
Management and Governance Cloud Environment Guide

AWS Well-Architected Framework

Management and Governance Cloud Environment Guide: AWS Well-Architected Framework

Copyright © 2022 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

Introduction	1
Journey to cloud-ready environments	2
Setup	2
Build and migrate	2
Operate	3
Manage and govern functions to interoperate	3
Manage and govern with a multi-account point of view	4
Controls	5
Interoperable functions	6
Implementation priorities	7
Define a multi-account strategy	7
Start with AWS Control Tower	7
Review and add preventive and detective controls	7
Select an aggregated view of your guardrails and findings	8
Create a base foundation of capabilities for each of your accounts	8
Control services	9
Integrated controls partners	10
Network connectivity	12
Interoperable functions	12
Implementation priorities	13
Plan your IP address space (IP address management – IPAM)	13
Design network connectivity	14
Define your VPC endpoint and DNS strategy	15
Establish network security	15
Establish network monitoring	16
AWS network connectivity management tools	16
Integrated network connectivity partners	18
Network orchestration	18
Gateway Load Balancer partners	18
AWS Transit Gateway partners, including SD-WAN solutions	18
Identity management	20
Interoperable functions	20
Implementation priorities	21
Establish a centralized identity provider for human identities	21
Define job functions and codify IAM roles	21
Continually collect, review, and refine permissions	21
Manage credential use	22
Source and distribute identity constructs with automation	22
AWS identity services	22
Integrated identity partners	23
Security management	25
Security architecture	25
Automated findings and campaigns	25
Security metrics	26
Security response management	26
Interoperable functions	27
Implementation priorities	27
Design a Well-Architected security environment	27
Choose security tools to match your enterprise needs	28
Analyze and model for threats	28
Automate incident management workflows, findings, and campaigns	28
Select, measure, and continually improve your security metrics	29
AWS security management services	29
Integrated security management partners	30

Service management	33
Provisioning and request management	33
Event and incident management	33
Problem management	34
Resource inventory management	34
Change management	34
Interoperable functions	35
Implementation priorities	35
Integrate provisioning processes with the ITSM tool suite	35
Enable event, incident, and problem management across your environment	36
Identify accounts, environments, and resources that require asset tracking	36
Align change request procedures and policies for rapid cloud deployment	36
Connect your ITSM system of record tooling to AWS	36
AWS service management tools	37
AWS Service Management Connectors	38
Integrated service management partners	38
Monitoring and observability	40
Interoperable functions	40
Implementation priorities	41
Collect, aggregate, and protect event and log data	41
Build capabilities to analyze and visualize log events and traces	41
Add detection and alerts for anomalous patterns across environments	41
Define, automate, and measure response and remediation	42
AWS observability tools	42
Integrated observability partners	43
Cloud Financial Management	45
Organize and report with user-defined methods	45
Manage billing and control costs	45
Use license management	45
Plan with flexible budgeting and forecasting	46
Select a unit metric to support your business	46
Optimize costs with pricing and resource recommendations	46
Interoperable functions	46
Implementation priorities	47
Enable Cloud Financial Management	47
Tag, track, and monitor resource costs across their lifecycle	47
Establish mechanisms for cost governance	47
Continually optimize for cost efficiency	48
AWS Cloud Financial Management services and tools	48
Integrated Cloud Financial Management partners	49
Sourcing and distribution	51
Interoperable functions	51
Implementation priorities	52
Implement a hub and spoke catalog model	52
Curate templates for reuse	52
Apply default parameters for reuse	53
Implement a lifecycle management and version distribution system	53
Identify a robust tagging strategy	53
Manage entitlements	54
Enable Private Marketplace	54
Integrate with procurement systems	54
AWS sourcing and distribution tools	54
AWS sourcing and distribution partners	56
Conclusion	57
Contributors	58
Document history	59
Notices	60

Definitions: Vocabulary	61
Application	61
Compliance	61
Governance	61
Guardrail	61
Management	61
Monitoring	62
Operations	62
Oversight	62
Resource	62
Supervision	62
AWS glossary	63

Management and Governance Cloud Environment Guide

Publication date: **November 22, 2021** ([Document history](#) (p. 59))

Customers of every size and industry type are moving to the cloud to increase their security, agility, cost efficiency, scalability, and ability to deploy more easily. Many customers are asking for our guidance to help them ensure their AWS environments meet those requirements. The Amazon Web Services (AWS) Well-Architected Management and Governance Cloud Environment Guide (M&G Guide) provides clear guidance for you to follow. This guidance includes answers to key questions, recommended guardrails, and identifying AWS services and solutions from AWS Partners to help you build development, test, and production workloads at scale regardless of the stage of cloud adoption you are in. This document can be read in its entirety or by individual section depending on the focus of the reader.

The [AWS Cloud Adoption Framework \(AWS CAF\)](#) helps you develop and run efficient and effective plans for your cloud adoption journey focused on people and processes. The M&G Guide builds on the CAF principles, offering prescriptive guidance with a focus on the technology. Cerner Corporation, a global health platform and technology company, completed an initiative in collaboration with AWS Professional Services to re-shape their cloud capabilities to support a growing, diverse, and global customer base:

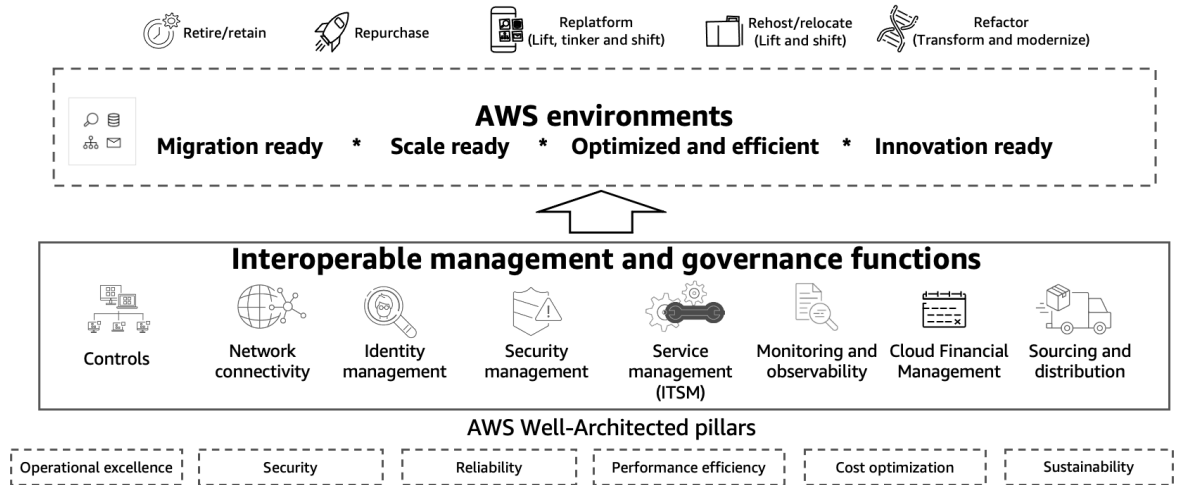
"Cerner has been operating about 50–100 accounts in AWS for years, but our strategy was very decentralized with regard to governance. We have business drivers to enable HITRUST compliance, which we combined with many other frameworks to create our Cerner Controls Framework (for example, compliance requirements). These requirements drove us to rethink how we applied governance principles in our cloud operating model through centralized governance. We could have moved faster and delivered value to our business in a more deliberate way had the AWS Management & Governance Cloud Environment Guide been available at the time. Every section would have created value for us as it really plants a flag in the ground to help distill AWS Best Practices from all the various people and places into a clear starting point."

- Eric Wright, Senior Director Cloud Engineering, Cerner Corporation

- Phil Brown, Director & Principal Engineer Cloud Engineering, Cerner Corporation

Based on our experiences from thousands of successful migrations, the M&G Guide helps decision makers, cloud, networking, and security architects configure their AWS environments to prepare for scale and evaluate if their environment is configured properly. The M&G Guide includes the following:

- Description of each of the management and governance functions.
- Information on how the functions interact and interoperate with each other to provide efficient management and governance.
- Detailed implementation priorities helping you to know what steps to take, and in what order.
- Recommended AWS services for each function.
- AWS Partner solutions available in [AWS Marketplace](#) that support multi-account environments and work with [AWS Control Tower](#).
- Implementation guidance as architectural diagrams, guides, and product videos.
- Aligned offerings and delivery kits from [AWS Professional Services](#).
- Turnkey complementary solutions and consulting services from [Built on Control Tower - AWS Partners](#).



How to prepare AWS environments

Journey to cloud-ready environments

A cornerstone of a successful, cost-efficient, secure, and compliant cloud strategy is to emphasize proactive management and governance. Incorporating best practices from the M&G Guide helps you grow with AWS, whether you are in the setup, migrate, or operate phase of the cloud journey. The progression along that journey includes the adoption of management and governance capabilities as you mature. The three phases of a typical cloud journey are described in the following sections.

Customer use case examples for this guide include those that are just getting started with AWS, those that are considering an expansion into a multi-account experience, or those that are planning an expansion, such as the launch of new applications or a data center migration.

Setup

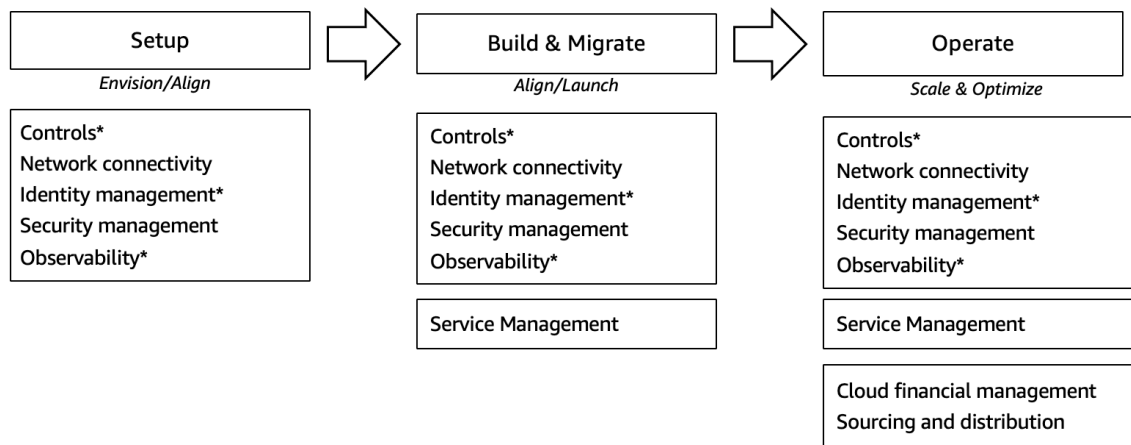
Getting started with AWS, you should configure identity management, logging, monitoring, observability, network connectivity to on-premises, and integrate security capabilities to their existing solutions. Gain a head-start on these capabilities by using AWS Control Tower to provision a landing zone embedded with controls. This is extended with basic network isolation, a base set of identities, and extending incident management and security capabilities to the new environments. In this phase, you begin building cloud-ready environments tuned to your enterprise needs. This lets you scale your management and governance functions alongside your workloads.

Build and migrate

In this phase, you want to extend and enhance your management and governance functions. This includes, extending network isolation boundaries, configuring further environment and workload-based controls, tuning change and incident management, and updating your observability to accommodate application-specific insights. Whether you are using a migration factory to quickly and efficiently migrate applications or workloads, or you are beginning to build out larger sets of applications or workloads, you should also add integration to your service management capabilities and enhance your security management tooling.

Operate

Evolving interoperability of the management and governance functions give you greater operational efficiency as you continue migrating, building, or modernizing your workloads. This phase typically includes the addition of full sourcing and distribution functions for your infrastructure templates or software solutions. Proactively using financial insights spanning across your workloads, accounts, and environments also position you for accelerating innovation activities.



*baseline functions provided by AWS Control Tower

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

Journey to AWS environments

Manage and govern functions to interoperate

The eight management and governance functions, supported by AWS services and AWS Partner solutions, should interoperate together in an informative relationship to help you manage and govern your environments at scale. Outputs from these functions are used to inform or integrate with other functions.

As an example, a report in [AWS Cost Explorer](#) shows a spike in month-to-month EC2 usage in a development environment. Further investigation leads you to discover that a new team has been launching `m5.16xlarge` instances, but using less than two percent of the CPU capacity. AWS Cost Explorer insights reveal that the development environment did not require the same instance size as test or production. As an input to controls, you can define a detective AWS Config rule in the development environment to alert on unapproved `m5.16xlarge` instances. In addition, you can permit builders to only self-service provision nano instance types in the development environment by using template constraints from AWS Service Catalog, or you can assign an AWS Organizations service control policy to restrict the instance types that can be launched in that environment.

With this interoperability example, you can tune the financial operations of your IT functions to automate cost controls, which permit you to continually evaluate mechanisms that can reduce your AWS costs. Although similar manual mechanisms might be effective, they are not as efficient as you scale further workloads on AWS. Throughout this M&G Guide, you will see the additive benefits of an interoperable and automated foundation of the proposed eight capabilities in your AWS environments.

Manage and govern with a multi-account point of view

AWS helps enable you to experiment, innovate, and scale more quickly, while providing flexible and secure cloud environments. An AWS account provides natural security, access, and billing boundaries for your AWS resources. The AWS account as a boundary helps you to achieve resource isolation as described in the [Security Pillar whitepaper](#). The Security Pillar specifically recommends the following best practices: separate workloads using accounts, secure AWS accounts, manage accounts centrally, set controls centrally, configure services and resources centrally.

The multi-account strategy prescriptive guidance provided in the [Organizing Your AWS Environment Using Multiple Accounts whitepaper](#) describes specific mechanisms to organize accounts. In addition, it describes how to apply a consistent set of controls so that you can efficiently manage your cloud assets. In AWS, accounts are a hard boundary. Account-level separation is recommended for isolating production workloads from development and test workloads. For instance, sandbox environments might need a different set of controls, network, change processes, and financial limits compared to other environments. Using this strategy helps you to centrally manage resources, permissions, and security standards across environments and accounts, improving your operational efficacy.

The M&G Guide complements the Security Pillar and the multi-account strategy to further define a set of eight foundational capabilities required to prepare your environments and operate efficiently in the AWS Cloud. You can start automating provisioning your accounts following this strategy with [AWS Control Tower](#). With this service you will provision a landing zone from your home Region, and deploy further accounts following your multi-account strategy.

The [Organizing Your AWS Environment Using Multiple Accounts whitepaper](#) recommends that you build a multi-account strategy using account boundaries to separate workloads. However, it is important to evaluate and plan your account management with automation and operational capacity in mind. That is, your accounts should employ the least privilege access, and provide boundaries to limit the effect of workload failures. Do not create more accounts than are feasible to operationally manage or scale. Furthermore, as you scale, consider reviewing your service quotas and deployment latencies when performing actions on a large number of accounts.

Controls

A *control* is a means of mitigating or detecting an issue that is a consequence of risk being realized, while *guardrails* are a technical implementation to meet those controls. More specifically, controls provide instruction for configuring resources to mitigate or address specific risks. We recommend you start their multi-account environment with AWS Control Tower, which offers predefined baseline preventive and detective guardrails that can be enabled at an environment, resource, account, or Organizational Unit (OU) level. Guardrails are an essential part of managing your AWS environments as they provide an automated way to deliver on policy intentions. Two kinds of guardrails exist: preventive and detective.

Preventive guardrails enforce specific policies to help ensure that your accounts operate in alignment to compliance standards, and disallow actions that lead to policy violations. Control what your AWS accounts can do by only permitting specific services, Regions, and service actions at the appropriate level. AWS Organizations provides service control policies (SCPs) to apply permission guardrails at the organization, organizational unit, or account level. For example, you can apply an SCP that restricts users from launching resources in Regions that you have not explicitly allowed. Or, you can create an SCP to [Disallow creation of access keys for the root user](#). This would help secure your AWS accounts by disallowing creation of access keys for the root user, thereby reducing risk of unrestricted access to all resources in the account.

Detective guardrails detect and alert on unexpected activity and noncompliance of resources within your accounts, such as policy violations. These are helpful in alerting when something requires remediation (either manual or automated). For example, you can create an AWS Config rule to [Detect whether public write access to Amazon S3 Buckets is allowed](#). This rule detects whether public write access is permitted to Amazon S3 buckets. You can use this alert to initiate remediation with a Systems Manager automation document, or a procedure outlined in your ITSM tools.

Selecting the right guardrails for your environments is an important step in managing and governing your resources across AWS. Managing configuration compliance for any IT service is typically required to ensure security (confidentiality, integrity, and availability) of your data. This includes reference to standards and regulatory requirements, individual policy definitions, risk management processes, remediation workflows, and exception procedures. To select the correct guardrails, we recommend building a portfolio from compliance frameworks, risk management processes, and AWS Best Practices to match the needs of your specific organization.

Compliance-based controls are often included in the compliance and framework specifications. As a reference, you can identify risk-based controls with guidance from the [National Institute of Standards and Technology \(NIST\) CyberSecurity Framework](#). The [NIST Risk Management Framework \(RMF\)](#) defines an approach for how to select controls, and the [Factor Analysis of Information Risk \(FAIR\)](#) defines a process for how to calculate your risk profile and measure risk reduction efforts related to controls.

We recommend aggregating the detective guardrails implemented through AWS Config Rules into conformance packs so that they can be easily provisioned across your AWS environments. A key feature of conformance packs is that they are immutable—individual rules cannot be changed outside of the pack in which they were deployed, regardless of access or account permissions. In addition, if the pack is deployed by an organization's management account, it cannot be modified by the organization's member accounts. This approach provides you with an additional level of security and certainty when managing compliance across your environments. It also enables aggregated reporting, as compliance summaries can be reported at the pack level. You can start with the [AWS Config conformance samples](#) we provide, and customize as you see fit. When using multiple conformance packs, determine if duplicate rules are being used as this might have cost implications across your environments.

AWS has provided a sample [set of Config Conformance Packs](#) that align to specific services and compliance frameworks. The sample templates, including those related to compliance standards and

industry benchmarks, are not designed to ensure your compliance with a specific governance standard, but rather are designed to help you form part of it. They cannot replace your internal efforts or ensure that you will pass a compliance assessment.

AWS Control Tower offers a simplified way to automate the provisioning of accounts that are preconfigured with baseline guardrails. Preventive guardrails deployed by AWS Control Tower are implemented via service control policies (SCPs). Detective guardrails deployed by AWS Control Tower are implemented using AWS Config Rules and AWS Lambda functions. In addition to the baseline guardrails found in SCPs and AWS Config Rules, guardrails can also be found in other M&G Guide capabilities. Some examples would be IAM policies, network security groups, NACLs, budget alarms, and constraints on AWS Service Catalog products.

BPX Energy, a BP company, used AWS Control Tower to establish their AWS environment with controls enabling them to deploy detective controls with AWS Config and preventive controls with AWS Organizations SCPs via AWS Control Tower. *“The key benefits of adopting AWS Control Tower included enhancing BPX Energy’s security posture, enabling enterprise governance at scale, and providing increased scalability.”* Grant Matthews, Chief Technology Officer, BPX Energy. Learn how BPX’s implementation further aligns to the controls function described in this M&G Guide by reviewing their [case study](#).

