Building a multicloud strategy in the financial services industry

AWS Prescriptive Guidance

AWS Prescriptive Guidance: Building a multicloud strategy in the financial services industry

Table of Contents

Building a multicloud strategy in the financial services industry

Chris Ibbitson, Dan MacKay, Rich McDonough, Ahmed Raafat, Aviad Tamir, and Matias Undurraga, Amazon Web Services

In the financial services industry, as organizations speed up their migrations to the cloud and transform their businesses, they frequently have to navigate the complexity of adopting multicloud operations. A multicloud environment can emerge from various needs, such as immediate data access requirements, regulatory concerns, or seeking differentiated capabilities from newly acquired companies that use different clouds. This paper discusses how financial institutions (FIs) can build a successful multicloud strategy in light of these challenges, and how AWS Cloud services can simplify the management and governance of diverse IT landscapes. Through AWS, FIs can more easily oversee their operations, make sure that their applications run smoothly, and analyze data across multiple platforms. They can balance innovation with operational efficiency to deliver better customer experiences, reduce costs, and innovate faster.

This paper extends the discussion in Proven practices for developing a multicloud strategy with guidance that is specific to the financial services industry. We recommend that you read that guide before reading this paper.

Cloud deployment strategies

AWS defines *cloud computing* as the on-demand delivery of IT resources over the internet with secure connectivity and pay-as-you-go pricing. Instead of owning and maintaining physical data centers and servers, you can access technology services—such as computing power, storage, and databases—on an as-needed basis from a cloud provider. Cloud computing allows FIs to avoid undifferentiated heavy lifting such as hardware procurement, maintenance, and capacity planning. When you adopt and deploy cloud solutions, you can choose from several options: single cloud, hybrid cloud, multicloud, and hybrid multicloud.

- **Single cloud** – This model uses only a single cloud service provider (CSP). Single cloud applications and workloads are implemented directly in the cloud, or they're hosted in another environment and then migrated to the cloud. These workloads might use lower-level infrastructure services (for example, virtual machines) from their cloud provider or they might also take advantage of higher-level, managed services (such as managed database services). In either case, this model adopts a single cloud provider and uses cloud services only from that provider.
- **Hybrid cloud** – This model distributes resources across an organization's own on-premises data center and at least one CSP. Typically, the purpose is to extend an organization's infrastructure into the cloud while maintaining private connectivity with existing internal systems that live on premises.
- **Multicloud** – In this model, an organization runs significant workloads on multiple CSPs, excluding software as a service (SaaS) solutions. An organization might choose this option as a deliberate strategy or as an unintentional result of individual teams, departments, or staff members having their own CSP preferences.
- **Hybrid multicloud** – In this model, hybrid and multicloud operations are not mutually exclusive, and many organizations use both operations. A discussion of this model is not within the scope of this paper.

Objectives for adopting a multicloud strategy

The adoption of a multicloud strategy is typically based on the following business drivers:

- **Mergers and acquisitions** – When an organization acquires, or merges with, companies that are using services or deploying workloads on different CSPs, it must optimize and harmonize the new IT landscape. This is time-consuming, because it involves diverse business priorities and requires integration effort. Adopting a multicloud strategy is key to successfully managing the complexity of varying systems and ensuring that standardized controls are in place.
- **Differentiated capabilities** – Some companies choose multiple CSPs to leverage the best capabilities and services for a given use case. Each cloud provider offers unique services that address specific requirements.
- **Line of business independence** – Centralized IT governance is often viewed as crucial for effective cloud deployment. However, allowing individual business units to choose their own cloud environments and provide their own governance can lead to optimized business outcomes, because different units might have requirements that are better served by specific cloud providers.
- **Contractual requirements** – An FI might have contractual requirements from their customers that mandate the use of a specific CSP.
- **Generating business value** – The key reason why an FI should look to operate in multiple CSPs is when it *adds value* for the business. An example would be where a key service that enables value to your business is available on only one CSP.

For more information about business drivers, see Tenet 1 in *Proven practices for developing a multicloud strategy*.

Anti-patterns for adopting a multicloud strategy

As outlined in the previous section, there are valid reasons for adopting a multicloud strategy. There are also misconceptions (or anti-patterns) that do not justify adopting a multicloud environment. These include:

- **Improved resilience** – Some FIs decide on a multicloud approach to reduce service disruption in the event that their primary CSP has an outage. Multicloud failover presumes that an application can be switched over to another cloud. However, this requires the company to maintain full portability between two CSPs, which adds complexity, risk, and additional work that make true portability nearly impossible. Adopting a multicloud strategy to improve resilience often adds complexity and is likely to reduce resilience instead of enhancing it.
- **Avoiding vendor lock-in** – Many firms cite the fear of lock-in—both contractual and technological—as a justification for pursuing a multicloud strategy. However, you should consider the potential cost of exiting one CSP and moving to another, along with the likelihood of needing to undertake that exit. You should also consider that the cloud is inherently more open (that is, it offers more open source technologies and standards) than traditional IT models. Avoiding lock-in doesn't require a multicloud approach. For more information, see the section Consider the advantages and disadvantages of vendor lock-in later in this paper.
- **Better pricing** – Traditional procurement approaches that use multiple suppliers to ensure price competitiveness do not adapt well to the pay-as-you-go purchasing, volume discounts, and price competition of the cloud. Companies often underestimate the operational complexity of managing multiple cloud environments and fail to properly evaluate whether the benefits justify this added complexity in competitive sourcing decisions.

For more information about these misconceptions, see Tenet 2 in *Proven practices for developing a multicloud strategy*.

Multicloud strategy recommendations

Adopting a multicloud approach requires balancing the need for security, resilience, and risk management with the need for flexibility and innovation. AWS guides organizations to adopt solutions that solve real-world business problems, and using multiple providers should similarly add value that outweighs potential additional costs and other challenges.

Cloud adoption provides a wealth of benefits that are essential to technology modernization, cost-effectiveness, and innovation. The technological and business changes involved in implementing workloads in the cloud require staff training and adjustments to core infrastructure, including networking, security, governance, and operations. If you are new to the cloud, we recommend that you start with a single provider to learn and achieve cloud benefits before you determine whether a multicloud strategy is right for your business. **It is typically more effective to master one cloud provider's operating model at one time.** This approach simplifies the upskilling of employees and minimizes the potential risk of cross-cloud dependencies and operational latency, while allowing you to establish consistent controls, policies, and playbooks.

Organizations that move to the cloud by concurrently adopting multiple cloud providers often regret that decision and the complexity it introduces.

