

Establishing Your Cloud Foundation on AWS

November 17, 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.

© 2021 Amazon Web Services, Inc. or its affiliates. All rights reserved.

Contents

- Introduction 1
- Capabilities2
- Working with the capabilities.....9
- Next steps 10
- Conclusion 10
- Contributors 10
- Further reading 11
- Document revisions11
- Appendix A: Capability structure and example..... 12
- Appendix B: Sample timeline 15

Abstract

The increasing breadth and depth of cloud services makes the cloud a powerful enabler of efficiency, agility, and rapid innovation. However, building a foundational AWS Cloud environment requires decisions across multiple AWS and partner products, services, and solutions. Customers are looking for guidance to help them set up and operate an environment that is compatible with their IT practices, and enables their builders and operators while adhering to governance requirements.

This whitepaper introduces a guided path approach to help customers build and evolve their AWS Cloud environment based on a consolidated set of definitions, use cases, guidance, and automations. The approach includes *people*, *process*, and *technology* considerations of establishing an AWS Cloud environment.

Introduction

The primary business drivers behind moving to the cloud include greater agility, innovation, and scale. When planning a cloud adoption strategy, the number of decisions that you need to make to stand up a production-ready cloud environment is significant. Decisions that are made early on can affect your ability to enhance and/or scale your environment in the future. This complexity has led customers to look for prescriptive guidance across the range of AWS services that can be used to create a foundational environment.

Establishing a cloud foundation on AWS requires guidance tailored to your business needs. Using a [capability-based approach](#), you can create an environment to deploy, operate, and govern your workloads. You can also enhance the capabilities to extend your environment as your requirements evolve and you deploy additional workloads to the cloud.

Building a foundational environment on AWS can be done with a standard, prescriptive set of capabilities across [different functional areas](#). These capabilities can be used as a structured way to quickly build or expand your AWS Cloud environment, and include use case scenarios and corresponding guidance.

You can adopt and implement capabilities according to your operational and governance needs. As your business requirements mature, the capability-based approach can be used as a mechanism to verify that your cloud environment is ready to support your workloads and scale as needed. This approach enables you to confidently establish your cloud environment for your builders and your business.

Capabilities

To support cloud adoption, AWS recommends that you have a foundational set of capabilities that enable you to deploy, operate, and govern your workloads.

A *capability* includes a definition, use case scenarios, opinionated guidance, and supporting automation to establish and operate a specific part of a cloud environment. Capabilities are components that can help you plan, implement, and operate your cloud environment, and include *people*, *process*, and *technology* considerations. Capabilities are designed to integrate into your overall technology environment.

In addition to technology implementation guidance, capabilities include operational guidance (for instance, notifications, event handling, and remediation, as well as team resource skills and processes) needed to stand up and operate each capability. For an example of what a capability should offer, refer to [Appendix A](#).

AWS has defined a set of 30 capabilities to help you establish a cloud foundation. One way to categorize these capabilities is by functional areas which can help you identify accountable owners and stakeholders in the development, operations, and governance of capabilities. Table 1 lists these capabilities organized by functional area. Each capability is listed under a single, *primary* functional area. However, most capabilities are also relevant to other functional areas.

Table 1 - Cloud Foundations capabilities by primary functional area

<u>Security</u>	<u>Central IT</u>	<u>Operations</u>	<u>Software Engineering</u>	<u>Network</u>	<u>Finance</u>
Identity Management & Access Control	Template Management	Support	Developer Experience & Tools	Network Security	Cloud Financial Management
Log Storage	Tagging	Rollout/Rollback	Application Security	Network Connectivity	Resource Inventory Management
Encryption & Key Management	Metadata Sorting/Searching	Backup			
Secrets Management	Service Onboarding	Disaster Recovery			
Data Isolation	Records Management				

<u>Security</u>	<u>Central IT</u>	<u>Operations</u>	<u>Software Engineering</u>	<u>Network</u>	<u>Finance</u>
Security Incident Response	Data De-Identification				
Forensics	Logging & Monitoring				
Patching	Governance				
Vulnerability & Threat Management	Audit and Assessment				
Workload Isolation Boundary	Change Management				

Each capability includes stages of maturity that enable you to implement based on where you are in your cloud journey, including your governance and operational requirements. As your cloud environment grows and matures, the *capabilities* can be enhanced to meet your new requirements.

