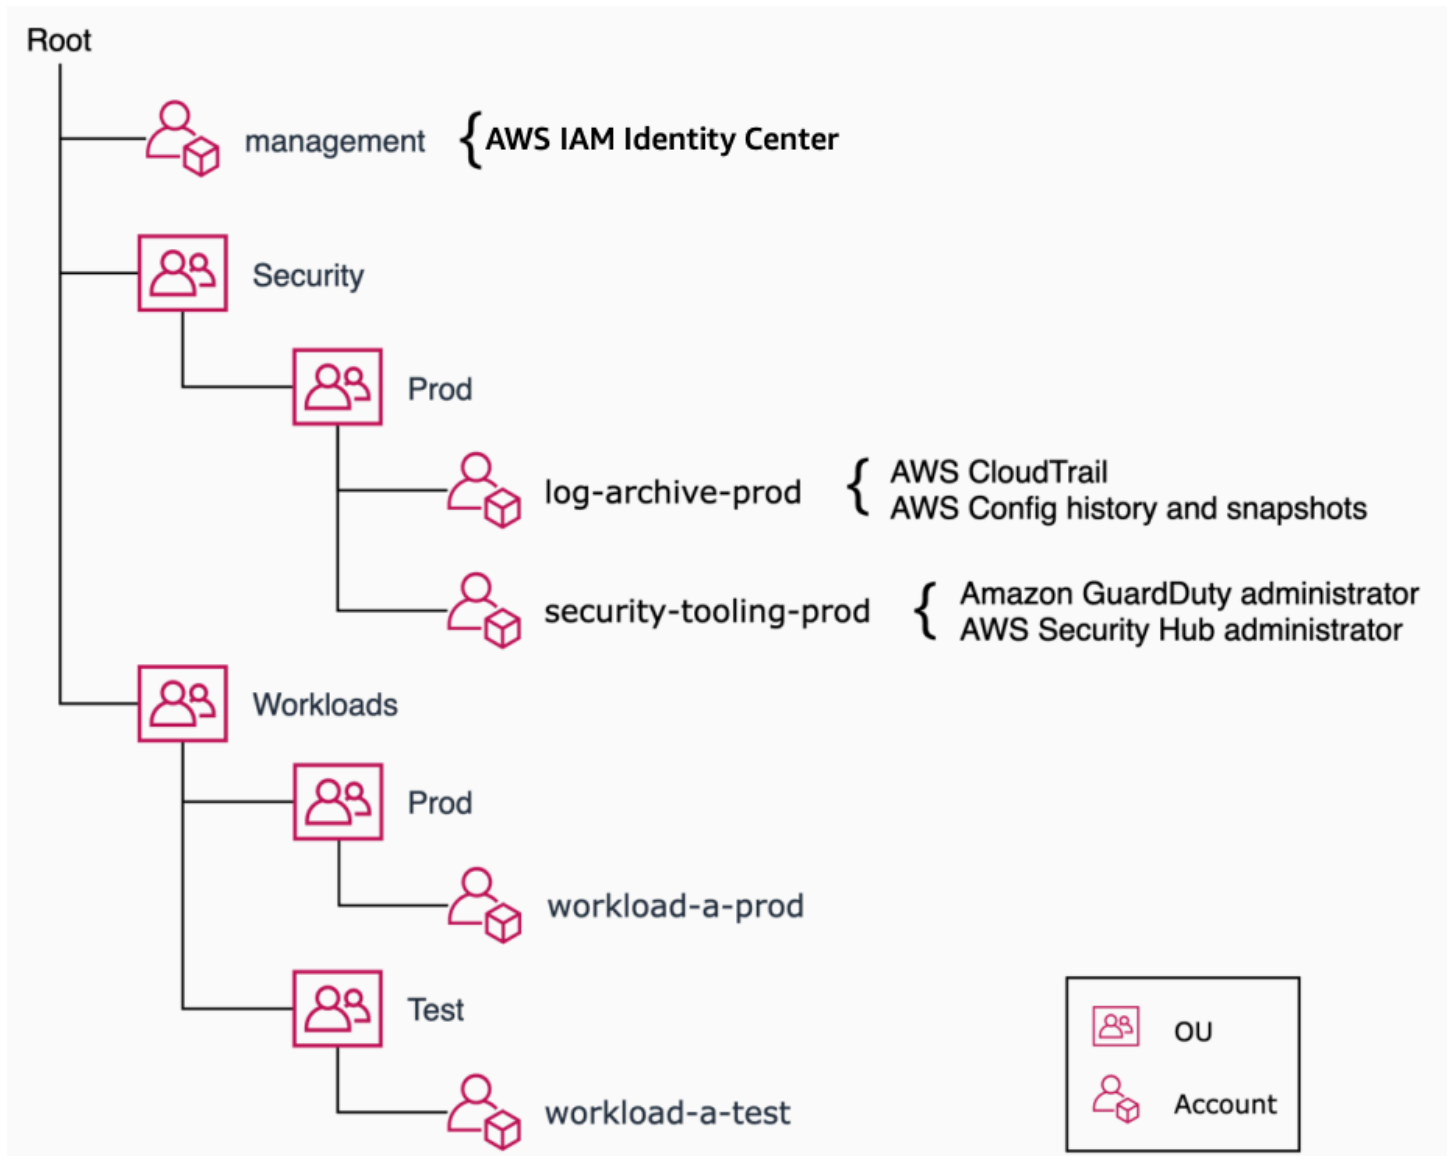
If you have a single account today and either have production workloads in place or are considering deploying production workloads, we recommend that you transition to the use of multiple accounts so that you can gain the [Benefits of using multiple AWS accounts](#). Refer to the [Transitional OU](#) for information on the considerations for moving accounts into a new overall AWS environment.

Production starter organization

This pattern represents a minimal starter environment in which the primary focus is on supporting a workload in a production environment.

For example, you might have an Amazon S3-backed on-premises backup solution that you simply need to test and deploy to production. Similarly, you might have a static web site that depends on Amazon CloudFront as a content delivery network (CDN) and uses a private bucket in Amazon S3 to manage the web content.

In these scenarios, you might not need sandbox and development environments. The following figure shows an example of this type of minimal starter production environment.



Example production starter organization

In this example, the organization's management account uses [AWS IAM Identity Center](#) to help provide your human users with federated access to the AWS accounts in your organization.

The [Security OU and accounts](#) contains a *log-archive-prod* account to act as the consolidation point in the organization for log data that is gathered from all of the accounts—not just other production environments—and primarily used by your security, audit, and compliance teams.

The Security OU also contains a *security-tooling-prod* account where you manage recommended security tools and service resources.

Since the capabilities provided in the *log-archive-prod* and *security-tooling-prod* accounts are expected to be of production quality, these accounts are contained in a Prod OU under the Security OU. The *-prod* suffix in these example account names emphasizes the production quality of their resources and workloads. The suffix is not intended to suggest that these accounts and their resources apply only to production accounts.

In future configurations, you can introduce non-production or test OUs and accounts associated with your Security OU. Where it's feasible to test changes inside the same organization, these non-production environments can help you develop and test changes for your production quality capabilities before promoting those changes to your production environments.