As FIs gain cloud experience and expand the size of their cloud teams, they might align different lines of business with specific cloud providers that best meet their business requirements. The following sections offer a structured approach for navigating the complexities of multicloud environments, and guide you through the essential steps to design and implement an effective multicloud framework.

- Establish a multi-cloud enabled CCoE
- Evaluate exit strategy requirements
- Accelerate and optimize operations with SaaS applications
- Establish security and governance for each CSP
- Establish a modern multicloud data and AI strategy
- Assess the viability of contiguous workloads across CSPs
- Consider the advantages and disadvantages of vendor lock-in
- Establish a workload placement strategy

Establish a multicloud-enabled CCoE

If you have decided to adopt a multicloud strategy, we recommend that you centralize your cloud leadership efforts through a transformation office or a Cloud Center of Excellence (CCoE). A CCoE develops and evangelizes an approach for implementing cloud technology at scale across an organization. For successful cloud adoption, design your CCoE to include representatives who can speak for the teams and departments involved. Start small and incrementally grow the CCoE to meet your needs as you progress through your cloud transformation. Your cloud provider representatives, such as your AWS account manager and solutions architect, can provide resources to guide you through the creation of your CCoE.

A CCoE speeds up your ability to establish subject matter expertise, achieve agreement, earn trust across your organization, and establish effective guidelines for meeting your mission requirements. There is no single organizational structure that works for every FI, but the following guidance will help you design your own CCoE. For more information, see Building your Cloud Operating Model on the AWS Prescriptive Guidance website.

We recommend that you build a single CCoE and then specialize within it for each cloud provider. Teams have favorite products, suppliers, software, architectural patterns, programming languages, and CSPs. Aligning your organization around these preferences can accelerate cloud adoption. When you add more providers, maintain a single CCoE that includes specialized teams. Creating separate CcoEs for each provider leads to competing engineering priorities and fragmented approaches across your IT organization.

What does a multicloud-enabled CCoE look like?

When a CCoE is fully operational, it reflects your business and represents a microcosm of your current IT and business functions. You might have dedicated architecture teams for each CSP or create shared teams that support multiple CSPs. The shape of your CCoE must be appropriate for your business model. The following diagram provides an example of organizational structure that's derived from the best experiences of AWS FSI multicloud customers.

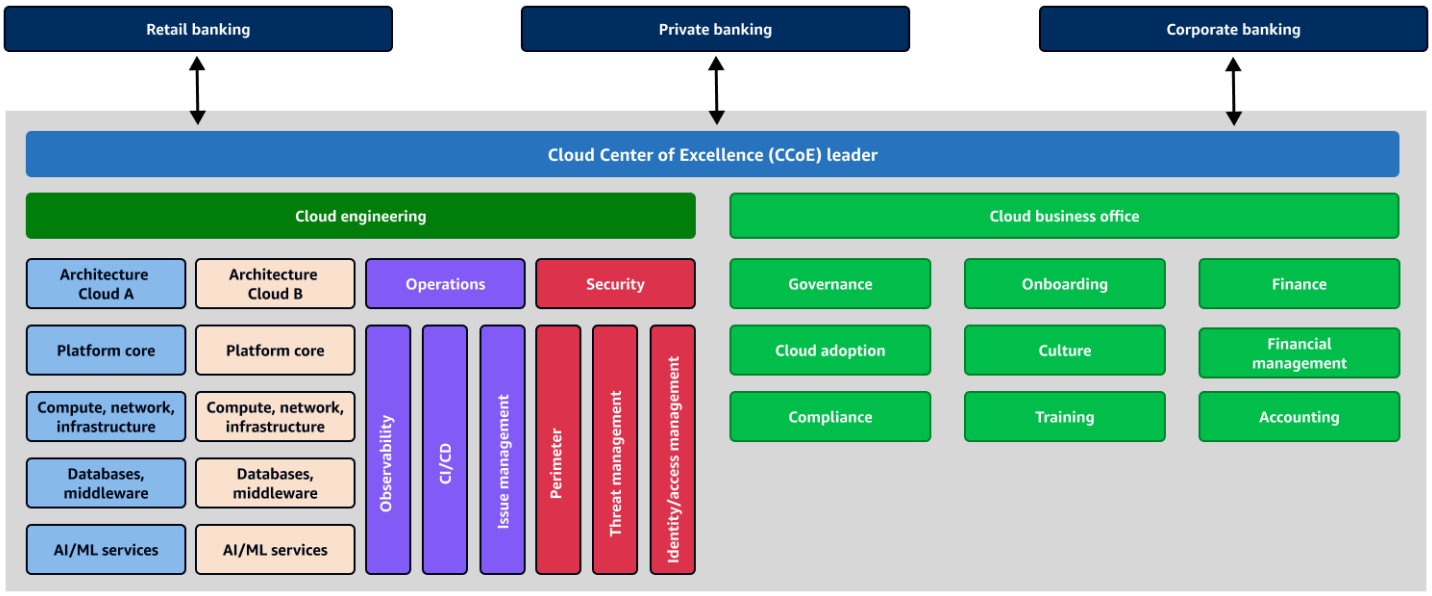

Best practices for multicloud CCoEs

- **Ensure critical capabilities in cross-cutting functions.** This includes security, governance, compliance, cloud financial management, and risk management. Your critical capabilities must not be unique to any one cloud provider, but driven from your business needs, which might exceed the capabilities of a single CSP.
- **Rotate staff between providers periodically.** Your engineering team will typically have uneven expertise across your CSPs. Rotating your staff between platforms to provide them with experience and exposure to new design patterns is an excellent way to speed up innovation. Ideally, rotate no more than 50% of your staff, and then rotate them back after some time, so they can share their learnings with the broader team.
- **Prioritize and reward certification programs.** Each CSP has its own certification program to improve and validate skills. Make sure that your staff has gained certifications across your CSPs and across multiple specialties. Create an incentive program for team members who earn multiple specialty certifications from multiple CSPs.
- **Embed security into your CCoE's inputs and outputs.** Your CCoE can create technical guardrails against future security threats because it authors your internal blueprints and provides guidance that shapes future projects. Create output goals for your CCoE, such as security escalations seen over time, and then measure changes to your workload security metrics from projects that use your CCoE's guidance. Compare security metrics between similar workloads across different CSPs.
- **Reward collaboration through cross-team metrics.** One approach to encourage solid engineering practices is to measure your operations teams on the number of code defects found

in production, and measure your developers on workload uptime. Apply the same approach to workloads in different CSPs: Ask your AWS teams to track performance metrics for workloads in another cloud, and vice versa. This cross-accountability encourages your teams to share successful strategies.

