



Building a strategy for single, hybrid, and multicloud in education

AWS Prescriptive Guidance



AWS Prescriptive Guidance: Building a strategy for single, hybrid, and multicloud in education

Copyright © 2026 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
Overview	1
Cloud deployment strategies	3
Single cloud	3
Hybrid cloud	3
Multicloud	3
Recommendations	4
Select a primary, strategic cloud provider	4
Establish a CCoE	6
Differentiate between SaaS applications and foundational cloud services	8
Establish security and governance requirements for each cloud service provider	11
Adopt cloud-native, managed services wherever possible and practical	14
Implement hybrid architectures when existing, on-premises investments incentivize continued use	18
Reserve multicloud only for workloads that can't meet their technical or business requirements through a single cloud provider	21
Example use cases	24
Virtual computer labs	24
Predicting student success	26
Identity federation and single sign-on	28
Cloud bursting for research computing	29
Next steps	33
Contributors	34
Further reading	35
Document history	36
Glossary	37
#	37
A	38
B	41
C	43
D	46
E	50
F	52
G	54

H	55
I	56
L	59
M	60
O	64
P	67
Q	69
R	70
S	73
T	77
U	78
V	79
W	79
Z	80

Building a strategy for single, hybrid, and multicloud in education

Amazon Web Services ([contributors](#))

September 2023 ([document history](#))

Educational institutions are seeking to support functions such as remote learning, research, student experience, data insights, and administration with the agility, cost savings, security, and resiliency that cloud computing offers. Many organizations are assessing hybrid and multicloud deployments as part of this digital transformation.

This paper provides prescriptive guidance on creating a single, hybrid, and multicloud technology and governance strategy, for executive leaders and decision-makers at educational institutions who are evaluating their cloud options. This guidance is based on our experience at AWS working with over 14,000 educational institutions of all sizes across the world—from primary and secondary schools through higher education.

Overview

As educational institutions digitally transform to deliver differentiated services and experiences to their students, parents, faculty, staff, and community, they face a multitude of technical decisions. Many organizations have already made the decision to adopt the cloud for increased agility, elasticity, resilience, security, and cost savings. Based on their existing relationships and investments across various teams, most organizations are using some combination of on-premises data centers, colocation facilities, and cloud providers. Given the availability of multiple cloud options, educational institutions must frequently decide from single, hybrid, and multicloud deployment models (defined in the section [Cloud deployment strategies](#)).

Multicloud, which is the use of services from at least two cloud service providers, is not uncommon for many institutions today. Your IT team might prefer one cloud provider, whereas other groups, departments, or individual users might choose or already be using alternate providers. Educational institutions that don't have a clear strategy to guide them to the appropriate cloud deployment model encounter many challenges. These include unnecessary complexity, increasing staff demands, inconsistent governance, and lowest common denominator approaches that limit them to the subset of basic capabilities that are common across providers. Each challenge stifles innovation and slows digital transformation.

Conversely, if you have a cloud strategy that guides you to use single, hybrid, and multicloud, you can meet your education mission requirements while realizing the benefits of the cloud in a way that is operationally sustainable for long-term success. For creating this strategy, we recommend the following:

- Select a primary, strategic cloud provider.
- Establish a Cloud Center of Excellence (CCoE).
- Differentiate between software as a service (SaaS) applications and foundational cloud services.
- Establish security and governance requirements for each cloud service provider.
- Adopt cloud-native, managed solutions wherever possible and practical.
- Implement hybrid architectures when existing, on-premises investments incentivize continued use.
- Reserve multicloud only for workloads that can't meet technical or business requirements through a single cloud provider.

These best practices are discussed in detail in the [Recommendations](#) section of this paper. Each recommendation is important, but your institution's priorities will depend on its stage of cloud adoption. For example, if you are just getting started with cloud adoption, focus on selecting a primary, strategic cloud provider, establishing a CCoE, and adopting cloud-native, managed solutions. If you are already using a single cloud provider, focus on establishing core security and governance requirements, and consider hybrid architectures when your existing data center investments incentivize continued use. If your organization is already using multiple cloud providers, focus on differentiating SaaS applications and reserving multicloud deployments to those rare workloads that truly require it.

Contents

- [Cloud deployment strategies](#)
- [Recommendations](#)
- [Example use cases](#)
- [Next steps](#)
- [Contributors](#)
- [Further reading](#)
- [Document history](#)

Cloud deployment strategies

AWS defines cloud computing as the on-demand delivery of IT resources over the internet with pay-as-you-go pricing. Instead of buying, 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 educational institutions 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 models: single cloud, hybrid cloud, and multicloud.

Single cloud

This model uses only a single cloud service provider. Single-cloud applications and workloads might be implemented directly in the cloud, or previously hosted in another environment and migrated to the cloud. These workloads might use lower-level infrastructure services from their cloud provider or also take advantage of higher-level, managed services. Regardless, this model adopts a single cloud provider and uses only cloud services from that provider.

Hybrid cloud

A hybrid cloud model distributes resources across an organization's own on-premises data center and at least one cloud service provider. Typically, the purpose of this model is to extend an organization's infrastructure into the cloud while maintaining private connectivity with existing internal systems that reside on premises.

Multicloud

A multicloud model distributes resources across, and uses services from, at least two cloud service providers. An organization might choose to be multicloud, but more often this is an unintentional result of individual teams, departments, or staff members having their own preferences for different cloud providers.

Recommendations

Now that you have a foundational understanding of single cloud, hybrid cloud, and multicloud, this section provides detailed recommendations for choosing a model.

- [Select a primary, strategic cloud provider](#)
- [Establish a CCoE](#)
- [Differentiate between SaaS applications and foundational cloud services](#)
- [Establish security and governance requirements for each cloud service provider](#)
- [Adopt cloud-native, managed services wherever possible and practical](#)
- [Implement hybrid architectures when existing, on-premises investments incentivize continued use](#)
- [Reserve multicloud only for workloads that can't meet their technical or business requirements through a single cloud provider](#)

Select a primary, strategic cloud provider

Cloud adoption provides a wealth of benefits that are essential to IT modernization, cost-effectiveness, and innovation. However, adopting cloud technologies beyond limited SaaS applications can introduce challenges that educational institutions must carefully plan through to avoid unnecessary cost and complexity. The technological and business changes involved in implementing workloads in the cloud require staff enablement and adjustments to core infrastructure, including networking, security, governance, and operations.

The best approach for addressing these challenges effectively, especially if your organization is in the early stages of its cloud journey, is to select a primary, strategic cloud provider to support the majority of your workloads. Begin with a focused adoption that's centered on that provider so that you can simplify and accelerate the realization of cloud benefits. Selecting a primary cloud provider is not an exclusive, irreversible decision. It enables your organization to evolve your cloud adoption iteratively. You can begin by focusing on a few services and then expand into other cloud services as and where needed, without delaying the overall benefits of the cloud. This approach maximizes your organization's ability to take advantage of a provider's capabilities, concentrate and develop employee skills and third-party partner relationships, and simplify vendor management.

We have seen customers embark on their cloud journey by trying to concurrently adopt multiple cloud providers but later regret that decision and the complexity it introduced. Gartner shares this insight in their article, [6 Steps for Planning a Cloud Strategy](#), in which step 2 is "Prioritize a primary provider in multicloud architectures."

Each cloud provider introduces different operating and support models, identity and access management, networking, operations, compliance capabilities, and more. **It is better to master one cloud provider's operating model at a time.** You can then incorporate additional cloud services iteratively and incrementally, where rationalized. Many factors can influence your decision to adopt a primary cloud provider, but use the following key questions to guide your choice.

- **What breadth and depth of services does the provider offer?**

Different cloud providers offer different services. At a minimum, make sure that your primary provider has the capabilities necessary to support all your functional requirements as well as your cross-cutting, operational needs such as security, governance, and automation. Select a provider that delivers these capabilities with a proven track record of innovation and operational excellence. Consider not only your applications, but also your data. Think about future data integration and transfer patterns to limit the cost, latency, and complexity of moving large amounts of data between providers. Choose a provider that has the greatest possible breadth and depth of services to satisfy your current application and data needs, and also to unlock new use cases that can meet your institution's needs as they change over time.

- **Can the provider support all your security and compliance needs?**

In education, security and compliance are critical to any technology deployment. Choose a cloud provider that is able to meet all your security and compliance needs. Tools such as [AWS Artifact](#) can help you evaluate providers by offering a central resource for on-demand access to security and compliance reports. Consider not only the security and compliance of the cloud provider's own infrastructure and services, but also how easy it is for you to architect secure, compliant solutions by using those services. Prefer a provider that offers some combination of prebuilt solutions, quick starts, and prescriptive guidance to accelerate your secure adoption of the cloud.

- **Does the provider have a robust partner network?**

No organization undergoes cloud transformation alone. To accelerate adoption, you should use the services and expertise of the cloud provider as well as their partner network. This network includes technology partners who provide software that runs on, integrates with, or supports cloud technology, as well as consulting partners who can help you design, build, run,

and manage your own applications in the cloud. You will find that many educational technology providers, independent software vendors (ISVs), consultants, and resellers that you already work with are members of the cloud provider's partner network. Prefer a cloud provider that has the most robust network of partners with vetted competencies. Having partners with proven industry and technical expertise is critical.

- **What support and enablement does the provider offer?**

To successfully adopt any new technology, you need mechanisms to request training and help, including best practice recommendations, configuration guidance, and break-fix problem resolution. Choosing a cloud provider that offers strong support and training options will set you up for success. Explore the provider's official support model and resources as well as any available third-party or community-based resources such as blogs, forums, videos, and how-to guides. Consider not only the provider's technical support programs, but also programs that focus on business and cultural transformation. For example, the [AWS Cloud Adoption Framework \(AWS CAF\)](#) helps organizations digitally transform by focusing on perspectives that include business processes and people, not just technology. Prefer a cloud provider that offers extensive training options and a proven, reliable support model and community.

Establish a CCoE

Consider evolving your cloud leadership function 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 evolve the CCoE to meet your needs as you progress through the transformation journey. Your primary 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 accelerates your ability to establish subject matter expertise, achieve buy-in, earn trust across your organization, and establish effective guidelines for meeting your mission requirements. There is no single organizational structure that works for every institution, but the following questions will help you design your own CCoE.

- **Who should you include in your CCoE?**

At its inception, a CCoE might include only a handful of early adopters and cloud champions. The CCoE might remain small, but it should evolve to include champions who can speak for both the business functions and the technical functions that are affected by cloud adoption.

Business functions include change management, stakeholder requirements, governance, training, procurement, and communications. These functions are usually represented by members of your institution's administrative and instructional teams. Technical functions include infrastructure, automation, operational tools, security, performance, and availability. These functions are usually represented by members of your institution's IT teams. The CCoE should also seek to involve vendors and partners, as necessary, to provide subject matter expertise. The CCoE is a living organization. Its membership, form, and function will likely change over time, and it might even disband at some point of future maturity.

- **How does the CCoE interact with its stakeholders?**

The CCoE is in service to other teams and is intended only to inform and enable successful cloud adoption. Look at embedding parts of the CCoE in various departments, schools, and functions. This enables access to a wider range of resources and faster internal feedback. Focus on building partnerships and open lines of communication between stakeholders early on to establish trust within the institution and break down organizational silos. The CCoE should have defined mechanisms for communicating with stakeholders, gathering feedback, and training users. The CCoE's success metrics should reflect such collaboration and communication. If a team is measured only on building technology, more technology will be built, but its use and outcomes will become an afterthought. Your metrics should instead measure things such as the number of teams that become self-sufficient through the CCoE's work, the number of times the CCoE is on the critical path for initiatives, the number of training events held, or the breadth of adoption of the CCoE's output. A well-constructed, trusted CCoE can be a stepping stone to a larger organizational transformation that is built on trust.

- **How should you establish a CCoE?**

Most organizations start their cloud adoption with specific, targeted pilot projects. Establish a CCoE as part of these projects. A good start is critical in defining the success of the whole journey.

- **Start with a business problem.** Technology for the sake of technology is a bad strategy. If you are experimenting with cloud technologies, identify a compelling business use case no matter how small it might seem. Then, work back from that use case to set clear goals on how technology can help. Do not implement the solution in a silo. Take constant inputs from business stakeholders before and during project implementation. All successful cloud projects rely on close collaboration with the institutional units that will use the technology.
- **Start small.** Choose a low-risk project that provides a two-way door. This means that the project is reversible and any mistakes can be quickly corrected. Pilot projects are all about

experimentation. Avoiding large-scale, high-risk projects gives you better control over implementation and results. It helps to target specific, definable problems instead of broad-based goals. For example, if automation is the ultimate goal, aim to automate specific tasks instead of entire jobs.

- **Define and measure the outcome.** Set clear metrics to assess the progress and performance of each project. Define the desired end state well in advance to avoid mismatched expectations among stakeholders. Work closely with business stakeholders and other leaders within the organization to define expectations and measurable gains. It is also important to translate the results into non-technical language. Talk in terms of institutional goals, such as how the project improved retention and reduced churn, how it lowered costs and increased the speed of delivery, and so on.
- **Start from the comfort zone.** Choose a project within a domain that your institution is familiar with. This way you can ensure that the project has meaningful, understandable goals with real impact. Such a project will build confidence and have greater long-term results for your organization. For example, if you already have expertise in data analytics, you can kickstart your cloud journey while leveraging your existing skill set by starting with an analytics project. Every institution has different expertise and needs to find its unique components to craft a successful digital transformation strategy.

Differentiate between SaaS applications and foundational cloud services

Most educational institutions have already adopted software as a service (SaaS) applications. SaaS provides your institution with a complete solution that is run and managed by the service provider. Common SaaS applications include productivity applications such as word processing and email, but SaaS options also exist for many mission-critical workloads such as enterprise resource planning (ERP), student information systems (SIS), and learning management systems (LMS). When your institution adopts a SaaS offering, your IT team doesn't have to think about how the service is maintained or how the infrastructure is managed—your users simply consume the service. This delivery model reduces the management burden on your IT staff. Many institutions choose to adopt a "SaaS first" approach in their IT strategy, especially if their IT teams lack the time, resources, or skill sets to sufficiently self-host the same application. Even if you have the resources to self-host, it might still be more cost-effective to adopt a SaaS solution and invest in other projects instead.

When you use SaaS applications, your IT team doesn't have to manage the underlying infrastructure, so where the vendor hosts the application (on-premises data center, your primary cloud provider, or an alternate cloud provider) becomes less important. After you choose a primary, strategic cloud provider, you might opt to use a SaaS offering that's hosted in another cloud provider or on premises, in the vendor's data center. Conversely, even if your SaaS applications are hosted in one cloud provider, you might choose a different primary, strategic cloud provider based on the strength of that provider for your non-SaaS workloads. The distinction between hosting environments is less important for SaaS than it is for self-hosted applications. However, you should still consider the following key questions when evaluating how SaaS fits in with the cloud as part of your IT strategy.