Both AWS Control Tower and AWS Security Hub continually evaluate all of your AWS accounts and workloads and provide dashboards so you can quickly identify areas of deviation from established guardrails. These insights can be used to improve and maintain your security posture across your AWS environments. For instance, AWS Control Tower applies a mandatory set of guardrails during the provisioning and management of your landing zone that indicate how your landing zone is compliant with best practices. AWS Security Hub provides a mechanism to deploy and categorize security-focused detective guardrails. This mechanism allows you to aggregate, organize, prioritize, and automate the remediation of the findings across your multi-account environment. There is an inclusive set of Security Hub standards that can be used to align to your specific compliance and security framework. These include [AWS Foundational Security Best Practices](#), the [CIS AWS Foundations Benchmark](#), and the [Payment Card Industry Data Security Standard \(PCI DSS\)](#). You can investigate findings via the AWS Security Hub integration with Amazon Detective, and you can build automated or semiautomated remediation actions using the Amazon EventBridge integration.

Review your use of detective guardrails to identify and remove duplicative detection efforts when using one or more of these frameworks. Also, as you use AWS services, remain aware of the inherent quotas being imposed. For example, AWS Control Tower describes its [limitations and service quotas within the service documentation](#). When you review these quotas, it is important to choose where to use preventive versus detective guardrails to work within the service quotas while still meeting your compliance needs.

Interoperable functions

The eight management and governance functions work together and interoperate to reduce complexity. Outputs from these functions are used to inform or integrate with other functions.

For controls, this includes:

- Inspect and protect out of band **Networking connectivity** changes.
- Access permissions and controls federated with your **Identity management** provider.
- Controls findings that initiate campaigns and playbooks in **Security management** operations.
- Integrated change, provisioning, and remediation capabilities with service level objectives for each control for your **Service management** framework.
- **Monitoring and observability** defined for both aggregated and granular views of each control and guardrail.
- Financial and process controls aligned to your **Cloud Financial Management** best practices.

- Infrastructure as code templates for your guardrails that are **Sourced** and **distributed** in a hub and spoke pattern for your multi-account strategy.

Implementation priorities

Using a centralized mechanism like AWS Control Tower to create accounts that are pre-configured for compliance can help you adjust to your changing scale needs. Having a multi-account strategy helps raise your security posture with the necessary separation of workloads and networks through logical and physical boundaries. As such, the following controls solutions should be prioritized:

Define a multi-account strategy

AWS recommends that you define a multi-account strategy that considers scale and operational efficiency concerns. This means that you should separate out your workloads into a logical pattern that best meets your operational needs. AWS provides [prescriptive guidance](#) that suggests you start with a foundational set of accounts to accommodate centralized and decentralized capabilities in your enterprise. You can centralize governance for distributed and autonomous teams using multiple AWS accounts, which lets you delineate at security, financial, and operational levels.

Start with AWS Control Tower

[Enable a landing zone using AWS Control Tower](#) in a new or existing management account. AWS Control Tower creates a secure, multi-account environment with an embedded set of default guardrails. AWS Control Tower automatically enables AWS Config in the AWS Control Tower Regions, with configuration history and snapshots delivered to an Amazon S3 bucket located in a centralized Log Archive account. It also provides the ability to add guardrails for each organizational unit (OU) in your AWS Organization. The landing zone includes a preconfigured security OU with an audit and log archive account provisioned. This includes guardrails to prevent unauthorized changes to the security baseline in your audit account. CloudTrail logs are encrypted (using AWS KMS) and enabled in all provisioned accounts with SCPs to prevent their modification. By default, a sandbox OU is provisioned for your use. Review the [best practices for organization unit](#) guidance to determine the multi-account strategy and OU structure that will support your unique enterprise needs.

Separating OUs by regulatory and SDLC environments is a commonly used pattern. Workload OUs are used for accounts that host your AWS resources to support your applications with the right policies applied. AWS Control Tower allows you to allow or deny the use of AWS Regions across your environments.

Review and add preventive and detective controls

AWS Control Tower uses AWS Organizations service control policies (SCPs) to provide preventive guardrails. SCPs define the guardrails or limits that IAM roles and users can have in the accounts located within the OU. Review the [strongly recommended](#) and [elective](#) detective guardrails AWS Control Tower provides, and choose which guardrails to apply. Use the AWS Security Foundations best practices in AWS Security Hub to identify controls that apply to your enterprise, and add any specific open standard controls required for your workloads. Create additional preventive controls as required, and group them by OUs to align them to your multi-account strategy.

Note

SCPs have [service quotas](#) in the size and number that can be applied. Carefully consider these quotas as you design your controls strategy.

Package your detective controls such that they can be deployed easily as you create or update your accounts. There are a variety of AWS Config Conformance Packs available to apply common sets of AWS Config rules to meet open standards and best practices. For instance, there is a sample pack that

includes the [Best Practices for the Well-Architected Security Pillar](#), which can provide a starting list of best practice rules to provision. Use these conformance packs to choose and add further guardrails to your environments.

Annotate and prioritize your detective guardrail findings so that they can be remediated in accordance with your security and compliance frameworks. Use automation to detect out of policy provisioning of resources. In addition, set and measure service level objectives alongside updating your runbooks and playbooks.

Select an aggregated view of your guardrails and findings

Centrally view the resource configuration and compliance data recorded in your observability findings. AWS Security Hub is a security and compliance service that provides security and compliance posture management, as a service. It uses AWS Config and AWS Config rules as its primary mechanism to evaluate the configuration of AWS resources. AWS Config rules can also be used to aggregate and evaluate resource configuration. Other AWS services, such as AWS Control Tower and AWS Firewall Manager, also provide an aggregated view of controls in their console view. Regularly review aggregated views of guardrails to alert you on any deviation of expected controls in your environments.

Create a base foundation of capabilities for each of your accounts

You can [provision AWS Control Tower accounts in an automated, batch fashion](#) by calling the AWS Service Catalog APIs. As you provision new accounts, use [Customizations for AWS Control Tower](#) to add in a base set of services and functions required for each account. This will include capabilities across each of the eight M&G Guide functions as well as tagging and support information. For example:

From a networking perspective, determine which VPC structure is provisioned and associate it to a central hub-spoke pattern. Add in any necessary network constructs that vary by account type (firewall, NAT gateway, etc.)

For identity management, after Control Tower is configured for single sign-on integrated with your federated user solution, you can provision additional roles and policies. These might include permissions boundaries to be distributed to member accounts.

Make sure that your monitoring and observability capabilities are updated as new accounts are provisioned. Workloads should be aligned to the environment logging strategies that describe which logs to locate where, and how to appropriately integrate log aggregation.

You might need to register new accounts with your security tools (SIEM, GuardDuty, Security Hub, etc.), or deploy security capabilities to specific accounts (XDR, CSPM, etc.).

As you create new accounts, it is important to align them with your service and incident management capabilities. New accounts should be integrated with your service management solution using native connectors configured to integrate those solutions with AWS. Update your playbooks and runbooks as appropriate.

[Tag policies](#) and tag libraries help create consistent tagging and can be used for common processes including Cloud Financial Management (CFM). Consider distributing new financial guardrails to detect deviations from expected budgets in spoke accounts. Allocate appropriate support levels for each account using the [AWS Support API](#).

Software assets managed in AWS Service Catalog as portfolios can be shared with users in one or more AWS accounts in a hub and spoke pattern. Using Private Marketplace and private offers, curate an assortment of third-party solutions and distribute them alongside your infrastructure as code templates.

Define which base set of resources should be directly provisioned or made available as a self-service model in each of your spoke accounts as they are created with solutions such as the Customization Framework for Control Tower.

Control services

The following AWS services can be used to help you follow the guidance provided by the M&G Guide:

[AWS Organizations](#) includes service control policies (SCPs) that you can use to provide centralized control over all accounts in your organization. You can configure an SCP to define a guardrail, or set a limit, on the actions that the account's administrator can delegate to the users and roles for the affected accounts. The administrator must still attach identity-based or resource-based policies to IAM roles, or to the resources in your accounts to actually grant permissions. The [effective permissions](#) are the logical intersection between what is allowed by the SCP and what is allowed by IAM and the resource-based policies.

[AWS Control Tower](#) complements AWS Organizations by implementing preventive and detective controls as you provision accounts. You can quickly set up and configure a new AWS environment, automate ongoing policy management, and view policy-level summaries of your AWS environments.

[AWS Security Hub](#) provides a single place that aggregates, organizes, and prioritizes your security alerts, or findings, from multiple AWS services. These include Amazon GuardDuty, Amazon Inspector, Amazon Macie, AWS Firewall Manager, AWS Systems Manager Patch Manager, AWS Config, AWS IAM Access Analyzer, as well as from many AWS Partner Network (APN) solutions.

[Amazon GuardDuty](#) is a threat detection service that continually monitors for malicious activity and unintended behavior to protect your AWS accounts, workloads, and data stored in Amazon S3. Amazon GuardDuty uses machine learning, anomaly detection, and integrated threat intelligence to identify and prioritize potential threats. GuardDuty analyzes tens of billions of events across multiple AWS data sources, such as AWS CloudTrail event logs, Amazon VPC Flow Logs, and DNS logs.

[Amazon Macie](#) gives you constant visibility of the data security and data privacy of your data stored in Amazon S3. Macie automatically and continually evaluates all of your S3 buckets and alerts you to any unencrypted buckets, publicly accessible buckets, or buckets shared with AWS accounts outside those you have defined in the AWS Organizations.

In [AWS Config](#), you can create and manage singular rules (detective controls), or group them as conformance packs. [AWS Config conformance packs](#) help you manage configuration compliance of your AWS resources at scale – from policy definition to auditing and aggregated reporting – using a common framework and packaging model. Additionally, AWS Config conformance packs enable you to simplify compliance reporting, as it is now reported at a new level - the pack level alongside the detailed view for each individual rule and resource level.

The [AWS Config Conformance Pack Sample Templates](#) help you create your own conformance packs with different or additional rules, input parameters, and remediation actions that suit your environment. The sample templates, including many related to compliance standards and industry benchmarks, are not designed to ensure your compliance with a specific governance standard. They cannot replace your internal efforts or ensure that you will pass a compliance assessment.

[AWS Audit Manager](#) helps you continually audit your AWS usage by simplifying how you assess risk and compliance with regulations and open standards. Audit Manager provides a fully customizable framework that automates evidence collection, simplifies the tracking of chain of custody for evidence, and manages evidence security and integrity.

If you would like support implementing this guidance, or assisting you with building the foundational elements prescribed by the M&G Guide, we recommend you review the offerings provided by [AWS Professional Services](#) or the AWS Partners in the [Built on Control Tower program](#).

If you are seeking help to operate your workloads in AWS following this guidance, [AWS Managed Services \(AMS\)](#) can augment your operational capabilities as a short-term accelerator or a long-term solution, letting you focus on transforming your applications and businesses in the cloud.

Integrated controls partners

The M&G Guide recommends you consider the following questions when choosing an AWS Partner solution for controls:

- Does it integrate with lifecycle events for AWS Control Tower?
- If controls are provided, are they updated on a regular basis?
- Does it support multiple AWS Regions?
- Can it be provisioned from an infrastructure as code template that is distributed from a service catalog?
- Does it integrate with an observability solution?
- Can changes be tracked automatically, or integrated to your service management tool?

The following controls partners have built integrations with AWS services, and are available to be provisioned from AWS Marketplace:

[Check Point CloudGuard](#) is a comprehensive cloud native security platform for visibility, workload protection, and posture management of cloud workloads and services. CloudGuard provides visualization of cloud assets, including network topology, and firewalls; comprehensive compliance management including automated continuous compliance to help assess and enforce regulatory requirements and security best practices; open-source auto-remediation to accelerate the resolution of dangerous misconfigurations and enforce compliance; automated reversion of unauthorized modifications to cloud accounts; and just-in-time privileged elevation with out-of-band authorization for IAM actions. Checkpoint findings are also integrated to AWS Security Hub.

[CloudCheckr CMx](#) is a unique, end-to-end governance solution that enables users to optimize security and monitor their compliance, while enacting self-healing automation to remediate security vulnerabilities and compliance gaps. CloudCheckr provides users with hundreds of security and performance optimization recommendations and dozens of options to fix security and resource utilization issues automatically anytime they are detected.

[Cutover](#) is a work orchestration and observability platform that allows teams to plan, orchestrate, and analyze complex workflows. It integrates with AWS Control Tower to accelerate your migration, drive effective governance, reduce risk, and help ensure standardization. The automation runbooks in Cutover work with existing toolsets to allow teams to achieve full visibility, control, and streamlined communications across their multi-account AWS environments.

[Flexera](#) offers a powerful policy engine that enables your cloud governance teams to manage and control cloud use with out-of-the-box and custom policies to automate governance of costs, operations, security, and compliance.

[Kion](#) is a comprehensive enablement software solution that delivers visibility and control of cloud workloads. [Kion](#) provides out-of-the box compliance checks to help enterprises auto-align with established standards like NIST and CIS, and delivers the flexibility to create custom checks. Auto-remediation and integrations with AWS Security Hub are also available. [Kion](#) allows enterprises to manage their cloud presence at scale with automation and orchestration, financial management, and continuous compliance.

[Palo Alto Networks Prisma Cloud](#) unifies Cloud Security Posture Management (CSPM) and workload protection (CWPP) into a single cloud native security platform. Continually monitor your environments and immediately enforce governance with hundreds of pre-built policies. Prisma Cloud ingests AWS APIs

and sources threat intelligence from over 30 feeds to provide comprehensive visibility. Risk-ranked alerts prevent remediation fatigue and one-click compliance reporting helps ease auditing across even the most complex distributed environments. Prisma findings are also integrated to AWS Security Hub.

[Sonrai Dig](#) is an enterprise cloud security platform providing complete visibility across all multi-account AWS environments. Dig's CSPM capabilities provide continuous, audit-based monitoring giving comprehensive visibility and control over the security posture of every cloud resource and identity. Detect drift and misconfigurations on identities, data stores, or a particular cloud resource to help ensure that compliance is baselined, monitored, and met.

[Trend Micro Cloud One - Conformity](#) is a cloud security posture management service that helps you fulfill your side of the shared responsibility model with continual security, compliance, and governance checks. With almost 1,000 cloud configuration checks out of the box that are mapped back to industry best practices, such as the AWS Well-Architected Framework, SOC2, NIST, CIS, PCI DSS, GDPR, and HIPAA, it provides a consistent approach to building cloud architectures that can scale over time. Infrastructure as code (IaC) template scanning also ensures deployment of the most secure and compliant templates aligned with industry best practices when building in the cloud.

Network connectivity

Workloads often exist in multiple locations or environments, both publicly accessible and private. Managing networks in AWS might require connecting many AWS-hosted VPCs from many accounts to specific enterprise networks, and to the internet. Your network strategy must allow for the interoperability of workloads while also aligning to your security architecture. The careful planning and management of your network design forms the foundation of how you provide isolation and resource boundaries within your workload. We think of network connectivity in three different groupings: connectivity between your on-premises network and your AWS environment, connectivity to and from the internet, and connectivity across your AWS environments—primarily between VPCs.

Where connectivity between VPCs is required, the M&G Guide recommends a hub and spoke model for your network design to connect to your existing environment. Intra-application connectivity requires multiple account network patterns that can be reusable for scale. Account types can include sandbox accounts that might require a separate network than the network used for your workload accounts. Regulatory requirements might require you to separate production data into distinct accounts and keep it separate from research and development activities in your sandbox accounts. To reinforce your data governance, you might restrict access using distinct network boundaries, along with specific controls. These boundaries could include controlling traffic with security groups and NACLs, implementation of firewalls, and implementing limited route configurations. Beyond data governance, your workload accounts might need further network refinement for regulated and non-regulated workloads.

Have a mechanism to enforce the use of non-overlapping private subnets when provisioning new accounts and VPCs in your multi-account framework. This automation should also encompass the definition of which network controls and patterns are implemented as you provision (and update) your AWS accounts and workloads. This automation would include definitions of which Regions are included and excluded from your network, as well as which mechanisms of access are allowed in your environments. Using AWS Control Tower, you can select a guardrail to detect if SSH or RDP is enabled for internet connections within your network, while specifically defining which Regions are allowed for the account and related VPC to operate. SSH and RDP traffic can also be restricted through security groups and NACLs.

Define and catalog your VPC in an infrastructure as code template such as AWS CloudFormation. Doing so will allow you to automate its provisioning as well as help with the necessary distributions of future version updates. AWS Control Tower provides a default VPC, or you can use the [Scalable VPC Architecture](#) from AWS Quick Starts as a building block for your own deployments. This template is also available within your console in the [AWS Service Catalog Getting Started Library](#).

Interoperable functions

The eight management and governance functions, supported by AWS services and AWS Partner solutions, work together and interoperate to reduce complexity. Outputs from functions are used to inform or integrate with other functions.

For network connectivity this includes:

- A specifically defined set of **Identity and permissions** to make changes to the networks.
- Provisioning each network with the appropriate **Controls** defined within the infrastructure as code template.
- Embedded integration with **Security management** including playbooks and runbooks.
- Integrated change, provisioning, and remediation capabilities for your networking capabilities for your **Service management** framework. This would include defining support escalation paths and dependencies, runbooks, and playbooks for each network as well.

- Complete **Monitoring and observability** with specific network logging, and identification of any necessary changes to the network design based on the behavior captured.
- Networking components should be included in the total cost of application management calculations for your business cases within **Cloud Financial Management**.
- Purchased third-party solutions or custom-built networking solutions **Sourced and distributed** across environments

Implementation priorities

Network connectivity is typically implemented in the early phases of a cloud journey. As you evolve your network strategy, the following items should be prioritized.

Plan your IP address space (IP address management – IPAM)

Similar to private networks, VPCs typically use private or RFC 1918 IPv4 space. However, you can also use publicly routable non-RFC 1918 CIDR blocks for your VPC. Carefully plan the IP address space that you will be allocating to your VPCs, particularly if you are using an IPv4 range. The best practice for IPv4 planning is to first allocate a non-overlapping and contiguous address block. Subdividing address space into subnets based on attributes like environment or AWS Region helps you to create separate network boundaries more easily. You might also consider other CIDR grouping approaches based on regulatory requirements or the sensitivity of their workloads. This approach can simplify routing, security policies and the ability to query logs.

In the following example, VPCs have been assigned a contiguous IP space aligned by VPC environment to simplify application of security groups and NACLs.

Example VPC CIDR ranges

	Dev VPCs	Test VPCs	QA VPCs	Prod VPCs
AWS Region 1	10.0.0.0/16	10.64.0.0/16	10.128.0.0/16	10.192.0.0/16
AWS Region 2	10.1.0.0/16	10.65.0.0/16	10.129.0.0/16	10.193.0.0/16

- 10.0.0.0/15 represents all Dev VPCs
- 10.64.0.0/15 represents all Test VPCs
- 10.128.0.0/15 represents all QA VPCs
- 10.192.0.0/15 represents all Prod VPCs

In the next example, VPCs have been assigned contiguous IP space aligned by AWS Region to simplify routing.