Evaluate exit strategy requirements

You might need a cloud provider exit strategy to comply with regulatory requirements, cloud concentration concerns, or your own corporate policies pertaining to the use of critical third-party service providers. An exit strategy outlines an organization's approach to move one or more workloads from a CSP to another environment.

Some key components of an exit strategy could include:

- Scope of the exit strategy (for example, the number of workloads)
- Objectives and success criteria for an exit
- Key risk indicators (KRIs) and thresholds that might necessitate an exit
- Target environment for workloads (for example, on premises or alternative service providers)
- Technical considerations (for example, whether migration to another environment requires refactoring)
- Operational considerations (for example, staffing and skills required to operate workloads in another environment)
- Legal considerations (for example, contractual commitments and termination rights)
- Key assumptions that underlie the exit strategy (for example, the resources available and the time it will take to complete the exit)
- Plans for testing the exit strategy and challenging key assumptions (for example, tabletop exercises and gamedays)
- Organizational responsibilities for development, maintenance, and governance of the exit strategy
- Data residency considerations

> **Note**
>
> To assist customers with a smooth transition, AWS offers free data egress for planned exits. For more information, see the AWS blog post Free data transfer out to internet when moving out of AWS.

Accelerate and optimize operations with SaaS applications

Most FIs have adopted software as as service (SaaS) applications for key parts of their business. Common SaaS applications include productivity applications such as word processing and email. SaaS options also exist for many mission-critical workloads, including payments processing, treasury, trading, anti-money laundering (AML), know your customer (KYC), learning management systems (LMS), and customer relationship management (CRM) workloads. When your FI adopts a SaaS offering, your IT team doesn't have to maintain the service or manage the infrastructure—your users consume the service without any IT involvement. Even if you have the resources to host the application yourself, it might be more cost-effective to adopt a SaaS solution and invest in other projects instead.

Knowing how and where your SaaS application is hosted, or choosing the right hosting environment, leads to lower data transfer costs, better integration, lower latency, and native backups. For example, Genesys is a call center application built on AWS. Integrating with this SaaS solution and using custom-built modules on AWS gives you the flexibility to define your users' experience on this application and to speed up implementation by taking advantage of the services and integrations offered by AWS.

As another example, services such as AWS Glue provide seamless, bidirectional data movement for Salesforce, and integrating AWS generative AI services complements the strengths of Salesforce. Consolidate and converge your SaaS landscape and understand how you will integrate with SaaS solutions to get the best experience for your users.

Many SaaS providers offer services that are hosted on another CSP's infrastructure. When you're deciding on workload placement, consider the latency and availability of these services and your access pattern. Some questions you should consider in this context include:

- Do I access these services asynchronously or through batch processing?
- Is near real-time, on-demand access a requirement?
- How latency-sensitive are my workload's requirements?

- Are my customers or internal users sensitive to whether my systems are colocated near our SaaS partners?

Align your SaaS providers and CSPs with your requirements. Your SaaS products should both speed up your business outcomes and reduce operational impact through cloud-native colocation where appropriate.

Establish security and governance for each CSP

Because of the shared responsibility model, FIs that adopt cloud services can reduce administrative burden by offloading some of the responsibility for infrastructure security to the CSP. FIs can benefit from purpose-built, cloud-native security services that offer features that are often unavailable, difficult to manage, or cost-prohibitive in an on-premises environment. A successful cloud security and governance strategy enables IT and security teams to focus on building systems that are secure by design, helps the FI rapidly adapt to evolving mission requirements, and provides staff with secure environments for innovation.

The best practice for security and governance is to make it easy for users to do the right thing. This is why automation, partner solutions, orchestration, and proper controls are critical to a multicloud operating model.

The following sections discuss best practices for establishing security and governance in a multicloud environment. For information about AWS services for the financial services industry, see Security, Compliance, and Governance for Financial Services on the AWS website.

Compliance networks

FIs must adhere to many regulatory requirements, and these are often specific to the markets and jurisdictions that the FI operates in. Regulations can include the Payment Card Industry Data Security Standard (PCI DSS), the European Union's Digital Operational Resilience Act (DORA), National Institute of Standards and Technology (NIST) compliance frameworks, and many others. Every regulation is distinct and might apply only to certain workloads. Consider the requirements for each workload and make sure that you can meet them in its cloud environment. Make sure that you understand your responsibilities compared with the cloud provider's responsibilities in each environment.

Supporting compliance without inhibiting innovation

We recommend that you embed security controls in self-service systems that enable users to rapidly deploy cloud environments with minimal intervention from IT teams. These self-service systems might take the form of service catalogs, where infrastructure can be vended with preset, vetted patterns, or common templates that are written by using infrastructure as code (IaC) products such as HashiCorp Terraform or the AWS Cloud Development Kit (AWS CDK).

You might choose to manage compliance for each provider separately or use custom-built or AWS Partner solutions that centralize management across providers.

Assessing and controlling cost and usage

Establish cost visibility and control mechanisms to gain insight into cloud services that are in use, who the cloud resources belong to, the purpose of those cloud resources, what potential cost savings can be achieved by optimizing consumption, and which geographies your workloads are deployed in. You can increase your firm's return on investment by partnering with your CSP to migrate and modernize mission-critical systems. You can negotiate enterprise-level agreements, benefit from volume pricing, and take advantage of the CSP's expertise. To manage costs across multiple providers, consider how you can aggregate and analyze cost and usage from each provider, either by using in-house processes and tooling, or by using AWS Partner solutions. Establish a core cloud financial operations (FinOps) function in your organization, and dedicate resources to evangelizing and implementing capabilities for cloud cost management and optimization.

Managing user permissions

You must identify the core needs of your users and confirm that there are appropriate mechanisms in place to grant and revoke access to cloud services. Different types of users require different types of access to cloud services. You should have tooling in place to centrally manage these identities in an automated way across all your cloud environments, and use established processes to identify, grant, and revoke permissions as roles and responsibilities change over time. The best practice is to use automation to review and report on privileges at least every 90 days. For information on identities and access, see Identity federation and single sign-on in the guide *Building a strategy for single, hybrid, and multicloud in education*.

Integrating new systems with your identity management solution

FIs should make it easy to integrate new systems with their identity management systems (for example, Microsoft Entra ID, Okta, or Ping Identity). This gives the organization the flexibility to support a variety of mission-critical functions by allowing stakeholders to buy and build systems that can easily be integrated into the identity management system. When you simplify the integration process, your stakeholders will be less likely to use their own access control measures, which might not enforce security best practices such as single sign-on, passkeys, and multi-factor authentication (MFA). Your identity management system should interoperate with the necessary systems through native integrations or industry-standard protocols, so it can support easy adoption across multiple CSPs.