Capabilities definitions

This section includes high-level definitions for each foundational capability organized by their primary functional area. For a deeper dive into a specific capability and what it includes, refer to [Appendix A](#).

Security

Security functional area capabilities include:

- **Identity Management & Access Control** enables your teams to efficiently build and centrally manage the access to your cloud platform environment. The capability enables you to structure your organization, organize your accounts, and set up access to your environment based on a least-privilege model.
- **Log Storage** enables you to securely collect and store environment logs centrally within tamper resistant storage. This capability enables you to later evaluate, monitor, alert, and audit access and actions performed on your AWS resources and events.

- **Data Isolation** enables you to limit access to data at rest and in transit so that data is only accessible to appropriate, authorized entities. This capability also includes the ability to detect misuse and/or unauthorized access, leak, and theft of data.
- **Encryption and Key Management** refers to the ability to centrally manage encryption keys for different workloads, and the ability to encrypt data at rest and in transit. Access to keys is provided based on least privilege, and usage is monitored to report any anomalies. This capability also includes different patterns of rotation based on requirements.
- **Secrets Management** applies to managing *secrets* (access credentials) such as passwords, access keys, other API keys, X.509, or SSH private keys. This capability includes storage, access control, access logging, revocation, and rotation aspects for managing secrets.
- **Security Incident Response** enables you to respond to a security incident. Based on decisions specified in policy, the response involves characterizing the nature of the incident and making changes (which may involve activities including restoration of operational status, identification and remediation of root cause, and gathering evidence pursuant to civil or criminal prosecution).
- **Forensics** involve the analysis of log data and evidentially-captured images of potentially compromised resources, to determine whether a compromise occurred (and if so, how). Outcomes of root cause analysis resulting from forensic investigations are typically used to produce and motivate the application of preventative measures.
- **Patching** is the ability to deploy sets of changes to update, fix, and/or enhance the operation and security properties of workloads. This includes addressing security vulnerabilities, bug fixes, and other related work. The scope of patching includes operating systems, applications, and any relevant software systems.
- **Vulnerability & Threat Management** is the ability to identify vulnerabilities that can affect the environment (availability, performance, or security). This capability enables you to assess the impact and scope (e.g., blast radius) of vulnerabilities and threats, and address/remediate them.

- **Workload Isolation Boundary** enables you to create and manage isolated environments to contain your newly created or migrated workloads. This approach reduces blast radius of vulnerabilities and threats, and eases the complexity of compliance by providing mechanisms to isolate access to resources.

Central IT

Central IT functional area capabilities include:

- **Template Management** is the ability to create and group reusable templates in a central repository to quickly deploy, manage, and update infrastructure, schemas, golden images, and resources across the environment. This capability includes the necessary processes to create, test, update, and validate the templates when required. These templates are pre-approved implementation patterns using already approved and onboarded AWS services, and are ready to be used by different teams based on requirements.
- **Tagging** enables you to group sets of resources by assigning metadata to cloud resources for a variety of purposes. These purposes include access control (e.g., ABAC), cost reporting, and automation (e.g., patching for select *tagged* instances). Tagging can also be used to create new resource constructs for visibility or control (e.g., grouping together resources that make up a microservice, application, or workload). Tagging is fundamental to providing enterprise-level visibility and control.
- **Metadata Sorting/Searching** is the ability to search and filter based on metadata applied to tagged resources within your environment. These resources can be accounts, or resources within these accounts
- **Service Onboarding** is the ability to review and approve AWS services for use based on consideration of internal, compliance, and regulatory requirements. This capability includes risk assessment, documentation, implementation patterns, and the change communication aspects of service consumption.
- **Records Management** enables you to set retention of data according to your internal policies and regulatory requirements, including how to transition data to archive before it is deleted. This data can include financial records, transactional data, audit logs, business records, personally identifiable information (PII), or other data subject to retention policies.