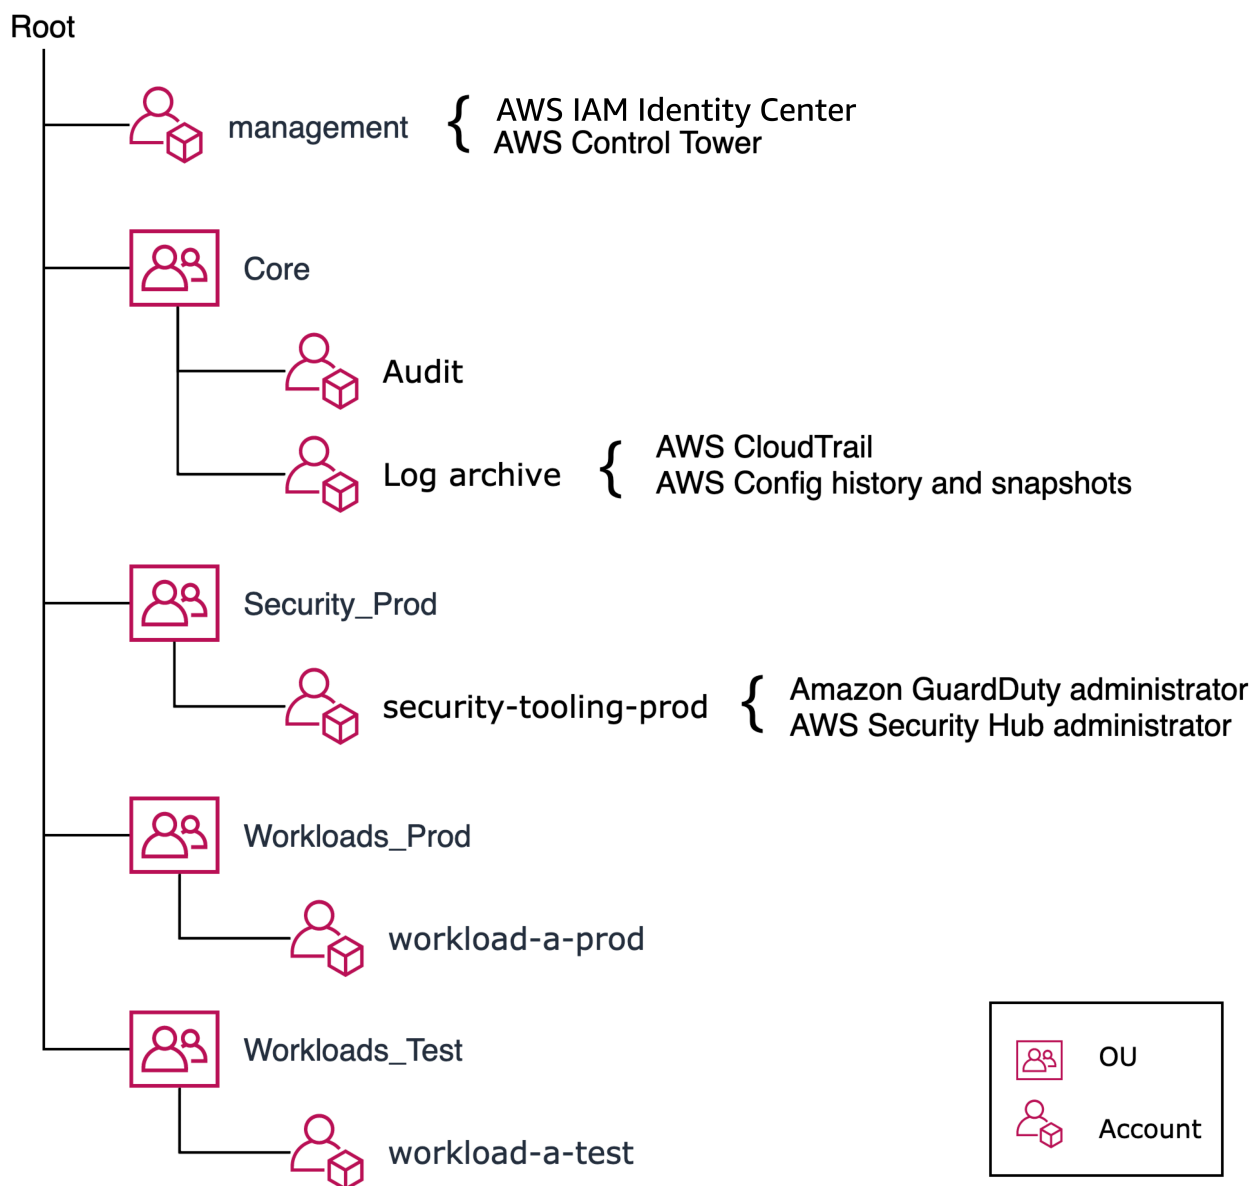
In cases where you cannot easily test certain changes that are foundational to your AWS environment in the same AWS organization, you might benefit from using a separate AWS organization to test such foundational changes. Refer to [Multiple AWS organizations](#) for more information.

A [Workloads OU](#), along with Prod and Test OUs, contains the *workload-a-test* and *workload-a-prod* accounts.

Production starter organization with AWS Control Tower

When you use [AWS Control Tower](#) to establish your AWS environment, it automatically creates the Audit and Log archive accounts under a Core OU. The Log archive account plays the same role as the Log archive account described in [Security OU and accounts](#). The Audit account is intended to provide your security team with cross-account access to other member accounts in your organization.

AWS Control Tower also automatically sets up IAM Identity Center in the organization's management account. The following figure shows this configuration.



Example production starter organization with AWS Control Tower

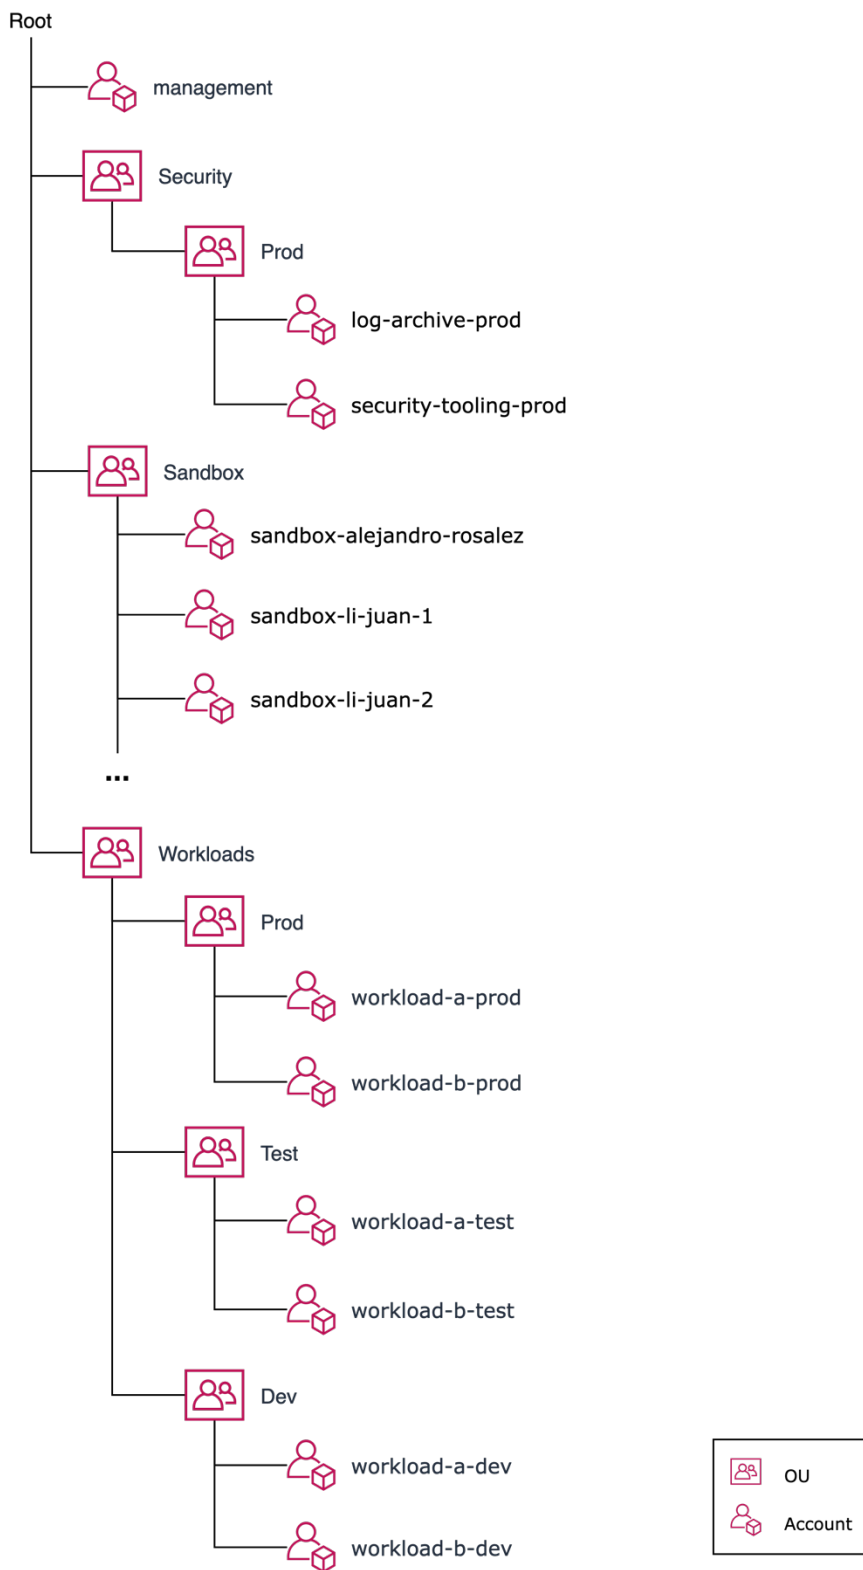
In this example, a [Security OU and accounts](#) is created by the AWS Control Tower Account Factory feature to contain a *security-tooling-prod* account where recommended security tools and service resources are managed.

Two [Workloads OUs](#) house the production and test environments for a workload.

Basic organization

The following example incorporates a security tooling environment for common security services, a second workload, and support for sandbox and development environments. Additions include:

- A [Sandbox OU](#) to contain a series of disconnected sandbox environment accounts.
- An additional workload in the form of *workload-b-prod*, *workload-b-test*, and *workload-b-dev* accounts.
- A Dev OU under the [Workloads OU](#) to contain development environment accounts associated with the workloads.