Effective incident detection and response

To detect and respond to cyberattacks effectively, use a dual approach:

- Focus your efforts on preventive measures and security controls that are automatically embedded in cloud environments, including adoption of CSP-specific services if necessary.
- Implement detection capabilities that help cyber incident responders detect, contain, and mitigate security breaches in a timely fashion, including the ability to isolate CSPs from one another if a breach occurs.

Perform regular simulations for attack scenarios, such as game days, tabletop exercises, and disaster recovery drills, so your people, processes, and technology are tested and ready across all your cloud environments. This includes simulating an attack on a single provider and analyzing the potential unintended impact to other providers.

Mapping data classifications to CSPs

Every FI has a data classification standard. Typically, this includes a data label such as *Personally Identifiable Information (PII)*, *Classified*, *Private*, or *Confidential*. These labels might include metadata, depending on the data and workload. The scope of discovery, audit, data access patterns, and controls will be mapped to your business's unique requirements. In a multicloud environment, establish mechanisms to ensure that each environment is vetted to process data based on its classification before data transmission. Prevent data processing in unauthorized or unintended environments that inadvertently expands the scope of compliance. For example, you might have a payment processing application on AWS that is PCI-compliant, but it might be associated with workloads in other CSPs that are not PCI-compliant.

Establish a modern multicloud data and AI strategy

A robust multicloud data strategy is essential for ensuring seamless data integration, stringent governance for compliance and resilience, and adherence to data sovereignty laws. It should include a cost-management framework that supports flexible scalability in line with organizational growth. By addressing these critical areas—integration, governance, resilience, sovereignty, and scalability—you can effectively optimize data assets across platforms. This ensures security and compliance, and enhances decision-making and operational efficiency.

When you build your multicloud data strategy, you should recognize that data management is the foundation of successful AI and generative AI applications. High-quality, well-managed data is both a differentiator and the cornerstone of innovative AI technologies. AI systems require diverse, accurate, and comprehensive datasets to train effectively. These systems can then create insights and generate value that would be impossible with traditional analytics. Working with data across multiple cloud service environments introduces additional considerations that are discussed in the following sections.

Integration and accessibility

The guiding principle of a multicloud data strategy is managing data that's spread across various platforms without sacrificing integrity and accessibility. You must develop a strategy that prioritizes interoperability and consistency. This includes implementing integration tools that connect disparate cloud services and provide a unified data view across clouds. A unified view facilitates an AI model's access to data, and ensures that teams and applications can easily find and use the data they need. We recommend that you create a unified data catalog that identifies data owners, custodians, and governance requirements while enabling a self-service experience for your data customers across your organization. AWS Partners such as Databricks, Snowflake, and Colibra have proven track records in successful data integration.

Data lineage

When you use multiple CSPs, you must have a complete understanding of your data lineage, that is, where your data comes from. You must document the sources and owners of your data, the results they produced, and how these were applied, especially if you're using generative AI. A data lineage strategy shows you and third parties what data you have and how it is used. For multicloud customers, we recommend that you focus on a federated governance model that includes automated lineage collection and data quality measures that cross all your cloud environments.

Governance and compliance

A comprehensive multicloud data strategy must enforce stringent data protection policies, privacy controls, and legal compliance across all cloud environments. This includes the application of consistent data encryption, anonymization, and adherence to regulations such as General Data Protection Regulation (GDPR) and PCI DSS. This governance supports AI systems by ensuring that the data they use is both secure and compliant, and by maintaining trust and integrity in AI outputs. We recommend that you set up proactive alarms for compliance violations by using automated tooling. This *continual compliance* approach should use cloud-native and ISV solutions to remove manual efforts from detective controls. These controls must operate in every cloud environment that you host sensitive data in.

Centralizing and analyzing security data

In a multicloud environment, security management is complex because of the variety of systems and data sources that are involved. Centralizing security data across multiple clouds helps enable better security by providing real-time analysis and response. This centralized approach requires tools that aggregate security alerts and automate the analysis of security data across CSPs. It simplifies the management of security alerts and helps enhance the detection of potential threats by applying consistent security policies and procedures across all cloud platforms.

Centralizing application monitoring

Central application monitoring in a multicloud environment is essential for maintaining performance and availability. It requires gathering and analyzing performance data from applications that are deployed across various cloud platforms. These solutions typically feature unified dashboards that display a wide range of operational metrics, logs, traces, and alarms. Centralized visibility enables administrators to monitor application health comprehensively and respond swiftly to any issues, and ensures consistent performance across all cloud environments. We recommend that you create centralized monitoring that includes cloud-native sources instead of relying exclusively on telemetry from your applications.

Performance optimization

Minimizing latency is essential to maintaining the performance of cloud-based applications and AI models. Optimizing data storage locations and employing edge computing are strategies for reducing latency. Implementing data caching strategies and choosing the appropriate data storage technologies significantly boost overall system performance, and facilitate quicker

AI processing and real-time analytics. We recommend that you maintain training data close to your ML workloads. You can exfiltrate trained models to other environments and operate them with lower latency for customer-facing workloads. Organizations typically use Retrieval Augmented Generation (RAG) to gain better AI insights from custom data sources, which must also be colocated with low-latency networks to their trained models.

Providing choices in AI/ML strategy

FIs often give their teams the flexibility to choose from a diverse array of machine learning (ML) tools and platforms. This choice gives organizations the ability to adopt the best technology that aligns with each project's specific requirements, whether they involve complex datasets, require real-time processing, or necessitate highly scalable solutions. For multicloud enterprises, we recommend that you let your teams self-select AI models and managed services that best meet their business goals from each CSP, instead of offering a single AI service on one cloud that is a suboptimal fit for the workload.

Optimizing multicloud operations with DataOps and MLOps

In multicloud operations, the challenges of managing cloud services require a strategic implementation of automation, data operations (DataOps), and machine learning operations (MLOps). These technologies and methodologies are essential for streamlining workflows, reducing complexity, and optimizing the management and operationalization of data and ML models.

DataOps improves the integration and collaboration of data flows between data managers and consumers. Implementing DataOps involves setting up agile data management processes that are similar to the agile methodologies in software development. This includes infrastructure as code (IaC), continuous integration and continuous delivery (CI/CD) pipelines for data, automated testing for data validity, and rapid iteration of data models. To improve your DataOps processes, we recommend that you create cross-functional teams that include data scientists, data engineers, and operations staff.