Example VPC CIDR ranges

	Dev VPCs	Test VPCs	QA VPCs	Prod VPCs
AWS Region 1	10.0.0.0/16	10.1.0.0/16	10.2.0.0/16	10.3.0.0/16
AWS Region 2	10.4.0.0/16	10.5.0.0/16	10.6.0.0/16	10.7.0.0/16

- 10.0.0.0/14 represents all VPCs in AWS Region 1
- 10.4.0.0/14 represents all VPCs in AWS Region 2

If insufficient IP space is a concern, consider [VPC sharing](#) to simplify IPv4 address allocation, and preserve scarce IP addresses. This approach gives you the ability to centralize control of network maintenance while still granting builders the ability to self-provision VPC based resources. AWS PrivateLink can also help alleviate IP exhaustion and overlap by enabling the creation of services in your VPCs that can be consumed through a PrivateLink endpoint with traffic flowing across Amazon's private network. Service consumers don't have to worry about overlapping IP addresses, arrange for VPC peering, or use a Transit Gateway. If exhaustion of IP space is still a concern, you can evaluate alternative solutions that rely on [Private NAT and TGW](#) to allow for communication between VPCs with overlapping CIDR ranges.

The IPv6 address space is much larger than the IPv4 address space, so it's not currently at risk of exhaustion. However, if you plan to [bring your own IPv6 space into AWS](#), it's still a good practice to structure it in a similar fashion to IPv4.

Continually review and refine your network isolation boundaries and perform impact analysis for any proposed network changes. VPC design would be incomplete without a scalable subnet design. As with other management and governance functions, consider the operational complexity when designing or refining the allocation of subnets and be mindful of allocating too many subnet tiers. Because individual subnets cannot span multiple Availability Zones (AZs), deploy workloads to multiple subnets across multiple AZs to allow for workload resiliency using capabilities like Auto Scaling groups, load balancers, and services that span AZs. Your subnet design will likely include a combination of public and private subnets. Public subnets are associated with a route table that has a route to an Internet Gateway, while private subnets do not have a route to an Internet Gateway and are typically associated with a NAT Gateway if they require internet access.

Design network connectivity

We think of network connectivity in three different areas:

- Connectivity between your on-premises network and your AWS environments
- Connectivity to and from the internet
- Connectivity across your AWS environments.

For on-premises connectivity, customers typically start with a VPN. They add additional VPNs or convert to Direct Connect and add resilience and bandwidth as time, maturity, and requirements progress. It is also common for customers to configure VPN over a Direct Connect connection to achieve consistent levels of throughput and encryption algorithms that protect data in transit. AWS Direct Connect also offers IEEE 802.1AE MAC Security Standard (MACsec) encryption for 10Gbps and 100Gbps Dedicated Connections at select locations to secure your high-speed, private connectivity to the cloud.

Decide whether to configure internet traffic in a centralized or distributed manner, depending on your enterprise needs. You might choose to centralize inbound or outbound internet traffic with a hub and spoke model using AWS Transit Gateway or an AWS Partner solution, or distribute internet traffic flows via appropriate VPCs in their environment. Establish internet connectivity by implementing internet gateways, public subnets as well as NAT gateways. Review the whitepaper on [building scalable multi-VPC architectures](#) to help you decide which pattern best fits your requirements.

For connectivity across your AWS environments, connect VPCs both within and across AWS Regions using a transit gateway hub and spoke model. [The Serverless Transit Network Orchestrator solution](#) automates the process of setting up and managing transit networks in distributed AWS environments. It creates a web interface to help control, audit, and approve (transit) network changes. This includes establishing peering connections between transit gateways to extend connectivity and build global networks spanning multiple AWS Regions.

Define your VPC endpoint and DNS strategy

To establish private connectivity from your VPC to supported AWS services, use [VPC Interface endpoints](#). An interface endpoint is an elastic network interface with a private IP address from the IP address range of your VPC subnets. Interface endpoints can be deployed across multiple AZs for resiliency. Interface endpoints serve as an entry point for traffic destined to a supported AWS service. In addition, add [endpoint policies](#) to control access from the endpoint to the specified service. Amazon Virtual Private Cloud documentation includes an updated [list of services](#) that support VPC Interface endpoints.

When deploying your VPC endpoints, consider two approaches. One approach is to centralize multiple endpoints in a single VPC reachable from other VPCs using AWS Transit Gateway. This approach allows you to lower the overall endpoint cost but also means that access policies and endpoint capacity would be shared between multiple VPCs. A second approach is to use interface endpoints for relevant services in each VPC. Access is localized and security policies and performance are scoped and consumed by a single VPC. It is important to consider that costs and operational complexity will rise with each additional VPC deployed.

Gateway VPC endpoints are available for Amazon S3 and Amazon DynamoDB and are recommended when accessing these services from within a VPC. Gateway VPC endpoints offer a more cost-effective alternative than the equivalent interface endpoints. For example, Gateway VPC endpoints don't have an associated data transfer or per-hour fee. Access to Gateway VPC endpoints is not directly accessible from your on-premises network and will require a proxy farm infrastructure to enable network connectivity.

A well-planned DNS strategy can help avoid complications as your AWS environments grow. If you maintain on-premises DNS capabilities, we recommend you design [hybrid DNS architectures](#) that use on-premises DNS infrastructure along with [Route 53](#) for any AWS based DNS requirements. Integrate DNS resolution with on-premises DNS environments using Route 53 Resolver Endpoints and [Forwarding rules](#). Use private hosted zones to hold information about how you want Amazon Route 53 to respond to DNS queries for a domain and its subdomains within one or more VPCs that you create with the Amazon VPC service. Establish distributed management of your [private hosted zones](#) by using Route 53 to associate your hosted zone to VPCs across your AWS accounts and Regions.

Establish network security

Securing your AWS network must align to your overall security strategy and follow the recommendations in the [Well-Architected security pillar](#). Understanding the risks that you're mitigating will help you apply appropriate network security controls for specific traffic flows. For instance, the [Security Reference Architecture](#) recommends a centralized network account that isolates inbound, outbound, and inspection VPCs. Network security should be designed to protect connectivity between your on-premises network and your AWS environment, to and from the internet, and across your AWS environments.

[AWS Shield](#) is a managed Distributed Denial of Service (DDoS) protection service that safeguards internet facing applications running on AWS and is offered in two tiers: standard and advanced. The standard plan, available to all AWS customers, is included for all tenants and defends against the most common, frequently occurring network and transport-layer DDoS attacks that target sites and applications. AWS Shield Advanced includes features such as additional capacity for large DDoS events, native integration with AWS WAF controls, historical reporting, assistance from the AWS DDoS Response Team, and some cost protection for charges incurred during an attack.

Configure [Amazon VPC Security Groups](#) to allow specific inbound and outbound traffic. In addition to security groups, you can also configure [stateless network ACLs](#) (NACLs) that operate at the subnet boundary. Configure security groups with more granular rules to govern access to specific applications or services. Use NACLs when security requirements require traffic be governed for an entire subnet.

Implement web application firewalls to help protect external facing web applications and APIs against common bugs and bots. These solutions can help block malicious application attacks like SQL injection, cross-site scripting (XSS), and others. These may include common threats such as OWASP Top 10 security

risks, Content Management Systems specific threats, or emerging Common Vulnerabilities and Exposures (CVE). AWS Solutions also provides templates and patterns in [AWS WAF Security Automations](#), and centralized [AWS WAF and VPC Security Group Management](#) to assist you in the deployment of AWS WAF controls in an automated manner.

Select [deployment models supported by AWS Network Firewall](#) that meet your specific use case. For each deployment model, you can have AWS Network Firewall chained together with other services (service chaining). For example, you can [chain AWS Network Firewall](#) and NAT gateway. Extend your security architecture as you scale to enable [Amazon Route 53 Resolver DNS Firewall](#) to block DNS queries made for known malicious domains and to allow queries for trusted domains. Adopt centralized management through AWS Firewall Manager to streamline operations across your multi-account framework. Save time by automating the process of provisioning a centralized AWS Network Firewall to inspect traffic between your Amazon VPCs with [AWS Network Firewall Deployment Automations for AWS Transit Gateway](#).

Consolidating AWS Partner virtual appliances with Gateway Load Balancer can reduce operational overhead and cost. Implement and consolidate AWS Partner security solutions such as intrusion detection and prevention, next-generation firewalls, and web application firewalls.

Establish network monitoring

Although network constructs are unique, you should pair network monitoring with the full breadth of your observability implementation, including specifying network metrics captured in [Amazon CloudWatch](#). For visibility into traffic patterns of your VPC, use [Amazon VPC Flow Logs](#). VPC Flow Logs provide metadata (IP addresses, ports, number of bytes transferred, etc.) about the networking flows (to and from interfaces) in your VPC. Collect VPC Flow Logs in a centralized S3 bucket for use with other log aggregation and analytics functions. When you need to perform content inspection, threat monitoring or troubleshooting, you can copy network traffic to specific monitoring appliances. For example, to capture the full packets, not just the metadata, use [Amazon VPC Traffic Mirroring](#) to replicate all traffic, or specific flows from an elastic network interface, to the destination of your choice.

Automate the monitoring of your AWS networks and identify where network access to your environments may be misconfigured by implementing tools like [VPC Reachability Analyzer](#) and [Amazon Inspector Network Reachability](#) from the [AWS Provable Security initiative](#). These tools let you implement detailed network security checks without having to install scanners and send packets. This will reduce complexity by providing automated monitoring and create a more efficient review process, especially across VPC peering connections and VPNs.

AWS network connectivity management tools

The following AWS services can be used to help you follow the guidance provided by the M&G Guide:

[Amazon VPC](#) is a service that lets you launch AWS resources in a logically isolated virtual network that you define. This can be done within one account, or within a multi-account strategy. You have complete control over this virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables and network gateways. You can use both IPv4 and IPv6 addresses for most resources in your virtual private cloud, helping to ensure secure and easy access to resources and applications.

[Amazon VPC IP Address Manager](#) (IPAM) is a VPC feature that makes it easier for you to plan, track, and monitor IP addresses for your AWS workloads. You can use the automated workflows in IPAM to more efficiently manage IP addresses.

For cloud-to-cloud connectivity, cloud-to-enterprise, and cloud-to-internet, we recommend using [AWS Transit Gateway](#) as a shared service in your multi-account strategy. Transit Gateway uses a hub and spoke pattern to simplify your network and provide a central point for network traffic inspection. Connections

of AWS accounts to a transit gateway can be deployed automatically by Control Tower Customizations and AWS Partners.

[AWS Direct Connect](#) establishes a dedicated network connection between your on-premises network and AWS. With this connection in place, you can create virtual interfaces directly to the AWS Cloud, bypassing your internet service provider. This can provide a more consistent network experience.

[AWS Virtual Private Network](#) solutions establish secure connections between your on-premises networks, remote offices, client devices, and the AWS global network. AWS VPN is comprised of two services: [AWS Site-to-Site VPN](#) and [AWS Client VPN](#). Each service provides a highly-available, managed, and elastic cloud VPN solution to protect your network traffic. AWS Site-to-Site VPN creates encrypted tunnels between your network and your Amazon Virtual Private Clouds or AWS Transit Gateways. For managing remote access, AWS Client VPN connects your users to AWS or on-premises resources using a VPN software client.

[AWS Transit Gateway Network Manager](#) reduces the operational complexity of managing a global network across AWS and on-premises. With Network Manager, you can set up a global view of your private network simply by registering your Transit Gateways and on-premises resources. Your global network can then be visualized and monitored via a centralized operational dashboard.

To provide preventive security for internet-to-cloud connectivity, we recommend implementation of [AWS Network Firewall](#). Network Firewall gives you granular visibility and control of your network traffic, enabling outbound domain filtering, and intrusion prevention through event driven logging, and the service automatically scales with network traffic to provide high availability protections without the need to set up or maintain the underlying infrastructure.

By deploying Network Firewall along with Transit Gateway, you can centrally inspect hundreds or thousands of VPCs and accounts and centrally configure and manage your network firewall, firewall policies, and rule groups.

[AWS Firewall Manager](#) is a security management service that helps you to simplify management of firewall rules across your accounts, easily deploy managed rules across accounts, meet compliance obligations of your existing and new application firewalls, and centrally deploy protections for your VPCs.

AWS automated reasoning provides tools that detect entire classes of misconfigurations, including both a VPC and network configuration tool. [VPC Reachability Analyzer](#) is a configuration analysis tool that enables you to perform connectivity testing between a source resource and a destination resource in your VPCs. When the destination is reachable, Reachability Analyzer produces hop-by-hop details of the virtual network path between the source and the destination. When the destination is not reachable, Reachability Analyzer identifies the blocking component. For example, paths can be blocked by configuration issues in a security group, network ACL, route table, or load balancer.

[Amazon Inspector Network Reachability](#) provides rules to analyze your network configurations to find security vulnerabilities of your EC2 instances. The findings that [Amazon Inspector](#) generates also provide guidance about restricting access that might not be secure. The Network Reachability rules package uses the latest technology from the [AWS Provable Security initiative](#). The findings generated by these rules show whether your ports are reachable from the internet through an internet gateway (including instances behind Application Load Balancers or Classic Load Balancers), a [VPC peering](#) connection, or a VPN through a virtual gateway. These findings also highlight network configurations that allow for potentially unwanted access, such as mismanaged security groups, ACLs, and internet gateways. These rules help automate the monitoring of your AWS networks and identify where network access to your EC2 instances might be misconfigured. By including this package in your assessment run, you can implement detailed network security checks without having to install scanners and send packets, which are complex and expensive to maintain, especially across VPC peering connections and VPNs.

If you would like support implementing this guidance, or assisting you with building the foundational elements prescribed by the M&G Guide, we recommend you review the offerings provided by [AWS Professional Services](#) or the AWS Partners in the [Built on Control Tower program](#).

If you are seeking help to operate your workloads in AWS following this guidance, [AWS Managed Services \(AMS\)](#) can augment your operational capabilities as a short-term accelerator or a long-term solution, letting you focus on transforming your applications and businesses in the cloud.

Integrated network connectivity partners

The M&G Guide recommends you consider the following questions when choosing an AWS Partner solution for network and connectivity:

- Does it support features you are considering using from Amazon VPC and Amazon EC2 instances?
- Does it integrate with AWS services such as AWS Firewall Manager, AWS Security Hub, AWS Transit Gateway, Amazon GuardDuty, Gateway Load Balancer, AWS WAF, and AWS Network Firewall?
- Does it support automatic scaling?
- Can it be provisioned from an infrastructure as code template that is distributed from a central catalog?
- Does it integrate with an observability solution? For instance, does it allow log aggregation across multiple instances, such as multiple firewalls and multiple routers?

Network orchestration

To set up and maintain cloud environments effectively, enterprises need network management solutions that scale in a multi-account environment to configure, manage, and coordinate AWS resources automatically. The following integrated network connectivity AWS Partners have provided solutions that align with the M&G Guide, and are available for entitlement in AWS Marketplace: Network orchestration solutions

Gateway Load Balancer partners

[Aviatrix – Cloud Network Platform](#) uses Gateway Load Balancer to scale and manage appliances that support GENEVE encapsulation. Gateway Load Balancer provides a high-performance connection to virtual appliances, reduces the need for source network address translation (SNAT), and allows you to add or remove appliances for scaling or in response to health checks without impacting existing sessions. The Aviatrix Controller automates attachment of Gateway Load Balancer, its associated Gateway Load Balancer endpoint, and all connected appliances to an Aviatrix Transit/FireNet Gateway.

[Cisco Systems – Cloud Application Centric Infrastructure \(Cisco Cloud ACI\)](#) drives networking automation in on-premises and AWS environments and allows you to have a consistent security posture and uniform operational processes across your hybrid cloud infrastructure.

[Palo Alto Networks – VM-Series](#) scales your traffic across multiple VM-Series firewalls using native AWS networking constructs to achieve higher throughputs – without the need for encrypted tunnels for east-west and outbound traffic inspection. VM-Series also reduces the number of firewalls needed to protect your AWS environments and consolidate your overall network security posture with centralized security management.

AWS Transit Gateway partners, including SD-WAN solutions

[Aviatrix](#) – With Aviatrix Secure Networking Platform AMI or Aviatrix software as a service (SaaS) listings, both available in AWS Marketplace, you can orchestrate the Transit Gateway in minutes without delving into the configuration detail required in each VPC and route table.

[Cisco](#) – Cisco SD-WAN offers automated connectivity provisioning to the most optimal AWS entry point for your data center, branch, and hub locations.

[Palo Alto Networks – Prisma SD-WAN](#) (formerly CloudGenix SD-WAN) is a cloud-delivered service that implements application-defined, autonomous SD-WAN to help you secure and connect your branch offices, data centers, and large campus sites without increasing cost and complexity.

Identity management

As you scale your use of the AWS Cloud, you need robust identity and permission management processes to help ensure that you follow the standard security advice of granting least privilege, or granting only the permissions required to perform a task. Robust identity management helps ensure that the right systems and people have access to the right resources under the right conditions. This also needs to be done while not overburdening operations capabilities with too many, too granular or too complex permission or identity constructs. The M&G Guide includes the recommendations of the [Security Pillar](#) for managing identities and access permissions across all cloud resources in your multi-account strategy in order to be migration ready, scale ready, and operating efficiently.

The [Security Pillar](#) describes the difference between human and machine identities. It also shows that centralized administration of human identities and access to your environments with an identity provider is a critical strategy to managing authentication and authorization across your enterprise. This is important for managing and governing, as it makes it easier to manage access across multiple applications and services because you are creating, managing, and revoking access from a single location. For example, if someone joins or leaves your organization, you can add or revoke that individual's access for all applications and services (including AWS) from one location. This aligns with ITIL best practices, and reduces the need for multiple credentials and provides an opportunity to integrate with existing human resources (HR) processes. In AWS, we consider machine identities distinctly from human identities. Machine identities (like service roles) still reside within AWS IAM and are designed to uphold the principle of least privilege, but are not managed by your identity provider.

Define access policies and mechanisms that include granting least privilege access, sharing resources securely, and reducing permissions continually (including the removal of unused permissions) with AWS IAM Access Analyzer. Review how permissions are actually being authored, validated, and used over time, so that you can remove unnecessary permissions in accordance with the principle of least privilege. This would include adding observability rules for "last accessed" data, such as a timestamp depicting when an identity policy or principal (such as a user or role) last used a service or performed an action from supported services. This enables you to more easily identify unused permissions and improve your security posture by removing the permissions that are not necessary for the user, group, or role to perform a specific task. Both AWS and AWS Partners provide tools for the creation, review, and revoking of permissions in an automated manner throughout your software development lifecycle (SDLC) or development, security and operations (DevSecOps) cycles.

Interoperable functions

The eight management and governance functions, supported by AWS services and AWS Partner solutions, work together and interoperate to reduce complexity. Outputs from functions are used to inform or integrate with other functions.

For identity management this includes:

- Specific identity **Controls** included within your preventive and detective mechanisms.
- **Network connectivity** designed as a complement to identity, forming a least privilege boundary for your environments.
- **Security management** with specific capabilities to remediate and address identity-related incidents.
- Using your **Service management** solution as the record of change for your identity constructs.
- Incorporating all Identity and access management activities across the **Monitoring and observability** functions so that they provide evidentiary findings for audit and compliance needs.