- **Data De-Identification** is the ability to anonymize subsets of data and information as they are stored and processed to reduce their sensitivity (for example, national ID numbers, trade data, healthcare information), and when required, preserving the underlying data format. This capability also includes the ability to tokenize data (such as credit card numbers, physical address, health care records) to reduce the need to access the underlying sensitive data.
- **Logging & Monitoring** is the ability to gather and aggregate security and operational data about system and application activities, including near-real-time analysis of data to identify anomalies, indicators of compromise, performance issues, and configuration changes.
- **Governance** is the ability to implement executive board policies that your AWS Cloud environment must adhere to. This policy includes the rules for your environment, defines risks, and informs alignment of internal policies. As your cloud foundation matures, a portion of this capability is embedded in all other capabilities to ensure adherence to governance requirements.
- **Audit & Assessment** is the gathering and organizing of documentary evidence to enable internal or independent assessment of your cloud environment, and activities within it, against standards (including information about who accessed what, when, and from where, and what changes happened). This capability allows you to validate assertions that all changes were performed in accordance with policy and via appropriate workflow mechanisms.
- **Change Management** enables you to deploy planned alterations to all configurable items that are in an environment within the defined scope, such as production and test. An approved change is an action which alters resource configuration that is implemented with a minimized and accepted risk to existing IT infrastructure.

Operations

Operations functional area capabilities include:

- **Support** is the ability to troubleshoot an environment, ask questions, submit tickets, integrate into existing ticketing systems, and escalate issues to an appropriate entity for a timely response depending on criticality and support level. Support may also require granting ability to access relevant resources to perform troubleshooting and remediation activities.

- **Rollout/Rollback** is the defined strategy to roll out application or configuration changes to the environment, or roll back these changes in case of failure. Application and configuration changes can include updated permissions, new policies, new or updated network configuration, new version of the application, or updated software development kits. These configuration changes can also include modifications to the orchestration framework that deploy these changes.
- **Backup** is the ability to make reliable copy of data in a reliable way for retrieval as needed to meet business and security goals, Recovery Point Objective (RPO), and Recovery Time Objective (RTO). Content to be backed up includes: orchestration framework data and configuration, application data, logs, and customer data.
- **Disaster Recovery** involves the use of automated mechanisms to resume processing of transactions hosted in one physical environment, in a different physical environment in the event that the physical environment where the transactions were originally being processed becomes unexpectedly unavailable.

Software Engineering

Software Engineering functional area capabilities include:

- **Developer Experience & Tools** includes the tools and processes required for developers to build and deploy workloads easily to the cloud. This capability spans from storing code, to building workflows, to promoting workloads from test to production environment.
- **Application Security** encompasses the protection of application software, and the detection of anomalous behavior in the context of the applications' interactions with clients. Threats to be addressed include unauthorized access, privilege escalation, and other application-level threats typically characterized in threat frameworks.

Network

Network functional area capabilities include:

- **Network Security** enables you to design and implement security policies and controls across different levels of the networking stack to protect your resources from external or internal threats to ensure confidentiality, availability, integrity, and usability. This capability includes prevention, detection, and blocking of anomalous network traffic based on monitoring of ingress/egress and lateral data movement.
- **Network Connectivity** enables you to design, build, and manage a secure and highly available network cloud infrastructure. This capability provides best practices and resources to automate network infrastructure build, configuration, and expansion.