- **Is the SaaS application highly available and scalable?**

Many vendors have already made the decision to adopt the cloud for their SaaS offerings. In doing so, the vendor is able to achieve the cloud benefits of increased availability and scalability. Furthermore, because the vendor can adopt the shared responsibility model of the cloud instead of managing and maintaining physical infrastructure, they can invest more time and resources into the delivery of new features. Because of these benefits, you should prefer providers that are cloud-first and offer cloud-hosted solutions.

- **Can the SaaS application meet your security requirements?**

When evaluating SaaS, it is important to know what data the application stores, how that data is used, and which security controls are in place to protect that data. Although you might not have direct control over data storage as you would in your own, self-hosted environment, you should ensure that the vendor has mechanisms and controls in place to handle your data appropriately. Be aware of which security features are built into the SaaS solution and which features require additional configuration. The cloud enables SaaS providers to build more available and scalable solutions, and they can also build more secure solutions because of the [shared responsibility model](#). You should prefer providers that are taking advantage of cloud security tooling and services as part of their solutions.

- **Who owns the SaaS application data and how can you access it?**

When you use SaaS, you trust the provider to properly handle your institution's data. Be sure to review the terms of service and service-level agreements for SaaS applications to understand contributing factors such as data ownership, availability, and durability. Evaluate the mechanisms for backing up or exporting your data; these are especially important in case you decide to switch providers or the provider ceases service.

- **Can your other services and self-hosted applications integrate with the SaaS application, regardless of the environment?**

When adopting a SaaS solution, it is easy to assume that services and applications that share the same hosting environment (that is, applications that use the same cloud provider or the same vendor's data center) will have a more seamless integration. However, most SaaS solutions today have broad support for API and third-party integrations, so don't limit yourself to solutions that are hosted in the same environment. If the necessary integrations exist, the solutions don't have to share the same underlying environment. For example, let's say that you're using a SaaS solution such as Google Drive or Microsoft OneDrive for cloud-based student file storage. To provide virtual desktops and application streaming to your students, you might determine that [Amazon WorkSpaces Applications](#) is the best fit for your requirements. Although these services run in different environments, WorkSpaces Applications has native integrations with Google Drive and Microsoft OneDrive, so your students can continue to use their existing storage.

- **Does the SaaS application support centralized identity management?**

To prevent your IT team from having to manage disparate identity stores and your users from having to remember multiple sets of credentials, make sure that your SaaS solutions support integration with your existing identity management or single sign-on solutions. Fragmented identity management decreases productivity and can lead to bad security practices such as privilege creep and weak passwords. If your desired SaaS solution doesn't support single sign-on or your existing identity store, evaluate whether the business value of adopting the solution outweighs the increased burden on users and staff.

- **How can you secure network communication with the SaaS application?**

In some cases, you might need a self-hosted application to communicate with a SaaS application. Typically, this communication will be through APIs that are secured with appropriate authentication and authorization mechanisms. However, depending on the hosting environments of the two applications, alternate or additional mechanisms might be required to simplify or secure that communication. For example, if you self-host an application with a cloud provider and need to integrate it with a SaaS application that's hosted on the same cloud provider, the vendor might provide several connection options. You might be able to use cloud-specific peering connections, private APIs, or private interfaces such as [AWS PrivateLink](#) to prevent that communication from traversing the public internet. Similarly, if your on-premises application has a dedicated network connection to a cloud provider through a service such as [AWS Direct Connect](#), you could use that same connection to communicate with SaaS applications that are hosted on the same cloud provider.

Establish security and governance requirements for each cloud service provider

Educational institutions have a variety of compliance, governance, and cybersecurity objectives that they must achieve. The risks of failing to meet these objectives can include institutional reputation loss, monetary fines, ransoms, sensitive data breaches, intellectual property theft, and degraded or complete loss of mission-critical functions. Because of the [shared responsibility model](#), institutions that adopt cloud services can reduce administrative burden by offloading some of the responsibility for infrastructure security to the cloud service provider. Furthermore, you 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 deployment. Examples include services such as [AWS WAF](#) for web application protection, [AWS Shield](#) for distributed denial of service (DDoS) protection, and [Amazon GuardDuty](#) for threat detection. A successful cloud security and governance strategy allows IT and security teams to focus on building systems that are secure by design, helps the institution rapidly adapt to evolving mission requirements, and provides faculty and researchers with secure environments for ground-breaking learning and innovation. To evaluate your security and governance requirements, consider the following key questions.

- **Which compliance frameworks must your workloads align to?**

Educational institutions must adhere to many compliance frameworks because of the multitude of stakeholders and workloads they support. These compliance frameworks include the Family Educational Rights and Privacy Act (FERPA), the Health Insurance Portability and Accountability Act (HIPAA), the Federal Risk and Authorization Management Program (FedRAMP), the Cybersecurity Maturity Model Certification (CMMC), the International Traffic in Arms Regulations (ITAR), the Criminal Justice Information Services (CJIS), and the Payment Card Industry Data Security Standard (PCI DSS). In some cases, such as with CMMC, research grant funding isn't released until the relevant workloads are certified as compliant. Each framework is unique and might apply only to a subset of workloads. Make sure that you know which workloads must adhere to which requirements and that you are able to achieve those requirements in each workload's environment. In cloud environments, make sure that you understand your responsibilities compared with the cloud provider's responsibilities. You should have the knowledge, resources, and skill sets that are necessary to achieve and maintain compliance.

- **Which mechanisms do you have in place to enforce compliance across multiple cloud providers without inhibiting innovation?**

If your academic institution is new to the cloud, we recommend that you select one primary strategic cloud service provider and focus on understanding how to architect, engineer, and operate cloud environments that are secure by design. Ideally, security controls that are automatically embedded within self-service systems allow users to rapidly deploy secure cloud environments with a minimum amount of intervention from IT teams. Focusing on a single provider limits the amount of resources and time you must invest to ensure security and compliance. The most successful institutions select a cloud service provider that can support the majority of compliance requirements, has a robust network of partners, offers prebuilt compliance solutions, and makes secure self-service automation available. If you must ensure security and compliance across multiple cloud providers, additional investment will be required to build the skill sets and resources to manage compliance for each environment. If each cloud provider uses a different foundational environment, or landing zone, you need to understand which compliance standards and requirements each landing zone can support, and this might determine whether certain workloads can be hosted on that provider. You might manage compliance for each provider separately or use custom-built or partner solutions that can centralize management across providers. [AWS Marketplace](#) provides turnkey solutions that can also meet your compliance requirements.

- **How can you assess and control cost and usage across multiple cloud providers?**

If your academic institution is new to the cloud, we recommend that you establish cost visibility and control mechanisms to gain insight into which cloud services are being used, who the cloud resources belong to, what the purpose of those cloud resources are, and what potential cost savings can be achieved by optimizing consumption. Institutions can achieve significant return on investment by partnering with their cloud service provider to migrate and modernize mission-critical systems, because they can negotiate enterprise-level agreements, benefit from volume pricing, and take advantage of the cloud service provider's expertise. If you must control cost and usage across multiple providers, consider how you can aggregate and analyze cost and usage from each provider, either with in-house processes and tooling or by using partner solutions. Many organizations are starting to identify cloud financial operations (FinOps) as a key function and dedicating resources to evangelizing and implementing capabilities for cloud cost management and optimization.

- **Do you have mechanisms in place to easily manage user permissions over time?**

We recommend that academic institutions understand core stakeholder needs when they first approach the cloud. Users of institutional systems include students, faculty, researchers, IT staff, administration, security, the general public, and third-party collaborators. You should identify

the core needs of these users and make sure that you have appropriate mechanisms in place to grant them access to cloud services. Different types of users require different types of access to cloud services. For example, students, faculty, and the general public need access to applications; IT staff, administrators, and security need access to cloud infrastructure; researchers and their third-party collaborators need access to secure research environments; faculty need access to secure teaching environments and might even want to provide students with hands-on access to cloud technologies. You should have tooling in place to [centrally manage these identities](#) in an automated way, and use established processes to identify, grant, and revoke permissions as roles and responsibilities change over time.

- **Do you have mechanisms in place to appropriately integrate new systems with your identity management solution?**

We recommend that academic institutions make it easy to integrate new systems with their identity management systems. This gives the institution the flexibility to support a variety of mission-critical functions by allowing stakeholders to procure and build systems that can easily be integrated into the identity management system. By simplifying the integration process, 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). Make sure that your identity management system can interoperate with the necessary systems through native integrations or industry-standard protocols.

- **Do you have mechanisms in place to enable effective incident detection and response?**

Educational institutions are frequently the target of cyberattacks and ransomware. To help detect and respond to such incidents effectively, we recommend a bifurcated approach:

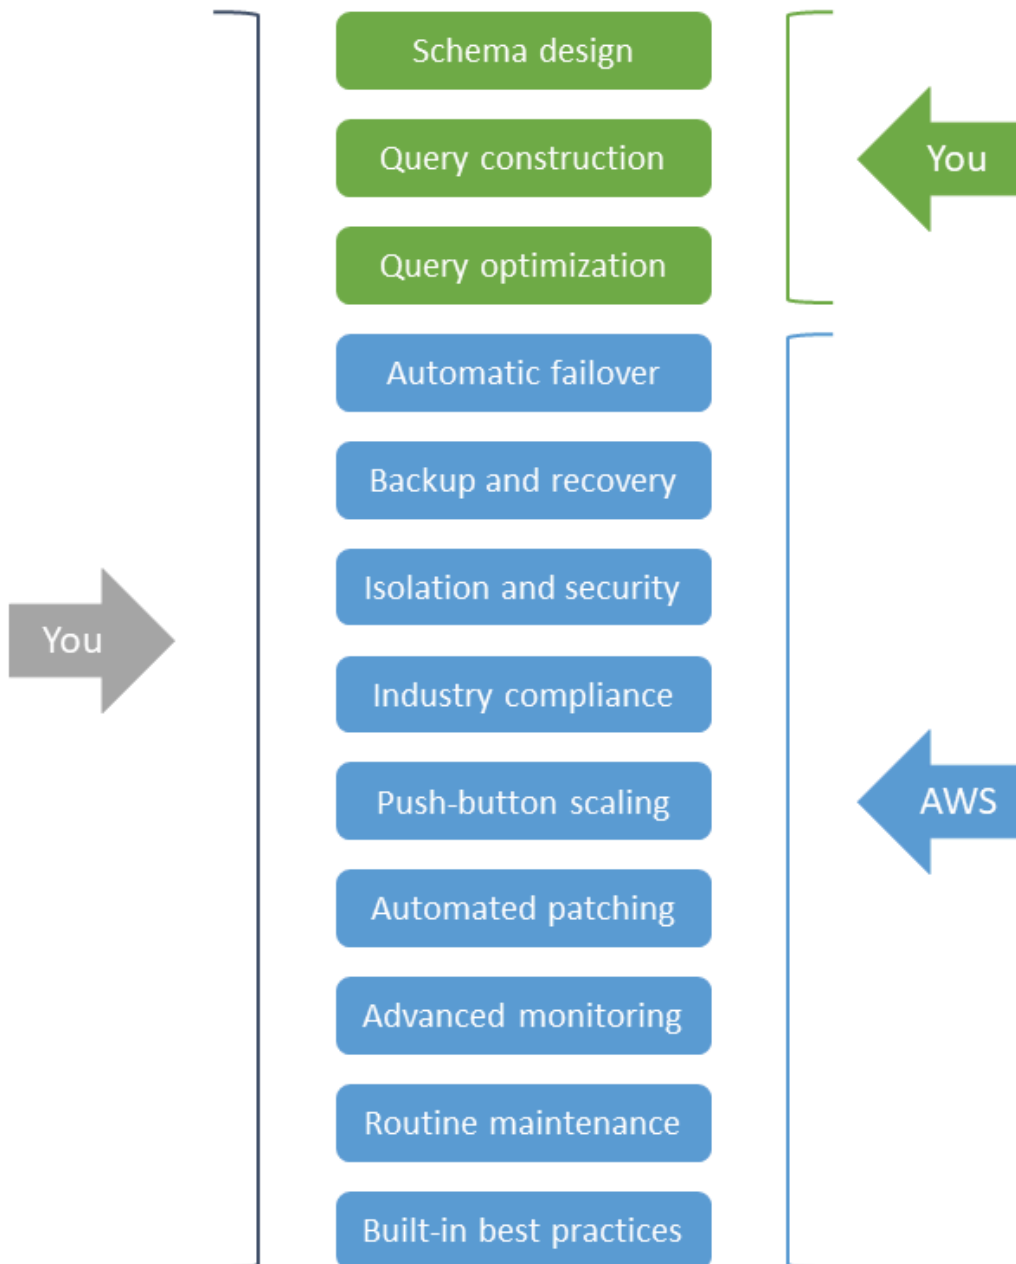
- Focus your efforts on preventative measures in the form of security controls that are automatically embedded in cloud environments.
- Implement detection capabilities that help cyberincident responders detect, contain, and mitigate security breaches in a timely fashion.

As with compliance, you must ensure that you have the resources, skill sets, and tools to detect, prevent, and respond to events in each environment. By focusing on a single, primary cloud provider, you can limit the resources that are required. Academic institutions that do not have a mature security operations team should look to independent software vendors, managed detection and response providers, and cybersecurity consultants for help in these areas.

Adopt cloud-native, managed services wherever possible and practical

When you initially consider how to take advantage of cloud services, using infrastructure services and development tools that your teams are familiar with might seem like the best path forward. However, selecting cloud-native managed services, especially serverless options, can greatly reduce cost, effort, and complexity.

Cloud-native, managed services eliminate many of the undifferentiated IT tasks that require time and effort from your staff that could be better spent on mission-focused activities. In addition, as providers improve the capabilities of their services, your solutions naturally inherit incremental improvements in efficiency, security, resilience, performance, and other characteristics. For example, a fully managed database service is a feature-rich relational database management system, but you don't have to provision and manage the underlying server and operating system that the database runs on. This eliminates administrative tasks that are typically required when you maintain a relational database in your own data center or on a self-managed virtual server that you provision in the cloud. The following diagram illustrates this difference.

Self-managed
database servicesFully managed
database services

The benefits of eliminating infrastructure management are clear when you compare any cloud-native managed service against a comparable self-managed approach. As a result, whenever you need to deploy components that your purchased or custom-developed applications will run on, you should use cloud-native, managed services to reduce time and effort.

When your team is responsible for building, deploying, or managing solutions in the cloud, use cloud-native, managed services to take full advantage of your cloud provider's differentiated capabilities and innovations. This strategy enables you to select, integrate, and deploy cloud services in a way that reduces the time and effort that these projects require, while increasing their resilience and security. For a successful cloud strategy, consider adopting these cloud-native *building blocks* when you migrate custom solutions to the cloud, develop new solutions in the cloud, or deploy licensed software on the cloud. When you evaluate options for cloud-native, managed services, consider the following key questions.