- Enabling **Cloud Financial Management** with identity management to provide specific cost and usage by defined roles and groups.
- As cloud assets are **Sourced and distributed**, defining identity and access policies in a manner that restricts controls the range of operations.

Implementation priorities

Having secure and scalable mechanisms to manage identities is a critical component of a cloud ready environment. As such, the following items should be prioritized.

Establish a centralized identity provider for human identities

Implementation of a centralized identity provider is a foundational capability for enterprises of all sizes and interwoven across all environments, systems, workloads, and processes. For workforce identities, restrict the use of IAM users and instead rely on an identity provider that enables you to manage identities in a centralized place. This makes it easier to manage access across multiple applications and services, because you are creating, managing, and revoking access from a single location. Use existing HR processes to manage creation, update, and removal of access to include your AWS environments. Federate access into your AWS environments by integrating the identity provider with a SAML 2.0 compliant SSO solution. Incorporate multi-factor authentication (MFA) in AWS for the root user, and use your identity provider MFA solution for other privileged roles.

Define job functions and codify IAM roles

Define the IAM roles to be granted to human and machine identities and strive to follow the principles of least privilege and separation of responsibilities. Verify that runbooks and playbooks reference identity constructs with sufficient permissions to run support activities (for example, emergency access). This might include "break glass" access in the event that your SSO solution becomes inaccessible. Optimizing your IAM permissions is a journey. Refine permissions over time and employ controls as an additional layer of protection while still enabling developer agility.

Consider that permissions will be variable by environment type. For instance, permissions defined for production accounts should be more restrictive than those defined in development or sandbox accounts. Use resource tags and IAM conditional statements to create more fine-grained access policies and apply permissions boundaries to allow safe delegation of administrative functions while protecting against privilege escalation. Attribute-based access control (ABAC) is an authorization strategy that defines permissions based on attributes. For AWS services that support tagging, ABAC policies can be designed to allow operations when the principal's tag matches the resource tag. ABAC is helpful in environments that are growing rapidly because it helps scale policy management with reusable attributes from your identity provider.

Continually collect, review, and refine permissions

Changes to identity roles and permissions are recorded in CloudTrail and detective guardrails should alert on deviations from your expected configuration state. With the centralized collection of events, you can use aggregated and pattern identification tools to review and refine permissions as required.

AWS Identity and Access Management (IAM) [access advisor](#) uses data analysis to help you set permission guardrails confidently by providing service last accessed information for your accounts, organizational units (OUs), and your organization managed by AWS Organizations. Use this feature to analyze service last accessed information and determine services not used and reduce permissions where appropriate.

Use [IAM Access Analyzer](#) to guide you to least privilege by helping you set, verify, and refine permissions. This includes identifying S3 buckets or IAM roles that are shared with an external entity outside of your organization or account. Establish a regular attestation process to help ensure permissions are still appropriate as personnel change roles within your organization. Review the IAM credential report for stale or unused account users and credentials.

Manage credential use

The M&G Guide recommends the use of IAM roles and temporary credentials. Use AWS Systems Manager to manage remote access to instances or on-premises systems using a pre-installed agent without the need for stored secrets. Reduce reliance on long-term credentials, and scan for hardcoded credentials in your infrastructure as code templates. In situations where you cannot use temporary credentials, use programmatic tools such as AWS Secrets Manager to automate credential rotation and management, such as application tokens and database passwords.

Source and distribute identity constructs with automation

Codify and version identity constructs such as roles, policies, and templates with infrastructure as code. Employ testing and linting to ensure coding standards are met within your continuous integration and continuous delivery (CI/CD) pipelines with tools like [cfn-guard](#). Use [IAM Access Analyzer](#) policy validation to check for findings that include security warnings, errors, general warnings, and suggested changes to your IAM policies. Where appropriate, deploy and remove identity constructs for temporary access to the environment in an automated manner and prohibit deployment by individuals using the console.

AWS identity services

Effective identity management is provided by AWS services, solutions, and AWS Partners that permit you to securely manage identities, resources, and permissions at scale. AWS identity services provide flexible options for where and how you manage your employee, partner, and customer identities. The following AWS services can be used to help you meet the prescribed benefits of the M&G Guide:

[AWS Identity and Access Management \(IAM\)](#) provides fine-grained access control across all of AWS. Using IAM, you can specify who can access which services and resources, and under which conditions. With IAM policies, you manage permissions to your workforce and systems to ensure least privilege permissions.

[AWS IAM Access Analyzer](#) guides you toward least privilege by helping you set, verify, and refine permissions. Policy validation with Access Analyzer helps you author secure and functional permissions with more than 100 policy checks. Policy generation with Access Analyzer makes it easier to apply fine-grained permissions by generating policies based on your access activity in AWS CloudTrail. Access Analyzer also continually monitors resources and generates public and cross-account findings to help you verify that existing access meets your intent.

[AWS IAM Identity Center \(successor to AWS Single Sign-On\) \(IAM Identity Center\)](#) helps you centrally manage access to multiple AWS accounts and business applications and provide users with single sign-on access to all their assigned accounts and applications from one place. With IAM Identity Center, you can manage access and user permissions to all of your accounts in AWS Organizations centrally. IAM Identity Center configures and maintains all the necessary permissions for your accounts automatically, without requiring any additional setup in the individual accounts. IAM Identity Center also includes built-in integrations to many business applications, such as Salesforce, Box, and Office 365.

[AWS Directory Service for Microsoft Active Directory](#), also known as AWS Managed Microsoft AD, enables your directory-aware workloads and AWS resources to use managed Active Directory in AWS. [AWS Managed Microsoft AD](#) is built on Microsoft Active Directory and does not require you to synchronize or

replicate data from your existing Active Directory to the cloud. You can use the standard Active Directory administration tools and take advantage of the built-in Active Directory features, such as group policy and single sign-on.

[AD Connector](#) is a directory gateway with which you can redirect directory requests to your on-premises Microsoft Active Directory without caching any information in AWS. AD Connector comes in two sizes, small and large. You can spread application loads across multiple Active Directory connectors to scale to your performance needs.

If you would like support implementing this guidance, or assisting you with building the foundational elements prescribed by the M&G Guide, we recommend you review the offerings provided by [AWS Professional Services](#) or the AWS Partners in the [Built on Control Tower program](#).

If you are seeking help to operate your workloads in AWS following this guidance, [AWS Managed Services \(AMS\)](#) can augment your operational capabilities as a short-term accelerator or a long-term solution, letting you focus on transforming your applications and businesses in the cloud.

Integrated identity partners

The M&G Guide recommends you consider at a minimum the following questions when choosing an AWS Partner solution for identity management:

- Does it integrate with a single sign-on provider such as AWS SSO?
- Does it support the System for Cross-domain Identity Management (SCIM) v2.0 standard for automating the exchange of user identity information?
- Does it support federated user and group mapping?
- Does it include a method for managing predefined permissions at scale such as AWS permission sets in IAM Identity Center?

Optimize identity management in a multi-account environment with a simplified single sign-on experience, user provisioning, and password management for your AWS environments. The following integrated identity AWS Partners have provided integrations that align to the M&G Guide, and are available for deployment from AWS Marketplace.

[CyberArk](#) helps organizations secure access to critical business applications and infrastructure, protect a distributed workforce, and accelerate business in the cloud. With CyberArk Identity Security Platform, enterprises can streamline access provisioning to AWS and give workers secure and frictionless access to all authorized AWS resources from any location, using any device. In addition to centralized management of end-user access to AWS environments, CyberArk provides comprehensive auditing and reporting capabilities to simplify access compliance.

[Kion](#) is a comprehensive enablement software solution that delivers visibility and control of cloud workloads. [Kion](#) provides integrations with identity providers to allow control over cloud federation and policy controls at an account and an organization level. [Kion](#) allows enterprises to manage their cloud presence at scale with automation and orchestration, financial management, and compliance.

[Okta](#) enables teams to securely and seamlessly manage AWS IAM Identity Center (successor to AWS Single Sign-On) (IAM Identity Center) entitlements at scale. After connecting Okta Identity Cloud to IAM Identity Center once, you can manage access to AWS centrally in IAM Identity Center, and enable end users to sign in using Okta to access all their assigned AWS accounts through AWS Organizations. This includes centralized reporting and auditing of end-user access across all apps and systems.

[OneLogin](#) cloud-based identity and access management enables IT teams to manage and provision access to AWS resources centrally. Whether you're newly migrating to AWS or an enterprise user,

integrating Control Tower with OneLogin helps ensure you can easily and securely scale your enterprise-wide environments and IAM permissions.

Ping Identity's PingOne Cloud Platform solution provides central authentication services to connect employees across any application, directory, and situation. By providing authentication for all end users and identities in customer environments, Ping can reduce authentication silos, and help your business increase agility. The result is a centrally-managed authentication hub that provides a highly-configurable, secure, and consistent experience for your workforce.

Sonrai Dig is an enterprise cloud security platform providing complete visibility across all multi-account AWS environments. Using Dig's Cloud Identity Entitlement Management (CIEM) capabilities, you can continually inventory your identities (people and non-people), compute their effective (end-to-end) permissions, enforce least privilege, and alert on any deviations as soon as they are detected.

Security management

Security plays a key role and is foundational to all functions of the M&G Guide. Security management is the process of setting up, measuring, and improving security processes and tools. The M&G Guide focuses on cloud-ready environments so that you are well prepared to host your workloads. We recommend following the security best practices described in the [Well-Architected Security Pillar](#) whitepaper for each type of workload you run on AWS. You will find that the same principles for well-architected workloads apply for how you effectively secure and manage cloud-ready environments. Specifically, the security pillar includes a comprehensive view of the best practices for management and governance of security capabilities, some of which are highlighted later in this guide. Further information on cloud adoption best practices that align with the Security Pillar can also be found in [AWS CAF](#).

To scale with AWS, it is important to continually address and refine your security capabilities alongside the rest of your management and governance functions. This includes the identification, management, and resolution of security issues and findings across all your environments. As your scale increases with AWS, it is essential to adapt your security management to the dynamic nature and ephemeral lifespan of cloud resources. This adaptation includes response mechanisms as well as ownership. In some cases, ownership of security might merge with, and in other cases require new, accountability and responsibility models.

The M&G Guide recommends standard ways to address AWS security across the eight management and governance functions. For instance, in the [Controls \(p. 5\)](#) section, we demonstrate the need for security controls to be included across your management and governance tooling. In this Security Management section, we outline security tools and functions that are equally important to operating and scaling efficiently. Each area of your cloud operations is responsible for implementing appropriate security controls. These should include capabilities to identify, protect, detect, respond, and recover from security issues and events.

Security architecture

The [AWS Security Reference Architecture \(AWS SRA\)](#) is a holistic set of guidelines for deploying the full complement of AWS security services in a multi-account environment that is aligned to the [Well-Architected Security Pillar](#). This overall architectural guidance complements detailed, service-specific recommendations, such as those found in the [AWS Security Documentation](#). For example, AWS SRA recommends complementing the security architecture implemented in your environments with a specific OU and account for security tooling. Where services support this, delegate administration of security-related services to the security tooling account. The security tooling account will then serve as a central pane of glass to the member accounts, providing insights for extended detection and response (XDR) activities. Where required, also provide for engineering and builder teams to create specialized or localized security capabilities that are specific to their workloads. Note that this reference architecture can be extended to include AWS Partner solutions following the same patterns.

Automated findings and campaigns

Following the prescriptive guidance in the [Controls \(p. 5\)](#) section of this guide, after you have detective controls in place across your multi-account strategy, deviations from the controls should result in security findings. A *finding* is a specific deviation from a control associated with a specific AWS account,

AWS Region, environment, or resource. For each detective mechanism you have, you should also have a clearly defined process in the form of a runbook or playbook to investigate. Tickets should be automatically created based on findings with information about the deviation, remediation guidance, and deadlines. Tickets are assigned to the resource, account, or environment owner.

A *campaign* is a way to aggregate issues around a particular control or set of controls and drive action towards remediation. Campaigns include the development of campaign metrics to measure progress. You can also use campaigns and tickets to drive action to have account owners put preventive controls in place.

Note that both campaigns and findings will need to be tuned along with your threat detection tools. This tuning will allow you to remove any noise created from false positives or negatives. In contrast, any patterns from campaigns or findings will need to be translated to additional controls.

Security metrics

A mature internal security metrics program is crucial for managing security in the cloud. In general, this is completed by following the guidance of “what gets measured, gets done”. After you have controls in place, security metrics are the primary way to assess whether your security posture is improving, and whether your controls are adequate. You should have metrics for each part of your security organization, and these metrics should be reviewed regularly to verify that they have the right level of organizational buy-in and attention. For example, mean time to identify (MTTI) root cause and mean time to respond (MTTR) provide insights into your security incident response effectiveness. Make sure that you have good processes and continuous improvement around capturing, reviewing, and remediating insights gained from them.

Security response management

Enterprises are mandated to protect their digital infrastructure from a wide range of threats and require in-depth visibility into their infrastructure and applications to make faster data-driven decisions. Enterprises need to take proactive actions to ensure timely threat intelligence. Security solutions must monitor workloads in real-time, identify security issues, and expedite root-cause analysis. Essential elements of these tools allow you to identify, prioritize, and mitigate threats, gain visibility into suspicious activities, and acknowledge risks. The Security Pillar outlines specific recommendations for building your workloads while thinking proactively about security. This is the foundation for helping ensure that you can respond effectively to security insights you are gathering.

Security management functions are responsible for analyzing and responding to security events. Where in the past this was done with human-powered processes, we recommend you automate these identification and remediation systems. This automation will help increase your security posture along with your ability to scale. At cloud scale, use automated workflows wherever possible to investigate events of interest and gather information on unexpected changes. Require that these workflows be tested in development environments to ensure operational resilience. Detect advanced security threats by combining monitoring from network, firewall, identity, control plane, vulnerability and patch management, workloads, and data protection processes with your existing threat detection capabilities. Threat detection can be used to determine the expected pattern of API calls per role, application, or service, and determine the levels that indicate an unexpected deviation. This activity will allow you to maximize your telemetry by layering behavioral analytics with your log analytics. The Security Pillar outlines how to build a comprehensive detective capability with options that include automated remediation and AWS Partner solutions. This capability is enabled through the [configuration of environments](#) with centralized analysis of logs, findings, and metrics. Automating aspects of your incident management process also improves reliability and increases the speed of your response, which creates an environment easier to assess in after-action reviews.

Interoperable functions

The eight management and governance functions, supported by AWS services and AWS Partner solutions, work together and interoperate to reduce complexity. Outputs from functions are used to inform or integrate with other functions. For security management this includes:

- **Controls** continually updated and tuned as a result of your security findings and patterns.
- Changes in the definition or behavior of your **Network connectivity** as part of your security findings and patterns, with automated remediation when applicable.
- Changes in the definition or behavior of your **Identity management** as part of your security findings and patterns, with automated remediation when applicable.
- A **Service management** framework integrated such that security incident response procedures and vulnerability management procedures (including security incident response findings and campaigns) are integrated with tooling from the service management framework.
- Centralized **Monitoring and observability** tools informing security management functions, with specific automated remediation.
- Unexpected changes in cost and spend patterns as part of your **Cloud Financial Management**, which should be visible and are included in your security findings and patterns, with automated remediation when applicable.
- Security tools which are **Sourced and distributed** with preconfigured security controls in a hub and spoke pattern across your environments.

Implementation priorities

For security management, we recommend that you deploy your security capabilities (XDR, CSPM, etc.) using the same mechanisms as the base foundation of capabilities for each of your accounts. In the [Controls \(p. 5\)](#) section, we recommend that you begin with assessing your risk posture and [developing a threat model](#), and thereafter selecting appropriate controls for each environment. In addition, you should set a foundation with specific security tools aligned to your environments and accounts, additional logging, and integration to your incident management and security analytics capabilities.

Design a Well-Architected security environment

Design capabilities with governance and security instrumentation in mind, following the best practices described in the [AWS Well-Architected Security Pillar](#). Your [security foundations](#) should include:

- Separating and securing your workloads across a multi-account strategy
- Identifying and validating control objectives based on your compliance requirements and risk assessments
- Recognizing and staying up to date with the latest security threats and vectors, recommendations, and effective controls
- Establishing secure baselines and templates for security mechanisms that are tested and validated continually as part of your build, pipelines, and processes
- Identifying and prioritizing risks using threat modeling
- Evolving the security posture of your workloads using new features and enhancements of AWS and AWS Partner services
- Enabling encryption at rest and in motion for cloud storage, databases and traffic that includes sensitive data in motion

Choose security tools to match your enterprise needs

Security monitoring tools should allow for granular security monitoring across infrastructure, applications, and workloads as well as provide aggregated views for pattern analysis. As with all other security management tools, it is important to extend your XDR tools to provide functions to assess, detect, respond, and remediate the security of your applications, resources, and environments on AWS. Using these tools with the interoperable functions of the M&G Guide can provide a mechanism for you to enable further use cases for compliance monitoring, incident response, DevSecOps integration, risk assessment and visualization. Cloud Security Posture Management (CSPM) tools can also be used to manage and remediate common vulnerabilities and exposures (CVEs) in your AWS environments. Use a vulnerability management solution that assesses infrastructure and applications for vulnerabilities or deviations from best practices, and produces a detailed list of findings prioritized by level of severity.

Analyze and model for threats

Implement continual monitoring and measurement against industry and security benchmarks. When designing your instrumentation approach, determine what types of event data and information will best inform your security management functions. This monitoring should encompass several attack vectors including service usage. Your security foundations should include a comprehensive secure logging and analytics capability across your multi-account environments that includes the ability to correlate events from multiple sources.

Prevent changes to this configuration with specific controls and guardrails. AWS Security Hub and AWS Partner tools provide dashboards across a multi-account environment and should be integrated with [event-triggering systems](#) in AWS for security and incident event management functions. Develop thresholds and metrics based on expected behavior of your environments. Use anomaly detection to identify unintended activities when thresholds are exceeded. Configure and monitor Amazon CloudWatch alarms for exceeded thresholds across IAM activity, resource creation, failed access attempts, policy and configuration changes, VPC-related changes (security groups, NACLs, gateways, and route tables), API calls, and activities in unapproved AWS Regions.

Develop a threat modeling practice to engage with business stakeholders, cloud infrastructure architects, compliance, application developers, security and other key stakeholders. The AWS Well-Architected Framework calls out threat modeling as a specific best practice within the Security Pillar, under the question [SEC 1: How do you securely operate your workload?](#) Preventive, detective, and responsive controls should be put in place as responses to both workload and environment level threats identified in threat modeling exercises.

Enable log aggregation as a foundation for your threat modeling and log analytics capabilities that is extended as new accounts or environments are created, updated, or deleted. Use XDR with multiple telemetry sources to identify if correlated events qualify as recordable incidents. Use the threat model as the basis for table top exercises, building incident response playbooks and runbooks, and develop automated testing. Codify your compliance objectives using AWS Config or AWS Partner products.