Finance

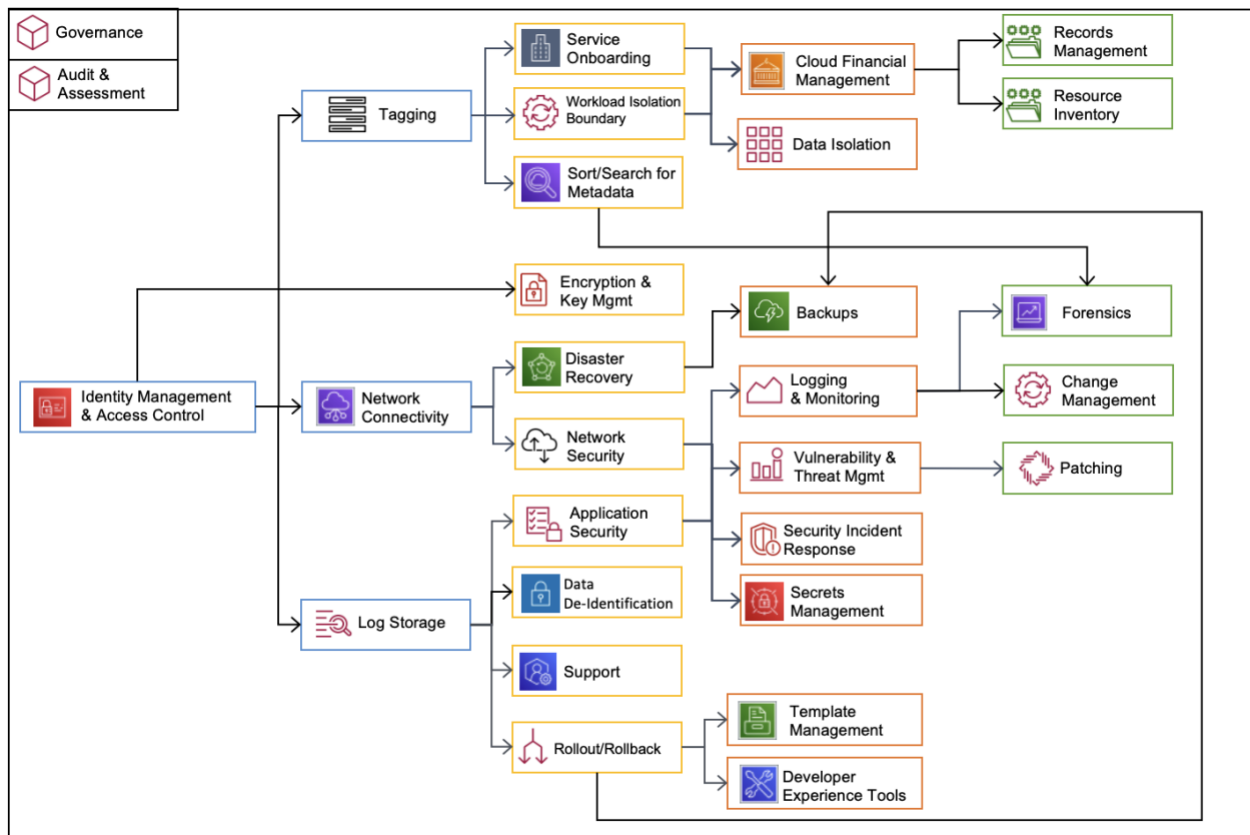
Finance functional area capabilities include:

- **Cloud Financial Management** provides the ability to manage and optimize your variable expense for cloud services. This capability includes near real-time visibility as well as cost and usage analysis to support decision making (e.g., spend dashboards, optimization, spend limits, chargeback, anomaly detection and response). This capability also includes budget and forecasting, to enable you to have a defined cost optimized architecture for your workloads, to select right pricing model, attributing cost of resources to the relevant teams. This enables you to track, notify, and apply cost optimization techniques across your environment and resources. Expense information is centrally managed and consumed, and access to critical stakeholders can be provided for targeted visibility and decision making.
- **Resource Inventory Management** enables visibility and configuration of cloud-based resources that make up an IT-level service or workload. Resources are tracked in the environment along with associated configurations via a system of record (e.g., CMDB for ITSM-managed environments) to enable a larger IT-level system of record for visibility and configuration management of all software, hardware, and firmware resources in the cloud environment.

Working with the capabilities

Each organization’s cloud adoption journey is unique. To successfully build your cloud environment, you need to understand your organization’s current state, the target state, and the transition required to achieve the target state. As you work on your plan to establish your environment, capabilities can help you drive the conversation and decisions across relevant stakeholders (identified by the functional areas for each capability).

The following graph shows a path that you can follow when planning your environment. It’s based on dependencies between capabilities, and can be used to create a project plan for the implementation of capabilities in your environment. In addition to the dependencies shown (via the arrows), some capabilities apply to the overall environment (for example, Governance, and Audit & Assessment).



Capability dependency guided path

Next steps

If you are still exploring the cloud, AWS recommends that you deploy a few proof-of-concepts (POCs) to demonstrate business value to your stakeholders. **If you are ready to start building a cloud environment** to host your workloads on the cloud, this set of defined capabilities can help you build your foundational cloud environment. Before getting started with your cloud adoption, AWS recommends that you complete the following activities, and reach out to your account team for more information:

- Review the list of capabilities and create a timeline for implementing capabilities, accounting for any dependencies.
- Identify the stakeholders in your organization that are responsible for each capability.
- Create an implementation plan and a timeline to build your cloud environment.

As your requirements change, to help you grow your presence in the AWS Cloud, you can use the defined capabilities to build your own approach using your own tools.