- **Do you need to focus more of your staff's time and effort on functionality that's core to your educational mission?**

Managing servers, even virtual ones, requires time and attention to ensure that they remain up to date with system software upgrades and patches. Using managed services that handle these tasks for you lets you direct IT staff time toward activities that align more directly to your institution's mission. For example, if you need to deploy containers, consider a serverless, managed service such as [AWS Fargate](#) so that you do not have to configure and maintain servers. By eliminating the need to procure, provision, and manage the underlying infrastructure, you are able to focus instead on delivering new functionality, optimizing performance, and improving user experience. Consider this benefit when you evaluate managed services against self-managed options.

- **What effort will it take for your team to adopt cloud-native managed services?**

There can be a learning curve to designing and implementing solutions with cloud-native, managed services, but these efforts will be repaid with reductions in cost, time, and complexity over the lifetime of a solution. Because of the on-demand, pay-as-you-go nature of cloud computing, cloud-native services enable you to quickly iterate and experiment in a more agile way while avoiding upfront investments. This leads to increased innovation and shorter project timelines. However, to realize these benefits effectively, consider what might be necessary to adopt and use the service, such as staff training on optimal usage patterns and code refactoring to accommodate service-specific APIs. Even if the service uses industry-standard or open source APIs, you might need to refactor or configure your application to handle feature disparity or version mismatches.

- **How do you currently deploy and manage infrastructure? Do you need to maintain that level of control?**

There are a variety of ways to host and manage infrastructure in the cloud, including using bare-metal hosts, virtual machines, managed container services, and serverless offerings. Even if you're currently using similar infrastructure, such as virtual machines or containers, in your on-premises environment, consider if an alternate approach would be suitable for certain workloads. For example, instead of running all applications on virtual machines, consider containerizing your applications and take advantage of managed container services such as [Amazon Elastic Container Service \(Amazon ECS\)](#). This might require refactoring, but you can use a tool such as [AWS App2Container](#) to simplify and assist with containerization. Taking this a step further, instead of deploying servers or containers for all components, consider fully serverless options. Serverless technologies feature automatic scaling, built-in high availability, and a pay-for-use billing model to increase agility and optimize costs. At the same time, they eliminate the need to manage servers and to plan for capacity. Serverless computing services such as [AWS Lambda](#) are core to serverless architectures. Lambda supports common programming languages and allows developers to focus on application code instead of managing infrastructure. Explore these options for each workload, and consider factors such as learning curve, management overhead, cost, and licensing.

- **Do you have to deploy and manage infrastructure for any licensed software?**

When you deploy and manage licensed software from independent software vendors (ISVs), it might seem logical to mimic your on-premises deployment with cloud infrastructure. For example, you might consider replacing on-premises virtual machines with cloud-hosted virtual machines. Although this is a viable option, consider whether you can replace any components of the architecture with cloud-native, managed services. For example, you might be able to replace a self-managed database server with a fully managed database service that reduces administrative burden while running the same database engine. Many ISVs already use cloud architectures that take advantage of managed services, and might even offer prebuilt templates to simplify deployment. Where possible, you should prefer ISVs that offer prescriptive guidance and support for cloud deployments. Before you deploy licensed software to the cloud, be sure to consult with your ISV to understand how cloud environment licensing might differ from on-premises licensing.

- **Are you concerned that using a managed service might result in vendor lock-in?**

Many cloud-native, managed services are built to support common industry standards and APIs. For example, analytics services such as [AWS Glue](#) and [Amazon EMR](#) are built on industry standard processing and storage frameworks such as Apache Spark and Apache Parquet. [AWS Lambda](#) natively supports Java, Go, Microsoft PowerShell, Node.js, C#, Python, and Ruby code.

[Amazon Relational Database Service \(Amazon RDS\)](#) supports multiple versions of common database engines, including SQL Server, Oracle, PostgreSQL, and MySQL. When services have proprietary APIs, native or partner solutions might be available to interact with the APIs by using common, cloud-agnostic protocols. For example, [Amazon Simple Storage Service \(Amazon S3\)](#) has a service-specific API for direct integration, but you can also interact with it by using standard storage protocols such as Network File System (NFS), Server Message Block (SMB), and Internet Small Computer Systems Interface (iSCSI) when you use [AWS Storage Gateway](#). You should still focus on choosing the cloud-native, managed service that best meets your needs while reducing operational overhead to the greatest extent, but you might prefer services that use or make available common industry standards and protocols.

Implement hybrid architectures when existing, on-premises investments incentivize continued use

Most educational institutions have invested in on-premises data centers of varying scale to host enterprise applications, data storage solutions, end-user computing (EUC) environments, and shared computing resources. All the resources in these data centers are subject to different refresh cycles, where you must consider future growth and provision enough capacity to accommodate peak scale, which might be necessary only a few times a year. As a result, resources often sit idle until the next refresh cycle. Planning for, budgeting, procuring, and deploying new hardware can take weeks, if not months or longer. This lengthy process stifles innovation and can delay learning and research.

Cloud computing solves many of these challenges. The cloud provides on-demand, pay-as-you-go IT resources, so you can more closely match current capacity with actual demands without large, upfront planning and investment. However, if you have already made a significant investment in on-premises hardware and resources, you should seek to utilize those resources efficiently and augment them as needed with cloud technology in a hybrid model.

A successful hybrid cloud strategy takes advantage of existing investments while providing greater agility, scalability, and reliability than those investments alone can support. The following considerations can help you get started.

- **When you must host a new workload, do you think about cloud first?**

How you use public and private cloud infrastructure together defines your hybrid cloud strategy. A cloud first approach doesn't mean that the cloud is the better choice for all your workloads.

However, when you plan for new workloads, evaluate the cloud as the first option, especially for workloads that require new technology or exceed the storage and compute capacity available on premises. Workloads that have transient, inconsistent usage patterns, need fast results, are easily portable, or require the newest hardware are ideal candidates for the scalability and elasticity of the cloud. Also, consider whether the workload would benefit from any cloud-native, managed services that are unavailable on premises, even if you do have available capacity.

- **Do you understand the TCO of your on-premises environment and partner with your CFO when making new investments?**

We recommend that you understand the true total cost of ownership (TCO) of maintaining your own on-premises data center. There are many hidden costs associated with owning and operating infrastructure on premises, including not just hardware, software, and support, but also facilities, utilities, insurance, and staff hours. These costs can negatively impact staff productivity, operational resilience, and business agility. Evaluate your current licensing structures and their renewal and maintenance periods as well. Partnering with your chief financial officer (CFO) can help you identify all hidden costs when you plan to make new investments. Some licenses might offer Bring Your Own License (BYOL) options in the cloud, or they might be more or less conducive to cloud services. Understanding the true TCO of your current infrastructure helps you prioritize cloud adoption for workloads that have the greatest impact on your organization's total TCO. Your AWS account team has tools readily available to help you better understand your on-premises TCO.

- **What infrastructure will you need to support hybrid deployments?**

To successfully adopt hybrid models, you will need foundational network, security, and infrastructure tooling. Make sure that you can maintain adequate network connectivity with your cloud provider. This could be through a combination of existing internet connectivity, virtual private networks (VPNs), dedicated connections such as Direct Connect, third-party connectivity providers, or [Internet2](#) and regional research and education networks. Make sure that you have unified identity and access management across your on-premises and cloud environments. Establish tools and processes to enforce consistent security, cost, and usage guardrails.

- **Is your IT staff ready to operate hybrid deployments?**

Cloud services can require specific skill sets that your team might not have. To limit the training and enablement necessary to upskill your IT staff for effective cloud adoption, consider whether the cloud provider offers any services that reuse and build upon existing skill sets across on premises and the cloud. For example, if you use and are familiar with Kubernetes, you might

consider using [Amazon Elastic Kubernetes Service \(Amazon EKS\)](#) or [Amazon EKS Anywhere](#). If you use and are familiar with NetApp, you might consider using [Amazon FSx for NetApp ONTAP](#). Similarly, also consider whether any existing partner solutions you use have native integrations or support for cloud environments.

- **Can you offload long-term storage or low-usage compute from on premises to the cloud?**

Cloud storage provides several cost-effective options for long-term data storage. For example, [Amazon Simple Storage Service \(Amazon S3\)](#) offers various storage tiers that are optimized for different use cases. If your institution is required to keep certain data for a long period of time, consider cold storage solutions such as [Amazon Glacier](#). Offloading this data into cloud storage can free up valuable high-performance, on-premises storage. Services such as [AWS Storage Gateway](#) make it easy for on-premises applications to access cloud storage tiers through standard protocols such as SMB, NFS, and iSCSI. Similarly, consider offloading any compute tasks that have infrequent or low usage. If you have on-premises servers that are dedicated to such tasks, you can instead use scalable cloud compute services, where resources are provisioned on demand and you pay only for what you use. Those low-cost, long-term storage and low-usage compute options also make the cloud ideal for backup and disaster recovery. You can use secure, durable, scalable storage and compute in the cloud to protect your data and quickly recover in case of a disaster without having to maintain the necessary storage and compute infrastructure yourself.

- **Do you have enough capacity on premises to experiment and innovate?**

The lack of elasticity and agility in fixed-size, on-premises environments can limit the services and technology available to your users. If you have strict refresh cycles, new workloads might have to wait until the next cycle for implementation. This operating model can limit experimentation and slow innovation. When you have a new or novel workload that needs to be tested, consider using scalable, elastic cloud services. Cloud resources can be provisioned and deprovisioned on demand and you pay only for what you use, so you can experiment and *fail fast* while minimizing organizational risk.

- **Do you have unique compliance or performance requirements that compel you to keep data on premises?**

Workloads with strict data residency or latency requirements might dictate that you keep data on premises or as close to your users as possible. For these use cases, you can prioritize the use of existing, on-premises resources. However, consider whether your cloud provider offers edge services or mechanisms to use cloud-based technology on premises. Edge services deliver

data processing, analysis, and storage closer to your own endpoints, and enable you to deploy tools outside of standard cloud provider data centers. For example, AWS offers services such as [AWS Local Zones](#) and [AWS Wavelength](#) to deploy applications in specific locations closer to end users. You can also bring cloud services and functionality into your existing data center with services such as [AWS Outposts](#), [AWS Storage Gateway](#), [Amazon ECS Anywhere](#), and [Amazon EKS Anywhere](#).

Reserve multicloud only for workloads that can't meet their technical or business requirements through a single cloud provider

Multicloud refers to the use of cloud services from multiple (two or more) cloud service providers. Having a multicloud strategy can offer certain benefits, such as the option to unlock the differentiated capabilities of multiple cloud providers or the ability to meet data sovereignty requirements that a single cloud provider might not be able to accommodate. However, for each provider that you use, make sure that you have the proper people, skills, training, and toolsets in place to use that provider effectively. Furthermore, if you want to use a multicloud strategy for a specific workload, you will need additional resources to integrate and interoperate the necessary services from each cloud provider. **We recommend that you consider multicloud only when the benefits outweigh the increased investment.** To determine whether you should choose a multicloud strategy, consider the following key questions.

- **Do you have the resources and skill sets to navigate services offered by different cloud providers?**

When multiple cloud providers offer various products and services, your staff needs essential skills to navigate each provider's capabilities. Using one cloud provider's services alone can require upskilling and training for your staff, depending on the services and features you are using. If you're considering a multicloud strategy, evaluate your existing resources to determine what additional skill sets you would need to use services from multiple cloud providers effectively. You might have to augment your staff or invest additional time and money in upskilling and training beyond what would be required for a single cloud provider. If you already have individual teams or users who are using different cloud providers, consider the organizational benefits of consolidating them onto a primary cloud provider on a case-by-case basis.

- **What additional overhead would a particular multicloud architecture introduce?**

A common driver for multicloud is the desire to use a specific managed service from one provider that has capabilities that can be differentiated from the services of another cloud provider. For example, you might want to use one cloud provider for your infrastructure needs and another provider's managed service for domain and directory services. However, even if that single managed service reduces administrative burden and simplifies the management of that architecture component, it could introduce additional overhead for other workloads, such as code refactoring, private connectivity needs, or manual integration work. Identify this additional overhead up front and make sure that it doesn't offset or eclipse the benefits your team stands to gain from the differentiated service.

- **How will you centralize monitoring and management across cloud providers?**

As you start to deploy applications and functionalities by using resources from different cloud providers, consider how you will tag, monitor, and manage such resources. Each provider will have their own tooling, which you might be able to extend into other environments. For example, you can use [Amazon CloudWatch](#) to monitor key metrics and logs, create alarms, and visualize your applications and infrastructure across single, hybrid, and multicloud environments. You can also use [AWS Systems Manager](#) to improve resource visibility and control, quickly diagnose and remediate operational issues, and automate processes such as updating and patching virtual machines across environments. If you have requirements that a provider's tools cannot support, you can explore partner solutions, but these could add additional cost or integration effort.

- **How can you manage infrastructure as code with automation when using different cloud providers?**

When you run resources in the cloud, automated provisioning and management of resources helps you manage various environments efficiently. The APIs and native automation tools vary across cloud providers. If possible, consider using a common set of orchestration and deployment tools that can accommodate different cloud provider resources. This provides greater flexibility and simplifies operations across multiple clouds. However, it might be simpler to use each provider's native automation separately and establish organizational processes to ensure appropriate usage.

- **Do you have compliance and regulatory requirements that each cloud provider must satisfy?**

You might have regulatory considerations that dictate how data should be stored and handled. Focus on standardizing policies (such as network traffic, storage, and security) that can be

applied automatically to each cloud environment across cloud providers. Consider how your applications will communicate with their data, and host them on the same provider. If your applications and their data are fragmented across providers, it will be difficult to ensure that you are meeting compliance and regulatory requirements. It is often best to have applications as close to data as possible to minimize network latency, maximize data throughput, and limit data egress while simplifying security and access controls.

- **Are you able to minimize TCO and maximize pricing discounts when you deploy applications across cloud providers?**

It is important to account for the total cost of ownership (TCO) when considering multicloud. Running your applications across multiple cloud providers can increase operational costs and administrative overhead to maintain and manage resources in each environment. Furthermore, spreading usage across multiple providers makes it more difficult to take advantage of a specific provider's volume pricing discounts or enterprise agreements. Take these factors into account when you determine whether the benefits of multicloud warrant the increased TCO.