MLOps focuses on the lifecycle management of ML models to ensure scalability, reproducibility, and maintainability. To implement MLOps, organizations adopt version control systems for models and datasets, automate model training and deployment processes, and continuously monitor the performance of deployed models to detect model and data drift. If performance degrades, they initiate model retraining. Tools, such as MLflow for model lifecycle management, TensorFlow Extended (TFX) for end-to-end ML pipelines, and Kubeflow for orchestrating scalable ML models on Kubernetes, help implement MLOps practices across environments.

We recommend that you prioritize both DataOps and MLOps for all flows between CSPs, and use automated pipelines that proactively deploy, monitor, and generate alerts when data quality is endangered. Do not rely on human intervention, and engineer solutions that scale through automation at every opportunity.

Data pipeline cost management

Managing costs effectively in a multicloud environment involves choosing the right mix of cloud services and pricing models. Using tools for monitoring and managing resource usage helps you allocate resources cost-effectively when you run cost-intensive AI computations and data storage across multiple clouds.

Adopt a strategy for zero-ETL, which is a set of integrations that eliminates the need to build extract, transform, and load (ETL) data pipelines, wherever possible. Traditionally, ETL is a major driver of data pipeline cost, because it often loads data that might not be required for all ML workloads. Zero-ETL features support a more cost-efficient multicloud strategy, and enable organizations to optimize their investments while using powerful AI and data analytics capabilities.

Assess the viability of contiguous workloads across CSPs

Distributing contiguous workloads across cloud providers introduces challenges and risks that require thoughtful consideration. To evaluate the viability of contiguous workloads, focus on these common considerations for financial services:

- **Latency**. The physical distance between data centers affects performance, particularly for real-time applications (such as real-time payments or trading solutions) that are prevalent in financial services. Weigh geographical diversity against potential performance issues that might affect user experience and operational efficiency.
- **Predictability of network routing**. Public internet connectivity between clouds lacks the predictability of dedicated network infrastructures. This can lead to inconsistent network performance and bottlenecks. Balance the convenience and cost savings of contiguous workloads against the need for dependable network performance.
- **Security**. Data transmission over the public internet requires strict security controls. Traffic between CSPs involves greater security risks than traffic within a single provider's network. When you adopt a multicloud strategy, you'll need to strengthen your security measures and possibly increase your cybersecurity investments to protect data in transit.

- **Complexity of billing**. Handling billing across multiple CSPs can add complexity because CSPs often have different cost models. This can complicate budget management and financial forecasting, and often necessitates advanced financial management tools.
- **Responsibility boundaries**. Maintaining clear boundaries for managing and securing workloads is challenging in multicloud environments. Each CSP has different policies, and service-level agreements (SLAs) can complicate IT governance. Businesses must define responsibility boundaries clearly to ensure comprehensive coverage of security, compliance, and performance.

You must consider these factors to make informed decisions that align with your strategic goals and operational needs. If you're planning a workload that spans multiple CSPs, follow the recommendations provided in the following sections.

Domain-based approaches in multicloud implementations

In financial services, most firms have workloads with significant data gravity—that is, the data is difficult to move between cloud environments. Financial systems have strong data gravity constraints because they require real-time transactions and strict operation sequences. Applications that don't have these constraints are easier to run across multiple clouds.

For example, compare your workloads with a workload that runs on a content delivery network (CDN). A CDN is a globally distributed system that has hundreds of points of presence (PoPs) and acts (broadly) as a pass-through cache for web assets. These systems do not need strong transactional consistency and are treated differently from financial systems that process credit card payments. The underlying technology and business mechanisms for each system are different.

Financial services workloads have the highest resilience and transactional consistency requirements. These requirements can sometimes conflict with operational resilience and concurrent operations goals if applied across multiple CSPs. Therefore, in some circumstances it might be advisable to keep each application or workload within a single CSP and bound its context within that CSP. If your workload requires concurrent operations across clouds, you should logically split or shard components according to failure domains. This approach removes strong dependencies between CSPs and helps you meet your service-level objectives (SLOs) and service-level agreements (SLAs).

When teams build applications that are distributed across CSPs, this domain-based approach is functionally similar to using a SaaS provider, although that provider is a team within your organization. In this scenario, the considerations discussed in the section Accelerate and optimize FI

operations with SaaS applications also apply to workloads that your organization operates across providers.

Hard dependencies and synchronous operations

Workloads that are spread across cloud environments can encounter similar design constraints as workloads that are spread between an on-premises data center and the cloud. When a single, logical operation requires synchronous or real-time API calls, splitting your workload might become logistically challenging and endanger SLAs, SLOs, and overall resilience. For example, if a customer initiates a credit card transaction that requires real-time calls to multiple physical locations, and each location is bound by its own SLAs and internal dependencies, a service interruption in any environment will endanger the entire transaction.

When you operate your system in multiple CSPs, you must turn hard dependencies across environments into soft (flexible), asynchronous dependencies, as illustrated in the following diagram. When API calls are part of your business process, use patterns such as circuit breakers and exponential backoff.