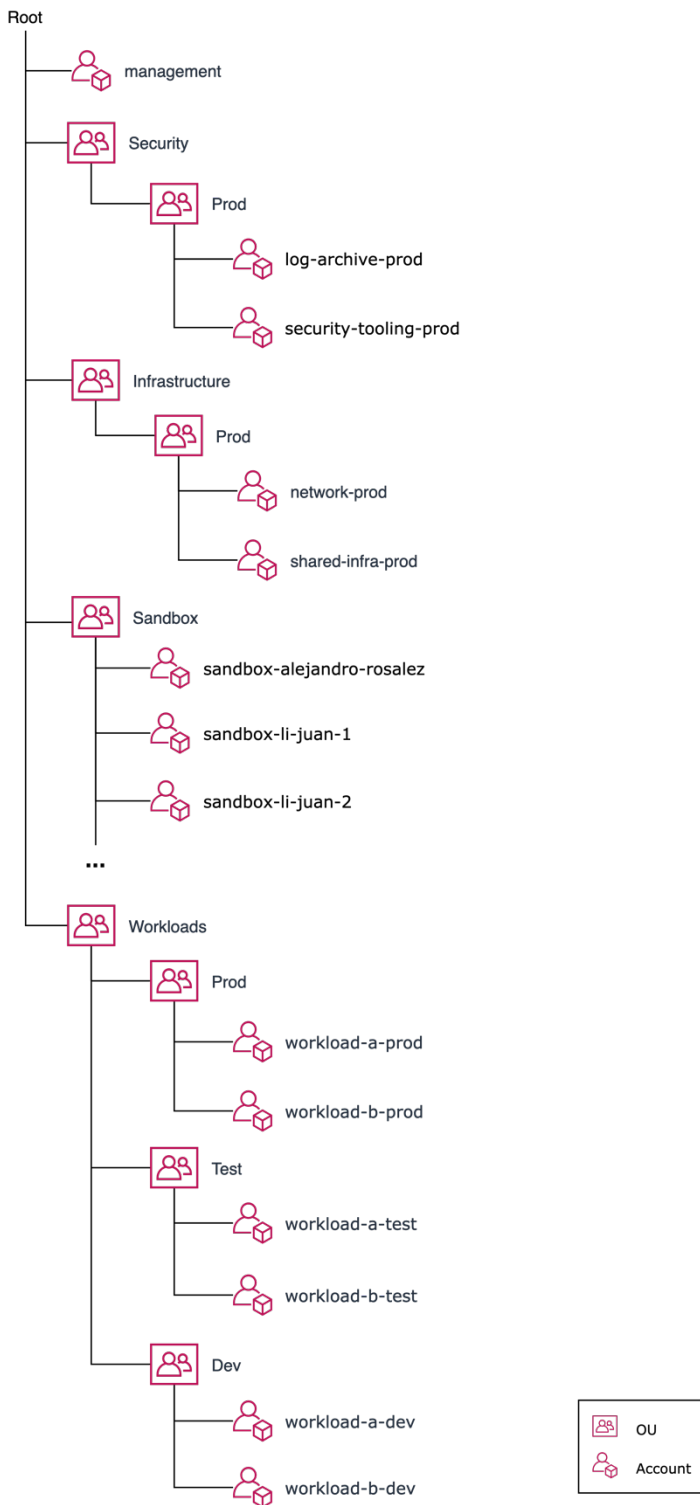


Example basic organization with multiple workloads

Basic organization with infrastructure services

This example includes support for common infrastructure resources. The additions include:

- An [Infrastructure OU and accounts](#) containing a Prod OU.
- A *network-prod* account to contain resources required to connect VPCs in the workload accounts to your on-premises network. For example, AWS Transit Gateway, AWS Site-to-Site VPN, and AWS Direct Connect resources.
- A *shared-infra-prod* account to contain common shared infrastructure services to be used by other accounts. For example, Amazon Route 53 resolver endpoints.




Example basic organization with infrastructure services

Basic organization with CI/CD as a separate function

This example incorporates support for CI/CD resources that are used to validate changes to their respective workloads and automate deployments to the test and production workload environments. The additions include:

- A [Deployments OU](#) and a child Prod OU.
- A set of *workload-a-cicd-prod* and *workload-b-cid-prod* accounts to contain the production quality CI/CD resources for each of the respective workloads.

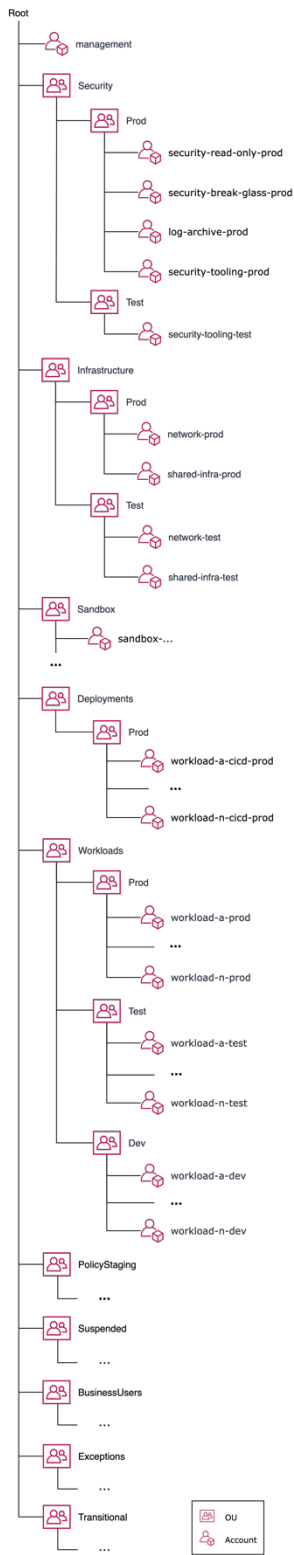


Example basic organization with CI/CD

Advanced organization

This example represents all of the recommended OUs and a greater number and diversity of workloads.

Additionally, Test child OUs are represented under the Security and Infrastructure OUs to support test environments for the security and infrastructure workloads.



Example advanced organization

Implementation

Topics

- [Getting started with your multi-account environment](#)
- [New customers](#)
- [Existing customers](#)
- [Available services](#)

Getting started with your multi-account environment

The earlier sections of this guide covered the benefits of using multiple accounts to organize your AWS environment, this section dives deeper into some implementation options you can use to build your multi-account environment.

Your environment might start with a small number of accounts, and as the number of accounts in your environment increase, you will need automations to manage your cloud environment. These services offer features to help you manage your multi-account environment at scale.

New customers

Managing your own operations

When starting on AWS, you should evaluate your business goals and determine your future operating model, and make a decision on which technology to use when building your multi-account environment. If you are planning on managing your own environment, start by evaluating [AWS Control Tower](#).

AWS Control Tower is a service built based on the guidance included in this paper, making it the preferred solution for new-to-the-cloud customers looking for a service-managed environment. AWS Control Tower offers a simplified experience and automatically deploys your initial environment, helping you manage the multi-account environment efficiently, leveraging other AWS services. For a complete list of these services, refer to the [AWS Control Tower](#) user guide.

Operations managed by AWS or AWS Partners

AWS and Partners can help you operate, update, monitor, or deploy your environment and your workload infrastructure, with [AWS Managed Services](#) and AWS Partners have offerings that can help with your operations.

Hybrid operations (mix of customer and AWS managed)

If your portfolio includes different operating models, you will need to consider a combination of the approaches described above. Proving the flexibility to decide what type of operations you will manage on your environment, and which operations you would like to be overseen. For example, you could manage the account vending process and the environment set up using AWS Control Tower, and use Amazon Managed Services for managing the operations related to moving to the cloud.

Existing customers

If you are not operating your current environment using AWS Organizations, or you are managing an AWS account structure that differs significantly from the strategies laid out in this whitepaper, consider examining your operational model and re-organizing your accounts, and consider operational and organizational changes to align with these strategies.

AWS has seen many challenges for customers because they have grown organically without employing these recommendations, or by implementing before these recommendations were available. These challenges include the following:

- Operational inefficiencies and difficulties, such as unintentionally encountering account quotas when operating many workloads in a small number of accounts due to lack of visibility of the shared quotas
- Use of service roles by humans, and exhaustion in the number of users due to the number of these roles needed to operate many workloads in a small number of accounts
- Unnecessary complexity due to the large number of accounts managed to operate one workload. This complexity can be part of root causes in incidents within security, availability, and other areas
- Lack of automation in the provisioning of accounts, causing large delays in implementing new workloads
- Challenges in managing the email contacts required for notifications about each account (due to lack of automation of account ownership attribution to the workload operators)

- Challenges in ensuring data governance and residency of data where customers have chosen to use a central account and region for storage, and use of this data in an effort to minimize the number of accounts they are operating

Embrace infrastructure as code

One of the foundational operational capabilities that will enable your ability to manage change effectively is to use infrastructure as code technologies and automation to build and deploy your workloads in AWS. This basic capability might be perceived as being part of your [Infrastructure OU](#), but these are foundational capabilities for all OUs.