Example use cases

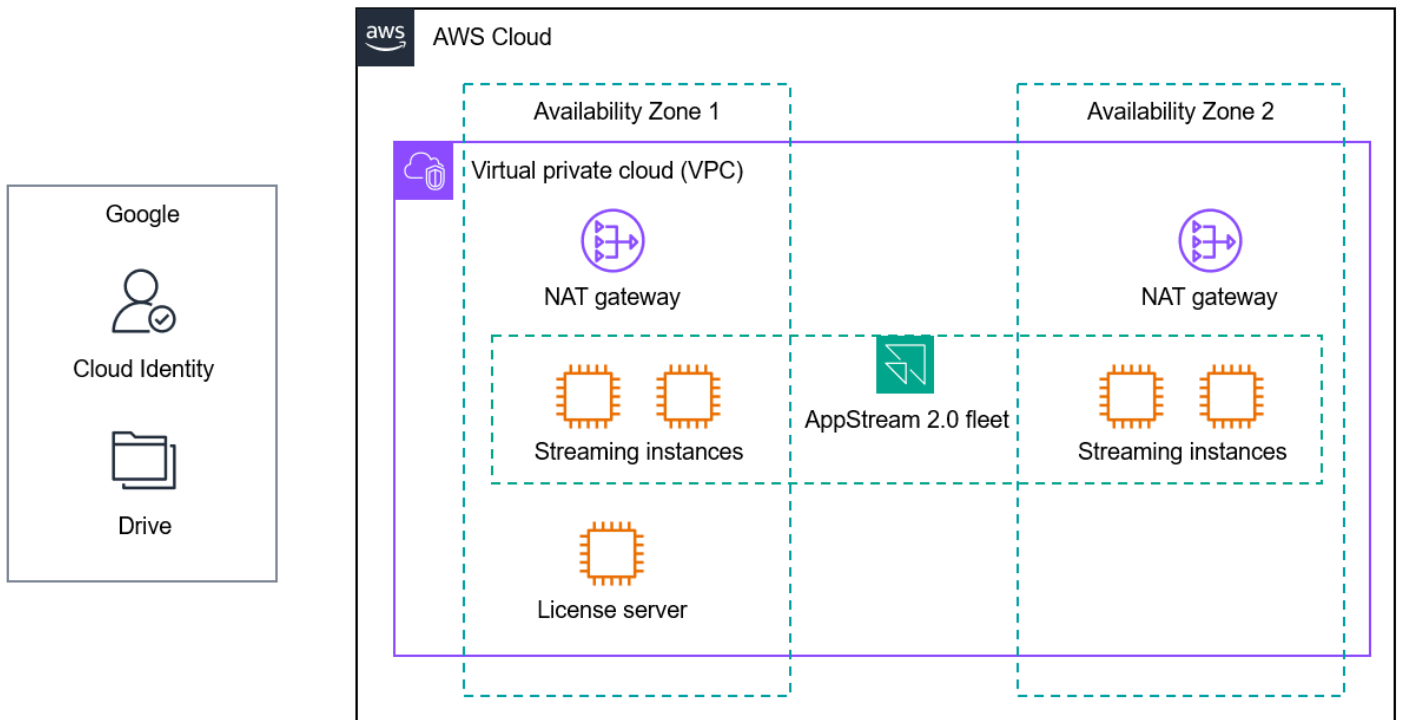
To better understand the application of these principles in different scenarios, let's discuss some example use cases. These use cases are based on how real-world educational institutions are adopting cloud services.

- [Virtual computer labs](#)
- [Predicting student success](#)
- [Identity federation and single sign-on](#)
- [Cloud bursting for research computing](#)

Virtual computer labs

Despite the popularity of web-based learning tools and the abundance of user devices such as laptops, Chromebooks, and tablets, most educational institutions maintain physical computer labs for resource-intensive or legacy applications. These computer labs are often necessities for science, technology, engineering, and math (STEM), career and technical education (CTE), media and art, engineering, and similar curricula. Schools can augment or replace physical computer labs with cloud-based virtual desktops or application streaming services to ensure that all students have access to the applications they need at any time, from any place, and on any device. This improves digital equity, enables remote learning, ensures a consistent user experience, and secures remote access while lowering cost.

In primary and secondary (K12) education, many US schools use [Amazon WorkSpaces Applications](#), a fully managed desktop and application streaming service, to deliver virtual computer labs to provide access to Adobe Creative Cloud, Autodesk software, STEM and CTE curricula such as Project Lead the Way (PLTW), and more. Many K12 organizations already manage student single sign-on and file storage through Google Workspace and Google Drive, which are SaaS applications. These institutions can set up single sign-on between Google Workspace and WorkSpaces Applications through SAML 2.0 federation. They can also configure native integration between WorkSpaces Applications and Google Drive so that students can use existing storage. The following diagram illustrates the WorkSpaces Applications deployment for this use case.



This architecture follows these recommendations:

- **Select a primary, strategic cloud provider.** This architecture uses cloud services from one primary cloud provider. Although it includes integration with SaaS applications that are not hosted on the same provider, those integrations are done through simple configurations. Cloud expertise and skill sets are necessary only to deploy and manage services from the primary cloud provider.
- **Differentiate between SaaS applications and foundational cloud services.** Google Workspace and Google Drive are not hosted on the same cloud provider as AppStream 2.0, but that is acceptable because this deployment provides the necessary integrations. Single sign-on enables centralized identity management and is securely configured through SAML 2.0. Enabling persistent cloud storage for students requires simple configuration changes in Google Drive and WorkSpaces Applications.
- **Establish security and governance requirements for each cloud service provider.** The services and integrations used in this architecture help meet an institution's security and governance requirements. Streaming traffic is encrypted. Federation through Google Workspace allows for centralized identity management. Network services such as [Amazon Virtual Private Cloud \(Amazon VPC\)](#) support the configuration of subnets, routing, and firewalls. You can filter content by using DNS configuration, agents, virtual appliances, or managed services such as Amazon

Route 53 Resolver DNS Firewall. You can use services such as [AWS Control Tower](#) to help ensure that the AWS account that hosts WorkSpaces Applications adheres to standard organizational guardrails and controls.


- **Adopt cloud-native, managed solutions wherever possible and practical.** WorkSpaces Applications is a managed service for desktop and application streaming. You can stream desktops and applications without worrying about provisioning, scaling, or maintaining servers. You install your applications, connect the appropriate identity, network, and storage solutions, and then centrally manage and stream those applications to your users. This eliminates much of the undifferentiated heavy lifting that would be required to manage your own virtual desktop streaming solution.

Predicting student success

A Midwest university in the US discovered that a handful of key activities for incoming first-year students was highly predictive of success, both in the student's first semester of classes and in attaining their degree. The university wanted to implement a system that watched for these activities to be completed, and when key deadlines approached or passed, they wanted to encourage students to complete these steps.

The SaaS learning management system (LMS) data was a key input for this solution, but its data proved to be challenging to access and process with the university IT team's data warehousing tools. In addition, the messages to students had to be sent through the school's cloud-based customer relationship management (CRM) system. To build a functional solution and assess the effectiveness of prompts to students, the university had to initiate messages through the CRM and gather data from it.

The university developed and deployed a solution into a single cloud environment. The solution is a mixture of cloud-native managed services, provisioned cloud servers, and integrations with on-premises systems and cloud-based SaaS applications. As the following diagram shows, the solution ingests data from the student information system (SIS), LMS, and CRM into a data lake. It uses this data to identify students who are in jeopardy of missing key activities, initiates messages to them through the CRM, and provides a dashboard to university leadership.



Amazon S3



AWS DMS



AWS Lambda



AWS Glue



Amazon SageMaker



Amazon Redshift



Amazon QuickSight

This architecture follows these recommendations:

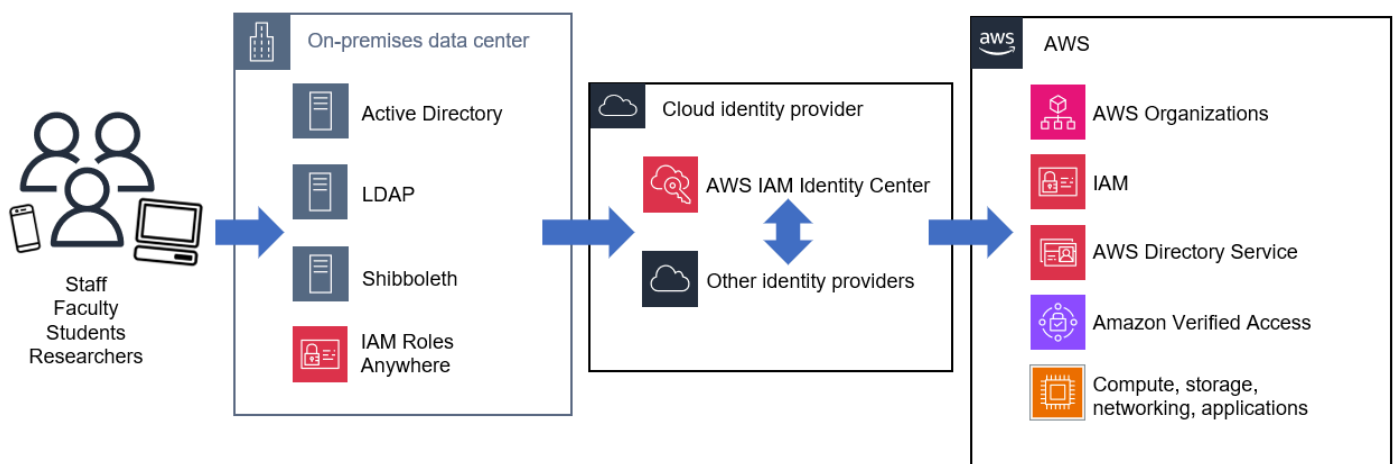
- **Select a primary, strategic cloud provider.** The university's strategic cloud provider houses the entire deployed solution. This enables the IT and business staff to focus on developing skills in a single, integrated set of cloud capabilities.
- **Differentiate between SaaS applications and foundational cloud services.** The university differentiates between SaaS applications and core cloud analytics services, and uses integrations with the SaaS applications to gather data and initiate the appropriate communications.
- **Establish security and governance requirements for each cloud service provider.** The university ensures that all components of the architecture are secure by enforcing guardrails and controls, including encryption in transit and at rest, to handle student data appropriately.

- **Adopt cloud-native, managed solutions wherever possible and practical.** Cloud-native managed services are used for data ingestion, storage, database, and extract, transform, and load (ETL) functionality, which reduces the time to develop the end-to-end data processing workflow.

Identity federation and single sign-on

Ensuring consistent identity management across core systems is key to successfully and securely adopting any technology. Educational institutions are increasingly adopting cloud-based identity and single sign-on solutions such as [AWS IAM Identity Center](#), Microsoft Entra ID (formerly Azure Active Directory), Okta, JumpCloud, OneLogin, Ping Identity, and CyberArk to simplify identity management, lower operational burden, and centrally enforce best practices such as multi-factor authentication and least privilege access.

Many of these institutions still maintain identity management and directory services such as Active Directory and Shibboleth for their on-premises environments. These can be integrated with cloud-based solutions to enable centralized identity management and single sign-on for your students, faculty, and staff. Cloud solution providers should have robust, easy-to-integrate identity management platforms that allow you to federate identities through cloud identity providers to your existing applications, your SaaS solutions, and cloud services. The following diagram shows an example architecture.



This architecture follows these recommendations:

- **Select a primary, strategic cloud provider.** This architecture uses AWS as the primary cloud provider. By integrating with a cloud identity provider and existing identity management

and directory services on premises, this architecture supports automated provisioning and management of access both to the primary cloud provider's services and to other applications and SaaS solutions. This ensures that security and governance requirements are met in a consistent, easy to manage way as more applications and services are added to the institution's technology portfolio.

- **Differentiate between SaaS applications and foundational cloud services.** This architecture integrates multiple types of cloud-based, SaaS, and on-premises identity systems to provide access to AWS Cloud services and other applications. Many cloud-based identity provider and single sign-on solutions are also SaaS applications, and they can use native integrations and standard protocols such as SAML to work across environments.
- **Establish security and governance requirements for each cloud service provider.** This architecture adheres to guidance on identity and access management issued by numerous security frameworks, including National Institute of Standards and Technology (NIST) Cybersecurity Framework (CSF), NIST 800-171, and NIST 800-53. Integrations with [AWS Organizations](#), [AWS Identity and Access Management \(IAM\)](#), and other [AWS security, identity, and compliance services](#) help provide secure, granular access controls based on group permissions.
- **Adopt cloud-native, managed services wherever possible and practical.** This architecture uses cloud-based, managed services for identity management and single sign-on. This decreases the time and energy spent on infrastructure management and makes it easier to maintain these critical systems.
- **Implement hybrid architectures when existing, on-premises investments incentivize continued use.** This architecture integrates existing, on-premises investments in infrastructure for hosting Active Directory, Lightweight Directory Access Control (LDAP), and Shibboleth workloads, and provides a path to eventually move core identity services into cloud-based infrastructure. Additionally, if your on-premises workloads need certificate-based access to AWS resources, you can use [AWS Identity and Access Management Roles Anywhere](#).

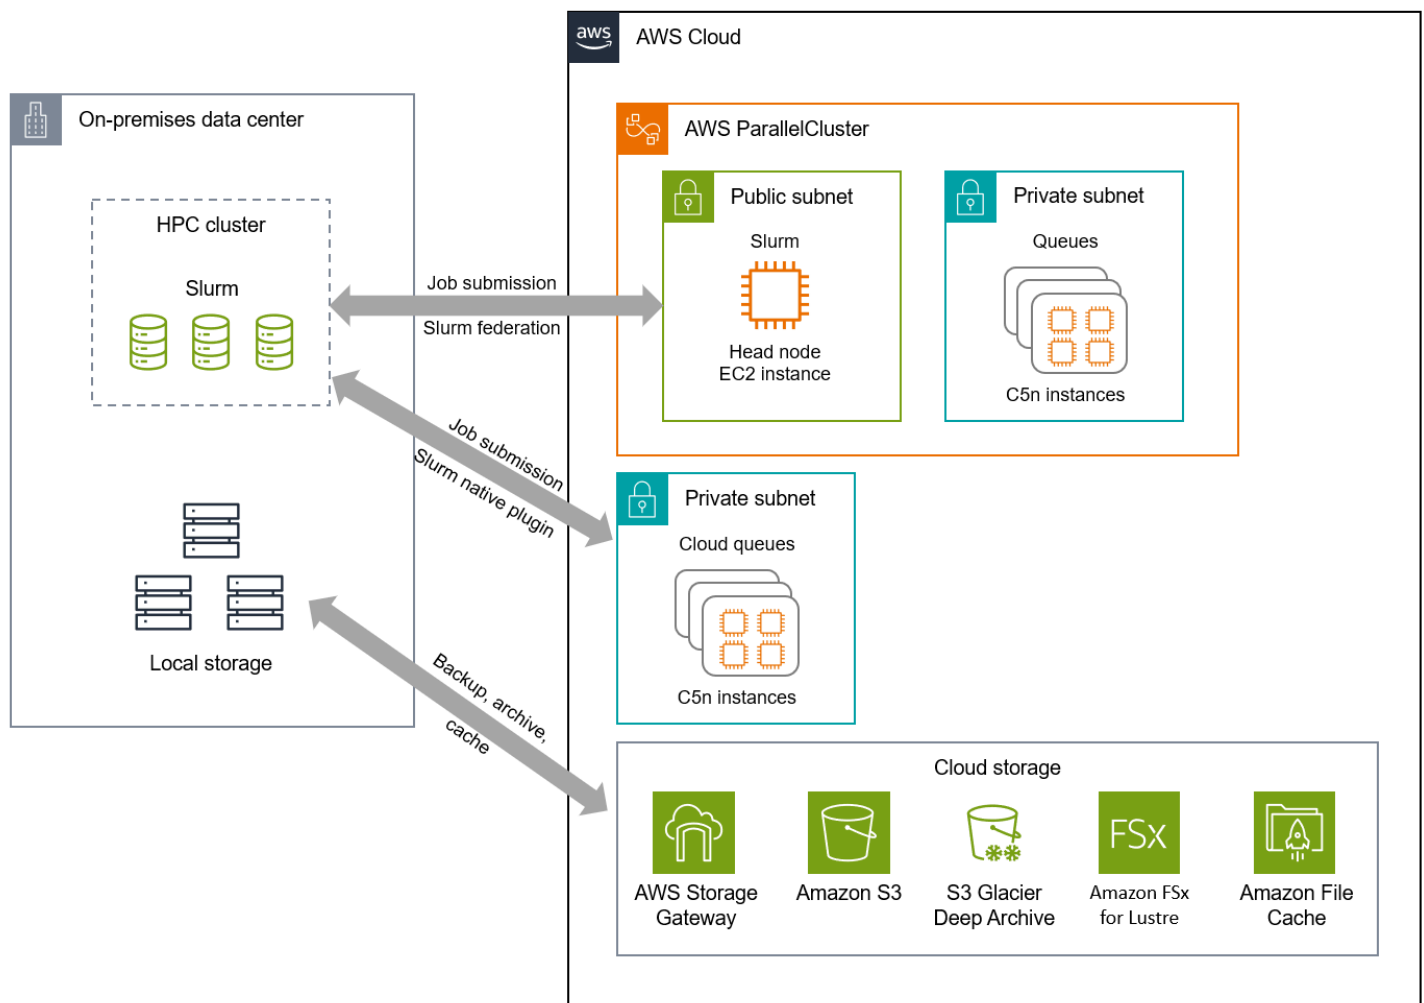
Cloud bursting for research computing

The research computing group at an R1 (Doctoral Universities – Very High Research Activity) research institution in the US had been running on-premises high performance computing (HPC) clusters with the Slurm scheduler for many years. Except for a few weeks of scheduled maintenance, the clusters were running at 80-95 percent utilization rate with most of their queues full.