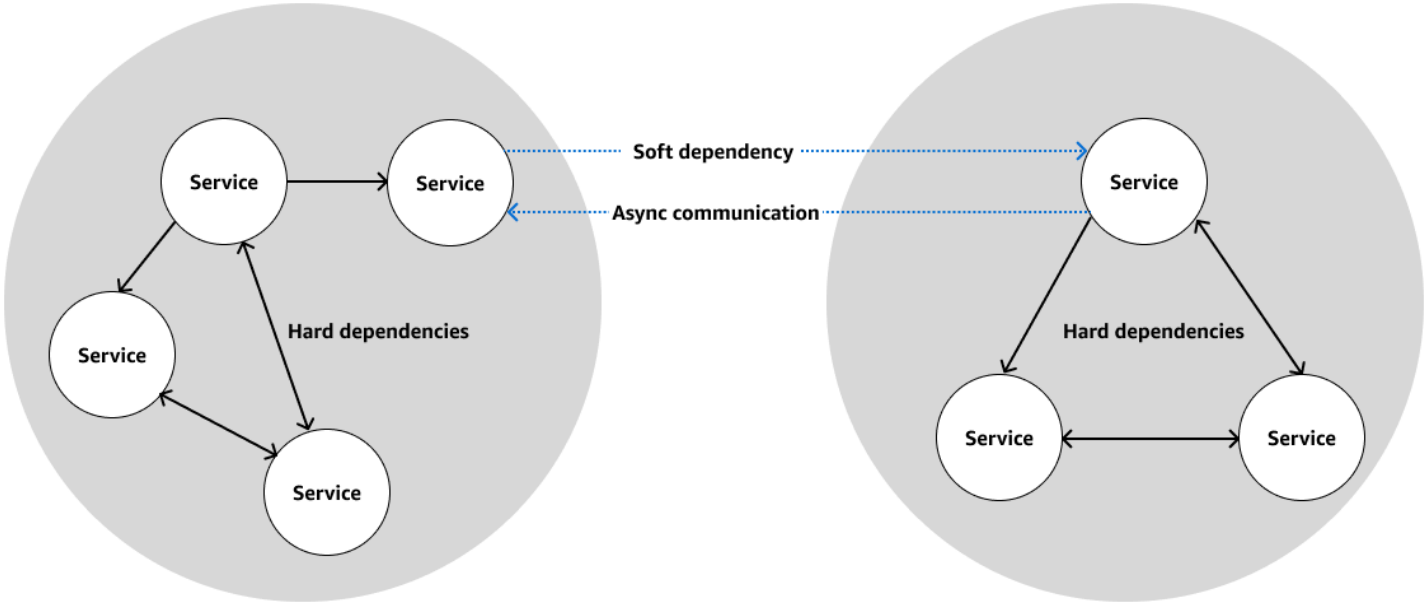

To transfer large amounts of data between cloud environments, use a bulk export approach. This approach enables faster data transfer than a series of synchronous API calls, simplifies data validation, and works with object storage across all cloud providers. One PUT operation is cleaner, faster, safer, easier to audit, and less expensive than millions (or billions) of PUT operations.

Impact on SLAs

Financial services firms must carefully assess the feasibility, practicality, and implications of multicloud operating models so they can meet their SLAs. Each CSP (and its services) has its own SLAs. If your workloads have stringent SLA requirements, consider the composite SLA between environments to make your workload design decisions.

For example, if your workload in CSP1 has an SLA of 99.99% and your workload in CSP2 has an SLA of 99.99%, the composite SLA is 99.98%. In addition to SLA considerations, you must communicate clearly with your stakeholders on how multicloud operations affect their cross-cloud workloads.

Consider the advantages and disadvantages of vendor lock-in

FIs often ask us how to design solutions that prevent vendor lock-in, and how they can use multicloud strategies to avoid lock-in.

Preventing vendor lock-in depends more on your organization's people and processes than on technology decisions alone. We recommend that you focus on creating IT processes that enable swift changes to your workload components, create domain-driven designs that include a microservices architecture where applicable, and use modern coding, testing, and deployment techniques.

These recommendations apply equally to your cloud infrastructure and your application development. Define your resources as code by using tools such as HashiCorp Terraform, the AWS Cloud Development Kit (AWS CDK), and Pulumi. By treating your infrastructure like software development, you can shift from long-term procurement to rapid testing and replacement.

When you choose a technology for long-term use (for example, a database for storing large amounts of data), consider the pros and cons of proprietary protocols compared with web standards such as REST. For some FIs, the value added by using a cloud-native product can be worth the perceived loss of short-term portability. Often the perceived lock-in from a technology decision is actually a valuable investment that creates rapid business value. For example, adopting Kubernetes could be perceived as a potential lock-in to an ecosystem but actually creates rapid business value by standardizing deployments.

Establish a workload placement strategy

When you define your workload placement strategy for a multicloud environment, first choose a primary CSP that meets the general needs of your organization. This provider should offer a robust set of services, security features that meet or exceed your needs, and cost efficiency. Your organization can then follow a clear process to select additional cloud providers based on specific criteria:

- **Organizational structure**. Consider how your organization is structured today (for example, by departments, business lines, or geographic locations), particularly if you need to comply with data sovereignty or other regulatory requirements. Also consider enterprise architecture, communications, operating procedures, and other influences from Conway's Law as inputs. Don't force workload placement or operations in ways that create friction with your existing culture or business operations.
- **Workload performance**. Consider how your applications are structured. For SaaS applications, determine their preferred CSP, if there is one. Different CSPs offer unique performance benefits for specific workloads. For example, some providers might provide an infrastructure that better supports applications or SaaS partners because of native optimizations or specialized services.
- **Security and compliance**. CSPs offer unique security features and compliance certifications. In a multicloud environment, all providers must meet your organization's security and regulatory requirements. Evaluate the security measures and compliance capabilities of each CSP to maintain a strong security posture across your multicloud infrastructure. Consider data residency and the physical fault isolation boundaries that each CSP supports.
- **Data management and integration**. Data management can become more complex in a multicloud environment. Evaluate how data will be stored, accessed, and integrated across different CSPs. You must have a cohesive data strategy that ensures data consistency, availability, and security across all cloud platforms. Federated data governance requires a strong *trust but verify* approach to data management, and adding metadata in other CSP environments can introduce unplanned scope expansion, which increases risk and potential exposure to unauthorized parties. Be prepared to maintain materialized views of your data in other clouds.
- **Cost management and optimization**. Multicloud environments can complicate cost tracking and optimization efforts, and many CSPs offer different cloud operating models and tools for assistance. Implement robust cloud financial management practices in your landing zone early in your adoption to monitor usage, control costs, and maximize the value of your cloud investments.

We recommend that you establish a multicloud governance framework that includes criteria for selecting additional CSPs. This framework should evaluate factors such as the cost, performance, security, and compliance capabilities of additional providers.

No single approach works best for all organizations, but successful multicloud implementations often use a scoreboard, such as in the following example, to evaluate and compare CSPs.

Capabilities	CSP1 score	CSP2 score
Landing zone alignment to organizational structure	5.00	5.00
Workload performance	4.00	5.00
Security	4.00	3.00
Data management	3.00	3.00
Cost optimization	4.00	3.00
Exit capabilities	4.00	3.00
Partner ecosystem	5.00	5.00
Operational excellence	5.00	4.00
Reliability	5.00	4.00
Sustainability	4.00	4.00
Total	**4.30**	**3.90**

However, this table isn't a fully developed example and doesn't reflect the weight or bias of each capability. For example, your business might require greater emphasis on exit capabilities compared with other capabilities. The key is to have a structured, systematic, repeatable, and deterministic method for selecting CSPs.

Each workload might require a minimum score per capability to meet external constraints. For example, payment processing applications might require a minimum score for reliability and de-emphasize cost as a consideration.

After you select and enable your CSPs—including setting up a secure landing zone, training your IT staff, and establishing a CCoE—follow a prescriptive selection process to determine which CSP will host each workload in production. In rare instances, the same application might be placed concurrently in multiple CSPs; however, this is not typical except for ISV and SaaS workloads. The placement strategy can be documented in a flowchart that addresses the considerations discussed in this paper. The following diagram provides an example.