Examine your operational model

Amazon has “commitment to operational excellence” as one of its four guiding principles of who we are. We have an operational model that enables rapid decisions and autonomy. The way you operate your business is a primary consideration in how you organize your environment. In many companies, this doesn’t change regularly, and changes should always be made with careful consideration of the business impact of the change. Your [Operating Model](#) could be aligned by business unit, regulatory restrictions, area of expertise, and/or organic growth. You might have a need for multiple operational models, for different businesses or business units. The ability for you to meet your goals using AWS is fundamentally affected by your ability to run your operational model. Refer to the [Patterns for organizing your AWS accounts](#) and tailor an approach that meets your needs.

Implement identity management and access controls and other security capabilities

The [Security OU and accounts](#) is a foundational OU. Using a central [identity provider](#) enables you to centrally manage both the identities of the people and services that will need to access resources in your environment, and the permissions associated with that access. To prevent a single source of compromise, consider having more than one identity provider. If you do not currently have an identity provider, consider [AWS IAM Identity Center](#). This also provides a central place to enforce multi-factor authentication and provide [federated](#) access into your environment. Evaluate your capabilities and operational model between your [identity management](#) and [Amazon S3](#) permissions management (also known as access controls).

Next, set up the security capabilities (for example: [centralized log storage](#), encryption and [key management](#), [secrets management](#), and [incident response](#) capabilities) in the appropriate

environment. You can then use these operational capabilities as you reorganize other parts of your operations. Decide if you want to have a security “read only” account to manage and use cross-account access in your environment, or use federated roles directly to access individual accounts in your environment.

The following capabilities can be built as you scale and deploy workloads, or in preparation for their use as you move workloads:

- The additional security capabilities of vulnerability and threat management (as part of security tooling)
- Data de-identification and data isolation (which is normally part of the individual workloads in your [Workloads OU](#).)
- Patching (which might use your [Deployments OU](#) to patch golden machine and container images)
- Forensics

Separate the production workload environment from non-production environment(s)

Next, consider implementing hard isolation between your production environment and your non-production environments by implementing them as separate accounts. This reduces your risk of exposing production data in non-production environments, but might increase the complexity of how you perform deployments. As a result, this change could affect many of your operational capabilities, including:

- [Network connectivity](#), in your [Infrastructure OU and accounts](#) and Workloads OU
- [Network security](#), in your Workloads and Security OUs
- [Application security](#), in your Workloads, Deployments, and Security OUs
- [Tagging](#), in all OUs
- Service onboarding and [cloud financial management](#), in all OUs
- [Rollout/rollback](#) and [change management capabilities](#), in your Deployments and Workloads OUs
- Developer experience and tools, in your [Sandbox OU](#), as well as in your Deployments and Workload OUs

It is not uncommon to have basic connectivity to and from your local networks and the internet shared between your production and non-production workloads. If you have deployment

tooling that is tightly integrated with your current environment, you can integrate the existing implementation as you move the workload, or you can implement the Deployments OU to also deploy to your existing environment while you move them to your new environment.

Networking considerations in a multi-account environment

The number of VPCs a customer operates is usually related to the number of accounts, regulatory requirements (such as the payment card industry (PCI), and compliant/non-PCI compliant), and staged environments (such as prod, dev, and test). With an increasing number of VPCs (account isolated or not), give careful consideration to cross-VPC connectivity management. This is an essential part of the customer's network operation. Additionally, IP address management becomes a key contributing factor to enabling scalability and future growth. You might consider designs that are built around IPv6 adoption that can be driven by the need to scale your network, or by a strategic initiative. Current recommendations for three specific areas in cross-VPC and hybrid connectivity can be grouped by:

- **Network connectivity** — Interconnecting VPCs, on-premises networks at scale, and choosing the right tool for the use case. For intra-AWS connectivity, [VPC peering](#), [AWS Transit Gateway](#), and [AWS PrivateLink](#) are just a [few options](#) which help with different use cases.
- **Network security** — Building [centralized](#) or distributed inspection points for accessing the internet and VPC-to-VPC traffic needs to account for the different options, such as [AWS Network Firewall](#), [AWS Gateway Load Balancer](#), and [AWS Web Application Firewall](#).
- **DNS management** — Resolving DNS within the AWS and hybrid environments at scale involves both right-sizing and choosing the deployment model that best suits the organization's scaling and growth goals. Inbound and outbound Route 53 Resolver endpoints can be centralized or distributed, depending on the operations and management models, and involve different needs across the AWS network environment.

Separate the workload environments to align with the operational organization units

To achieve agility and autonomy, your environment will separate into organizational boundaries. Align your workload environments (both production and non-production) with these organizational boundaries to ensure you enable further agility and autonomy. As you grow and expand, it is common to create new boundaries and move workloads to ensure you continue to have these business capabilities. This could further affect your operational capabilities, such as [backups and disaster recovery](#), [support](#), template management, records management, and sorting and searching

via metadata in your workloads OU. Consider centralizing the management of these capabilities to simplify the realignment. You can use [tag policies](#) as a means to classify the workloads. You can [use AWS Backup to back up to a designated location](#) in your environment and allow you to [centralize protection and monitoring](#) of your backups.

Create the additional organizational units to enable other capabilities

You should consider creating an [Exceptions OU](#). Also, consider creating a [Policy Staging OU](#). As your business might either acquire or divest parts of your business, consider having a [Transitional OU](#) to enable an area to change the policies to align with new requirements. As you deprecate accounts, you will want a [Suspended OU](#) to contain these accounts until you are comfortable having them permanently removed. You should consider creating a [Business Users OU](#) to enable business users and teams who need access to manage AWS resources directly, rather than management by the Workload OU.

Other considerations for implementing these changes

The reorganization will require movement of accounts, or migrations of workloads between accounts if you have deployed workloads in accounts that don't align with your desired operational model. There are mechanisms to [move accounts between organizations and organizational units](#), but if your new approach indicates some workloads in an account should belong to one organizational unit and other workloads belong to another organizational unit, you will have misalignment. To align the accounts with the operational model, you will have to migrate workloads between existing accounts, or to new accounts. The methods and capabilities to migrate between accounts are similar to [accomplish these migrations](#).

Available services

AWS Organizations

AWS Organizations provides the underlying infrastructure and capabilities for customers to build and manage their multi-account environments. Using AWS Organizations, customers can automate AWS account creation and management; govern access within the organization to AWS services, resources, and Region using preventative guardrails; centrally manage policies across multiple AWS accounts (SCPs, Tag Policies, AWS Backup, ML opt-out); configure multi-account capabilities for AWS services (such as AWS Config, AWS CloudTrail, AWS CloudFormation, GuardDuty, Amazon Macie, and IAM Identity Center); share resources across accounts; and consolidate their bill.

AWS has the following resources available for help you establish your multi-account environment using AWS Organizations:

- [AWS Organizations features](#)
- [Organizations supported multi-account services](#)
- [Organization Quotas](#)

AWS Control Tower

AWS Control Tower provides a simplified way to set up and govern a secure, multi-account AWS environment based on the guidance in this paper. AWS Control Tower automates the creation of your multi-account environment using AWS Organizations, instantiating a set of initial accounts and with some default guardrails and configurations for the environment. Although AWS Control Tower reduces flexibility, it also provides automations to manage your cloud environment efficiently.

Note

The OU structure that AWS Control Tower initially deploys is slightly different from the guidance in this paper, review [AWS multi-account environment with Control Tower](#) for a detailed implementation and mapping between the AWS Control Tower implementation and the guidance offered in this paper.

AWS has the following resources available for help you establish your multi-account environment using AWS Control Tower:

- [Appendix E: How does AWS Control Tower establish your multi-account environment?](#)
- [Getting started with AWS Control Tower](#)
- [AWS Control Tower Quotas](#)

AWS Managed Services

AWS Managed Services (AMS) uses AWS services and a growing library of automations, configurations, and run books, to provide an end-to-end operational solution for both new and existing AWS environments. AMS covers the people element of operating the technology that AWS services provide.

For additional information, refer to [AWS Managed Services features](#).

Conclusion

If you are in the early stages of adopting AWS, you can use these best practices to start implementing an AWS environment structure that is sufficient to meet your initial needs. As your adoption of AWS expands and your requirements increase, you can be confident that your AWS environment can be expanded to meet those needs without requiring significant restructuring.