The increasing number of research activities at the institution introduced capacity and capability challenges. A few high-profile researchers were always performing long-running simulations on certain queues, which increased the wait time for other users. Newly hired faculty needed to run large numbers of weather simulations to build a novel artificial intelligence and machine learning (AI/ML) model for weather forecasting, but they required more capacity than was available. The research computing group was also getting more requests for the latest graphics processing units (GPUs) to train machine learning models. Even with funding for new GPUs, the team would need to wait months to get approval for expanding rack space in the data center.

Many researchers were unwilling to delete old data, so local storage capacity was also a challenge. A more scalable, long-term storage option was needed to free up valuable, high-performance storage on premises.

The cloud addresses these challenges with hybrid compute and storage solutions that let you *burst* research computing into the cloud when on-premises capacity isn't enough. The following architecture diagram illustrates a few compute and storage bursting approaches, using tools such as [AWS ParallelCluster](#) and [AWS Storage Gateway](#).



This architecture follows these recommendations:

- **Select a primary, strategic cloud provider.** This architecture uses one primary cloud provider to avoid being restricted by the least common denominator approach. This way, the institution can take advantage of the innovation and native compute and storage services that the primary cloud provider offers. The research computing team can focus on optimizing workloads in the environment provided by the primary cloud provider, not how to work in different cloud environments.
- **Establish security and governance requirements for each cloud service provider.** Each service and tool used in this architecture can be configured to meet the research computing team's security and governance requirements, including private connectivity, data encryption in transit and at rest, activity logging, and more.
- **Adopt cloud-native, managed services wherever possible and practical.** This architecture provides the ability to use managed storage and compute services as well as tools to simplify

cluster management. This way, the research computing team doesn't have to worry about managing clusters or underlying infrastructure on their own, which can be complex and time-consuming.

- **Implement hybrid architectures when existing, on-premises investments incentivize continued use.** This architecture allows the institution to continue using its on-premises resources and take advantage of the cloud to increase capacity and expand computing power on demand. With the cloud, the institution can right-size the compute type to maximize price-performance and access the latest technology to promote innovation without a large upfront investment in additional on-premises hardware.

Next steps

Selecting the right deployment model for cloud workloads requires careful consideration. Use the recommendations outlined in this paper to guide your decision-making and to avoid common pitfalls such as unnecessary complexity, increasing staff demands, inconsistent governance, and lowest common denominator approaches. By following these best practices, you can accelerate your cloud adoption to meet and exceed your institutional goals more effectively.

Remember to select a primary, strategic cloud provider, and establish a Cloud Center of Excellence (CCoE) to help drive organizational maturity to 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, reserve multicloud for only those workloads that truly require it.

AWS is well positioned to help you manage single, hybrid, and multicloud environments. Your institution 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 can 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 system for cloud implementation to start as a proof of concept. This will help you define the cloud foundation or framework to validate any assumptions, and will also enable future cloud implementations.
3. Engage your [AWS account team](#) to discuss your cloud implementation goals. The AWS account team can help provide clarifications, suggest approaches, identify dependencies, and also work with your teams to map out your journey from initial concept to implementation.

Contributors

Contributors to this guide include:

- Kevin Arand, Senior Manager, Solutions Architecture, Education, AWS
- Kevin McCandless, Senior Solutions Architect, K-12 Education, AWS
- Craig Jordan, Principal Solutions Architect, Education, AWS
- Jesse Roberts, Principal Solutions Architect, SLG & K-12 Education, AWS
- Jianjun Xu, Principal Solutions Architect, Education, AWS
- Josh Badal, Senior Solutions Architect, Education, AWS
- Raj Chary, Senior Solutions Architect, Education, AWS

Further reading

For additional information, refer to:

- [AWS Architecture Center](#)
- [Public Sector Cloud Transformation](#)
- [AWS Cloud Adoption Framework \(AWS CAF\)](#)
- [AWS Solutions for Hybrid and Multicloud](#)

Document history

The following table describes significant changes to this guide. If you want to be notified about future updates, you can subscribe to an [RSS feed](#).

Change	Description	Date
Initial publication	—	September 15, 2023

AWS Prescriptive Guidance glossary

The following are commonly used terms in strategies, guides, and patterns provided by AWS Prescriptive Guidance. To suggest entries, please use the **Provide feedback** link at the end of the glossary.

Numbers

7 Rs

Seven common migration strategies for moving applications to the cloud. These strategies build upon the 5 Rs that Gartner identified in 2011 and consist of the following:

- Refactor/re-architect – Move an application and modify its architecture by taking full advantage of cloud-native features to improve agility, performance, and scalability. This typically involves porting the operating system and database. Example: Migrate your on-premises Oracle database to the Amazon Aurora PostgreSQL-Compatible Edition.
- Replatform (lift and reshape) – Move an application to the cloud, and introduce some level of optimization to take advantage of cloud capabilities. Example: Migrate your on-premises Oracle database to Amazon Relational Database Service (Amazon RDS) for Oracle in the AWS Cloud.
- Repurchase (drop and shop) – Switch to a different product, typically by moving from a traditional license to a SaaS model. Example: Migrate your customer relationship management (CRM) system to Salesforce.com.
- Rehost (lift and shift) – Move an application to the cloud without making any changes to take advantage of cloud capabilities. Example: Migrate your on-premises Oracle database to Oracle on an EC2 instance in the AWS Cloud.
- Relocate (hypervisor-level lift and shift) – Move infrastructure to the cloud without purchasing new hardware, rewriting applications, or modifying your existing operations. You migrate servers from an on-premises platform to a cloud service for the same platform. Example: Migrate a Microsoft Hyper-V application to AWS.
- Retain (revisit) – Keep applications in your source environment. These might include applications that require major refactoring, and you want to postpone that work until a later time, and legacy applications that you want to retain, because there's no business justification for migrating them.

- **Retire** – Decommission or remove applications that are no longer needed in your source environment.

A

ABAC

See [attribute-based access control](#).

abstracted services

See [managed services](#).

ACID

See [atomicity, consistency, isolation, durability](#).

active-active migration

A database migration method in which the source and target databases are kept in sync (by using a bidirectional replication tool or dual write operations), and both databases handle transactions from connecting applications during migration. This method supports migration in small, controlled batches instead of requiring a one-time cutover. It's more flexible but requires more work than [active-passive migration](#).

active-passive migration

A database migration method in which the source and target databases are kept in sync, but only the source database handles transactions from connecting applications while data is replicated to the target database. The target database doesn't accept any transactions during migration.

aggregate function

A SQL function that operates on a group of rows and calculates a single return value for the group. Examples of aggregate functions include SUM and MAX.

AI

See [artificial intelligence](#).

AIOps

See [artificial intelligence operations](#).

anonymization

The process of permanently deleting personal information in a dataset. Anonymization can help protect personal privacy. Anonymized data is no longer considered to be personal data.

anti-pattern

A frequently used solution for a recurring issue where the solution is counter-productive, ineffective, or less effective than an alternative.

application control

A security approach that allows the use of only approved applications in order to help protect a system from malware.

application portfolio

A collection of detailed information about each application used by an organization, including the cost to build and maintain the application, and its business value. This information is key to [the portfolio discovery and analysis process](#) and helps identify and prioritize the applications to be migrated, modernized, and optimized.

artificial intelligence (AI)

The field of computer science that is dedicated to using computing technologies to perform cognitive functions that are typically associated with humans, such as learning, solving problems, and recognizing patterns. For more information, see [What is Artificial Intelligence?](#)

artificial intelligence operations (AIOps)

The process of using machine learning techniques to solve operational problems, reduce operational incidents and human intervention, and increase service quality. For more information about how AIOps is used in the AWS migration strategy, see the [operations integration guide](#).

asymmetric encryption

An encryption algorithm that uses a pair of keys, a public key for encryption and a private key for decryption. You can share the public key because it isn't used for decryption, but access to the private key should be highly restricted.

atomicity, consistency, isolation, durability (ACID)

A set of software properties that guarantee the data validity and operational reliability of a database, even in the case of errors, power failures, or other problems.

attribute-based access control (ABAC)

The practice of creating fine-grained permissions based on user attributes, such as department, job role, and team name. For more information, see [ABAC for AWS](#) in the AWS Identity and Access Management (IAM) documentation.

authoritative data source

A location where you store the primary version of data, which is considered to be the most reliable source of information. You can copy data from the authoritative data source to other locations for the purposes of processing or modifying the data, such as anonymizing, redacting, or pseudonymizing it.

Availability Zone

A distinct location within an AWS Region that is insulated from failures in other Availability Zones and provides inexpensive, low-latency network connectivity to other Availability Zones in the same Region.

AWS Cloud Adoption Framework (AWS CAF)

A framework of guidelines and best practices from AWS to help organizations develop an efficient and effective plan to move successfully to the cloud. AWS CAF organizes guidance into six focus areas called perspectives: business, people, governance, platform, security, and operations. The business, people, and governance perspectives focus on business skills and processes; the platform, security, and operations perspectives focus on technical skills and processes. For example, the people perspective targets stakeholders who handle human resources (HR), staffing functions, and people management. For this perspective, AWS CAF provides guidance for people development, training, and communications to help ready the organization for successful cloud adoption. For more information, see the [AWS CAF website](#) and the [AWS CAF whitepaper](#).

AWS Workload Qualification Framework (AWS WQF)

A tool that evaluates database migration workloads, recommends migration strategies, and provides work estimates. AWS WQF is included with AWS Schema Conversion Tool (AWS SCT). It analyzes database schemas and code objects, application code, dependencies, and performance characteristics, and provides assessment reports.

B

bad bot

A [bot](#) that is intended to disrupt or cause harm to individuals or organizations.

BCP

See [business continuity planning](#).

behavior graph

A unified, interactive view of resource behavior and interactions over time. You can use a behavior graph with Amazon Detective to examine failed logon attempts, suspicious API calls, and similar actions. For more information, see [Data in a behavior graph](#) in the Detective documentation.

big-endian system

A system that stores the most significant byte first. See also [endianness](#).

binary classification

A process that predicts a binary outcome (one of two possible classes). For example, your ML model might need to predict problems such as "Is this email spam or not spam?" or "Is this product a book or a car?"

bloom filter

A probabilistic, memory-efficient data structure that is used to test whether an element is a member of a set.

blue/green deployment

A deployment strategy where you create two separate but identical environments. You run the current application version in one environment (blue) and the new application version in the other environment (green). This strategy helps you quickly roll back with minimal impact.

bot

A software application that runs automated tasks over the internet and simulates human activity or interaction. Some bots are useful or beneficial, such as web crawlers that index information on the internet. Some other bots, known as *bad bots*, are intended to disrupt or cause harm to individuals or organizations.

botnet

Networks of [bots](#) that are infected by [malware](#) and are under the control of a single party, known as a *bot herder* or *bot operator*. Botnets are the best-known mechanism to scale bots and their impact.

branch

A contained area of a code repository. The first branch created in a repository is the *main branch*. You can create a new branch from an existing branch, and you can then develop features or fix bugs in the new branch. A branch you create to build a feature is commonly referred to as a *feature branch*. When the feature is ready for release, you merge the feature branch back into the main branch. For more information, see [About branches](#) (GitHub documentation).

break-glass access

In exceptional circumstances and through an approved process, a quick means for a user to gain access to an AWS account that they don't typically have permissions to access. For more information, see the [Implement break-glass procedures](#) indicator in the AWS Well-Architected guidance.

brownfield strategy

The existing infrastructure in your environment. When adopting a brownfield strategy for a system architecture, you design the architecture around the constraints of the current systems and infrastructure. If you are expanding the existing infrastructure, you might blend brownfield and [greenfield](#) strategies.

buffer cache

The memory area where the most frequently accessed data is stored.

business capability

What a business does to generate value (for example, sales, customer service, or marketing). Microservices architectures and development decisions can be driven by business capabilities. For more information, see the [Organized around business capabilities](#) section of the [Running containerized microservices on AWS](#) whitepaper.

business continuity planning (BCP)

A plan that addresses the potential impact of a disruptive event, such as a large-scale migration, on operations and enables a business to resume operations quickly.

C

CAF

See [AWS Cloud Adoption Framework](#).

canary deployment

The slow and incremental release of a version to end users. When you are confident, you deploy the new version and replace the current version in its entirety.

CCoE

See [Cloud Center of Excellence](#).

CDC

See [change data capture](#).

change data capture (CDC)

The process of tracking changes to a data source, such as a database table, and recording metadata about the change. You can use CDC for various purposes, such as auditing or replicating changes in a target system to maintain synchronization.

chaos engineering

Intentionally introducing failures or disruptive events to test a system's resilience. You can use [AWS Fault Injection Service \(AWS FIS\)](#) to perform experiments that stress your AWS workloads and evaluate their response.

CI/CD

See [continuous integration and continuous delivery](#).

classification

A categorization process that helps generate predictions. ML models for classification problems predict a discrete value. Discrete values are always distinct from one another. For example, a model might need to evaluate whether or not there is a car in an image.

client-side encryption

Encryption of data locally, before the target AWS service receives it.