Automate incident management workflows, findings, and campaigns

The [Security Pillar](#) outlines how to build a comprehensive detective capability with options that include automated remediation and AWS Partner integrations. This capability is enabled through the [configuration of environments](#) with centralized analysis of logs, findings, and metrics. Typical automation might include AWS Lambda function “responders” that react to specific changes in the environment, orchestrating automatic scaling, isolating suspect system components, deploying just-in-time investigative tools, and creating workflow and ticketing to shut down and learn from a closed loop organizational response. Each account, application, or resource should be provisioned with a baseline configuration aligned with your security operations. This includes provisioning specific security tools,

which also align to your observability requirements. Develop remediation processes which allow you to isolate cloud resources for forensic analysis.

Select, measure, and continually improve your security metrics

Follow the guidance of “*what gets measured, gets done*”. Implement metrics for each part of your security organization and review regularly to verify you have the right level of organization buy-in and attention. Measure the performance of your security operations along with the threats themselves. Include metrics around your security operations paired with metrics around security campaigns, findings, and tools. For example, mean time to identify (MTTI) root cause and mean time to respond (MTTR) provide insights into your security incident response effectiveness. Drive operational insights and reviews to continually improve your threat modeling, threat detection, incident management, and response and remediation capabilities.

AWS security management services

The following AWS services can be used to help you meet the prescribed benefits of the M&G Guide:

[AWS Security Hub](#) is a cloud security posture management service that performs security best practice checks, aggregates alerts, and enables automated remediation. AWS Security Hub aggregates, organizes, and prioritizes your findings from multiple AWS services as well as from AWS Partner solutions, enabling you to quickly assess the security posture across your AWS accounts. AWS Security Hub runs automated configurations and compliance checks based on open standards, such as CIS Benchmarks, NIST frameworks, and AWS Foundational Security Best Practices.

[Amazon GuardDuty](#) is a threat detection service that continually monitors for malicious activity and unintended behavior to protect your AWS accounts, workloads, and data stored in Amazon S3. Amazon GuardDuty uses machine learning, anomaly detection, and integrated threat intelligence to identify and prioritize potential threats. GuardDuty analyzes tens of billions of events across multiple AWS data sources, such as AWS CloudTrail event logs, Amazon VPC Flow Logs, and DNS logs.

Both AWS Security Hub and Amazon GuardDuty have the concept of an *administrator* and *member* account. The administrator account can view the aggregated findings of all member accounts within a Region. You should delegate administration of Security Hub and GuardDuty to the security audit account provisioned by AWS Control Tower.

[AWS Security Hub Automated Response and Remediation](#) is a solution that uses AWS Security Hub to provide a ready-to-deploy architecture and a library of automated playbooks. The solution creates an AWS Service Catalog portfolio of predefined security response and remediation actions called playbooks. Individual playbooks are deployed in the Security Hub primary account. Each playbook contains the necessary custom actions, AWS Identity and Access Management (IAM) roles, Amazon CloudWatch Events, Systems Manager automation documents, AWS Lambda functions, and AWS Step Functions needed to start a remediation workflow within a single AWS account, or across multiple accounts.

[Amazon Detective](#) automatically collects log data from your AWS resources and uses machine learning, statistical analysis, and graph theory to build a linked set of data that enables you to easily conduct faster and more efficient security investigations.

[AWS Control Tower](#) implements centralized logging and audit accounts that use AWS CloudTrail and Amazon CloudWatch. This is done using AWS Config for detective guardrail enablement, and SCPs from AWS Organizations for preventive controls.

[AWS Systems Manager](#) allows you to create automated responses to security misconfigurations via specific automation documents, with patch management functions.

Using [automated reasoning technology](#) (the application of mathematical logic to help answer critical questions about your infrastructure), AWS is able to identify opportunities to improve your security posture. We call this *provable security* providing higher assurance in security of the cloud and in the cloud. Automated reasoning capabilities include [IAM Access Analyzer](#), [VPC Reachability Analyzer](#), [Amazon CodeGuru](#), [Amazon S3 Block Public Access](#), and Amazon Inspector network reachability.

If you would like support implementing this guidance, or assisting you with building the foundational elements prescribed by the M&G Guide, we recommend you review the offerings provided by [AWS Professional Services](#) or the AWS Partners in the [Built on Control Tower program](#).

If you are seeking help to operate your workloads in AWS following this guidance, [AWS Managed Services \(AMS\)](#) can augment your operational capabilities as a short-term accelerator or a long-term solution, letting you focus on transforming your applications and businesses in the cloud.

Integrated security management partners

The M&G Guide recommends you consider the following questions when choosing an AWS Partner solution for security management functions:

- Is the solution from an AWS Security Competency Partner?
- Does the solution support multi-account, and work across all your required AWS Regions?
- Are security findings aligned to your controls surfaced with appropriate remediation steps? Is this auditable?
- Does the AWS Partner incorporate new threat vectors, maintain and manage their own findings, and add them to the operations tools on a regular basis?
- Does the solution provide analysis and troubleshooting tools for security operations teams?

To help improve the security posture across a multi-account environment, you need to implement security functions, such as vulnerability assessment, firewalls, and intrusion prevention. AWS Marketplace offers integrated software solutions for AWS Control Tower that help enterprises secure diverse workloads and provide broader visibility into assets, events and vulnerabilities.

[Alert Logic Managed Detection and Response \(MDR\)](#) is always on, providing protection across your entire organization through five key elements: intelligence driven by data and humans, a scalable MDR platform, security experts named to your account, security insights at your fingertips, and protection tailored to each asset in your environments.

[Aqua Security SaaS](#) provides a SaaS-based, cloud security posture management (CSPM) solution for AWS Control Tower. Aqua CSPM continually audits your AWS accounts for security risks and misconfigurations. This is performed across hundreds of configuration settings and compliance best practices, enabling consistent, unified multi-account security. It also provides self-securing capabilities to help ensure your cloud accounts do not drift out of compliance by applying a policy-driven approach.

[Cloud Custodian](#) is a tool that unifies the dozens of tools and scripts most enterprises use for managing their public cloud accounts into one open source tool. It uses a stateless rules engine for policy definition and enforcement, with metrics, structured outputs and detailed reporting for clouds infrastructure. Cloud Custodian's integration with Security Hub allows it to both send findings to Security and receive findings for response and remediation actions.

[CrowdStrike Falcon Endpoint Protection](#) uses advanced artificial intelligence (AI), machine learning, behavioral protection, kernel level visibility and proactive threat hunting to identify potential attacks in real time. For enterprises who are adopting or migrating to cloud workloads, CrowdStrike Falcon Endpoint Protection provides comprehensive visibility and breach protection allowing you to rapidly adopt and secure technology across any workload.

[ExtraHop Reveal\(x\) 360](#) provides multi-layered visibility, threat detection, and investigation in AWS via integrations with Amazon VPC Traffic Mirroring for packet-level visibility and VPC Flow Logs for broad coverage. ExtraHop is an AWS Security Competency Partner and offers a free trial of Reveal(x) 360. To learn more, see Reveal(x) 360 in the AWS Marketplace.

[Logz.io](#) AI-Powered ELK-as-a-Service is a cloud-native observability platform providing unified monitoring, troubleshooting, and security for distributed cloud environments. Intelligent log analytics help engineers and businesses resolve incidents faster and simplify cloud security. Logz.io's analytics and optimization tools help businesses reduce overall logging expenses and identify production and security incidents in real time.

[Palo Alto Prisma Cloud](#) provides cloud security posture management (CSPM) and cloud workload protection (CWP) as a single pane of glass for comprehensive visibility and control. Securely provision automated account registrations, continual governance, and enterprise-wide management of multiple AWS accounts in just a few clicks. Prisma Cloud also extends cloud automation to integrated Lambda serverless remediation and manages it through a common policy and governance framework.

[Prowler](#) is a security assessment tool that gives customers direct insights into the security best practices of their AWS infrastructure. Customers can run Prowler to continuously monitor their security status. The main differentiators between Prowler and other existing services or solutions are the number of checks that are included out-of-the-box; no configuration needed to get insights; and no direct cost associated to its use. Prowler's checks follow guidelines from the CIS Amazon Web Services Foundations Benchmark and performs additional checks related to GDPR, PCI, and HIPAA. Prowler supports natively sending findings to AWS Security Hub.

[Qualys](#) The Qualys integration with AWS Security Hub provides customers the ability to consume security and compliance findings about their AWS Instances and accounts within the AWS Security Hub console. Customers have access to critical vulnerabilities, missing patches, open ports, as well as the compliance to CIS, PCI, NIST, HIPAA, and security policies of their Instances and AMIs. Customers can also assess misconfigurations of VPCs, Security Groups, Amazon S3, and IAM against the CIS Benchmark. The Qualys integration with AWS Security Hub allows customers to prioritize their risks and automate remediation using services, such as AWS Lambda.

[Rapid7](#) InsightVM, a vulnerability assessment solution, uses the power of the Insight platform to provide visibility across your modern ecosystem, prioritize risk using attacker analytics, and remediate or contain threats with SecOps agility. With InsightVM, vulnerabilities are discovered in real time and prioritized actionably. By integrating InsightVM with AWS Security Hub, vulnerabilities detected in a business's Amazon EC2 instances are automatically sent to AWS Security Hub for a holistic view of its cloud security posture. With additional vulnerability context from InsightVM, businesses can prioritize its team's security tasks more efficiently and reduce measurable risk in its AWS Cloud.

[Sonrai Dig](#) is an enterprise cloud security platform providing complete visibility across all multi-account AWS environments. Built on our patented graph, Dig combines platform (CSPM), identity (CIEM), and data (Cloud DLP) controls, delivering speed and security where it matters in your cloud apps. Maturity Modeling effectively addresses alert fatigue by providing workload/environment context, while our Governance Automation Engine automates workflow, remediation, and prevention capabilities across cloud and security teams improving operational efficiency and ensuring end-to-end security.

[Splunk Cloud](#)'s integration into AWS Control Tower allows administrators to automatically configure and set up AWS services. Data from AWS CloudTrail, AWS Config, and other sources can be incorporated into your Splunk deployment using Kinesis Data Firehose and Splunk HTTP Event Collector (HEC). With Splunk Cloud, you can automatically collect data from newly vended AWS Accounts and dashboards and alert compliance with AWS Control Tower guardrails.

[Sumo Logic Cloud-Native Machine Data Analytics](#) pulls in critical operational data across services and accounts to give a unified view of AWS environments. Easily navigate from overview dashboards into account, Region, Availability Zone, or service-specific views. Intuitive navigation across logs and metrics data ensures that teams can quickly resolve issues, minimize downtime, and improve system availability.

The Sumo Logic Continuous Intelligence Platform automates the collection, ingestion, and analysis of application, infrastructure, security, and IoT data to derive actionable insights.

[Sysdig Secure](#) helps cloud and security teams detect and respond to threats, and manage cloud configurations, permissions, and compliance. Integration with AWS CloudTrail enables customers to protect existing and newly enrolled AWS accounts via AWS CloudTrail logs. Sysdig detects anomalous activity across AWS workloads with out-of-the-box policies based on open source Falco.

[Tenable Vulnerability Management](#) for Modern IT, Tenable.io provides the most accurate information about assets and vulnerabilities in your IT environments. Available as a cloud-delivered solution, Tenable.io features the broadest vulnerability coverage, intuitive dashboard visualizations for rapid analysis, and seamless integrations that help you maximize efficiency and increase effectiveness.

[Trend Micro Cloud One - Workload Security](#) is purpose-built for server, cloud, and container environments, providing visibility across your entire hybrid cloud. Automatically protect against vulnerabilities, malware, and unauthorized changes with a wide range of powerful and intelligent capabilities. Workload Security automatically integrates with the DevOps toolchain and includes a rich set of REST APIs, which facilitate deployment, policy management, health checks, and compliance reporting.

Service management

The IT service management (ITSM) framework enables enterprises to align the relationship between people, process, and tooling needed through the lifecycle of IT services. The service management framework is also used to create evidentiary support for compliance and risk audits, Cloud Financial Management (CFM) capabilities, and business service requests. Enterprises also use ITSM tools to track business approvals, capture service issue resolutions, inventory technical assets, identify customer technical inquiries, and capture data points to make business decisions. These ITSM tools not only handle the daily operations for business services and applications (incidents/tickets, and CMDB transparency) but also enable everyday workflow and approvals of business requests (for example, facilities, HR, marketing, etc.). Integrating your service management framework to managing and governing your cloud capabilities will increase your operational excellence and agility.

The M&G Guide recommends five capabilities as a baseline for your service management framework within your AWS environments:

- Provisioning and request management
- Event and incident management
- Problem management
- Resource inventory management
- Change management

Provisioning and request management

Provisioning procedures help plan, implement, and maintain a stable technical infrastructure to support organizational business processes. Provisioning focuses on repeatable, standardized, approved, and curated templates to ensure resilient, cost effective, scalable resources. It enables enterprises to transition to a mindset of “infrastructure as code.”

Request management helps in maintaining the curated templates as AWS Service Catalog items. Fulfillment to enterprise end users for any of the AWS services and infrastructure is ensured by AWS Service Catalog through an automated workflow-driven process. The M&G Guide recommends integrating your provisioning, request, and distribution processes with your ITSM tool suite.

Event and incident management

Event and incident management enables enterprises to control and restore environments and data. Event management helps in understanding what is currently happening, detect events, assess potential impact, and determine the appropriate control action. Event management provides the ability to detect and interpret environment issues, and initiate appropriate response and remediation. It is a basis for operational monitoring and control and an entry point for many service operation activities. Automation should be implemented where necessary, based on operations data and metrics. Analyzing event, incident, and operations metrics will support continual service improvement activities of service assurance. This analysis is used as inputs for organizational SLAs.

Incident management restores normal service operation and minimizes adverse business impacts on operations. Combining trend metrics with the identification of common or adverse patterns in service

designs, can also help inform service availability design and reporting calculations. The M&G Guide recommends you enable an issue management mechanism across your AWS accounts. Integrating AWS events with other ITSM processes, such as incident and change management, can also increase your ability to scale.

Problem management

Problem management focuses on identifying and resolving underlying issues (root cause) in the production environment that can lead to incidents. Problems are the underlying causes of incidents. Initially, problem management enables you to resolve the root causes of incidents to minimize impact and prevent them from happening again. Over time, problem management enables you to predict similar incidents using trend analysis and helps you proactively correlate incidents.

The main focus of problem management is root-cause analysis (RCA) with the goal of identifying why an incident occurred and defining measures so that similar incidents don't happen to resources such as applications, infrastructure, and procedures. At the core, problem management capabilities include RCA, incident analysis, knowledge management, collaboration, and reporting. The M&G Guide recommends that you extend and update existing incident and problem management capabilities with specific roles and responsibilities, support escalation paths, and standard operating procedures for your AWS environments.

Resource inventory management

Resource inventory management provides the ability to define and control the components of services and infrastructure, and maintain accurate configuration records. The configuration management database (CMDB) ensures assets required to deliver services are properly controlled, and that accurate and reliable information about those assets is continuously available. The goal of configuration management is to define and control service components and maintain accurate configuration records. CMDB provides a single source of truth of resources and their relationships. The CMDB enables resource transparency for:

- Compliance with corporate governance
- Audit support
- Visibility into service assets and their dependencies
- Cost optimization
- Effective change (impact analysis) and release management
- Faster incident and problem resolution

The CMDB ensures that systems configuration management is ubiquitous and scalable. As AWS adoption progresses and more applications are deployed and running on AWS, the complexity and interdependence might become challenging. The M&G Guide recommends using hierarchical configuration management tools to help manage configurations across account, environment, stack, application, and versions.

Change management

Change management provides the ability to request, prioritize, authorize, and approve, schedule, and implement changes to assets. This helps provide a balanced approach to modify IT services

while minimizing the risk to production environments. The evidentiary controls included with change management functions allow for ease of audit and compliance reporting. Distributing your infrastructure as code in your multi-account framework should be part of change management processes and approval. This basis facilitates the automation of changes and provides for the documentation, review, and storage of changes in configuration management tools. The M&G Guide recommends that you develop an iterative approach for integrating change management with automation and distribution functions.

Interoperable functions

The eight management and governance functions, supported by AWS services and AWS Partner solutions, work together and interoperate to reduce complexity. Outputs from these functions are used to inform or integrate with other functions.

For service management this includes:

- Resources, applications, and accounts provisioned from ITSM tools with embedded **Controls**.
- **Network connectivity** designs, including boundaries and isolation, provisioned, updated, and recorded with ITSM tools.
- Aligned **Identity and access management** across ITSM and AWS tooling.
- Incorporating **Security management** runbooks and playbooks with your service management framework.
- Aggregating **Monitoring and observability** findings alongside incident management frameworks within ITSM.
- ITSM capabilities and configuration libraries aligned to **Cloud Financial Management** reporting and insights.
- **Sourcing and distribution** requests initiated from ITSM to support change and incident management.

Implementation priorities

Implementation priorities for service management include incorporating or establishing policies and procedures to account for cloud services. It is a best practice to involve your operations teams early in a preparation phase before migrating production workloads using iterative production readiness criteria and checklists. Use of ITSM tools and integration with AWS services can be accomplished in an iterative manner. For example, initially tooling and integration can be implemented manually and progressing to full automation in later phases. The M&G Guide recommends that you implement service management framework processes in a phased approach that establishes a cloud operational baseline and connects with your system of record declared ITSM tools.

Integrate provisioning processes with the ITSM tool suite

Infrastructure as code templates are the cornerstone of distribution, but also enable full service management. Integrating your provisioning and distribution processes with your ITSM tool suite is an essential step. Prioritize common requests for self-service in your ITSM tool. When creating a service template naming convention across your accounts and workloads, establish a lifecycle management process using approvals and standard workflow processes. Build templates with operational metadata (tags) and parameters to be populated in alignment with the configuration management system. Sensitive data should not be used for tag keys or values. Ensure that template access is enforced across both distribution and ITSM tooling permissions.

Enable event, incident, and problem management across your environment

Enable an issue management mechanism such as ticketing across your AWS accounts. Integrate event management with other ITSM processes, such as incident and change management. Identify service owners, dependencies, and third-party integrations required to scale effectively with updated event store and sourcing patterns. Extend existing roles, procedures, and governance activities to accommodate cloud scale. This extension includes: incident and problem management roles and responsibilities, support escalation paths, and standard operating procedures. Establish remediation runbooks for common issue patterns to improve mean time to repair. Use game-day scenarios to validate support procedures. Analyze service trends to help provide recommendations and improve designs of your applications, resources, and environments on AWS.

Identify accounts, environments, and resources that require asset tracking

Identify the accounts, environments, and resources that require asset tracking for compliance. Update registration in the CMDB as part of the account and asset provisioning and decommissioning processes. Track standards (regulatory, enterprise, security, financial, etc.) and compliance of required resources within AWS by creating integrations to your ITSM tooling that will enable a federated view of your AWS services and resources.

Align change request procedures and policies for rapid cloud deployment

Align or create change request types in your policies and procedures that allow rapid deployment of resources. Determine which service templates (infrastructure as code) can be deemed as pre-approved changes. Determine how continuous integration and continuous delivery (CI/CD) pipelines will be accounted for in your change procedures. Align service and resource change requests through your ITSM tools.

Connect your ITSM system of record tooling to AWS

Implementing [AWS Service Management Connectors](#) is composed of three key steps: configuring AWS native services, configuring ITSM tooling, and validating your configuration and connectivity.