If you already have an AWS environment in place, you can use these best practices to assess its current state. By doing so, you can determine if you're fully realizing the benefits of using multiple OUs and AWS accounts and, if necessary, you can make plans to enhance your current environment.

In either case, your [AWS sales team](#) and [AWS Partner Network](#) (APN) partners are ready to help you apply these best practices to meet your business needs.

Contributors

Contributors to this document include:

- George Rolston, Sr. Solutions Architect, Amazon Web Services
- Todd Gruet, Sr. Solutions Architect, Amazon Web Services
- Paul Bayer, Principal Consultant, Amazon Web Services
- Sam Elmalak, Principal Solutions Architect, Amazon Web Services
- Ilya Epshteyn, Senior Manager, Identity Solutions, Amazon Web Services
- Christopher Kampmeier, Senior Solutions Architect, Amazon Web Services
- Tomas Riha, Senior Solutions Architect, Amazon Web Services
- Dave Walker, Principal Solutions Architect, Security and Compliance, Amazon Web Services
- Alex Torres, Sr. Solutions Architect, Amazon Web Services
- Rodney Lester, Principal Solutions Architect, Amazon Web Services
- Brandon Wu, Senior Security Solutions Architect, Amazon Web Services
- Nathan Case, Security Strategist, Amazon Web Services
- Brian Mycroft, Chief Technologist, Amazon Web Services
- Jason DiDomenico, Sr Solutions Architect, Amazon Web Services

Additional resources

For additional information, see:

- [Management and Governance on AWS](#)
- [Security, Identity, and Compliance on AWS](#)
- [AWS Best Practices for Security, Identity, and Compliance](#)
- [AWS Well-Architected](#)
- [AWS IAM Permissions Guardrails](#)
- [Establishing your Cloud Foundation on AWS](#)

Appendix A: Relation to AWS Well-Architected

[AWS Well-Architected](#) helps cloud architects build secure, high-performing, resilient, and efficient infrastructure for their applications and workloads. Based on six pillars — operational excellence, security, reliability, performance efficiency, cost optimization, and sustainability — AWS Well-Architected provides a consistent approach for customers and partners to evaluate architectures, and implement designs that can scale over time.

The best practices for organizing your AWS environment addressed in this guide augment and support the best practices represented in the following sections of the Well-Architected pillars.

Topics

- [Operational Excellence Pillar](#)
- [Security Pillar](#)
- [Reliability Pillar](#)
- [Cost Optimization Pillar](#)

Operational Excellence Pillar

- [Organization - Operating Model](#) – Use of multiple AWS accounts and OUs enable you to support multiple operating models within a common overall AWS environment.
- [Prepare – Mitigate Deployment Risks](#) – Use of separate AWS accounts for deployment-oriented operations help secure those operations and isolate them from the targeted test and production workload environments.

Security Pillar

- [Security Foundations - AWS Account Management and Separation](#) – Use of separate accounts for your test and production workloads helps maintain isolation between those environments.
- [Identity and Access Management – Identity Management](#) – Use of a centralized identity provider and federated access helps you more efficiently manage human access across your accounts.
- [Identity and Access Management – Permissions Management](#) – You can define permission guardrails for your AWS environment and provide least privilege access to identities that need access to your accounts.

- [Configure service and application logging \(SEC04-BP01\)](#) – Your security operations team can benefit from the centralization of logs generated across your accounts in support of analysis and detection requirements.

Reliability Pillar

- [Foundations - Manage Service Quotas and Constraints](#) – By isolating workloads in their own accounts, you can more easily manage service quotas for those workloads.

Cost Optimization Pillar

- [Practice Cloud Financial Management – Cost-Aware Processes](#) – You should build in cost awareness from the start of your cloud adoption journey.
- [Expenditure and Usage Awareness – Governance](#) – Your account structure can help you isolate different workloads for fiscal and billing purposes.
- [Expenditure and Usage Awareness – Monitor Cost and Usage](#) – Since costs are allocated by default at the account level, you can distinguish costs across accounts.

Appendix B: Worksheet for mapping workload environment purposes to hosting environment types

You can use the worksheet in this appendix to help you determine where workload environments for a given purpose will be deployed in your overall AWS environment. This worksheet, the worksheet in [Appendix C](#), and the [Organizing workload-oriented OUs](#) section are intended to be used together.

Combined, these worksheets and guidance can help your team quickly iterate on and arrive at an informed view of your own approach to organizing your workload-oriented OUs. They can also help inform how you might identify and scope your workload-oriented accounts.

For example, as you complete the worksheets in Appendix B and Appendix C, you should gain a better understanding of:

- How you expect to position AWS sandbox environments in relation to work done on your corporate desktops and in your AWS development environments.
- How you expect to generally divide and position your set of AWS development, test, and production environments in relation to your SDLC. Coupled with these best practices, this knowledge should help you refine your set of workload-oriented OUs and help inform the overall scoping of your workload-oriented accounts.

Topics

- [Use the results for internal documentation](#)
- [How to use this worksheet](#)
- [Descriptions of example purposes of workload environments](#)
- [Descriptions of example workload hosting environment types](#)
- [Example worksheet](#)
- [Empty worksheet](#)

Use the results for internal documentation

Later, as you establish and expand your standards documentation for using AWS, you might find value in publishing the resulting customized table and descriptions internally. Your teams might

find value in being able to quickly understand where workload environments for a given purpose are meant to be positioned in your AWS environment.

How to use this worksheet

1. Review [Descriptions of example purposes of workload environments](#) and [Description of example types of workload hosting environments](#). Note how your current on-premises standards and your initial expectations for working in AWS are similar to and different from the examples.
2. Review the [Example worksheet](#) to gain a sense of typical mappings of workload environment purposes to workload hosting environment types. Note where your expectations align with and differ from the examples. Mark where you expect to position workloads of a given purpose.
3. Make a copy of the [Empty worksheet](#), expand and modify the table by adding, removing, and/or changing rows and columns based on your needs. You can revisit and refine your initial assignments after you complete the worksheet in Appendix C.
4. Modify the example descriptions to suit your needs.
5. Use [Appendix C](#) to complete a related worksheet to help you document the key attributes of each of your own list of workload hosting environment types.
6. Update and refine the table and descriptions, based on your initial review of both this worksheet and the worksheet in [Appendix C](#).

Descriptions of example purposes of workload environments

This section provides descriptions of each of the example purposes of workload environments shown in the [Example worksheet](#). These examples and descriptions are meant to act as a reference and starting point for you to modify and extend to suit your needs.

Self-paced learning and experimentation workload environments

Builders learning and experimenting on their own, ideally with access to a wide variety of technologies.

Development workload environments

Work done by builder teams up and down the stack to develop and support their workloads and services. Examples of code includes proprietary application code, test scripts, test data, data models, machine learning models, security roles and policies, and Infrastructure as Code (IaC).

Generally, it's most efficient for builders to carry out development tasks locally on their corporate desktops.

Static analysis, build, and unit testing workload environments

Teams naturally develop, test, and carry out static analysis, builds of draft artifacts, and unit tests in their corporate desktop and/or cloud development environments. More formally controlled executions of these processes including builds of formal candidate artifacts typically occur through either continuous integration (CI) jobs or in early CI stages of continuous delivery (CD) pipelines.

CI jobs and CD pipelines workload environments

CI jobs and CD pipelines are developed and tested by teams in their development environments before they are potentially formally validated in a test environment. Given the common requirements to formally control creation of candidate artifacts and orchestration of validation stages, formal execution of CI jobs and CD pipelines typically occurs in production environments.

Smoke testing workload environments

Smoke testing is often used to quickly determine the overall success or failure of either deploying or releasing a change to any workload environment. A smoke test is an initial indication of the success or failure of a deployment or release. Unlike many other types of system level testing, smoke testing is also typically used against production workloads to initially validate that release of a change was generally successful.

Development integration testing workload environments

Development integration testing is the early integration testing of code and configurations that typically occur prior to formal types of system testing. This is done so that basic integration issues are detected early in the lifecycle.

Production-like system testing workload environments

System level testing is the end-to-end testing of services and applications that occur in production-like environments prior to changes being promoted for use in production. In the spirit of *shifting left* so that issues are found earlier in the lifecycle, builder teams can take advantage of automation and the on-demand and elastic nature of the cloud to perform early system level testing in their development environments.

Common examples of system level testing include, but are not limited to:

- **Functional acceptance testing** – Verification that the service meets the business requirements. In mature situations, this testing is often largely automated.
- **User acceptance testing** – Validation from a user’s perspective or through a proxy, such as a product owner, that the service meets their expectations.
- **Exploratory testing** – Functionality testing by humans in an ad hoc manner.
- **Performance testing** – Systems level testing to ensure a service is able to meet the expected performance goals with the supporting capacity.
- **Workload-recovery testing** – Testing the ability of your workloads to recover from component-level and workload-level failures in the face of sustained failures.
- **Penetration testing** – Also known as *pen* testing, testing services to identify security vulnerabilities.
- **Multi-region testing** – Testing that a workload is deployed in an active-active mode across multiple AWS Regions.
- **Disaster recovery / business continuity testing** – Business testing to ensure that services can be restored when in the event of a large-scale disaster.