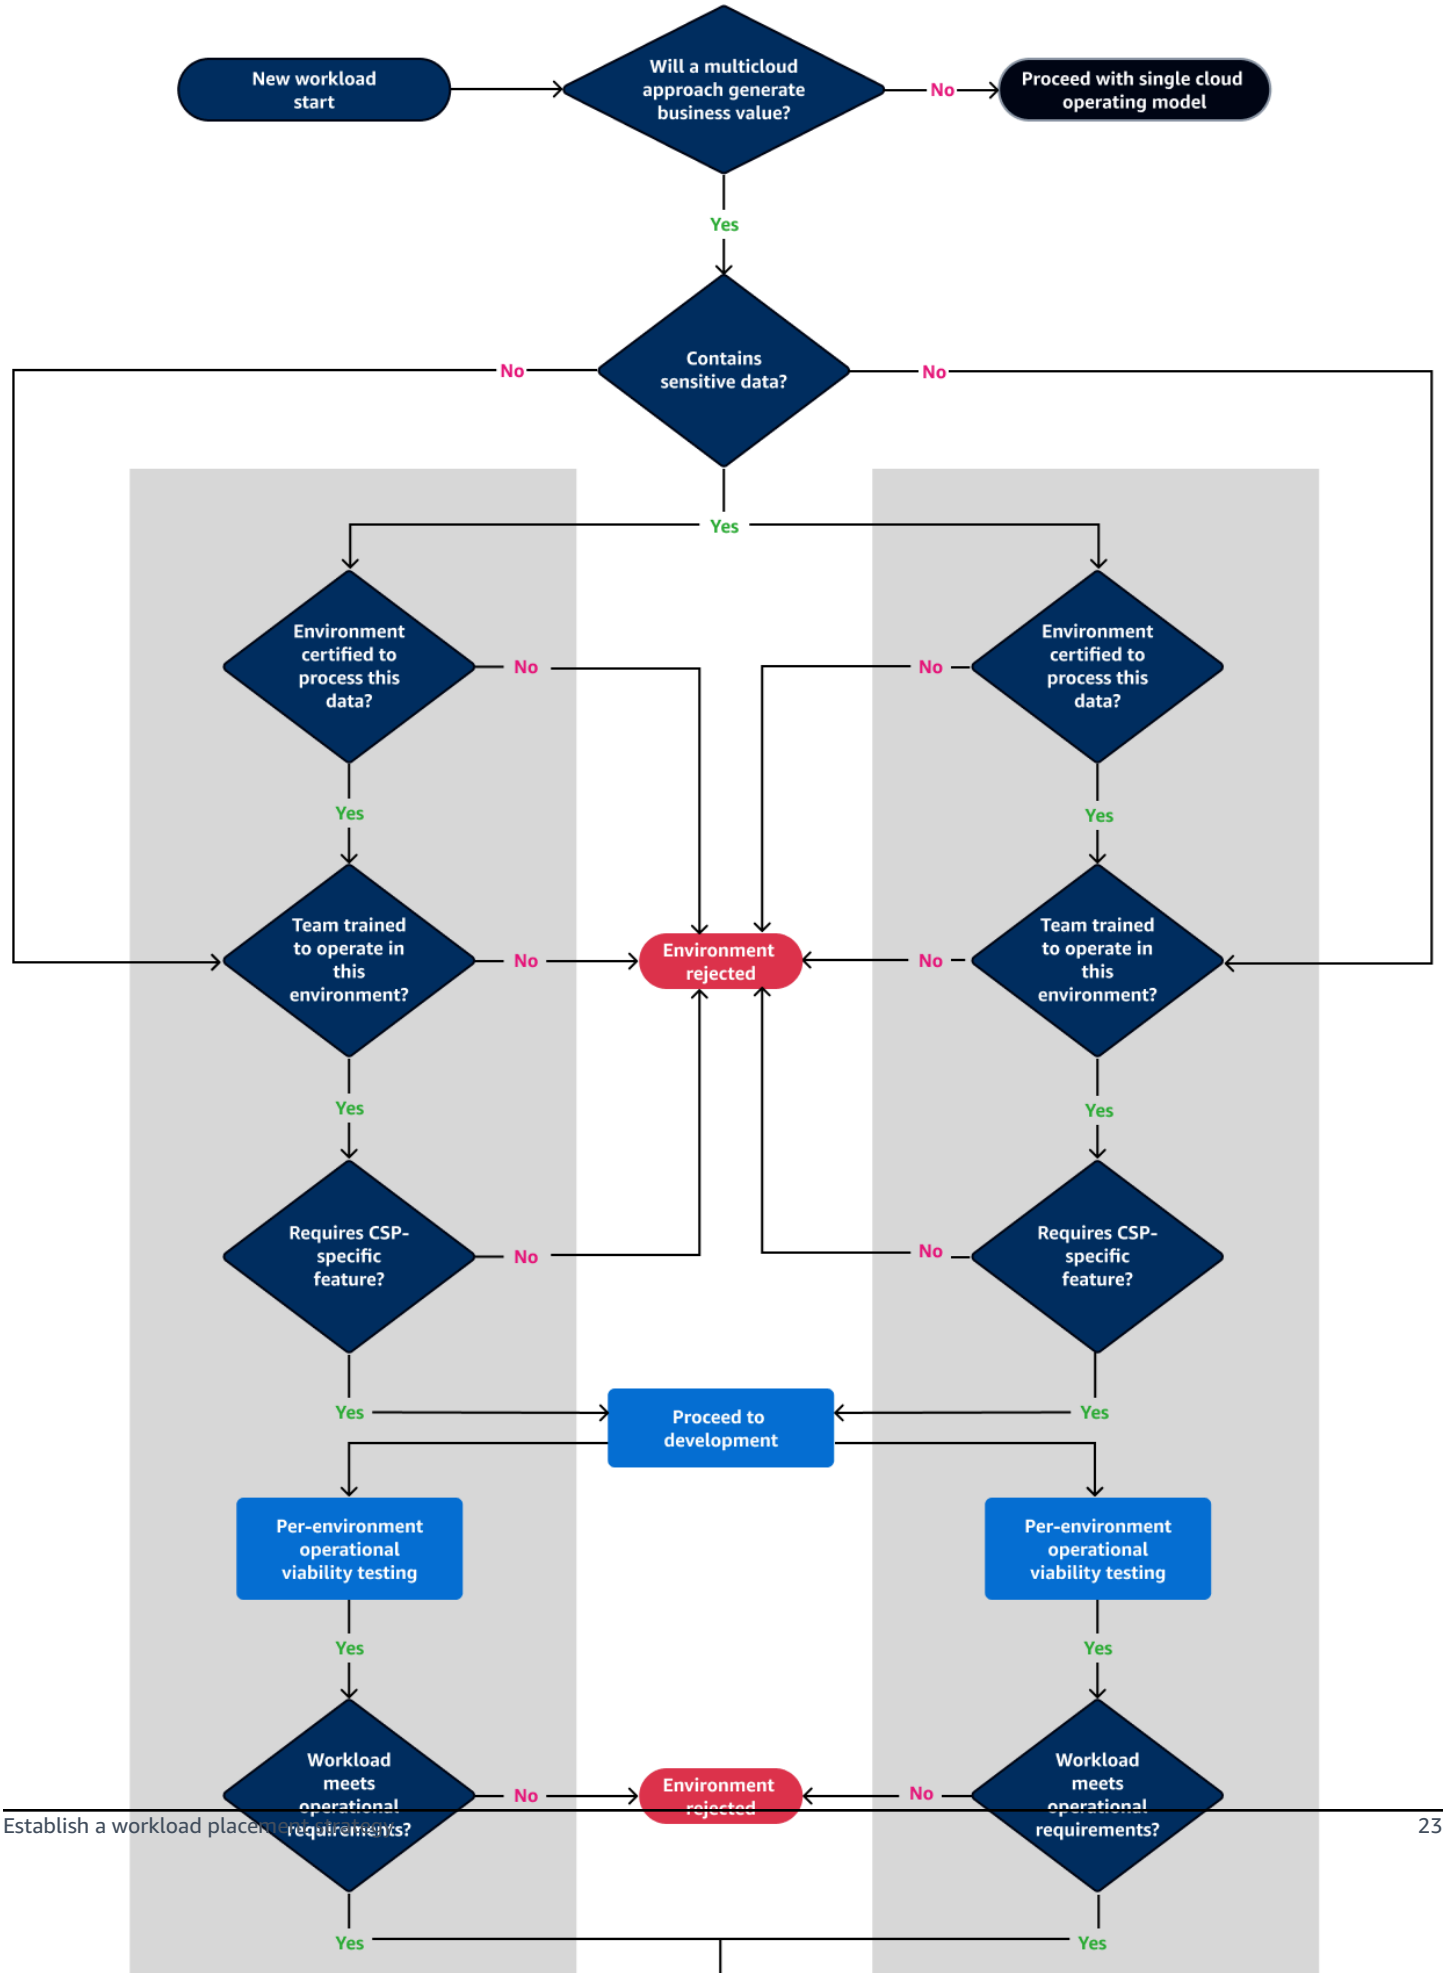
New workload start
Will a multicloud approach generate business value?
No
Proceed with single cloud operating model
Yes
Contains sensitive data?
No
No
Yes
Environment certified to process this data?
No
Environment certified to process this data?
No
Yes
Yes
Team trained to operate in this environment?
No
Environment rejected
No
Team trained to operate in this environment?
Yes
Yes
Requires CSP-specific feature?
No
No
Requires CSP-specific feature?
Yes
Proceed to development
Yes
Per-environment operational viability testing
Per-environment operational viability testing
Yes
Yes
Workload meets operational requirements?
No
Environment rejected
No
Workload meets operational requirements?
Yes
Yes