1. Configuring AWS management and governance services

AWS Service Management Connectors enable integration features for AWS Service Catalog, AWS Config, AWS Systems Manager Automation, AWS Systems Manager OpsCenter, and AWS Security Hub. AWS Service Management Connectors requires baseline configurations and permissions to these services. For more information on these specific requirements, refer to the following documentation:

- [AWS Service Management Connector for ServiceNow](#)
- [AWS Service Management Connect for Jira Service Management](#)

2. Configuring ITSM Tool

AWS Service Management Connectors enable integration for ServiceNow and Atlassian Jira Service Management. The connector requires you to download the Connector plugin (scoped app) from the respective ITSM tool platform. For more information on these specific requirements, refer to the:

- [\[ServiceNow Store App\] AWS Service Management Connector for ServiceNow](#)
- [\[Atlassian Marketplace App\] AWS Service Management Connect for Jira Service Management](#)

3. Validating configurations

Once the AWS and ITSM Tooling configurations are complete, the final step is to validate that AWS and the respective ITSM tool connected successfully and the intended service management actions are enabled. For more information on these validation actions, refer to the [public documentation](#) for the desired ITSM tool.

AWS service management tools

The AWS Management & Governance product suite allows you to enable, provision, and operate AWS resources to determine the health and predictability of your cloud workloads. The following AWS services can be used to help you meet the prescribed benefits of the M&G Guide, establish a cloud operational baseline, and align to your ITSM solution implementation:

[AWS Systems Manager](#) is a management service that helps you automatically collect software inventory, apply operating system patches, create system images, and configure Windows and Linux operating systems. These capabilities help you define and track system configurations, prevent drift, and maintain software compliance of your Amazon EC2 and on-premises configurations. By providing a management approach that is designed for the scale and agility of the cloud but extends into your on-premises data center, Systems Manager makes it easier for you to seamlessly bridge your existing infrastructure with AWS.

[AWS Systems Manager Explorer](#) is a customizable dashboard providing key insights and analysis into the operational health and performance of your AWS environments. Systems Manager Explorer aggregates operational data from across AWS accounts and AWS Regions to help you prioritize and identify where action might be required.

[AWS Systems Manager Automation](#) allows you to safely automate common and repetitive IT operations and management tasks. With Systems Manager Automation, you can use predefined runbooks, or you can build, run, and share wiki-style automated playbooks to enable AWS resource management across multiple accounts and AWS Regions. The runbooks can also be used to remediate issues such as AWS Systems Manager OpsCenter OpsItems.

[AWS Systems Manager OpsCenter and Incident Manager](#) provide an issue management mechanism that you can enable across your AWS accounts. This service provides a central location where operations engineers and IT professionals can view, investigate, and resolve operational issues related to any AWS resource. OpsCenter aggregates and standardizes operational issues, referred to as OpsItems, while providing contextually-relevant data that helps with diagnosis and remediation.

[AWS Systems Manager Change Manager](#) simplifies the way you request, approve, implement, and report on operational changes to your application configuration and infrastructure in the AWS Cloud and on premises. With Change Manager, you can use pre-approved change workflows to help avoid unintentional results when making operational changes. Change Manager helps you safely implement changes, while detecting schedule conflicts with important business events and automatically notifying impacted approvers. Using Change Manager's change reports, you can monitor progress and operational changes across your organization, providing improved visibility and accountability.

[AWS Config](#) is a service that enables detective controls to assess, monitor, and evaluate the configurations of supported AWS resources. AWS Config monitors and records AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. With AWS Config, you are able to not only track the relationships among resources and quickly review the history of the resource's configuration but you can also identify the compliance of resources based on defined config rules. Use [AWS Config](#) to view status, compliance, and the relationships of your provisioned AWS resources. [Getting started with AWS Config](#) entails turning on recording and establishing the right detective controls based on your governance and compliance requirements.

[AWS Security Hub](#) is a service that gives you a comprehensive view of your security alerts and security posture across your AWS accounts. With Security Hub, you have a single place that aggregates, organizes, and prioritizes your security alerts, or findings. Security Hub findings can also enable your organization to create incidents within ITSM tooling via integrations depending on the finding's severity level.

[AWS Service Catalog](#) allows you to centrally manage commonly deployed AWS services and provisioned software products. The curated products are vetted and enable end users to request services and resources as needed without having direct permissions enabling segregation of duty. AWS Service Catalog also helps your organization achieve consistent governance and compliance requirements, while enabling users to quickly deploy only the approved AWS services they need.

AWS Service Management Connectors

The M&G Guide recommends using AWS-supplied service management connectors that enable you to access AWS services and features in familiar ITSM tooling, such as ServiceNow and Atlassian. By using your existing service management tools to provide governance, your organization can accelerate its migration and adoption of AWS at scale.

The AWS Service Management Connector for ServiceNow enables ServiceNow end users to provision, manage, and operate AWS resources natively through ServiceNow. With the connector, ServiceNow administrators can:

- Provide pre-approved, secured, and governed AWS resources to end users through AWS Service Catalog.
- Run automation playbooks through AWS Systems Manager.
- Track resources in the CMDB powered by AWS Config seamlessly on ServiceNow with the AWS Service Management Connector.
- Define new resource types based on ServiceNow CMDB tables and synchronize these with AWS Config custom resources.
- Configure syncing AWS Security Hub findings to ServiceNow incidents or problems.

The AWS Service Management Connector for Jira Service Management allows Jira Service Management end users to provision, manage, and operate AWS resources natively through Atlassian's Jira Service Management. Jira Service Management administrators can provide pre-approved, secured, and governed AWS resources to end users through AWS Service Catalog, create and manage operational items through AWS Systems Manager OpsCenter, run automation playbooks through AWS Systems Manager Automation and track resources in a configuration item view powered by AWS Config seamlessly on the Jira Service Management with the AWS Service Management Connector.

If you would like support implementing this guidance, or assisting you with building the foundational elements prescribed by the M&G Guide, we recommend you review the offerings provided by [AWS Professional Services](#) or the AWS Partners in the [Built on Control Tower program](#).

If you are seeking help to operate your workloads in AWS following this guidance, [AWS Managed Services \(AMS\)](#) can augment your operational capabilities as a short-term accelerator or a long-term solution, letting you focus on transforming your applications and businesses in the cloud.

Integrated service management partners

The M&G Guide recommends you consider the following questions when choosing an AWS Partner solution for service management:

- Does it provide ITSM process enablement?
- Does it allow users to relate ITSM processes (that is, relating incidents to change requests)?
- Does it enable business workflows and approvals?
- Does it allow for configuration, customization, and integration to other systems and platforms?
- Does it enable the ability to create reports and dashboards?
- Does it enable self-service for business and service requests?

[Atlassian Jira Service Management](#) is service management software for modern IT teams. Jira Service Management request types and projects enable self-service for developers and end users to order IT services based on request fulfillment approvals and workflows.

[BMC Helix ITSM](#) (formerly BMC Remedy) is a service management tooling that uses emerging technologies to integrate IT service support functions.

[ServiceNow](#) is an enterprise service management platform that places a service-oriented lens on the activities, tasks, and processes that enable day to day work life in a modern work environment.

Monitoring and observability

Like security, monitoring and observability are required for all teams who operate and administer cloud applications and services. As described in the [Operational Excellence Pillar whitepaper](#), your teams must define, capture, and analyze operations metrics to gain visibility into workload events so that you can take appropriate action. In the management layer, this also means understanding operational metrics as you provide guardrails, network, security, and identity services in your management platform.

All of your teams, whether responsible for many cloud environments or a single application, must be able to understand the health of their operations easily. Your teams will want to use metrics based on operations outcomes to gain useful insights. You should use these metrics to make informed decisions, and as key inputs into each of the eight M&G Guide capabilities. AWS makes it easier to bring together and analyze your operations logs so that you can generate metrics, know the status of your operations, and gain insight from operations over time. These activities are supported centrally when you provide an observability solution for consumption, storage, analysis, and presentation of operational data for analysis.

As described in [Responding to Events](#), you should anticipate both planned operational events (such as, sales promotions, deployments, and failure tests) and unplanned ones (such as, surges in utilization and component failures). Use simulations, custom runbooks, and playbooks, and iterate to deliver consistent results when you respond to alerts. Defined alerts should be owned by a role or a team that is accountable for the response and escalations. You will also want to know the business impact of your system components and use this to target efforts when needed. Perform a root cause analysis (RCA) after events, and then introduce necessary changes and controls to prevent recurrence of failures or document workarounds.

In many enterprises, technical teams share integrated systems to monitor the services or infrastructure they manage. Shared observability systems bring together all the performance data for an entire organization, enabling teams to visualize the connections between services and components, collaborate with real-time data, and quickly identify the source of performance or security issues.

Observability systems collect data directly from applications, and AWS logging and service metric capabilities. AWS provides several services that can help increase your monitoring and observability posture. These services include [AWS CloudTrail](#), [Amazon CloudWatch](#), [Amazon Managed Service for Prometheus](#), [VPC Flow Logs](#), [AWS X-Ray traces](#), [Amazon EventBridge events](#), [Amazon Managed Grafana](#), [Elastic Load Balancing](#), and [AWS Network Firewall](#).

Interoperable functions

The eight management and governance functions, supported by AWS services and AWS Partner solutions, work together and interoperate to reduce complexity. Outputs from these functions are used to inform or integrate with other functions. For monitoring and observability this includes:

- Incorporating complementary **Controls** to observe changes and highlight them in the observability tools.
- **Network capabilities** that have VPC Flow Logs archived with the central infrastructure log archives and included in the log aggregation tools.
- Access to observability tooling defined by **Identity management** with changes to configuration recorded.
- **Security management** with observability by design, and specific systems to alert for changes in observability practices.
- **Service management** frameworks integrated to observability with operational tooling such as patch management and change and incident management.

- **Cloud Financial Management** with observability measures to alert for changes (including outliers in both upper and lower spend) in incurred and forecasted costs.
- **Sourcing and distribution** for both custom solutions and purchased solutions with specific logging integrated with your observability design.

Implementation priorities

Collect, aggregate, and protect event and log data

After you have provisioned your multi-account framework with [AWS Control Tower](#), you will have enabled the centralized collection of observable metrics and events to a log archive account, using CloudTrail. This collection uses a dedicated and encrypted Amazon S3 bucket, in a dedicated account, with access restricted. Encryption keys should be rotated on a regular basis to increase the security posture of the log archive. Use log aggregation to increase your visibility at scale. Use a service control policy to prevent changes to log configurations.

Use [AWS Systems Manager Quick Setup](#) with policies defined at the organization level, to deploy the [CloudWatch agent](#) to EC2 instances across your environments. This will enable system-level metrics to be aggregated alongside your other log data. Feed events into an event management or SIEM platform that has been adapted for AWS environments via API integration. Logs, metrics, and traces should be collected across the following observability categories:

- **Control plane observability**—Enable CloudTrail logging to capture API call activity. As accounts are provisioned from AWS Control Tower, a service control policy will be provisioned which prevents changes to the CloudTrail configuration and log archive account.
- **Network observability**—Monitor and track network events and behaviors including network firewalls, network intrusion detection and prevention, load balancers, AWS WAF, proxy tools, and network flow data collection and monitoring. Track events and behaviors related to access controls (for example, security groups and firewall services) and monitor network activity with Amazon VPC Flow Logs and packet inspection with Amazon VPC Traffic Mirroring.
- **Workload observability** (including distributed tracing within your application observability solutions for serverless, container, storage, and database workloads)—Track events and behaviors at scale as workloads communicate within the cloud environment as a whole, in addition to the local application logs on individual systems.

Build capabilities to analyze and visualize log events and traces

Build capabilities to interactively search and analyze your local and centralized log data. As you scale with AWS, you will need to include the ability to index and visualize your log insights and metrics. Correlate logs and performance metrics across different types of data collection to drive meaningful conclusions and insights. Use rules to effectively respond to security events or patterns identified in your logs. Develop a nearly continuous monitoring strategy to scale your observability capabilities as you migrate and grow solutions on AWS.

Add detection and alerts for anomalous patterns across environments

Proactively assess environments for known vulnerabilities and add detection for anomalous patterns of events and activities. Monitor for unusual activity or behavior related to users and workloads using tools

such as [Amazon GuardDuty](#), [Amazon CloudWatch ServiceLens](#), and [Amazon CloudWatch dashboards](#). Start with patterns or indicators of unintended account usage or permissions including any login activity to cloud management consoles, any changes, or attempted changes to important cloud objects and data, and any creation, deletion, or modification of credentials or cryptographic keys. Detect incidents and patterns of denials of access, unidentified network traffic, atypical increases in cloud services costs, and unusual application traffic behavior. Configure Amazon CloudWatch alarms, GuardDuty, and SIEMs to initiate alerts and notifications using [Amazon Simple Notification Service](#) (Amazon SNS). Identify anomalous behavior with [Amazon DevOps Guru](#), [AWS X-Ray Insights](#), and [Amazon CloudWatch Contributor Insights](#).

Define, automate, and measure response and remediation

Establish expected behavior thresholds paired with business metrics to understand KPIs for workloads and environments. Determine appropriate incident and response actions to pursue. Use SIEM solutions to monitor workloads in real-time, identify security issues, and expedite root-cause analysis.

Automations can be initiated by [several different triggers](#), such as EventBridge, State Manager associations, and maintenance windows. By using triggers, you can run automations because of a specific event or on a scheduled basis. Events can be derived from pattern matching using Amazon CloudWatch alerts or SIEM. Take advantage of security orchestration, automation, and response platforms (SOAR) while pairing with responses created from recorded events with tools like AWS Lambda. Maintain a process to continually improve mean time to identify (MTTI) root cause and mean time to respond (MTTR) to problems. Establish and measure goals to reduce the time to detect, identify, and remediate issues. This can also be done in conjunction with post-mortem or lessons learned procedures that align with your existing software development lifecycle or management practices.

AWS observability tools

The following AWS services can be used to help you meet the prescribed benefits of the M&G Guide:

[AWS CloudTrail](#) provides event history of your AWS API activity, including actions taken through the AWS Management Console, AWS SDKs, command line tools, and other AWS services that you specifically enable. By default, AWS Control Tower uses AWS CloudTrail where it is enabled as a multi-account guardrail control, and stores control plane logs in a centralized account. Use the central account to store and analyze all trails.

[Amazon CloudWatch](#) is a monitoring and observability service built for DevOps engineers, developers, site reliability engineers, and IT managers. CloudWatch provides you with data and actionable insights to monitor your applications, respond to system-wide performance changes, optimize resource utilization, and get a unified view of operational health. CloudWatch collects monitoring and operational data as logs, metrics, and events, providing you with a unified view of AWS resources, applications, and services that run on AWS and on-premises servers. CloudWatch should be used to integrate AWS service, resource, and application logs.

With [AWS X-Ray](#), you can understand how your application and its underlying services are performing to identify and troubleshoot the root cause of performance issues and errors. X-Ray provides an end-to-end view of requests as they travel through your application, and shows a map of your application's underlying components. You can use X-Ray to analyze both applications in development and in production, from simple three-tier applications to complex microservices applications consisting of thousands of services.

To visualize, query, and correlate your metrics, logs, and traces at scale, and to provide a deeper analysis of your observability data, we recommend [Amazon Managed Grafana](#). Developed in collaboration with

Grafana Labs, Amazon Managed Grafana manages the provisioning, setup, scaling, and maintenance of Grafana servers, decreasing the need for you to manage the underlying infrastructure. Based on open source Grafana with enhanced features such as single sign-on support, Amazon Managed Grafana enables you to query, visualize, alert on, and understand your observability metrics, logs, and traces no matter where the data is stored, such as querying container metrics stored in [Amazon Managed Service for Prometheus](#).

Amazon Managed Service for Prometheus is a fully managed, Prometheus-compatible service that enables you to securely ingest, store, and query metrics from container environments. Amazon Managed Service for Prometheus scales on demand, collecting and accessing performance and operational data from container workloads on AWS and on premises. With Amazon Managed Service for Prometheus, you can use the open source Prometheus query language (PromQL) to monitor the performance of containerized workloads without having to manage the underlying infrastructure. Amazon Managed Service for Prometheus automatically scales as your workloads grow or shrink, and uses AWS security services to enable fast and secure access to data. You can use Amazon Managed Service for Prometheus to collect and query metrics from AWS container services including Amazon Elastic Kubernetes Service (EKS) and Amazon Elastic Container Service (Amazon ECS), via AWS Distro for OpenTelemetry or Prometheus servers as the collection agents.

[Amazon OpenSearch Service](#) (successor to Amazon Elasticsearch Service) is a distributed, open-source search and analytics suite used for a broad set of use cases, such as real-time application monitoring, log analytics, and website search. Amazon OpenSearch Service provides a highly scalable system for providing fast access and response to large volumes of data with an integrated visualization tool, OpenSearch Dashboards, that makes it easy for users to explore their data. Like Elasticsearch and Apache Solr, OpenSearch Service is powered by the Apache Lucene search library. OpenSearch Service and OpenSearch Dashboards were originally derived from Elasticsearch 7.10.2 and Kibana 7.10.2.

If you would like support implementing this guidance, or assisting you with building the foundational elements prescribed by the M&G Guide, we recommend you review the offerings provided by [AWS Professional Services](#) or the AWS Partners in the [Built on Control Tower program](#).

If you are seeking help to operate your workloads in AWS following this guidance, [AWS Managed Services \(AMS\)](#) can augment your operational capabilities as a short-term accelerator or a long-term solution, letting you focus on transforming your applications and businesses in the cloud.

Integrated observability partners

The M&G Guide recommends you consider the following questions when choosing an AWS Partner solution for observability:

- Does it continually monitor security risk? Does it provide or implement configuration changes across cloud environments?
- Does it offer threat detection, logging, and reports that align to your specific enterprise standards or regulatory compliance needs?
- Does it provide automation to address issues ranging from cloud service configurations to security settings as they relate to governance, compliance, and security for AWS resources?
- Does it highlight over-allocation of permissions and permissive traffic policies?
- Does it allow for inter-operability between observability and automation?

The following integrated monitoring and observability AWS Partners have provided integrations that align to the M&G Guide, and are available for entitlement in AWS Marketplace:

[AppDynamics](#) is designed for production and pre-production environments, and gives you visibility into your entire application topology from a single pane of glass. It allows you to monitor and manage:

- End-to-end performance of complex distributed applications with Application Performance Management
- Real user monitoring and browser synthetic monitoring
- Insights from correlating server and data base performance with application performance
- Real-time business awareness into IT operations, customer experience, and business outcomes with transaction, log, browser, and mobile analytics

[Datadog](#) collects and unifies data streaming from complex AWS environments, with a one-click integration for pulling in metrics and tags from over 70 AWS services. You can deploy the Datadog Agent directly on your hosts and compute instances to collect metrics with greater granularity—down to one-second resolution. And with Datadog's out-of-the-box integration dashboards, you get not only a high-level view into the health of your infrastructure and applications but also a deeper visibility into individual services such as AWS Lambda and Amazon EKS.

[Dynatrace](#) provides software intelligence to simplify cloud complexity. With automatic and intelligent observability at scale, it delivers precise answers about the performance of cloud platform environments. It seamlessly integrates with AWS Control Tower and securely governs AWS accounts as soon as they are created. A smart baselining capability adapts dynamically and monitors the performance of your environments in real time.