Stable shared test workload environments

Teams who own shared services are sometimes responsible for managing stable test instances of their workloads that are loaded with test data in support of development and integration testing.

Resiliency testing workload environments

Also known as *chaos engineering*, this scenario is where failures are purposely introduced into the environment to ensure that the workloads respond in the expected manner. More advanced scenarios include running such tests against production workloads.

Demo workload environments

Internal teams demonstrating services and capabilities to internal and/or external customers throughout the lifecycle.

External pre-release workload environments

Sometimes referred to as *alpha*, *beta*, or *early access*, these workloads are typically accessed by external customers. Given the external access nature of these workloads, it’s common for these workloads to be contained in either production or test environments.

Production workload environments

Formally managed and monitored workloads that are deployed to distinct production environments.

Descriptions of example workload hosting environment types

The following descriptions provide more detail for the example workload hosting environment types. These examples and descriptions are meant to act as a reference and starting point for you to modify and extend depending on your needs.

Each workload hosting environment type is not intended to represent a single account in your overall AWS environment. Rather, a workload hosting environment type is meant to help you think about the overall categorization of and distinction between the accounts in which your workloads reside. You might end up having numerous accounts of any given hosting environment type.

For example, if there are significantly different security and operations needs of the environment types that you identify, then those environments might help you refine your overall OU structure for workloads.

For a set of example attributes for each of the example hosting environment types, refer to [Appendix C](#).

Corporate desktop environments

Depending on the degree of access your builders have on their corporate desktops, they might perform much of their day-to-day, non-cloud work locally on their desktops. As you adopt AWS, many of your builder teams will start creating and managing AWS resources and workloads in their sandbox and development AWS environments.

Sandbox environments

Sandbox environments are environments allocated to individuals in which they have significant degrees of administrative access so that they can learn and dive deep into a wide variety of AWS services and other technologies. Typically, sandbox environments do not have access to your internal networks, services, and data.

Foundational guardrails are used to mitigate the risk of this extensive access. For more information, refer to [Sandbox OU](#).

Development environments

An extension of your corporate desktops that enable your builder teams to carry out formal development tasks both with and in AWS that cannot occur only on their local corporate desktops.

Development environments are often allocated on a team and/or individual basis. Depending on the permissions granted to your builder teams, varying degrees of self-paced learning, experimentation, and prototyping can also occur in development environments.

Foundational guardrails and shared infrastructure services that are treated as production stable resources are typically used to ensure the proper degrees of access control and stability of development environments.

Data-oriented development environments

A hybrid of the attributes of development and production environments. Data scientists and data engineers need an environment in which they can develop and perform early forms of testing machine learning and other algorithms that require access to production data.

Given the hybrid nature of this environment type, the guardrails and controls for data development are typically a blend of those used to support development and production environments.

Test environments

A type of workload hosting environment in which changes are formally validated before they are promoted for release to production environments. Depending on how you view this overall stage of your SDLC and the distinct security and operational requirements of it, you might end up dividing this overall type into multiple types of environments.

Typically, test environments are secured and managed in much the same manner as production environments so that validation efforts occur in production-like environments and issues are detected prior to release to production.

Production environments

Highly secured, formally managed hosting environments in which production data and workloads reside.

Example worksheet

The following example worksheet represents a typical mapping of workload purposes to types of workload hosting environments. This example is meant to spur discussion of and comparison to both your current on-premises environments and your expectations for working on AWS.

Table 1 — Example worksheet

Workload environment purpose	Corporate desktops	Sandbox	Development	Data-oriented development	Test	Production
Self-paced learning/early experimentation	•	•	•	•		
Development	•		•	•		
Static analysis, building, unit testing	•		•	•		
CI jobs and CD pipelines			•	•	•	•
Smoke testing			•	•	•	•
Development integration testing			•	•	•	

Workload environment purpose	Corporate desktops	Sandbox	Development	Data-oriented development	Test	Production
Production-like system testing			•	•	•	
Stable shared test			•	•	•	
Resiliency testing			•	•	•	•
Demo	•	•	•	•	•	•
External pre-release (Alpha, Beta)					•	•
Production						•

Empty worksheet

Use this worksheet to help you describe your expectations for where certain purposes of workloads are expected to occur across your AWS workload environments. Copy and modify the rows and column headers in the following worksheet as needed.

Table 2 — Empty worksheet

Workload environment purpose	Corporate desktops	Sandbox	Development	Data-oriented development	Test	Production
Self-paced learning/early experimentation						
Development						
Static analysis, building, unit testing						
CI jobs and CD pipelines						
Smoke testing						
Development integration testing						
Production-like system testing						
Stable shared test						

Workload environment purpose	Corporate desktops	Sandbox	Development	Data-oriented development	Test	Production
Resiliency testing						
Demo						
External pre-release (Alpha, Beta)						
Production						

Appendix C: Worksheet for identifying attributes of workload hosting environments

Use the following worksheet to help you identify and refine your expectations for common security and operational attributes across your set of AWS environments. This worksheet helps you plan how to structure your overall AWS environment based on your requirements.

Topics

- [How to use this worksheet](#)
- [Descriptions of example attributes](#)
- [Example worksheet](#)
- [Empty worksheet](#)

How to use this worksheet

1. Follow the instructions in [Appendix B: Worksheet for mapping workload environment purposes to hosting environment types](#).
2. Review [Descriptions of example attributes](#) located after the worksheets below. Note how your current on-premises standards and your initial expectations for working in AWS are similar to and different from the example attributes.
3. Review the [Example worksheet](#) example attributes to get a sense of typical attributes of workload hosting environment types. Note where your expectations align with and differ from the examples. Modify the example attributes in each cell based on your expectations.
4. Make a copy of the [Empty worksheet](#). Expand and modify the table by adding, removing, and/or changing rows and columns based on your needs.
5. Mark where you expect to position workloads of a given purpose. Include learnings from the completion of the worksheet in [Appendix B](#).
6. Modify the example attribute descriptions to suit your needs.
7. Iterate on the worksheet from [Appendix B](#) based on the outcome of completing the following worksheet.
8. Publish and share internally a draft of the resulting tables and descriptions so that your cross-functional teams can further evolve them.

Descriptions of example attributes

The following descriptions elaborate on the meaning of each of the example attributes of workload hosting environment types. These examples and descriptions are meant to act as a reference and starting point for you to modify and extend to suit your needs.

Note

In the table of example attributes of workload hosting environment types, you'll see a close alignment between test and production environment types. This is a typical approach used by organizations to help ensure that testing occurs in production-like environments prior to changes being promoted to production environments.

If this degree of alignment is important to you, then it can factor into how closely you align your production OU to test environment OUs, both in terms of SCPs and overall security baselines.

Owners/tenants

These are the teams or individuals accountable for the environment and the workloads deployed to it.

Sandbox, development, and data-oriented development

Although sandbox environments are typically allocated on an individual basis and development environments are often allocated on a team basis, your requirements might vary. For example, in support of short duration events, such as company hackathons, you might need to allocate sandbox environments to teams. Similarly, you might have the need for individual owned development environments if team owned environments don't meet your needs.

Regardless of whether a team or an individual owns an environment, you should consider applying virtually the same attributes across a given type of environment.

Test and production

The ownership and tenancy model typically depends on the operational model you use. For example, in a more traditional centralized operational model, your operations team might own the production and test environments and possibly the workloads residing in them.

In a more federated DevOps style operating mode, individual application and data services teams might own their own production and test environments and the workloads residing in them.

Tolerance for extended outages

This refers to the tolerance of the business for extended outages of a given type of environment. For example, if access to a particular type of environment is down for several hours, what's the impact to the business?

Generally, organizations treat the environments in which they perform their formal day-to-day work as needing to have production qualities in terms of stability and access.

Corporate desktops

An inability to use corporate desktops for any appreciable length of time is considered to be a significant impact in most organizations.

Sandbox

In our example definition of sandbox environments, they are not typically viewed as having a great deal of importance to everyday formal tasks.

Development and data oriented development

To the extent that development and early testing of code for a variety of foundation and business workloads occurs through development environments, extended outages of development environments can have a significant impact on your business.

Test

Similar to development, your ability to validate important changes will be severely impacted due to extended outages of your test environments.

Internet access

This refers to the extent to which access to and from the internet is allowed.