Cloud Center of Excellence (CCoE)

A multi-disciplinary team that drives cloud adoption efforts across an organization, including developing cloud best practices, mobilizing resources, establishing migration timelines, and leading the organization through large-scale transformations. For more information, see the [CCoE posts](#) on the AWS Cloud Enterprise Strategy Blog.

cloud computing

The cloud technology that is typically used for remote data storage and IoT device management. Cloud computing is commonly connected to [edge computing](#) technology.

cloud operating model

In an IT organization, the operating model that is used to build, mature, and optimize one or more cloud environments. For more information, see [Building your Cloud Operating Model](#).

cloud stages of adoption

The four phases that organizations typically go through when they migrate to the AWS Cloud:

- Project – Running a few cloud-related projects for proof of concept and learning purposes
- Foundation – Making foundational investments to scale your cloud adoption (e.g., creating a landing zone, defining a CCoE, establishing an operations model)
- Migration – Migrating individual applications
- Re-invention – Optimizing products and services, and innovating in the cloud

These stages were defined by Stephen Orban in the blog post [The Journey Toward Cloud-First & the Stages of Adoption](#) on the AWS Cloud Enterprise Strategy blog. For information about how they relate to the AWS migration strategy, see the [migration readiness guide](#).

CMDB

See [configuration management database](#).

code repository

A location where source code and other assets, such as documentation, samples, and scripts, are stored and updated through version control processes. Common cloud repositories include GitHub or Bitbucket Cloud. Each version of the code is called a *branch*. In a microservice structure, each repository is devoted to a single piece of functionality. A single CI/CD pipeline can use multiple repositories.

cold cache

A buffer cache that is empty, not well populated, or contains stale or irrelevant data. This affects performance because the database instance must read from the main memory or disk, which is slower than reading from the buffer cache.

cold data

Data that is rarely accessed and is typically historical. When querying this kind of data, slow queries are typically acceptable. Moving this data to lower-performing and less expensive storage tiers or classes can reduce costs.

computer vision (CV)

A field of [AI](#) that uses machine learning to analyze and extract information from visual formats such as digital images and videos. For example, Amazon SageMaker AI provides image processing algorithms for CV.

configuration drift

For a workload, a configuration change from the expected state. It might cause the workload to become noncompliant, and it's typically gradual and unintentional.

configuration management database (CMDB)

A repository that stores and manages information about a database and its IT environment, including both hardware and software components and their configurations. You typically use data from a CMDB in the portfolio discovery and analysis stage of migration.

conformance pack

A collection of AWS Config rules and remediation actions that you can assemble to customize your compliance and security checks. You can deploy a conformance pack as a single entity in an AWS account and Region, or across an organization, by using a YAML template. For more information, see [Conformance packs](#) in the AWS Config documentation.

continuous integration and continuous delivery (CI/CD)

The process of automating the source, build, test, staging, and production stages of the software release process. CI/CD is commonly described as a pipeline. CI/CD can help you automate processes, improve productivity, improve code quality, and deliver faster. For more information, see [Benefits of continuous delivery](#). CD can also stand for *continuous deployment*. For more information, see [Continuous Delivery vs. Continuous Deployment](#).

CV

See [computer vision](#).

D

data at rest

Data that is stationary in your network, such as data that is in storage.

data classification

A process for identifying and categorizing the data in your network based on its criticality and sensitivity. It is a critical component of any cybersecurity risk management strategy because it helps you determine the appropriate protection and retention controls for the data. Data classification is a component of the security pillar in the AWS Well-Architected Framework. For more information, see [Data classification](#).

data drift

A meaningful variation between the production data and the data that was used to train an ML model, or a meaningful change in the input data over time. Data drift can reduce the overall quality, accuracy, and fairness in ML model predictions.

data in transit

Data that is actively moving through your network, such as between network resources.

data mesh

An architectural framework that provides distributed, decentralized data ownership with centralized management and governance.

data minimization

The principle of collecting and processing only the data that is strictly necessary. Practicing data minimization in the AWS Cloud can reduce privacy risks, costs, and your analytics carbon footprint.

data perimeter

A set of preventive guardrails in your AWS environment that help make sure that only trusted identities are accessing trusted resources from expected networks. For more information, see [Building a data perimeter on AWS](#).

data preprocessing

To transform raw data into a format that is easily parsed by your ML model. Preprocessing data can mean removing certain columns or rows and addressing missing, inconsistent, or duplicate values.

data provenance

The process of tracking the origin and history of data throughout its lifecycle, such as how the data was generated, transmitted, and stored.

data subject

An individual whose data is being collected and processed.

data warehouse

A data management system that supports business intelligence, such as analytics. Data warehouses commonly contain large amounts of historical data, and they are typically used for queries and analysis.

database definition language (DDL)

Statements or commands for creating or modifying the structure of tables and objects in a database.

database manipulation language (DML)

Statements or commands for modifying (inserting, updating, and deleting) information in a database.

DDL

See [database definition language](#).

deep ensemble

To combine multiple deep learning models for prediction. You can use deep ensembles to obtain a more accurate prediction or for estimating uncertainty in predictions.

deep learning

An ML subfield that uses multiple layers of artificial neural networks to identify mapping between input data and target variables of interest.

defense-in-depth

An information security approach in which a series of security mechanisms and controls are thoughtfully layered throughout a computer network to protect the confidentiality, integrity, and availability of the network and the data within. When you adopt this strategy on AWS, you add multiple controls at different layers of the AWS Organizations structure to help secure resources. For example, a defense-in-depth approach might combine multi-factor authentication, network segmentation, and encryption.

delegated administrator

In AWS Organizations, a compatible service can register an AWS member account to administer the organization's accounts and manage permissions for that service. This account is called the *delegated administrator* for that service. For more information and a list of compatible services, see [Services that work with AWS Organizations](#) in the AWS Organizations documentation.

deployment

The process of making an application, new features, or code fixes available in the target environment. Deployment involves implementing changes in a code base and then building and running that code base in the application's environments.

development environment

See [environment](#).

detective control

A security control that is designed to detect, log, and alert after an event has occurred. These controls are a second line of defense, alerting you to security events that bypassed the preventative controls in place. For more information, see [Detective controls](#) in *Implementing security controls on AWS*.

development value stream mapping (DVSM)

A process used to identify and prioritize constraints that adversely affect speed and quality in a software development lifecycle. DVSM extends the value stream mapping process originally designed for lean manufacturing practices. It focuses on the steps and teams required to create and move value through the software development process.

digital twin

A virtual representation of a real-world system, such as a building, factory, industrial equipment, or production line. Digital twins support predictive maintenance, remote monitoring, and production optimization.

dimension table

In a [star schema](#), a smaller table that contains data attributes about quantitative data in a fact table. Dimension table attributes are typically text fields or discrete numbers that behave like text. These attributes are commonly used for query constraining, filtering, and result set labeling.

disaster

An event that prevents a workload or system from fulfilling its business objectives in its primary deployed location. These events can be natural disasters, technical failures, or the result of human actions, such as unintentional misconfiguration or a malware attack.

disaster recovery (DR)

The strategy and process you use to minimize downtime and data loss caused by a [disaster](#). For more information, see [Disaster Recovery of Workloads on AWS: Recovery in the Cloud](#) in the AWS Well-Architected Framework.

DML

See [database manipulation language](#).

domain-driven design

An approach to developing a complex software system by connecting its components to evolving domains, or core business goals, that each component serves. This concept was introduced by Eric Evans in his book, *Domain-Driven Design: Tackling Complexity in the Heart of Software* (Boston: Addison-Wesley Professional, 2003). For information about how you can use domain-driven design with the strangler fig pattern, see [Modernizing legacy Microsoft ASP.NET \(ASMX\) web services incrementally by using containers and Amazon API Gateway](#).

DR

See [disaster recovery](#).

drift detection

Tracking deviations from a baselined configuration. For example, you can use AWS CloudFormation to [detect drift in system resources](#), or you can use AWS Control Tower to [detect changes in your landing zone](#) that might affect compliance with governance requirements.

DVSM

See [development value stream mapping](#).

E

EDA

See [exploratory data analysis](#).

EDI

See [electronic data interchange](#).

edge computing

The technology that increases the computing power for smart devices at the edges of an IoT network. When compared with [cloud computing](#), edge computing can reduce communication latency and improve response time.

electronic data interchange (EDI)

The automated exchange of business documents between organizations. For more information, see [What is Electronic Data Interchange](#).

encryption

A computing process that transforms plaintext data, which is human-readable, into ciphertext.

encryption key

A cryptographic string of randomized bits that is generated by an encryption algorithm. Keys can vary in length, and each key is designed to be unpredictable and unique.

endianness

The order in which bytes are stored in computer memory. Big-endian systems store the most significant byte first. Little-endian systems store the least significant byte first.

endpoint

See [service endpoint](#).

endpoint service

A service that you can host in a virtual private cloud (VPC) to share with other users. You can create an endpoint service with AWS PrivateLink and grant permissions to other AWS accounts or to AWS Identity and Access Management (IAM) principals. These accounts or principals can connect to your endpoint service privately by creating interface VPC endpoints. For more information, see [Create an endpoint service](#) in the Amazon Virtual Private Cloud (Amazon VPC) documentation.

enterprise resource planning (ERP)

A system that automates and manages key business processes (such as accounting, [MES](#), and project management) for an enterprise.

envelope encryption

The process of encrypting an encryption key with another encryption key. For more information, see [Envelope encryption](#) in the AWS Key Management Service (AWS KMS) documentation.

environment

An instance of a running application. The following are common types of environments in cloud computing:

- development environment – An instance of a running application that is available only to the core team responsible for maintaining the application. Development environments are used to test changes before promoting them to upper environments. This type of environment is sometimes referred to as a *test environment*.
- lower environments – All development environments for an application, such as those used for initial builds and tests.
- production environment – An instance of a running application that end users can access. In a CI/CD pipeline, the production environment is the last deployment environment.
- upper environments – All environments that can be accessed by users other than the core development team. This can include a production environment, preproduction environments, and environments for user acceptance testing.

epic

In agile methodologies, functional categories that help organize and prioritize your work. Epics provide a high-level description of requirements and implementation tasks. For example, AWS CAF security epics include identity and access management, detective controls, infrastructure security, data protection, and incident response. For more information about epics in the AWS migration strategy, see the [program implementation guide](#).

ERP

See [enterprise resource planning](#).

exploratory data analysis (EDA)

The process of analyzing a dataset to understand its main characteristics. You collect or aggregate data and then perform initial investigations to find patterns, detect anomalies, and check assumptions. EDA is performed by calculating summary statistics and creating data visualizations.

F

fact table

The central table in a [star schema](#). It stores quantitative data about business operations. Typically, a fact table contains two types of columns: those that contain measures and those that contain a foreign key to a dimension table.

fail fast

A philosophy that uses frequent and incremental testing to reduce the development lifecycle. It is a critical part of an agile approach.

fault isolation boundary

In the AWS Cloud, a boundary such as an Availability Zone, AWS Region, control plane, or data plane that limits the effect of a failure and helps improve the resilience of workloads. For more information, see [AWS Fault Isolation Boundaries](#).

feature branch

See [branch](#).

features

The input data that you use to make a prediction. For example, in a manufacturing context, features could be images that are periodically captured from the manufacturing line.

feature importance

How significant a feature is for a model's predictions. This is usually expressed as a numerical score that can be calculated through various techniques, such as Shapley Additive Explanations (SHAP) and integrated gradients. For more information, see [Machine learning model interpretability with AWS](#).

feature transformation

To optimize data for the ML process, including enriching data with additional sources, scaling values, or extracting multiple sets of information from a single data field. This enables the ML model to benefit from the data. For example, if you break down the "2021-05-27 00:15:37" date into "2021", "May", "Thu", and "15", you can help the learning algorithm learn nuanced patterns associated with different data components.

few-shot prompting

Providing an [LLM](#) with a small number of examples that demonstrate the task and desired output before asking it to perform a similar task. This technique is an application of in-context learning, where models learn from examples (*shots*) that are embedded in prompts. Few-shot prompting can be effective for tasks that require specific formatting, reasoning, or domain knowledge. See also [zero-shot prompting](#).

FGAC

See [fine-grained access control](#).

fine-grained access control (FGAC)

The use of multiple conditions to allow or deny an access request.

flash-cut migration

A database migration method that uses continuous data replication through [change data capture](#) to migrate data in the shortest time possible, instead of using a phased approach. The objective is to keep downtime to a minimum.

FM

See [foundation model](#).

foundation model (FM)

A large deep-learning neural network that has been training on massive datasets of generalized and unlabeled data. FMs are capable of performing a wide variety of general tasks, such as understanding language, generating text and images, and conversing in natural language. For more information, see [What are Foundation Models](#).

G

generative AI

A subset of [AI](#) models that have been trained on large amounts of data and that can use a simple text prompt to create new content and artifacts, such as images, videos, text, and audio. For more information, see [What is Generative AI](#).

geo blocking

See [geographic restrictions](#).

geographic restrictions (geo blocking)

In Amazon CloudFront, an option to prevent users in specific countries from accessing content distributions. You can use an allow list or block list to specify approved and banned countries. For more information, see [Restricting the geographic distribution of your content](#) in the CloudFront documentation.

Gitflow workflow

An approach in which lower and upper environments use different branches in a source code repository. The Gitflow workflow is considered legacy, and the [trunk-based workflow](#) is the modern, preferred approach.

golden image

A snapshot of a system or software that is used as a template to deploy new instances of that system or software. For example, in manufacturing, a golden image can be used to provision software on multiple devices and helps improve speed, scalability, and productivity in device manufacturing operations.

greenfield strategy

The absence of existing infrastructure in a new environment. When adopting a greenfield strategy for a system architecture, you can select all new technologies without the restriction

of compatibility with existing infrastructure, also known as [brownfield](#). If you are expanding the existing infrastructure, you might blend brownfield and greenfield strategies.

guardrail

A high-level rule that helps govern resources, policies, and compliance across organizational units (OUs). *Preventive guardrails* enforce policies to ensure alignment to compliance standards. They are implemented by using service control policies and IAM permissions boundaries. *Detective guardrails* detect policy violations and compliance issues, and generate alerts for remediation. They are implemented by using AWS Config, AWS Security Hub CSPM, Amazon GuardDuty, AWS Trusted Advisor, Amazon Inspector, and custom AWS Lambda checks.