[ExtraHop Reveal\(x\) 360](#) provides multi-layered visibility, threat detection, and investigation in AWS via integrations with Amazon VPC Traffic Mirroring for packet-level visibility and VPC Flow Logs for broad coverage. ExtraHop is an AWS Security Competency Partner and offers a free trial of Reveal(x) 360. To learn more, see Reveal(x) 360 in the AWS Marketplace.

[New Relic One](#) includes a Telemetry Data Platform to ingest, analyze, and alert on your metrics, events, logs, and traces, full-stack observability to quickly visualize and troubleshoot your entire software stack in one connected experience, and applied intelligence to automatically detect anomalies, correlate issues, and reduce alert noise.

[Splunk Cloud](#) enables you to search, monitor, and analyze machine data from various sources to gain valuable intelligence and insights across your entire organization.

[Sysdig Monitor](#) provides real-time, deep visibility into rapidly changing AWS Cloud and container environments. You can resolve issues faster using granular data derived from Linux system calls enriched with cloud and Kubernetes context along with Prometheus metrics. With Sysdig, cloud teams can optimize costs by visualizing capacity utilization across regions, services, and clusters.

Cloud Financial Management

Managing cloud finance requires evolving your existing finance processes to establish and operate with cost transparency, control, planning, and optimization for your AWS environments. Cloud Financial Management (CFM) involves more than just reining in costs. It is about how to embrace the agility, innovation, and scale of AWS to maximize the value that the cloud provides to your business.

Applying traditional, static waterfall planning, IT budgeting, and cost assessment models to dynamic cloud usage can create risks, lead to inaccurate planning, and result in less visibility. Ultimately, this results in a lost opportunity to effectively optimize and control costs and realize long-term business value. To avoid these pitfalls, actively manage costs throughout the cloud journey, whether you are building applications natively in the cloud, migrating your workloads to the cloud, or expanding your adoption of cloud services.

CFM solutions help transform your business through cost transparency, control, forecasting, and optimization. These solutions can also help enable a cost-conscious culture that drives accountability across all teams and functions. Finance teams can see where costs are coming from, run operations with minimal unexpected expenses, plan for dynamic cloud usage, and save on cloud expenses while teams scale their adoptions on the cloud. Sharing this with engineering teams can provide necessary financial context for their resource selection, use, and optimization.

Organize and report with user-defined methods

To understand your AWS costs and optimize spending, you need to know where those costs are coming from. This requires a deliberate structure for your accounts and resources, to enable finance to track spending flows and ensure that teams are accountable for their portion of the bottom line. The M&G Guide recommends appointing a dedicated owner or team to develop, obtain stakeholder buy-in, monitor, and actively design and implement the cost allocation model to drive accountability and cost-conscious cloud consumption. Will you charge cloud and internal costs out to business function or product teams (internal chargeback)? Or, will you make the costs visible (show-back model)? The former drives accountability, but can be perceived as a tax. The latter requires less overhead to administer but may not drive as much accountability for costs.

Manage billing and control costs

Establish guardrails and set governance to help ensure that expenses stay in line with budgets. It is critical to establish basic governance policies to guide permissions and accessibility as related to cost control. Customers who are successful doing this have centralized ownership through designated teams, such as a Cloud Center of Excellence (CCoE), or a Cloud Business Office (CBO). These teams help design and implement governance mechanisms and drive best practices company-wide.

Use license management

Cloud Financial Management includes a perspective on vendor license management. License management validates compliance of your purchased assets across AWS. Aligning license management

capabilities with your financial management can help you understand a complete cost picture and make appropriate procurement decisions as described in [Sourcing and distribution \(p. 51\)](#).

Plan with flexible budgeting and forecasting

Once you've established visibility and cost controls, plan, and set expectations for spending on cloud projects. The tools and capabilities described in the M&G Guide are designed to give you the flexibility to build dynamic forecasting and budgeting processes, and help you stay informed on whether costs are adhered to, or exceed budgetary limits. They also help you act quickly in response to negative variances in forecasted spend, and mitigate risks of overspending and failing to meet the return-on-investment target.

Select a unit metric to support your business

[Unit metrics](#) allow you to normalize your cost and usage information to a common measure, and tie them back to your business outcome. These normalized metrics bring consistency, fairness, and clarity to your IT planning and evaluation cycle. You can use the unit metric to gauge how efficient your team uses technology resources, and you can also use it to forecast how much you need to invest as your business grows. The unit metric is a straightforward tool that helps you get buy-in and tell your IT value story inside your organization.

The objective of a unit metric is to present incremental cost or incremental consumption in terms of a unit of the demand driver. A demand driver is a factor that is correlated to AWS spend or AWS resource consumption. The quantity of AWS resources consumed and the cost of using those resources are directly impacted by increases or decreases in the demand driver. To learn more about this topic, refer to the [Unit metrics blog](#).

Optimize costs with pricing and resource recommendations

Optimizing costs begins with having a well-defined strategy for your new cloud operating model. This should start as early as possible in your cloud journey, setting the stage for a cost-conscious culture reinforced by the right processes and behaviors. The M&G Guide recommends focusing on selecting the right purchase model and matching capacity with demand.

Interoperable functions

The eight management and governance functions, supported by AWS services and AWS Partner solutions, work together and interoperate to reduce complexity. Outputs from these functions are used to inform or integrate with other functions. For Cloud Financial Management this includes:

- Setting specific financial **Controls** and organizing them across your environments.
- Measuring cost and usage for your **Network connectivity** and adjusting assets to optimize costs accordingly.
- Granting access to financial reporting tools and real-time distribution of reports with **Identity management**.
- Using **Security management** to detect large variances in spend and respond accordingly.

- Using **Service management** tools to integrate AWS costs and forecast for chargeback or showback and provisioning.
- Using **Monitoring and observability** to incorporate aggregated financial findings from AWS and an enterprise portfolio perspective.
- Transparent financial controls: budget, cost, and forecast for the **Sourcing and distribution** of cloud resources.

Implementation priorities

The M&G Guide recommends that you implement your Cloud Financial Management capabilities with transparency in mind. This includes enabling your builder teams to see the financial impact of their cloud usage for the resources they provision, as well as to define specific controls related to the financial governance of your resources.

Enable Cloud Financial Management

Configure detailed information sources including [Billing and Cost Management](#) tools to create the reporting your organization needs. Regularly review (minimally on a monthly basis) the cost and usage by different dimensions to understand cost drivers. Establish organizational metrics, such as a [unit metric](#) to identify cost attribution categories as you scale. If required, ensure that your cost reporting includes all costs (labor, licensing, infrastructure, and more) to create the total cost of application management (TCAM).

Tag, track, and monitor resource costs across their lifecycle

A consistent and [well-designed tagging strategy](#) is required to manage and track costs across your AWS environments. Once resources in your environments are tagged, you must activate both [AWS-generated tags](#) and [user-defined tags](#) separately to use them in your cost reporting and analysis tools. Enforce [tag options](#) using distribution and preconfigured infrastructure as code templates for governance. Use [tag policies](#) to enforce and maintain consistent tags across your organization and resources.

Track resources over their lifetime and design your workloads to gracefully handle resource termination as you automatically identify and decommission non-critical or low utilization resources. Analyze the design, architecture, and all components of each workload or application for cost effectiveness, including license costs. Use [Managed entitlements](#) to track and help ensure that you have compliance with your established agreements while avoiding unexpected true-up bills for exceeding license limits. Determine if the component and resources will be running for extended periods (for commitment discounts), or dynamic and transiently running (for Spot or On-Demand Instances). Implement the appropriate [pricing models](#) for all components of your applications sourced from AWS Marketplace.

Establish mechanisms for cost governance

Create policies and mechanisms that define how resources are managed by your organization. The policies should cover cost aspects of resources and workloads, including creation, modification, and decommissioning over the resource lifetime. Create an obsolescence plan and defined retention period with lifecycle policies for resources as they are provisioned. Implement account structure, groups, and roles to help allocate costs and control who can create, modify, or decommission instances and resources in each group. Identify any new controls that could support a more efficient cost spend. Update your distribution of infrastructure as code templates in [AWS Service Catalog](#) so that cost is transparent and only approved instance sizes are available in a self-service manner across your multi-account framework. Enforce tagging of resources as they are provisioned to ensure effective cost governance.

Continually optimize for cost efficiency

Review historic spend patterns to detect cost spikes (one-time or recurring) or continual cost increases, assuming 14–30 days of historical spend. Implement mechanisms to periodically identify and [right-size instances based on current workload metrics](#) and characteristics. This can be evaluated using AWS Cost Explorer, AWS Trusted Advisor, and AWS Compute Optimizer, along with AWS Partner tools, such as VMware CloudHealth, Apptio Cloudability, and CloudCheckr. Cost efficiencies can also be achieved with Compute Savings Plans, Reserved Instances, Spot Instances for ephemeral workloads, and Amazon CloudFront Security Savings Bundle. Continually reviewing cost metrics can help to identify over purchased or underutilized savings mechanisms. For example, you can optimize your storage costs with S3 Intelligent-Tiering, Amazon S3 Glacier, or implementing lifecycle policies and purge processes. Centralize redundant or shared infrastructure to optimize costs. Manage demand and supply resources dynamically by implementing scheduled or automatic scaling, buffering, or throttling. Review new EC2 instance types as they are released to take advantage of a better price-performance ratio.

AWS Cloud Financial Management services and tools

The following AWS services can be used to help you meet the prescribed benefits of the M&G Guide:

The [AWS Cost & Usage Report](#) contains a comprehensive set of AWS cost and usage data, including additional metadata about AWS services, pricing, Reserved Instances, and Savings Plans. You should use this information to inform and create controls.

[AWS Cost Explorer](#) is a tool that enables you to view and analyze your costs and usage. You can explore your usage and costs using the main graph, the Cost Explorer cost and usage reports, or the Cost Explorer RI reports. You can view data for up to the last 12 months, forecast how much you're likely to spend for the next 12 months, and get recommendations for what Reserved Instances to purchase. You can use Cost Explorer to identify areas that need further inquiry and see trends that you can use to understand your costs.

AWS uses [cost allocation tags](#) to organize your resource costs on your cost allocation report, which makes it easier for you to categorize and track your AWS costs. AWS provides two types of cost allocation tags: AWS-generated tags and user-defined tags. AWS (or AWS Partners) defines, creates, and applies the AWS-generated tags for you, and you define, create, and apply user-defined tags.

[AWS Cost Anomaly Detection](#) is an AWS cost management feature that uses machine learning to continually monitor your cost and usage to detect unusual spends.

[AWS Budgets](#) allows you to set custom budgets to track your cost and usage on a wide variety of use cases. With AWS Budgets, you can choose to be alerted by email or Amazon SNS notification when actual or forecasted cost and usage exceed your budget threshold, or when your actual RI and Savings Plans utilization or coverage drops below your desired threshold. With AWS Budgets actions, you can also configure specific actions to respond to cost and usage statuses in your accounts, so that if your cost or usage exceeds or is forecasted to exceed your threshold, actions can be run automatically or with your approval to reduce unintentional over-spending. AWS Budgets integrates with multiple AWS services, such as AWS Cost Explorer, so that you can easily view and analyze your cost and usage drivers, AWS Chatbot, so you can receive budget alerts in your designated Slack channel or Amazon Chime room, and AWS Service Catalog, so you can track costs on your approved AWS portfolios and products.

[AWS Cost Categories](#) is an AWS cost management feature that enables you to group cost and usage information into meaningful categories based on your needs. You can create custom categories and map your cost and usage information into these categories based on the rules defined by you using various dimensions such as account, tag, service, charge type, and even other cost categories.

[Tag policies](#) are a type of policy that can help you standardize tags across resources in your organization's accounts. In a tag policy, you specify tagging rules applicable to resources when they are tagged. [AWS Resource Groups Tag Editor](#) allows you to add tags to—or edit or delete tags of—multiple AWS resources at once. With Tag Editor, you can search for the resources that you want to tag, and then manage tags for the resources in your search results.

[AWS License Manager](#) enables management of your software licenses from vendors across AWS and on-premises environments. AWS License Manager lets administrators define and enforce licensing rules that mirror the terms of their licensing agreements and prevent breaches. Portfolio administrators gain control and visibility of all their licenses with the AWS License Manager dashboard integrated with AWS Organizations and reduce the risk of non-compliance, misreporting, and additional costs due to licensing overages. Independent software vendors (ISVs) can also use AWS License Manager to easily distribute and track licenses.

[AWS Compute Optimizer](#) recommends optimal AWS resources for your workloads to reduce costs and improve performance by using machine learning to analyze historical utilization metrics. Compute Optimizer helps you choose optimal configurations for three types of AWS resources: Amazon EC2 instances, Amazon EBS volumes, and AWS Lambda functions, based on your utilization data.

[AWS Application Cost Profiler](#) provides you the ability to track the consumption of shared AWS resources used by software applications and report granular cost breakdown across tenant base.

[AWS Billing Conductor](#) simplifies the show-back and charge-back workflows for AWS Solution Providers and Enterprise customers. Using AWS Billing Conductor, you can customize your AWS pricing and monthly billing report according to the billing relationship between you and your end customers. A pro forma billing and cost and usage report is available for viewing and cost allocation purposes.

If you would like support implementing this guidance, or assisting you with building the foundational elements prescribed by the M&G Guide, we recommend you review the offerings provided by [AWS Professional Services](#) or the AWS Partners in the [Built on Control Tower program](#).

If you are seeking help to operate your workloads in AWS following this guidance, [AWS Managed Services \(AMS\)](#) can augment your operational capabilities as a short-term accelerator or a long-term solution, letting you focus on transforming your applications and businesses in the cloud.

Integrated Cloud Financial Management partners

The M&G Guide recommends you consider the following questions when choosing an AWS Partner solution for Cloud Financial Management:

- Does it help track spending at desired granularity and trace it back to approved initiatives, or allocate costs to the right business unit or project?
- Does it provide guardrails to control and govern cloud spending, which helps prevent unexpected or unexplainable costs, compliance, and security risks?
- Does it help estimate future costs and create financial predictability?
- Does it help analyze the implications of different AWS services, Availability Zones, or pricing models that can improve unit costs?

The following integrated Cloud Financial Management partners have provided integrations that align to the M&G Guide, and are available for entitlement in AWS Marketplace:

[CloudCheckr CMx](#) by NetApp is a unique, end-to-end governance solution that enables users to gain total visibility into their AWS environments, optimize costs, and perform cost allocation and chargebacks. CloudCheckr users can reduce their monthly cloud spend by 30% or more by acting on hundreds of

optimization recommendations. The platform has tailored dashboards and analytics to address reserved capacity purchase options, including Savings Plans and Reserved Instances. CloudCheckr CMx features an advanced cost query engine that surfaces the most detailed level of information available to help AWS users properly analyze consumption and associated costs.

[Flexera](#) cloud cost optimization simplifies cloud cost management and gives enterprises full visibility into cloud usage and costs. Powerful functionality enables cloud governance teams to work collaboratively with business units and cloud resource owners to report, manage, and optimize cloud spend.

[Kion](#) is a comprehensive enablement software solution that delivers visibility and control of cloud workloads. [Kion](#) provides insights to enable planning and reporting, allocates proper budgeting, and prevents overspending. [Kion](#) allows enterprises to manage their cloud presence at scale with automation and orchestration, financial management, and continuous compliance.

[Spot](#) by NetApp's solutions allows end users to use Amazon EC2 Spot Instances and Reserved Instance capacity without operational overhead and complexity. Spot by NetApp automates and optimizes your AWS infrastructure delivering SLA-backed availability and performance at the lowest possible cost. Machine learning and application-driven scaling enables you to run workloads of various sizes, providing an optimal blend of Savings Plans, Reserved, Spot, and On-Demand Instances.

Sourcing and distribution

Your sourcing and distribution strategy defines how you procure and deploy software and distribute infrastructure as code in a hub-and-spoke model across your cloud environment. Integration between sourcing systems, like AWS Marketplace, and your procurement system helps centralize governance for your software purchasing. With this integration, you can also use your existing workflows for procurement approval. AWS Marketplace provides for this integration using Commerce XML (cXML), an open standard communication protocol. With this feature, builders can find, buy, and deploy solutions, where IT administrators and procurement teams streamline approvals and spend directly from their procurement systems. We recommend that you further simplify your software procurement to work with your distribution systems so that software can be provisioned across your environments. This can be achieved by distributing infrastructure as code templates via a hub and spoke model.

Infrastructure as code templates are the cornerstone for agility in the cloud. These templates allow you to rapidly iterate and provision workloads and environments to meet your evolving customer needs. In the same manner, a consistent application of governance controls should be used to help meet ongoing and changing compliance requirements for internal enterprise standards and controls, as well as regulatory compliance frameworks. This applies to how you source software and distribute your templates across your multi-account strategy. Governance functions should be implemented proactively in order to verify you scale without introducing workflow bottlenecks. Having resources preconfigured for compliance (either with internal or external standards) allows you to reuse and scale your cloud assets without introducing manual review and reaction processes.

After you have obtained your software, or built your infrastructure as code templates, you then centrally manage and distribute these services, applications, resources, and metadata (tags). We recommend creating and recording these assets into a catalog, and distributing access to a curated assortment from the catalog in a hub and spoke manner. Meaning, all templates are stored at a top-level repository, and then replicated and shared as required in spoke repositories with permissions granted as required. This helps ensure that not only have you preconfigured and created immutable infrastructure as code assets, but that by curating the assortment you are introducing efficiencies as well for your builder teams. Providing the right template, preconfigured for governance controls, at the right time, in the right accounts, helps ensure that your teams can self-service any provisioning (including updating and shutting down) they require on an as-needed basis.

For example, you might have a collection of infrastructure as code templates in your hub catalog that have been preconfigured with Amazon S3, Amazon EC2, and Amazon RDS. In your member accounts, you would select the appropriate parameters for each AWS service in such a way that each team or end user would only see the templates that they require, with the preconfigured options for each parameter defined for their specific use. Your tagging strategy should be well defined and enforced at the Organization level, and a reusable repository of tag options should be available during the self-service provisioning of resources from the central catalog. This will help you achieve consistent governance, while also enabling users to quickly provision the approved assets with the right tags included. This self-service model is a core component of operating efficiently. With the ability to provision resources preconfigured for compliance, your development teams will be empowered, and able to move at their own pace or agility.

Interoperable functions

The eight management and governance functions, supported by AWS services and AWS Partner solutions, work together and interoperate to reduce complexity. Outputs from functions are used to inform or integrate with other functions. For sourcing and distribution this includes:

- Using **Controls** to enforce which resources, applications or accounts can be provisioned.
- **Network connectivity** defined as infrastructure as code and managed in the central catalog.

- Configuring sourcing and distribution with least privilege access with **Identity management**
- configuring **Service management** to utilize the available connectors for the service catalog and procurement systems.
- Managing and procuring **Security management** tools using centralized capabilities
- Logging, sourcing, provisioning, and distribution of each resource, application, account, and environment into centralized **Monitoring and observability** solutions.
- Providing **Cloud Financial Management** transparency across sourcing and distribution functions.