Sandbox and development outbound access

Typically, users of sandbox and development environments benefit from the ability to download packages from the internet. For example, access to public language platform and OS-native

package repositories, including npms, rpms, PyPi modules, etc. Additionally, sandbox and development environments typically benefit from having access to public web services to support experimentation, development, and early testing.

In either case, you might implement filters on outbound requests to the internet to reduce the risk of accessing malicious or unauthorized software packages and services.

Sandbox and development inbound access

Sandbox environments might have the ability to contain workloads that are accessible from the internet. Given that they contain internal data and Intellectual Property (IP), inbound access from the internet is not typically allowed for development environments. In this respect, your AWS development environments might be similar to your corporate desktops.

Test and production internet access

Outbound and inbound internet access with test and production environments is typically more controlled than in lower environments. For example, you might require defined connectivity for each workload that performs outbound requests to the internet or receives inbound requests.

Internal network access

This refers to the extent to which workloads in an environment are allowed to connect to other internal services whether they reside in your on-premises or on AWS environments.

Sandbox

Given the broad administrative access and bidirectional connectivity with the internet, sandbox environments are typically inhibited from accessing internal data and services.

Development

Development environments, like corporate desktops, are typically allowed to connect to your source code management, artifact management, CI/CD, and deployment/release automation services. These environments also typically have connectivity to shared infrastructure services, such as DNS resolution and user authentication and identity management instances populated with test data.

Data

This refers to the overall type of data allowed in each environment.

Sandbox

Unlike corporate desktops and development environments, sandbox environments are typically intended to house only public data. Due to the risk of data loss and unintended exposure, intellectual property (IP) and internal data are typically not allowed in sandbox environments.

IP often includes proprietary source code, configuration data, artifacts (such as applications and binary packages), non-public test data, business, and operational data.

Development and test

Use of sensitive production data is typically prohibited in these environments.

Data oriented development

A variation of a development environment is one that supports the needs of data scientists and data engineers. These teams typically need direct human access to cloud environments to develop and iterate on machine learning (ML) models. Because these workloads need access to production sensitive data, additional guardrails and controls are typically applied to these forms of development environments.

Third-party software and cloud services

This refers to the extent to which third-party and overall cloud services are allowed in each environment. Typically, your existing standards for acceptable use of third-party software and cloud services should extend to your AWS environments.

Degree of access

This refers to the extent to which people interacting with an environment have access to resources in the environment.

Foundation team development

Given the nature of their development work, your cloud foundation or cloud platform team of cross-functional network, security, and other responsibilities typically needs greater write access to foundational resources in their development environments.

Lifespan of resources

This refers to how long workloads and supporting resources are expected to live in a given environment. Even in development environments, teams should be encouraged to automate important configurations as much as feasible. Automation enables teams to recreate resources and workloads on demand so that they have less dependency on hand-built configurations.

Direct human write access to workload resources

This refers to whether human users have the ability to directly create, update, and delete workload resources.

Corporate desktop, sandbox, and development

Across corporate desktop, sandbox, and development environments, individuals and team often need to perform direct manual manipulation of cloud resources in early stages of experimentation. Later, as certain cloud resource and workload configurations mature, investment in automation often makes sense.

Production

In support of workloads in production environments, as a general recommendation, changes made to production should be automated. In these cases, human user write access to workload resources should not typically be necessary and not allowed.

Direct read access to production resources and data could be allowed on a least privileged basis to support audit and operational troubleshooting scenarios.

Break-glass scenarios should be supported to enable temporary write access with strict auditing under exceptional conditions.

Automated workload provisioning

This refers to the extent to which automated workload provisioning applies to a type of environment.

Corporate desktops

A modest degree of automated workload provisioning can be developed and tested on local corporate desktops. To the extent that [type 2 hypervisors](#) for running virtual machines locally

and/or containers are supported on corporate desktops, builders can carry out some degrees of environment automation testing on their corporate desktops. However, unless an AWS service API is available locally, workload automation that depends on AWS service APIs needs to occur in their AWS development environments.

Sandbox

Because sandbox environments are not typically connected to your source code management systems and don't usually contain internal data, there might be limited value in developing workload automation in sandbox environments. Reuse of open source or other third-party automation is typical in sandbox environments.

Development

Teams typically use a mix of manual and automated means to initially experiment and perform early iterations on their cloud resource configurations. Once teams mature their desired configurations, they typically invest in infrastructure as code (IaC) to automate workload configurations. Your standard source code management systems are typically used to house the IaC automation files.

Formal change management for workloads

This refers to whether formal change management processes are required to apply changes to workloads.

Corporate desktop, sandbox, and development

Formal change management does not typically apply to workloads contained in these environments.

Test

Minimally, we recommend that you at least test your change management processes in test environments.

Production

Typically, some form of change management process applies to all changes made to workload in your production environments.

Degree of centrally managed foundation

This refers to the extent to which there's a centrally managed set of foundation resources in a given environment.

Sandbox

In sandbox environments, given the experimental and wide-ranging administrative access, there might be little if any centrally managed foundation resources other than common enterprise guardrails (see next section).

Development

In development environments, you can provide centrally-managed foundation services to teams. By doing so, you can remove undifferentiated heavy lifting from their day-to-day work.

Test and production

Similar to development, you can relieve individual teams from needing to manage underlying foundation services.

Common enterprise guardrails

Across all of your AWS environments, you'll benefit from applying a set of enterprise guardrails. The makeup and depth of these guardrails might vary across your types of environments depending on the level of risk to the business and the degree of access allowed to owners of the environments.

For example, AWS API logging using AWS CloudTrail and cloud resource configuration recording via AWS Config is typically a common guardrail applied across all environments.

Example worksheet

This example worksheet represents a typical mapping of attributes to types of workload hosting environments. This example is meant to spur discussion and comparison to both your current on-premises conventions and your expectations for working on AWS.

Table 3 — Example worksheet - Attributes of workload hosting environment types

Attribute	Corporate desktops	Sandbox	Development	Data-oriented Development	Test	Production
Owners / tenants	Individual	Individual	Team	Team	Same as Production	Depends on operating model
Tolerance to extended outages	Low	High	Low to medium	Low to medium	Low	Extremely low
Internet access	Outbound requests subject to proxying and filtering controls No inbound requests	Outbound and inbound requests	Outbound requests subject to proxying and filtering controls No inbound requests	Given the presence of production data, outbound requests will likely be more controlled than development environments. No inbound requests	Same as production	Workload-specific outbound and inbound requests Proxy-based access to external services

Attribute	Corporate desktops	Sandbox	Development	Data-oriented Development	Test	Production
Internal network access	Shared development and infrastructure services Other development workloads Corporate services	No connectivity to corporate and data center services No connectivity to shared development and infrastructure services	Shared development and infrastructure services Other development workloads No access to business production services	Shared development and infrastructure services Access to defined production data sources	Shared development and infrastructure services Other test services	Shared infrastructure services Other production services
Data	Intellectual property (IP) Test data	Public data only Public test data (no Intellectual property)	IP Test data No access to production data	IP Production data	IP Test data Typically avoid use of production sensitive data unless sanitized	IP Production data

Attribute	Corporate desktops	Sandbox	Development	Data-oriented Development	Test	Production
Third-party software and cloud services	Installation of approved software	<p>Access to broad set of AWS services</p> <p>Installation of approved software</p>	<p>Access to enterprise standardized AWS services</p> <p>Controlled access to AWS services undergoing standardization for the purposes of testing</p> <p>Installation of approved software</p>	<p>Access to enterprise standardized AWS services</p> <p>Controlled access to AWS services undergoing standardization for the purposes of testing</p> <p>Installation of approved software</p>	<p>Access to enterprise standardized AWS services</p> <p>Access to AWS services undergoing standardization</p> <p>Installation of approved software</p>	<p>Access to enterprise standardized AWS services</p> <p>Installation of approved software.</p>

Attribute	Corporate desktops	Sandbox	Development	Data-oriented Development	Test	Production
Degree of access	Limited OS configuration	Wide ranging administrative cloud resource write access Some limits of modifying foundation resources	Wide ranging access including write access to development and test workload-specific IAM service roles and policies Some limits of modifying foundation resources	Given the presence of production data, likely more limited access to cloud resource write access than in development	Same as production	Least privileged access Strictly controlled access based on operating model that is in effect Service-to-service access based on authorization.
Lifespan of resources	Up to builder to manage	Temporary	Up to owning teams to manage	Up to owning teams to manage	Same as production	Depends on business need
Direct human write access to workload resources	Yes	Yes	Yes	Yes	Same as production	No

Attribute	Corporate desktops	Sandbox	Development	Data-oriented Development	Test	Production
Automated workload provisioning	Limited	Limited	Mix of manual and automated	Mix of manual and automated	Same as production	Yes
Formal change management for workloads	No	No	No	No	Same as production	Yes
Degree of centrally managed foundation	As appropriate for corporate desktops	Sufficient to ensure overall security	Typical foundation resources centrally managed	Typical foundation resources centrally managed	Same as production	Typical foundation resources centrally managed
Common enterprise guardrails	Desktop specific	Yes Guardrails to prevent write access to baseline security monitoring services and configuration	Yes Guardrails to prevent write access to foundation resources	Yes Guardrails might be a hybrid of those used for development and production environments	Same as production	Yes

Empty worksheet

Copy and modify the rows and column headers in the following worksheet and note how each attribute relates to a given environment type.