H

HA

See [high availability](#).

heterogeneous database migration

Migrating your source database to a target database that uses a different database engine (for example, Oracle to Amazon Aurora). Heterogeneous migration is typically part of a re-architecting effort, and converting the schema can be a complex task. [AWS provides AWS SCT](#) that helps with schema conversions.

high availability (HA)

The ability of a workload to operate continuously, without intervention, in the event of challenges or disasters. HA systems are designed to automatically fail over, consistently deliver high-quality performance, and handle different loads and failures with minimal performance impact.

historian modernization

An approach used to modernize and upgrade operational technology (OT) systems to better serve the needs of the manufacturing industry. A *historian* is a type of database that is used to collect and store data from various sources in a factory.

holdout data

A portion of historical, labeled data that is withheld from a dataset that is used to train a [machine learning](#) model. You can use holdout data to evaluate the model performance by comparing the model predictions against the holdout data.

homogeneous database migration

Migrating your source database to a target database that shares the same database engine (for example, Microsoft SQL Server to Amazon RDS for SQL Server). Homogeneous migration is typically part of a rehosting or replatforming effort. You can use native database utilities to migrate the schema.

hot data

Data that is frequently accessed, such as real-time data or recent translational data. This data typically requires a high-performance storage tier or class to provide fast query responses.

hotfix

An urgent fix for a critical issue in a production environment. Due to its urgency, a hotfix is usually made outside of the typical DevOps release workflow.

hypercure period

Immediately following cutover, the period of time when a migration team manages and monitors the migrated applications in the cloud in order to address any issues. Typically, this period is 1–4 days in length. At the end of the hypercure period, the migration team typically transfers responsibility for the applications to the cloud operations team.

I

IaC

See [infrastructure as code](#).

identity-based policy

A policy attached to one or more IAM principals that defines their permissions within the AWS Cloud environment.

idle application

An application that has an average CPU and memory usage between 5 and 20 percent over a period of 90 days. In a migration project, it is common to retire these applications or retain them on premises.

IIoT

See [Industrial Internet of Things](#).

immutable infrastructure

A model that deploys new infrastructure for production workloads instead of updating, patching, or modifying the existing infrastructure. Immutable infrastructures are inherently more consistent, reliable, and predictable than [mutable infrastructure](#). For more information, see the [Deploy using immutable infrastructure](#) best practice in the AWS Well-Architected Framework.

inbound (ingress) VPC

In an AWS multi-account architecture, a VPC that accepts, inspects, and routes network connections from outside an application. The [AWS Security Reference Architecture](#) recommends setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

incremental migration

A cutover strategy in which you migrate your application in small parts instead of performing a single, full cutover. For example, you might move only a few microservices or users to the new system initially. After you verify that everything is working properly, you can incrementally move additional microservices or users until you can decommission your legacy system. This strategy reduces the risks associated with large migrations.

Industry 4.0

A term that was introduced by [Klaus Schwab](#) in 2016 to refer to the modernization of manufacturing processes through advances in connectivity, real-time data, automation, analytics, and AI/ML.

infrastructure

All of the resources and assets contained within an application's environment.

infrastructure as code (IaC)

The process of provisioning and managing an application's infrastructure through a set of configuration files. IaC is designed to help you centralize infrastructure management, standardize resources, and scale quickly so that new environments are repeatable, reliable, and consistent.

industrial Internet of Things (IIoT)

The use of internet-connected sensors and devices in the industrial sectors, such as manufacturing, energy, automotive, healthcare, life sciences, and agriculture. For more information, see [Building an industrial Internet of Things \(IIoT\) digital transformation strategy](#).

inspection VPC

In an AWS multi-account architecture, a centralized VPC that manages inspections of network traffic between VPCs (in the same or different AWS Regions), the internet, and on-premises networks. The [AWS Security Reference Architecture](#) recommends setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

Internet of Things (IoT)

The network of connected physical objects with embedded sensors or processors that communicate with other devices and systems through the internet or over a local communication network. For more information, see [What is IoT?](#)

interpretability

A characteristic of a machine learning model that describes the degree to which a human can understand how the model's predictions depend on its inputs. For more information, see [Machine learning model interpretability with AWS](#).

IoT

See [Internet of Things](#).

IT information library (ITIL)

A set of best practices for delivering IT services and aligning these services with business requirements. ITIL provides the foundation for ITSM.

IT service management (ITSM)

Activities associated with designing, implementing, managing, and supporting IT services for an organization. For information about integrating cloud operations with ITSM tools, see the [operations integration guide](#).

ITIL

See [IT information library](#).

ITSM

See [IT service management](#).

L

label-based access control (LBAC)

An implementation of mandatory access control (MAC) where the users and the data itself are each explicitly assigned a security label value. The intersection between the user security label and data security label determines which rows and columns can be seen by the user.

landing zone

A landing zone is a well-architected, multi-account AWS environment that is scalable and secure. This is a starting point from which your organizations can quickly launch and deploy workloads and applications with confidence in their security and infrastructure environment. For more information about landing zones, see [Setting up a secure and scalable multi-account AWS environment](#).

large language model (LLM)

A deep learning [AI](#) model that is pretrained on a vast amount of data. An LLM can perform multiple tasks, such as answering questions, summarizing documents, translating text into other languages, and completing sentences. For more information, see [What are LLMs](#).

large migration

A migration of 300 or more servers.

LBAC

See [label-based access control](#).

least privilege

The security best practice of granting the minimum permissions required to perform a task. For more information, see [Apply least-privilege permissions](#) in the IAM documentation.

lift and shift

See [7 Rs](#).

little-endian system

A system that stores the least significant byte first. See also [endianness](#).

LLM

See [large language model](#).

lower environments

See [environment](#).

M

machine learning (ML)

A type of artificial intelligence that uses algorithms and techniques for pattern recognition and learning. ML analyzes and learns from recorded data, such as Internet of Things (IoT) data, to generate a statistical model based on patterns. For more information, see [Machine Learning](#).

main branch

See [branch](#).

malware

Software that is designed to compromise computer security or privacy. Malware might disrupt computer systems, leak sensitive information, or gain unauthorized access. Examples of malware include viruses, worms, ransomware, Trojan horses, spyware, and keyloggers.

managed services

AWS services for which AWS operates the infrastructure layer, the operating system, and platforms, and you access the endpoints to store and retrieve data. Amazon Simple Storage Service (Amazon S3) and Amazon DynamoDB are examples of managed services. These are also known as *abstracted services*.

manufacturing execution system (MES)

A software system for tracking, monitoring, documenting, and controlling production processes that convert raw materials to finished products on the shop floor.

MAP

See [Migration Acceleration Program](#).

mechanism

A complete process in which you create a tool, drive adoption of the tool, and then inspect the results in order to make adjustments. A mechanism is a cycle that reinforces and improves itself as it operates. For more information, see [Building mechanisms](#) in the AWS Well-Architected Framework.

member account

All AWS accounts other than the management account that are part of an organization in AWS Organizations. An account can be a member of only one organization at a time.

MES

See [manufacturing execution system](#).

Message Queuing Telemetry Transport (MQTT)

A lightweight, machine-to-machine (M2M) communication protocol, based on the [publish/subscribe](#) pattern, for resource-constrained [IoT](#) devices.

microservice

A small, independent service that communicates over well-defined APIs and is typically owned by small, self-contained teams. For example, an insurance system might include microservices that map to business capabilities, such as sales or marketing, or subdomains, such as purchasing, claims, or analytics. The benefits of microservices include agility, flexible scaling, easy deployment, reusable code, and resilience. For more information, see [Integrating microservices by using AWS serverless services](#).

microservices architecture

An approach to building an application with independent components that run each application process as a microservice. These microservices communicate through a well-defined interface by using lightweight APIs. Each microservice in this architecture can be updated, deployed,

and scaled to meet demand for specific functions of an application. For more information, see [Implementing microservices on AWS](#).

Migration Acceleration Program (MAP)

An AWS program that provides consulting support, training, and services to help organizations build a strong operational foundation for moving to the cloud, and to help offset the initial cost of migrations. MAP includes a migration methodology for executing legacy migrations in a methodical way and a set of tools to automate and accelerate common migration scenarios.

migration at scale

The process of moving the majority of the application portfolio to the cloud in waves, with more applications moved at a faster rate in each wave. This phase uses the best practices and lessons learned from the earlier phases to implement a *migration factory* of teams, tools, and processes to streamline the migration of workloads through automation and agile delivery. This is the third phase of the [AWS migration strategy](#).

migration factory

Cross-functional teams that streamline the migration of workloads through automated, agile approaches. Migration factory teams typically include operations, business analysts and owners, migration engineers, developers, and DevOps professionals working in sprints. Between 20 and 50 percent of an enterprise application portfolio consists of repeated patterns that can be optimized by a factory approach. For more information, see the [discussion of migration factories](#) and the [Cloud Migration Factory guide](#) in this content set.

migration metadata

The information about the application and server that is needed to complete the migration. Each migration pattern requires a different set of migration metadata. Examples of migration metadata include the target subnet, security group, and AWS account.

migration pattern

A repeatable migration task that details the migration strategy, the migration destination, and the migration application or service used. Example: Rehost migration to Amazon EC2 with AWS Application Migration Service.

Migration Portfolio Assessment (MPA)

An online tool that provides information for validating the business case for migrating to the AWS Cloud. MPA provides detailed portfolio assessment (server right-sizing, pricing, TCO

comparisons, migration cost analysis) as well as migration planning (application data analysis and data collection, application grouping, migration prioritization, and wave planning). The [MPA tool](#) (requires login) is available free of charge to all AWS consultants and APN Partner consultants.

Migration Readiness Assessment (MRA)

The process of gaining insights about an organization's cloud readiness status, identifying strengths and weaknesses, and building an action plan to close identified gaps, using the AWS CAF. For more information, see the [migration readiness guide](#). MRA is the first phase of the [AWS migration strategy](#).

migration strategy

The approach used to migrate a workload to the AWS Cloud. For more information, see the [7 Rs](#) entry in this glossary and see [Mobilize your organization to accelerate large-scale migrations](#).

ML

See [machine learning](#).

modernization

Transforming an outdated (legacy or monolithic) application and its infrastructure into an agile, elastic, and highly available system in the cloud to reduce costs, gain efficiencies, and take advantage of innovations. For more information, see [Strategy for modernizing applications in the AWS Cloud](#).

modernization readiness assessment

An evaluation that helps determine the modernization readiness of an organization's applications; identifies benefits, risks, and dependencies; and determines how well the organization can support the future state of those applications. The outcome of the assessment is a blueprint of the target architecture, a roadmap that details development phases and milestones for the modernization process, and an action plan for addressing identified gaps. For more information, see [Evaluating modernization readiness for applications in the AWS Cloud](#).

monolithic applications (monoliths)

Applications that run as a single service with tightly coupled processes. Monolithic applications have several drawbacks. If one application feature experiences a spike in demand, the entire architecture must be scaled. Adding or improving a monolithic application's features also becomes more complex when the code base grows. To address these issues, you can

use a microservices architecture. For more information, see [Decomposing monoliths into microservices](#).

MPA

See [Migration Portfolio Assessment](#).

MQTT

See [Message Queuing Telemetry Transport](#).

multiclass classification

A process that helps generate predictions for multiple classes (predicting one of more than two outcomes). For example, an ML model might ask "Is this product a book, car, or phone?" or "Which product category is most interesting to this customer?"

mutable infrastructure

A model that updates and modifies the existing infrastructure for production workloads. For improved consistency, reliability, and predictability, the AWS Well-Architected Framework recommends the use of [immutable infrastructure](#) as a best practice.

O

OAC

See [origin access control](#).

OAI

See [origin access identity](#).

OCM

See [organizational change management](#).

offline migration

A migration method in which the source workload is taken down during the migration process. This method involves extended downtime and is typically used for small, non-critical workloads.

OI

See [operations integration](#).

OLA

See [operational-level agreement](#).

online migration

A migration method in which the source workload is copied to the target system without being taken offline. Applications that are connected to the workload can continue to function during the migration. This method involves zero to minimal downtime and is typically used for critical production workloads.

OPC-UA

See [Open Process Communications - Unified Architecture](#).

Open Process Communications - Unified Architecture (OPC-UA)

A machine-to-machine (M2M) communication protocol for industrial automation. OPC-UA provides an interoperability standard with data encryption, authentication, and authorization schemes.

operational-level agreement (OLA)

An agreement that clarifies what functional IT groups promise to deliver to each other, to support a service-level agreement (SLA).

operational readiness review (ORR)

A checklist of questions and associated best practices that help you understand, evaluate, prevent, or reduce the scope of incidents and possible failures. For more information, see [Operational Readiness Reviews \(ORR\)](#) in the AWS Well-Architected Framework.

operational technology (OT)

Hardware and software systems that work with the physical environment to control industrial operations, equipment, and infrastructure. In manufacturing, the integration of OT and information technology (IT) systems is a key focus for [Industry 4.0](#) transformations.

operations integration (OI)

The process of modernizing operations in the cloud, which involves readiness planning, automation, and integration. For more information, see the [operations integration guide](#).

organization trail

A trail that's created by AWS CloudTrail that logs all events for all AWS accounts in an organization in AWS Organizations. This trail is created in each AWS account that's part of the

organization and tracks the activity in each account. For more information, see [Creating a trail for an organization](#) in the CloudTrail documentation.

organizational change management (OCM)

A framework for managing major, disruptive business transformations from a people, culture, and leadership perspective. OCM helps organizations prepare for, and transition to, new systems and strategies by accelerating change adoption, addressing transitional issues, and driving cultural and organizational changes. In the AWS migration strategy, this framework is called *people acceleration*, because of the speed of change required in cloud adoption projects. For more information, see the [OCM guide](#).

origin access control (OAC)

In CloudFront, an enhanced option for restricting access to secure your Amazon Simple Storage Service (Amazon S3) content. OAC supports all S3 buckets in all AWS Regions, server-side encryption with AWS KMS (SSE-KMS), and dynamic PUT and DELETE requests to the S3 bucket.

origin access identity (OAI)

In CloudFront, an option for restricting access to secure your Amazon S3 content. When you use OAI, CloudFront creates a principal that Amazon S3 can authenticate with. Authenticated principals can access content in an S3 bucket only through a specific CloudFront distribution. See also [OAC](#), which provides more granular and enhanced access control.

ORR

See [operational readiness review](#).

OT

See [operational technology](#).

outbound (egress) VPC

In an AWS multi-account architecture, a VPC that handles network connections that are initiated from within an application. The [AWS Security Reference Architecture](#) recommends setting up your Network account with inbound, outbound, and inspection VPCs to protect the two-way interface between your application and the broader internet.

P

permissions boundary

An IAM management policy that is attached to IAM principals to set the maximum permissions that the user or role can have. For more information, see [Permissions boundaries](#) in the IAM documentation.

personally identifiable information (PII)

Information that, when viewed directly or paired with other related data, can be used to reasonably infer the identity of an individual. Examples of PII include names, addresses, and contact information.

PII

See [personally identifiable information](#).

playbook

A set of predefined steps that capture the work associated with migrations, such as delivering core operations functions in the cloud. A playbook can take the form of scripts, automated runbooks, or a summary of processes or steps required to operate your modernized environment.

PLC

See [programmable logic controller](#).