Implementation priorities

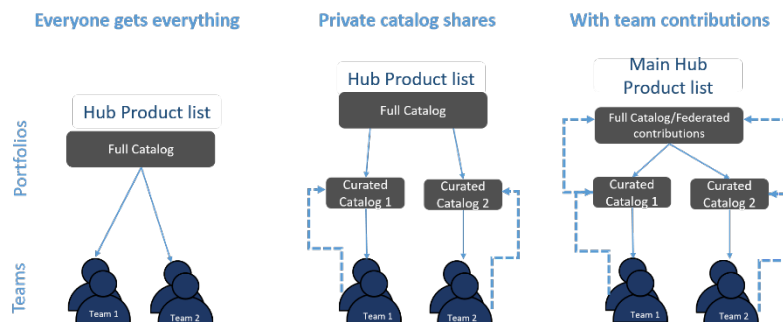
Implement a hub and spoke catalog model

In order to plan for scale, having the ability to reuse infrastructure as code templates across your environments, Regions, and accounts is a base requirement. The hub and spoke model for the distribution of infrastructure as code templates should be implemented alongside the templates sourced from AWS Partners. This will provide for administrative isolation between workloads, while also allowing for an expedited reuse of templates.

However, planning a centralized hub (catalog) for all templates might also slow down agility by adding unnecessary bottlenecks if not carefully implemented. Incorporating the concept of a federated hub and spoke architecture, allows for disparate builder teams to create and manage their own catalog repository (increasing their agility), while at the same time providing a mechanism for them to contribute to broader reuse by sharing their templates in the main hub catalog as well.

This distributed model is an effective mechanism for scale, as it allows builder teams to contribute to the broader repositories any innovations or new use cases, while still following the governance model of approvals in the main hub. Planning what shared resources to put in a central catalog (for instance, either Private Marketplace, or AWS Service Catalog), how builders can contribute and have their templates reviewed for the central catalog, as well as how the central catalog can be curated for specific use by each builder team will be important.

Depending on your requirements, you may elect to implement more than one version of the hub and spoke model. For example, for templates that are regulated or have strict enterprise requirements, the deviation from the approved template may not be allowed for higher-level environments. In this case, a mix of the restricted catalog distribution might be necessary for higher environments, where a team contribution approach would make sense for lower-level environments.



Distribution models for catalog sharing

Curate templates for reuse

Once you have codified your solutions as infrastructure as code templates and defined your hub and spoke model, you will need to define two categories of templates for each of your spoke accounts:

Provisioned/Enforced and *Available to Consume*. *Provisioned/Enforced* are a set of templates that will be provisioned directly from the management account into each member account as foundational capabilities. *Available to Consume* are a set available for builders to browse and self-service provision.

For example, templates can be provisioned directly to spoke accounts through stacks, stack sets, Lambda functions, or Step Functions. Alternatively, you can distribute a self-service curated assortment to the spoke accounts using AWS Service Catalog. Curating templates that are specific for your builders without overloading a catalog with all available templates can allow for faster discovery and reuse of primary templates.

Apply default parameters for reuse

Implement infrastructure as code templates that allow default parameters to be pre-selected for use by your builders. This enables them to align to governance without having to evaluate the details of each parameter option. Where possible, set default values for parameters, sparing end users (builder teams) the need to choose or possibly make incorrect choices. This approach exposes only what is needed for setup. For example, AWS Service Catalog implements this with a constraint capability that allows for a choice of parameters to vary based on the builder team. This is done so that it is pre-configured when the builder team self-service provisions a template.

Implement a lifecycle management and version distribution system

Maintain infrastructure as code template versions throughout their development lifecycle. This is done for the templates, and the services provisioned from the templates. Using the hub and spoke model you implemented for your catalog, you can define if a forced update is required at a spoke level, or if concurrent versions should be available for self-service provisioning, and which versions need to be marked for obsolescence. Using a hub and spoke catalog will also help allow the audit and distribution of new versions as required.

Identify a robust tagging strategy

Tags are key-value textual metadata that can be applied to most resources, and are a critical component for successfully implementing your environments in a scalable, efficient manner. Tags are most commonly used for:

- Organizing resources in the AWS Management Console
- Allocating costs
- Business-related metadata
- Enabling automation and operations support (for example, backups, and the ongoing management of environments)
- Indicating resource risk profiles and compliance requirements (for example, regulatory, and data privacy)

[Centralizing your tags into an enforced library](#) of metadata will help you expedite their proper application while provisioning and updating resources. Adding tagging policies in your controls functions can help supplement this enforcement. AWS Resource Groups and the Resource Groups Tagging API enable programmatic control of tags. This makes it easier to centrally manage, search, and filter tags and resources. In addition, AWS Service Catalog provides a TagOption library where the tag-option key-value pairs are stored. The TagOption library makes it easier to enforce a consistent taxonomy, the proper tagging of resources, and defines user-selectable options. Meaning, with only a defined set of choices available, misconfigured or tags in error are less likely to occur.

As the catalog is curated for the hub and spoke model, these tags can also be used as required parameter values and enforced with constraints on your templates. In addition, with the TagOption library from AWS Service Catalog, you can deactivate TagOptions and retain their associations to portfolios or products, and reactivate when needed. This approach not only helps maintain library integrity, it also allows you to manage TagOptions that might be used intermittently, or under special circumstances. With AWS Service Catalog TagOptions, auto-tags are also built into each product and portfolio.

Manage entitlements

You should enable controls that allow only authorized users and workloads to consume a license within vendor-defined limits. This helps reduce the risk of costly audits and unexpected licensing tune-ups.

Use [AWS License Manager](#) to manage software licenses from vendors, such as Microsoft, SAP, Oracle, and IBM, across your AWS and on-premises environments. This service allows you to define rules based on your licensing agreements to prevent license violations. You can set notification alerts along with license rules to help ensure that you are meeting license requirements. [Managed entitlements](#) enable you to distribute, activate, and track software license entitlements acquired in AWS Marketplace through AWS License Manager.

Enable Private Marketplace

Private Marketplace serves as a curated catalog of purchased products (software, data, and professional services) and should be implemented in a hub and spoke pattern (management-member) such that spoke accounts are limited to subscribing to only the approved software. This product governance is done to control software costs and streamline legal and contractual reviews. Create a Private Marketplace at the management account level to serve as the primary hub.

For the member accounts, determine whether the software can operate centrally as a shared service, or is required to be available in individual accounts. If the software is required in a set of member accounts, you will need to individually subscribe and associate the software you've curated in the Private Marketplace to enable use and self-service provisioning. Once an association is completed, an account is governed by the Private Marketplace. The users in the member accounts will be limited to the software that is curated in their specific Private Marketplace experience. After the subscription is set for the member accounts, the same catalog distribution capabilities can be used to provision self-service products (for example, using AWS Service Catalog).

Integrate with procurement systems

Complement your existing procurement processes, with integration to AWS Marketplace. This is done by extending your procurement systems (Coupa or SAP Ariba) to Private Marketplace so your users can use the existing procurement and approval processes to obtain software. Create the appropriate IAM-managed permissions, use AWS Marketplace to generate the necessary information to configure your procurement solution, and finally configure your procurement solution to complete the integration. For example, you can [set up a punchout](#), attach purchase orders to your AWS invoices, and then align your procurement processes to use the standard provisioning solutions.

AWS sourcing and distribution tools

The following AWS services can be used to help you meet the prescribed benefits of the M&G Guide:

[AWS Service Catalog](#) allows enterprises to create and manage catalogs of IT services that are approved for use on AWS. These IT services can include everything from virtual machine images, servers, software, and databases to complete multi-tier application architectures. AWS Service Catalog allows you to

centrally manage deployed IT services and your applications, resources, and metadata. This helps you achieve consistent governance and meet your compliance requirements, while enabling users to quickly deploy only the approved IT services they need.

[AWS Marketplace](#) provides listings across categories like security, networking, storage, and observability from independent software vendors that make it easy to find, test, buy, and deploy software that runs on AWS.

[AWS Private Marketplace](#) controls what products users in your AWS account, such as business users and engineering teams, can get from AWS Marketplace. Built on top of AWS Marketplace, it enables your administrators to create and customize curated digital catalogs of approved independent software vendors (ISVs) and products that conform to their in-house policies. Users in your AWS account can find, buy, and deploy approved products from your Private Marketplace, and ensure that all available products comply with your organization's policies and standards.

[Managed entitlements](#) enable you to distribute, activate, and track software license entitlements acquired in [AWS Marketplace](#) through [AWS License Manager](#). Your administrators can use AWS License Manager to automate the distribution and activation of software entitlements to your users and workloads across accounts in your AWS Organization. Managed entitlements also provide built-in controls that allow only approved users and workloads to consume licenses.

You can subscribe to, track, and manage software licenses at scale, as third-party products purchased in AWS Marketplace create managed entitlements in AWS License Manager. Once subscribed to a third-party Amazon Machine Image (AMI), container, or machine learning product or dataset in AWS Marketplace, you can automate the distribution of license entitlements across multiple accounts set up in AWS Organizations. This removes the need for each account to manually subscribe to each product. With a centralized administrative AWS account, you can govern license access by distributing entitlements to individual accounts on an as needed basis. This reduces the potential for duplicative buying or license overuse.

[AWS Solutions Library](#) offers a collection of cloud-based solutions for dozens of technical and business problems, vetted for you by AWS. You can use patterns from AWS Solutions Constructs if you want to build your own well-architected application, explore our collection of AWS Solutions Reference Architectures as a reference for your project, browse the portfolio of AWS Solutions Implementations for applications that you can automatically deploy directly into your AWS account, or choose an AWS Solutions Consulting Offer if you want help from an AWS Partner with deploying, integrating, and managing a solution.

[AWS Quick Starts](#) are automated reference deployments built by AWS Solutions Architects and AWS Partners. Quick Starts help you deploy popular technologies on AWS based on AWS best practices for security and high availability. These accelerators reduce hundreds of manual procedures into just a few steps so that you can build your production environment in minutes and start using it immediately. Each Quick Start includes AWS CloudFormation templates that automate the deployment and a guide that describes the architecture and provides deployment instructions.

[AWS Service Catalog Getting Started Library](#) provides a library of well-architected product templates so that you can get started quickly. You can copy any of the products in our Getting Started Library portfolios to your own account, then customize them to suit your needs. This library exists within the AWS Management Console to allow you an easy way of provisioning in your multi-account environments.

If you would like support implementing this guidance, or assisting you with building the foundational elements prescribed by the M&G Guide, we recommend you review the offerings provided by [AWS Professional Services](#) or the AWS Partners in the [Built on Control Tower program](#).

If you are seeking help to operate your workloads in AWS following this guidance, [AWS Managed Services \(AMS\)](#) can augment your operational capabilities as a short-term accelerator or a long-term solution, letting you focus on transforming your applications and businesses in the cloud.

AWS sourcing and distribution partners

The M&G Guide recommends you consider the following questions when choosing an AWS Partner solution for integrated sourcing solutions:

- Does the procurement solution support using cXML?
- Does it integrate with IAM Identity Center?
- Does it support notifications or other alerting mechanisms?
- Does it provide workflow integration where approvals are recorded for compliance evidence?

The following integrated sourcing and distribution partners have provided integrations that align to the M&G Guide, and are available to be entitled in AWS Marketplace.

[AWS Marketplace](#) integrates with procurement systems through an open standard communication protocol, Commerce XML (cXML). With this [procurement system integration](#), builders can find, buy, and deploy from thousands of solutions quickly, and IT administrators can streamline approvals and manage spend directly from their procurement system.

[Coupa](#) platform provides greater visibility into, and control over, how companies spend money, helping you maximize your spend under management, achieve cost savings, and drive profitability.

[SAP Ariba](#) is a cloud-based innovative solution that allows suppliers and buyers to connect and do business on a single platform. It improves the overall vendor management system of an organization by providing less-costly ways of procurement and making business simple. Ariba acts as supply chain, procurement service, and contract management.

Conclusion

The M&G Guide recommends a foundational set of eight functions that are needed to configure and build cloud-ready environments. We believe that by following the recommendations outlined in this guide, you will be able to set up your cloud environment for scalability and create what we call migration-ready environments. Deploying the eight capabilities covered in this guide in an interoperable manner will also enable you to realize cost-efficiencies while scaling workloads on AWS. This will help you improve your cloud value, and more importantly increase the speed to achieving value for your customers.

We have learned from customers migrating thousands of applications to AWS that they achieve success through complementing scaling with the progressive adoption of management and governance capabilities. The M&G Guide provides prescriptive guidance on how to accelerate value from these phases with recommended AWS M&G services and Technology Partner solutions powered by AWS services.

As you launch production workloads in your environments, we recommend that in addition to this guide, you also evaluate your AWS environments using the AWS Well-Architected Framework, which describes in detail how you can monitor, manage, and operate production workloads on AWS. As you work towards building and deploying production workloads on AWS, we recommend reviewing the Well-Architected Lens whitepapers as well.

For support implementing this guidance, or for assistance with building the foundational elements described in this guide, we recommend that you review the [Cloud Environment Workshop](#) and the offerings provided by [AWS Professional Services](#), including the [AWS Control Tower Partners](#) program.

If you are seeking help to operate your workloads in AWS following this guidance, [AWS Managed Services \(AMS\)](#) can help you to use AWS services using a growing library of automations, configurations and runbooks. AMS can augment your operational capabilities as a short-term accelerator or a long-term solution, letting you focus on transforming your applications and businesses in the cloud. AMS provides an operating model for your AWS fleet using detective guardrails, monitoring, security, and incident management best practices for your workloads and environments. AMS is available with two operations plans that offer specific sets of features with differing levels of service, technical capabilities, requirements, price and restrictions. AMS Accelerate helps you operate the day-to-day infrastructure management of your AWS environments; AMS Advanced extends AMS Accelerate to also include additional services such as landing zone management, infrastructure changes and provisioning, access management and endpoint security. You can also extend these plans with additional capabilities using [Operations on Demand](#). Choosing a plan to meet your specific scale and environments can help you to proactively accelerate your scale with AWS.

Contributors

The following organizations and individuals contributed to this document:

- AWS Well-Architected Management & Governance Guidance
 - Benjamin Andrew: Principal, Product Management -Technical, Amazon Web Services
 - Nam Le: Senior Partner Solutions Architect Control Services and Marketplace, Amazon Web Services
 - Jim McDonald: Senior M&G Specialist, Amazon Web Services
 - Belinda Quick: Principal, Product Management -Technical, Amazon Web Services
 - Mahdi Sajjadpour: Worldwide BD Leader for AWS Control Tower, AWS Service Catalog, Amazon Web Services
 - Tim Honychurch: Principal, Management & Governance, Amazon Web Services
- Controls
 - Bryan Miller: Principal, Control Services, Amazon Web Services
- Network Connectivity
 - Tom Adamski: Principal Networking Solutions Architect, Amazon Web Services
- Identity Management
 - Chris Mercer: Security Solutions Architect, Identity, Amazon Web Services
- Monitoring and Observability
 - Rich McDonough: Senior WW CloudOps Specialist Solutions Architect, Amazon Web Services
 - Bobby Hallahan: Senior WW CloudOps Specialist Solutions Architect, Amazon Web Services
- Security Management
 - Reef Dsouza: Senior Solutions Architect, Amazon Web Services
- Cloud Financial Management
 - Lisa Harnett: Product Marketing Manager, Amazon Web Services
 - Shankar Ramachandran: Principal Solutions Architect, Amazon Web Services
- Service Management
 - MaSonya Scott: Principal, Control Services, Amazon Web Services
 - Chandra Chappa: Service Management Specialist, Amazon Web Services
- Sourcing and Distribution
 - Raphael Sack: Principal, Control Services, Amazon Web Services
 - Murphy Tiggelaar: Principal Product Manager, Marketplace Catalog Services, Amazon Web Services

Document history

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

Change	Description	Date
Minor update (p. 23)	Added CyberArk as an integrated identity partner.	November 17, 2022
Minor update (p. 59)	Minor changes throughout.	November 2, 2022
Minor update (p. 16)	Added Amazon VPC IP Address Manager (IPAM) to network connectivity management tools.	May 24, 2022
Minor update (p. 59)	Corrected broken links and adjusted document path.	May 11, 2022
Minor update (p. 59)	Added Sysdig as an integrated observability and security management partner.	May 6, 2022
Minor update (p. 59)	Added ExtraHop Reveal(x) 360 as an integrated observability and security management partner.	April 20, 2022
Minor update (p. 48)	Added AWS Billing Conductor.	March 22, 2022
Document name change (p. 59)	The name of this document has changed from the <i>Management and Governance Lens</i> to the <i>Management and Governance Cloud Environment Guide</i> .	March 21, 2022
Minor update (p. 42)	Added Amazon OpenSearch Service to AWS observability tools.	February 14, 2022
Minor update (p. 1)	Minor updates to diagram.	January 31, 2022
Minor update (p. 59)	Minor changes.	January 6, 2022
Major update (p. 59)	Updated with new guidance throughout.	November 22, 2021
Major update (p. 59)	Management and Governance Lens first published.	April 28, 2021
Initial publication (p. 59)	Preview of Management and Governance Lens published.	December 3, 2020

Note

To subscribe to RSS updates, you must have an RSS plug-in enabled for the browser you are using.

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.

Definitions: Vocabulary

Application

An AWS application is the top-level node in a hierarchy of related cloud resource abstractions. Applications can include one or more of these abstractions (such as CloudFormation stacks and resource groups), which consist of one or more running AWS resources.

Compliance

In general, compliance means conforming to a defined rule, such as a policy, regulation, law, or specification. Organizations strive to ensure that they are aware of and take appropriate steps to identify the full scope of compliance requirements, and establish processes to ensure expectations are met and can be evidenced.

Governance

Governance focuses on the *what*, with emphasis on defining and achieving an organization's strategy. Governance activities emphasize how decisions are made, by whom, and accountability for the resulting decisions.

Key outputs of governance include:

- organizational strategy defined in formal policies and standards
- funding
- clear ownership and accountability for all aspects of implementation, ongoing assessment, and oversight

Guardrail

Guardrails define the boundaries that limit activity aligned with organizational control requirements. This would include standards in technology and technology resources, application architecture, operational requirements, and security requirements.

Management

Management activities focus on the *how*, looking to optimize processes to achieve organization's vision. Management is responsible for the execution and delivery of the strategy, with responsibility for allocation of resources to run and oversee day-to-day operations.

Key outputs of management include:

- Defined operational processes and tooling

- Budget and expense management
- Resource allocation and talent management
- Team structure and direction
- Performance monitoring (human and environment)

Monitoring

Monitoring is the systematic process of collecting and analyzing information to form an opinion.

Operations

Operations are activities taken to create, monitor, modify, expand, and remove applications and resources.

Oversight

Oversight is the responsibility over cloud resources and activities to validate compliance with organizational requirements.

Resource

Resource refers to cloud-based resources and are defined as cloud services, infrastructure and objects that make up part of a cloud platform supporting a variety of application types. This can refer to compute services such as Amazon EC2 instances, a database service such as Amazon RDS, cloud network services such as gateways and load balancers, and storage services such as Amazon S3 and Amazon Elastic File System (Amazon EFS).

Supervision

Supervision is the activity of directing, managing, or overseeing operational activities.

AWS glossary

For the latest AWS terminology, see the [AWS glossary](#) in the *AWS General Reference*.