Table 4 — Attributes of workload hosting environment types

Attribute	Corporate desktops	Sandbox	Development	Data-oriented development	Test	Production
Owners / tenants						
Tolerance to extended outages						
Internet access						
Internal network access						
Data						
Third-party software and cloud services						
Degree of access						

Attribute	Corporate desktops	Sandbox	Development	Data-oriented development	Test	Production
Lifespan of resources						
Direct human write access to workload resources						
Automated workload provisioning						
Formal change management for workloads						
Degree of centrally managed foundation						
Common enterprise guardrails						

Appendix D: Multiple AWS Regions

If you plan to use multiple AWS Regions, keep the following considerations in mind as you design your overall AWS environment.

Topics

- [Geographic scopes of data protection](#)
- [Performance considerations](#)
- [Log management](#)

Geographic scopes of data protection

If you use different AWS Regions that are in the same geographic scope defined by the data protection requirements applicable to your workloads, you can use the same IAM IdP(s) to federate to all accounts in live, disaster recovery, or load balanced live environments. You can replicate databases between environments using appropriate mechanisms, such as [Amazon DynamoDB global tables](#) or [Amazon RDS read replicas](#). In such circumstances, it is also possible for you to distribute core elements of your foundational AWS environment such that the log archive bucket is in one Region and assets in other accounts in other Regions log cross-Region to it.

You should carefully consider whether the data protection requirements applicable to your workload differ across countries, or are subject to data sovereignty requirements or export control. This might impact your ability to make cross-Region data transfers. (Note that cross-Region data transfers incur networking costs.)

Performance considerations

There are also performance considerations to keep in mind for certain workloads. Some services are by their nature per-Region, which makes it more sensible for you to deploy such workloads with all assets in the same Region. For example, AWS KMS keys cannot be exported from a Region, and use of a KMS key in another Region is likely going to add latency to an application. We therefore recommend using AWS KMS in the same Region, unless specific governance policies, regulatory or corporate, mandate otherwise.

Close collaboration between your security and architecture teams and your workload owning teams is important to properly using KMS. Your design of how Amazon S3 objects, EBS volumes,

and other data are encrypted and potentially replicated across Regions should factor in low latency when required.

Where cross-account replication of these assets is required, Amazon S3 Cross-Region Replication (CRR) enables on-the-fly re-encryption of an object with an AWS KMS key in the destination Region. Multi-Region duplication of AWS KMS keys for the decryption of cross-Region copied EBS volumes can be achieved using the techniques covered in [Busy Engineer's Document Bucket](#).

Log management

When logs are generated, we recommend that you implement secondary controls to filter them before they are passed outside a compliance scope boundary associated with an account, or are passed cross-Region. If your logs contain sensitive data, this approach helps ensure that such sensitive data cannot escape your defined compliance scope boundary using AWS logging capabilities.

Although AWS CloudTrail has built-in cross-account logging capability and AWS Config can aggregate configuration and compliance data across accounts and Regions, it might be more appropriate for you to aggregate logs in an account. You can use AWS Lambda functions or similar to filter the logs before sending them to another Region for aggregation into a multi-Region logging archive.

Appendix E: How does AWS Control Tower establish your multi-account environment?

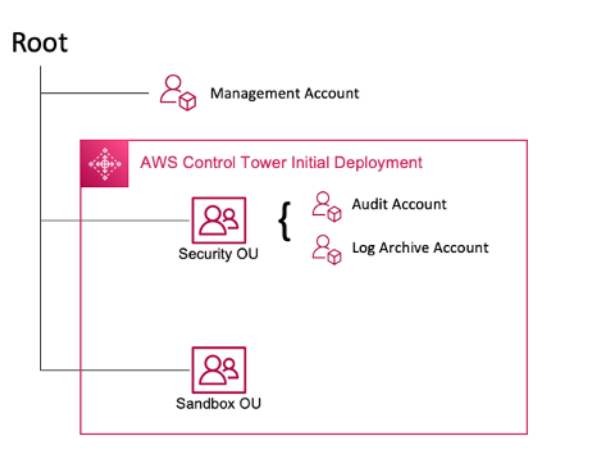
Topics

- [Establish your multi-account environment with AWS Control Tower](#)
- [Next steps for setting up your multi-account environment](#)

Establish your multi-account environment with AWS Control Tower

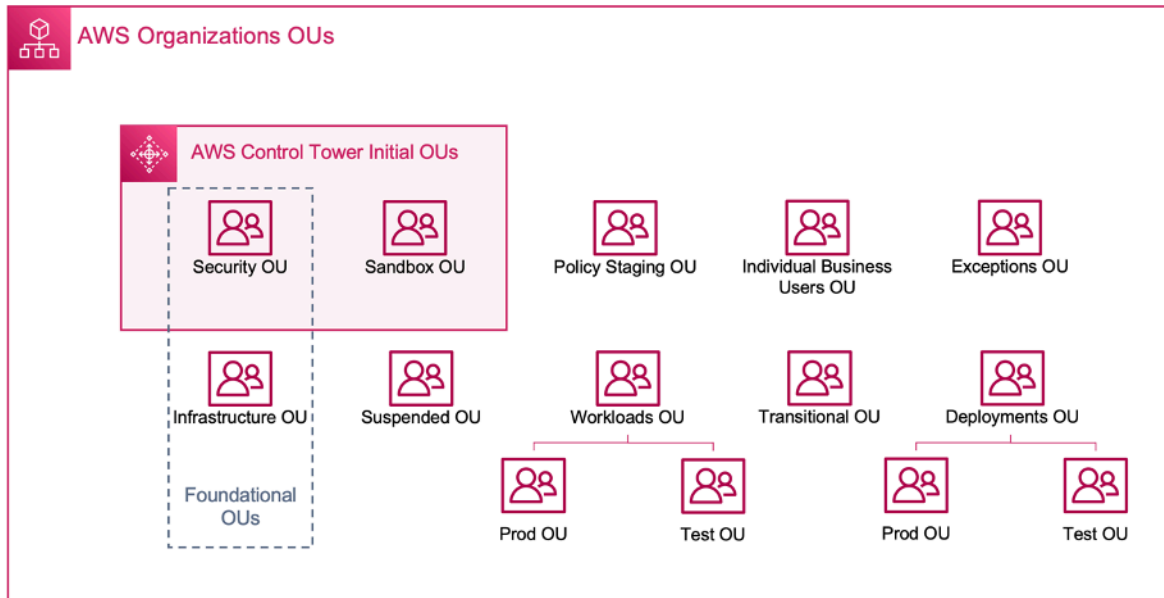
When you set up your multi-account environment using AWS Control Tower, it creates two OUs.

- **Security OU** - Within this OU, AWS Control Tower creates two accounts:
 - Log Archive
 - Audit (*This account corresponds to the Security Tooling account discussed previously in the guidance.*)
- **Sandbox OU** - This OU is the default destination for accounts created within AWS Control Tower. It contains accounts in which your builders can explore and experiment with AWS services, and other tools and services, subject to your team's acceptable use policies.



AWS Control Tower allows you to create, register, and manage additional OUs to expand the initial environment to implement the guidance.

The following diagram shows the OUs initially deployed by AWS Control Tower. You can expand your AWS environment to implement any of the recommended OUs included in the diagram, to meet your requirements.



OUs initially deployed by AWS Control Tower

Next steps for setting up your multi-account environment

To get started with AWS Control Tower, visit the [Getting Started with AWS Control Tower](#) documentation page. We recommend that you review the pre-requisites and next steps required to establish your multi-account environment on AWS.

For complete guidance on establishing your multi-account environment, you can review the guidance included in this whitepaper.

Document history

To be notified about updates to this whitepaper, subscribe to the RSS feed.

Change	Description	Date
Updated	Updated for technical accuracy.	March 28, 2024
Updated	Updated break-glass guidance and included disaster recovery guidance.	March 15, 2023
Updated	Updates to include the latest best practices for managing a multi-account environment.	July 26, 2022
Updated	Updated guidance for existing customers getting started with their multi-account environment.	March 31, 2022
Updated	Updated guidance for establishing a multi-account environment.	July 19, 2021
Initial release	Whitepaper first published.	March 18, 2021

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

© 2023 Amazon Web Services, Inc. or its affiliates. All rights reserved.