Conclusion

This whitepaper introduces a capability-based approach to establishing the foundation for your AWS environment, and helps you identify the relevant stakeholders needed to make important decisions along your journey. The defined *capabilities* in this paper are based on current AWS best practices, and the experience of thousands of customers that have built their foundational environment on the AWS Cloud.

Contributors

Contributors to this document include:

- Alex Torres, Solutions Developer, Amazon Web Services
- Eamonn Faherty, Principal Solutions Developer, Amazon Web Services
- Fabian Labat, Sr. Solutions Architect, Amazon Web Services
- Glenn Dasmalchi, Sr. Manager Tech Leader, Amazon Web Services
- Sam Elmalak, Tech Leader, Amazon Web Services

Further reading

For additional information, refer to:

- [AWS Architecture Center](#)
- [AWS Whitepapers & Guides](#)

Document revisions

Date	Description
November 17, 2021	First publication

Appendix A: Capability structure and example

Capability structure

Definition

The definition includes a high-level description of what the capability will help you enable in your cloud environment.

Scenarios

Scenarios are a set of use cases that expand the capability definition, and detail what parts of your environment the guidance included in the capability solves. Each capability provides a baseline, which establishes the minimum requirement for the capability, and can be expanded and customized to add additional scenarios based on your requirements or as your business needs mature and your AWS presence grows.

Guidance

This section outlines prescriptive recommendations for how the capability should be built in your environment to implement the included scenarios. It also includes responsible stakeholders, and a description of how the capability will work in the overall environment. Additionally, this section includes the people and recommended skillsets necessary to successfully establish the capability in your environment.

Implementation guidance

Each capability provides prescriptive, opinionated AWS guidance to establish the capability in your environment. Runbooks are included to help you operate the capability efficiently in your environment using AWS services.

Capability example - Log Storage

Definition

The Log Storage capability enables you to securely collect and store your environment logs centrally within an immutable storage. This will enable you to evaluate, monitor, alert, and audit access and actions performed on your AWS resources and objects.

Scenarios

- Your cloud team wants to log *individual user access* to resources, and what systems are accessed and actions taken (*Individual user access* also includes access by system administrators and system operators).
- Your cloud team wants to set controls to prevent modification of the related logs.
- Your cloud team wants to set controls to prevent unauthorized access to logs.
- Your cloud team wants to generate logs that can show if inappropriate or unusual activity has occurred.
- Your cloud team wants to store logs in near real-time for resiliency during a determined period of time (matching your governance requirements).
- Your cloud team wants the stored logs to be encrypted at rest.

Guidance

The Log Storage capability primary mapping is to the **Security Functional Area**. This means the **Security team** should be responsible for implementing this capability.

When establishing your capability, the builders owning the implementation will need to receive inputs from the owners of additional functional areas to ensure the proper interlock of the functions in the cloud environment. The list of secondary functional areas required are:

- Operations
- Central IT

Having a separated Log Storage allows you to establish a secure location where the logs become the source of truth to show what is happening in your environment related to security and operations. As your environment expands to accommodate your business needs, centrally aggregating the information will enable you to later build monitoring and observability capabilities, to monitor in near real-time what is happening across your environment.

The Log Storage must be secured, built for resilience, to avoid tampering with the logs, and only accessed by controlled, automated, and monitored mechanisms, based on least privilege access by role. The following controls need to be implemented around the Log Storage to protect the integrity and availability of the logs and their management process. The logs delivered to Log Storage should be encrypted, and the

encryption key access and permissions should also be based on least privilege permissions.

- **Detective controls** should be implemented to alert and remediate the collection of permissions used on the log storage, and to actively monitor access to the logs within the Log Storage.
- **Preventive controls** should be implemented to protect from changes to your configuration and access in your Log Storage, and restricting permissions on your Log Storage.

The Log Storage should also have retention policies, establishing a lifecycle for your logs based on your governance and data retention policy requirements (for example, automatically archiving infrequent access or delete the logs over time to reduce the cost while meeting retention requirements).

Appendix B: Sample timeline

In this section you can find a sample timeline that includes all 30 capabilities that are needed to meet your requirements when establishing a foundational cloud environment on AWS. Enable your teams to work with the capabilities, and start building an environment which initially allows you to deploy experiments and then workloads. As you scale or your business needs evolve, you can assess your established capabilities, and enhance them as necessary to meet your requirements.