We recommend that you follow a phased approach with checkpoints to develop your workload placement strategy. Remember that multicloud is a strategy, not a goal. It requires balancing business value, technological feasibility, product velocity, and cost.

Develop a clear business case for using multiple providers. Do not proceed without documented justification. Next, map each critical capability for your cloud environments against your business requirements. A "no" response to any critical capability question disqualifies the workload. A multicloud approach is feasible only when a workload meets all your business and technical requirements.

Next steps and resources

Selecting the right deployment model for multicloud workloads requires careful consideration. Use the recommendations outlined in this paper to guide your decision-making and avoid common pitfalls such as unnecessary complexity and inconsistent governance. By following these best practices, you will accelerate your cloud adoption to meet and exceed your firm's goals more effectively.

Establish a Cloud Center of Excellence (CCoE) to drive organizational maturity and ensure your long-term success. Differentiate between SaaS applications and foundational cloud services, and identify core security and governance requirements for each. Whenever possible, adopt cloud-native, managed services, and implement hybrid architectures when your existing data center investments incentivize continued use. Lastly, use multicloud approaches when they add business value that outweighs complexity and cost.

AWS can help you manage single, hybrid, and multicloud environments. Your FI can use AWS management and observability solutions such as AWS Systems Manager, AWS Config, and Amazon CloudWatch to simplify and centralize the management and monitoring of your infrastructure and applications, regardless of your environment. With data and analytics services such as Amazon Athena, AWS Glue, and AWS DataSync, you gain insights from all your data, wherever it is stored. Hybrid solutions such as AWS Outposts, AWS Wavelength, and AWS Snow Family let you bring AWS infrastructure and services to wherever they are needed. Tools such as Amazon EKS Distro help you build self-managed Kubernetes clusters on AWS, on premises, or on other clouds.

As you define your cloud strategy, consider these next steps:

1. Review the AWS Cloud Adoption Framework (AWS CAF) to identify and prioritize transformation opportunities, evaluate and improve your cloud readiness, and iteratively evolve your transformation roadmap.
2. Identify a cloud implementation project as a proof of concept. This will help you define the cloud foundation or framework to validate any assumptions and to enable future cloud implementations.
3. Engage your AWS account team to discuss your cloud implementation goals. The AWS account team can provide clarifications, suggest approaches, identify dependencies, and work with your teams to plan your project from initial concept to implementation. The account team can also connect you to multicloud experts at AWS to help address your specific multicloud needs.

For additional information, see:

- AWS Architecture Center
- AWS Cloud Adoption Framework (AWS CAF)
- AWS Solutions for Multicloud

Contributors

Authors and contributors to this paper include:

- Chris Ibbitson, Principal Solutions Architect, AWS Financial Services
- Ahmed Raafat, Senior Manager, AI/ML Specialist Solutions Architects, AWS Worldwide Specialist Organization
- Aviad Tamir, Solutions Architect, AWS Global Sales
- Dan MacKay, Principal FSI Compliance Specialist, AWS Financial Services
- Matias Undurraga, Enterprise Technologist, AWS Cloud and AI Innovation
- Rich McDonough, Principal Technologist, Multicloud Program

Document history

The following table describes significant changes to this guide.

Change	Description	Date
Initial publication	—	November 6, 2025