PLM

See [product lifecycle management](#).

policy

An object that can define permissions (see [identity-based policy](#)), specify access conditions (see [resource-based policy](#)), or define the maximum permissions for all accounts in an organization in AWS Organizations (see [service control policy](#)).

polyglot persistence

Independently choosing a microservice's data storage technology based on data access patterns and other requirements. If your microservices have the same data storage technology, they can encounter implementation challenges or experience poor performance. Microservices are more

easily implemented and achieve better performance and scalability if they use the data store best adapted to their requirements.

portfolio assessment

A process of discovering, analyzing, and prioritizing the application portfolio in order to plan the migration. For more information, see [Evaluating migration readiness](#).

predicate

A query condition that returns `true` or `false`, commonly located in a `WHERE` clause.

predicate pushdown

A database query optimization technique that filters the data in the query before transfer. This reduces the amount of data that must be retrieved and processed from the relational database, and it improves query performance.

preventative control

A security control that is designed to prevent an event from occurring. These controls are a first line of defense to help prevent unauthorized access or unwanted changes to your network. For more information, see [Preventative controls](#) in *Implementing security controls on AWS*.

principal

An entity in AWS that can perform actions and access resources. This entity is typically a root user for an AWS account, an IAM role, or a user. For more information, see *Principal* in [Roles terms and concepts](#) in the IAM documentation.

privacy by design

A system engineering approach that takes privacy into account through the whole development process.

private hosted zones

A container that holds information about how you want Amazon Route 53 to respond to DNS queries for a domain and its subdomains within one or more VPCs. For more information, see [Working with private hosted zones](#) in the Route 53 documentation.

proactive control

A [security control](#) designed to prevent the deployment of noncompliant resources. These controls scan resources before they are provisioned. If the resource is not compliant with the control, then it isn't provisioned. For more information, see the [Controls reference guide](#) in the

AWS Control Tower documentation and see [Proactive controls](#) in *Implementing security controls on AWS*.

product lifecycle management (PLM)

The management of data and processes for a product throughout its entire lifecycle, from design, development, and launch, through growth and maturity, to decline and removal.

production environment

See [environment](#).

programmable logic controller (PLC)

In manufacturing, a highly reliable, adaptable computer that monitors machines and automates manufacturing processes.

prompt chaining

Using the output of one [LLM](#) prompt as the input for the next prompt to generate better responses. This technique is used to break down a complex task into subtasks, or to iteratively refine or expand a preliminary response. It helps improve the accuracy and relevance of a model's responses and allows for more granular, personalized results.

pseudonymization

The process of replacing personal identifiers in a dataset with placeholder values. Pseudonymization can help protect personal privacy. Pseudonymized data is still considered to be personal data.

publish/subscribe (pub/sub)

A pattern that enables asynchronous communications among microservices to improve scalability and responsiveness. For example, in a microservices-based [MES](#), a microservice can publish event messages to a channel that other microservices can subscribe to. The system can add new microservices without changing the publishing service.

Q

query plan

A series of steps, like instructions, that are used to access the data in a SQL relational database system.

query plan regression

When a database service optimizer chooses a less optimal plan than it did before a given change to the database environment. This can be caused by changes to statistics, constraints, environment settings, query parameter bindings, and updates to the database engine.

R

RACI matrix

See [responsible, accountable, consulted, informed \(RACI\)](#).

RAG

See [Retrieval Augmented Generation](#).

ransomware

A malicious software that is designed to block access to a computer system or data until a payment is made.

RASCI matrix

See [responsible, accountable, consulted, informed \(RACI\)](#).

RCAC

See [row and column access control](#).

read replica

A copy of a database that's used for read-only purposes. You can route queries to the read replica to reduce the load on your primary database.

re-architect

See [7 Rs](#).

recovery point objective (RPO)

The maximum acceptable amount of time since the last data recovery point. This determines what is considered an acceptable loss of data between the last recovery point and the interruption of service.

recovery time objective (RTO)

The maximum acceptable delay between the interruption of service and restoration of service.

refactor

See [7 Rs](#).

Region

A collection of AWS resources in a geographic area. Each AWS Region is isolated and independent of the others to provide fault tolerance, stability, and resilience. For more information, see [Specify which AWS Regions your account can use](#).

regression

An ML technique that predicts a numeric value. For example, to solve the problem of "What price will this house sell for?" an ML model could use a linear regression model to predict a house's sale price based on known facts about the house (for example, the square footage).

rehost

See [7 Rs](#).

release

In a deployment process, the act of promoting changes to a production environment.

relocate

See [7 Rs](#).

replatform

See [7 Rs](#).

repurchase

See [7 Rs](#).

resiliency

An application's ability to resist or recover from disruptions. [High availability](#) and [disaster recovery](#) are common considerations when planning for resiliency in the AWS Cloud. For more information, see [AWS Cloud Resilience](#).

resource-based policy

A policy attached to a resource, such as an Amazon S3 bucket, an endpoint, or an encryption key. This type of policy specifies which principals are allowed access, supported actions, and any other conditions that must be met.

responsible, accountable, consulted, informed (RACI) matrix

A matrix that defines the roles and responsibilities for all parties involved in migration activities and cloud operations. The matrix name is derived from the responsibility types defined in the matrix: responsible (R), accountable (A), consulted (C), and informed (I). The support (S) type is optional. If you include support, the matrix is called a *RASCI matrix*, and if you exclude it, it's called a *RACI matrix*.

responsive control

A security control that is designed to drive remediation of adverse events or deviations from your security baseline. For more information, see [Responsive controls](#) in *Implementing security controls on AWS*.

retain

See [7 Rs](#).

retire

See [7 Rs](#).

Retrieval Augmented Generation (RAG)

A [generative AI](#) technology in which an [LLM](#) references an authoritative data source that is outside of its training data sources before generating a response. For example, a RAG model might perform a semantic search of an organization's knowledge base or custom data. For more information, see [What is RAG](#).

rotation

The process of periodically updating a [secret](#) to make it more difficult for an attacker to access the credentials.

row and column access control (RCAC)

The use of basic, flexible SQL expressions that have defined access rules. RCAC consists of row permissions and column masks.

RPO

See [recovery point objective](#).

RTO

See [recovery time objective](#).

runbook

A set of manual or automated procedures required to perform a specific task. These are typically built to streamline repetitive operations or procedures with high error rates.

S

SAML 2.0

An open standard that many identity providers (IdPs) use. This feature enables federated single sign-on (SSO), so users can log into the AWS Management Console or call the AWS API operations without you having to create user in IAM for everyone in your organization. For more information about SAML 2.0-based federation, see [About SAML 2.0-based federation](#) in the IAM documentation.

SCADA

See [supervisory control and data acquisition](#).

SCP

See [service control policy](#).

secret

In AWS Secrets Manager, confidential or restricted information, such as a password or user credentials, that you store in encrypted form. It consists of the secret value and its metadata. The secret value can be binary, a single string, or multiple strings. For more information, see [What's in a Secrets Manager secret?](#) in the Secrets Manager documentation.

security by design

A system engineering approach that takes security into account through the whole development process.

security control

A technical or administrative guardrail that prevents, detects, or reduces the ability of a threat actor to exploit a security vulnerability. There are four primary types of security controls: [preventative](#), [detective](#), [responsive](#), and [proactive](#).

security hardening

The process of reducing the attack surface to make it more resistant to attacks. This can include actions such as removing resources that are no longer needed, implementing the security best practice of granting least privilege, or deactivating unnecessary features in configuration files.

security information and event management (SIEM) system

Tools and services that combine security information management (SIM) and security event management (SEM) systems. A SIEM system collects, monitors, and analyzes data from servers, networks, devices, and other sources to detect threats and security breaches, and to generate alerts.

security response automation

A predefined and programmed action that is designed to automatically respond to or remediate a security event. These automations serve as [detective](#) or [responsive](#) security controls that help you implement AWS security best practices. Examples of automated response actions include modifying a VPC security group, patching an Amazon EC2 instance, or rotating credentials.

server-side encryption

Encryption of data at its destination, by the AWS service that receives it.

service control policy (SCP)

A policy that provides centralized control over permissions for all accounts in an organization in AWS Organizations. SCPs define guardrails or set limits on actions that an administrator can delegate to users or roles. You can use SCPs as allow lists or deny lists, to specify which services or actions are permitted or prohibited. For more information, see [Service control policies](#) in the AWS Organizations documentation.

service endpoint

The URL of the entry point for an AWS service. You can use the endpoint to connect programmatically to the target service. For more information, see [AWS service endpoints](#) in *AWS General Reference*.

service-level agreement (SLA)

An agreement that clarifies what an IT team promises to deliver to their customers, such as service uptime and performance.

service-level indicator (SLI)

A measurement of a performance aspect of a service, such as its error rate, availability, or throughput.

service-level objective (SLO)

A target metric that represents the health of a service, as measured by a [service-level indicator](#).

shared responsibility model

A model describing the responsibility you share with AWS for cloud security and compliance. AWS is responsible for security *of* the cloud, whereas you are responsible for security *in* the cloud. For more information, see [Shared responsibility model](#).

SIEM

See [security information and event management system](#).

single point of failure (SPOF)

A failure in a single, critical component of an application that can disrupt the system.

SLA

See [service-level agreement](#).

SLI

See [service-level indicator](#).

SLO

See [service-level objective](#).

split-and-seed model

A pattern for scaling and accelerating modernization projects. As new features and product releases are defined, the core team splits up to create new product teams. This helps scale your organization's capabilities and services, improves developer productivity, and supports rapid

innovation. For more information, see [Phased approach to modernizing applications in the AWS Cloud](#).

SPOF

See [single point of failure](#).

star schema

A database organizational structure that uses one large fact table to store transactional or measured data and uses one or more smaller dimensional tables to store data attributes. This structure is designed for use in a [data warehouse](#) or for business intelligence purposes.

strangler fig pattern

An approach to modernizing monolithic systems by incrementally rewriting and replacing system functionality until the legacy system can be decommissioned. This pattern uses the analogy of a fig vine that grows into an established tree and eventually overcomes and replaces its host. The pattern was [introduced by Martin Fowler](#) as a way to manage risk when rewriting monolithic systems. For an example of how to apply this pattern, see [Modernizing legacy Microsoft ASP.NET \(ASMX\) web services incrementally by using containers and Amazon API Gateway](#).

subnet

A range of IP addresses in your VPC. A subnet must reside in a single Availability Zone.

supervisory control and data acquisition (SCADA)

In manufacturing, a system that uses hardware and software to monitor physical assets and production operations.

symmetric encryption

An encryption algorithm that uses the same key to encrypt and decrypt the data.

synthetic testing

Testing a system in a way that simulates user interactions to detect potential issues or to monitor performance. You can use [Amazon CloudWatch Synthetics](#) to create these tests.

system prompt

A technique for providing context, instructions, or guidelines to an [LLM](#) to direct its behavior. System prompts help set context and establish rules for interactions with users.

T

tags

Key-value pairs that act as metadata for organizing your AWS resources. Tags can help you manage, identify, organize, search for, and filter resources. For more information, see [Tagging your AWS resources](#).

target variable

The value that you are trying to predict in supervised ML. This is also referred to as an *outcome variable*. For example, in a manufacturing setting the target variable could be a product defect.

task list

A tool that is used to track progress through a runbook. A task list contains an overview of the runbook and a list of general tasks to be completed. For each general task, it includes the estimated amount of time required, the owner, and the progress.

test environment

See [environment](#).

training

To provide data for your ML model to learn from. The training data must contain the correct answer. The learning algorithm finds patterns in the training data that map the input data attributes to the target (the answer that you want to predict). It outputs an ML model that captures these patterns. You can then use the ML model to make predictions on new data for which you don't know the target.

transit gateway

A network transit hub that you can use to interconnect your VPCs and on-premises networks. For more information, see [What is a transit gateway](#) in the AWS Transit Gateway documentation.

trunk-based workflow

An approach in which developers build and test features locally in a feature branch and then merge those changes into the main branch. The main branch is then built to the development, preproduction, and production environments, sequentially.

trusted access

Granting permissions to a service that you specify to perform tasks in your organization in AWS Organizations and in its accounts on your behalf. The trusted service creates a service-linked role in each account, when that role is needed, to perform management tasks for you. For more information, see [Using AWS Organizations with other AWS services](#) in the AWS Organizations documentation.

tuning

To change aspects of your training process to improve the ML model's accuracy. For example, you can train the ML model by generating a labeling set, adding labels, and then repeating these steps several times under different settings to optimize the model.

two-pizza team

A small DevOps team that you can feed with two pizzas. A two-pizza team size ensures the best possible opportunity for collaboration in software development.

U

uncertainty

A concept that refers to imprecise, incomplete, or unknown information that can undermine the reliability of predictive ML models. There are two types of uncertainty: *Epistemic uncertainty* is caused by limited, incomplete data, whereas *aleatoric uncertainty* is caused by the noise and randomness inherent in the data.

undifferentiated tasks

Also known as *heavy lifting*, work that is necessary to create and operate an application but that doesn't provide direct value to the end user or provide competitive advantage. Examples of undifferentiated tasks include procurement, maintenance, and capacity planning.

upper environments

See [environment](#).

V

vacuuming

A database maintenance operation that involves cleaning up after incremental updates to reclaim storage and improve performance.

version control

Processes and tools that track changes, such as changes to source code in a repository.

VPC peering

A connection between two VPCs that allows you to route traffic by using private IP addresses. For more information, see [What is VPC peering](#) in the Amazon VPC documentation.

vulnerability

A software or hardware flaw that compromises the security of the system.

W

warm cache

A buffer cache that contains current, relevant data that is frequently accessed. The database instance can read from the buffer cache, which is faster than reading from the main memory or disk.

warm data

Data that is infrequently accessed. When querying this kind of data, moderately slow queries are typically acceptable.

window function

A SQL function that performs a calculation on a group of rows that relate in some way to the current record. Window functions are useful for processing tasks, such as calculating a moving average or accessing the value of rows based on the relative position of the current row.

workload

A collection of resources and code that delivers business value, such as a customer-facing application or backend process.

workstream

Functional groups in a migration project that are responsible for a specific set of tasks. Each workstream is independent but supports the other workstreams in the project. For example, the portfolio workstream is responsible for prioritizing applications, wave planning, and collecting migration metadata. The portfolio workstream delivers these assets to the migration workstream, which then migrates the servers and applications.

WORM

See [write once, read many](#).

WQF

See [AWS Workload Qualification Framework](#).

write once, read many (WORM)

A storage model that writes data a single time and prevents the data from being deleted or modified. Authorized users can read the data as many times as needed, but they cannot change it. This data storage infrastructure is considered [immutable](#).

Z

zero-day exploit

An attack, typically malware, that takes advantage of a [zero-day vulnerability](#).

zero-day vulnerability

An unmitigated flaw or vulnerability in a production system. Threat actors can use this type of vulnerability to attack the system. Developers frequently become aware of the vulnerability as a result of the attack.

zero-shot prompting

Providing an [LLM](#) with instructions for performing a task but no examples (*shots*) that can help guide it. The LLM must use its pre-trained knowledge to handle the task. The effectiveness of zero-shot prompting depends on the complexity of the task and the quality of the prompt. See also [few-shot prompting](#).

zombie application

An application that has an average CPU and memory usage below 5 percent. In a migration project, it is common to retire these applications.