GET STARTED

The Enterprise Guide to Multi-Cloud

vmware

Welcome to the World of Multi-Cloud Computing

Across the globe, major social and economic forces are accelerating enterprise transformation. The need for business continuity, the constant demand for remote services, and the ongoing mission to improve customer experience are key drivers of change. It's no surprise that at the heart of this transformation is the cloud.

Organizations of all industries and sizes are increasing cloud adoption to meet the demands of today and tomorrow. Whether to optimize what it takes to manage an on-premises data center, enable development teams to build modern and agile applications, or to leverage best-of-breed functionalities, the cloud provides the speed, flexibility, and scalability that organizations require to succeed in our rapidly changing world.

However, the breadth of offerings across the cloud presents a myriad of choices to the business. Teams struggle to balance the needs of new applications with existing infrastructure and legacy applications, and the pressures of reliability, security, and governance are compounded with a growing portfolio of inconsistent architectures, infrastructure requirements, cloud providers, services, tools, and processes.

Ultimately, the challenge organizations face is not whether to adopt cloud but how to map their requirements to the most suitable environment for both current and future needs. Therefore, nearly every business is on a multi-cloud journey. A future-ready, multi-cloud operating model, enabled by a multi-cloud platform, is the most efficient and cost-effective pathway to modernize with minimum risk. You can unlock the power of cloud to rapidly migrate apps, scale resources up or down based on demand, deliver resources for distributed work initiatives, and drive app modernization strategies.

This *Enterprise Guide to Multi-Cloud* is designed to help organizations understand the value you can expect from an effective multi-cloud strategy. We dive into the primary multi-cloud use cases and service models, and the most common challenges that can stand in the way of success. From there, we outline the ideal approach to multi-cloud and what is needed for organizations to achieve a successful multi-cloud operating model that hinges on choice, speed, and control.

What Is Multi-Cloud?

The definition of multi-cloud is a difficult one and leads to debates about "what is" and "what is not" multi-cloud. Is hybrid multi-cloud? Is consuming multiple public clouds multi-cloud? What about edge? Is it multi-cloud without a private cloud?

All these questions are relevant when trying to understand multi-cloud. So before we dive into the advantages or challenges of a multi-cloud strategy, let's first come to an understanding of what it means.

A brief history of cloud computing

To start, let's go back to the inception of cloud. Today, many people associate the beginning of cloud with Amazon Web Services (AWS), but in fact, its foundation goes back several decades.

In the 1960s and 70s, computers were huge, expensive and often sat idle for long periods of time. Businesses realized they could sell excess capacity to make a profit. For instance, Dartmouth Time-Sharing System (DTSS), originally an on-campus network, connected students to a GE-635 mainframe computer via standard telephone lines using acoustic couplers. The $16 per month student fee paid for the computer's capacity to be expanded. In 1977, DTSS was extended beyond the campus to colleges too small to have their own computers.

While DTSS was developing and expanding, so were other seeds of cloud computing. The development of packet switching technology, internetworking, and the internet protocol suite (TCP/IP) resulted in the launch of the World Wide Web in the 1980s, which—having originally been intended for government and government body use only—was subsequently made available for commercial use, and soon after, the first internet service providers (ISPs) launched. It was also around this time that virtualization technology started to gain momentum.

Beginning as a way to divide system resources between different applications, virtualization has since evolved and serves as the foundation for cloud computing as we know it today. Virtualization enables organizations to create multiple simulated environments or dedicated resources from a single, physical hardware system.

Organizations leverage virtualization technology on their own on-premises data centers to create what we refer to as a private cloud. With a private cloud, an organization is essentially transforming its traditional infrastructure to become an on-demand, self-service platform for internal users to easily request and access IT resources.

This strategy became the business model for what we know today as public cloud. Amazon, who had acknowledged the need for an internal infrastructure to support the e-commerce platforms it was building for other retailers, launched its Elastic Compute Cloud (EC2) service in 2006 as part of its AWS business segment. With the introduction of EC2, customers could take advantage of on-demand compute services to run their applications without hosting or managing the servers themselves. Instead, AWS owns and handles the infrastructure, and services are made available to customers via the internet.

Shortly thereafter, other tech giants followed suit with their own public cloud offerings, such as Microsoft (Azure), Google (Google Cloud Platform), IBM, Oracle, Alibaba, and more. Today, the complexity and sheer number of cloud computing services that customers have access to is immense, and the competition among cloud service providers is more intense than ever.

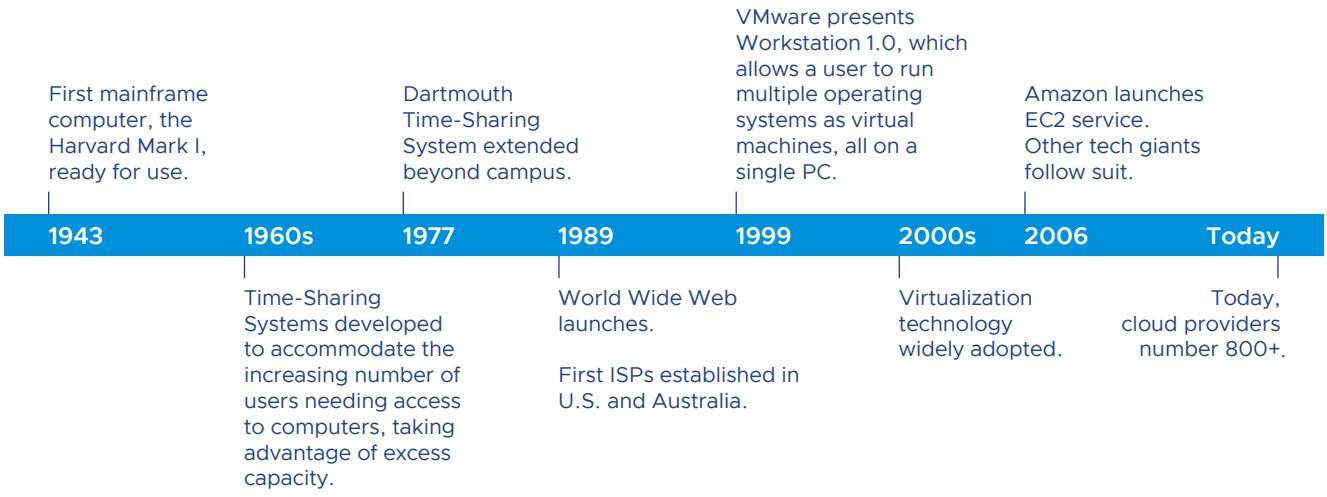

Cloud deployment models

As defined by the National Institute of Standards and Technology (NIST), cloud computing is "a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."

Expanding on this definition, NIST lists five essential characteristics of a cloud computing environment.

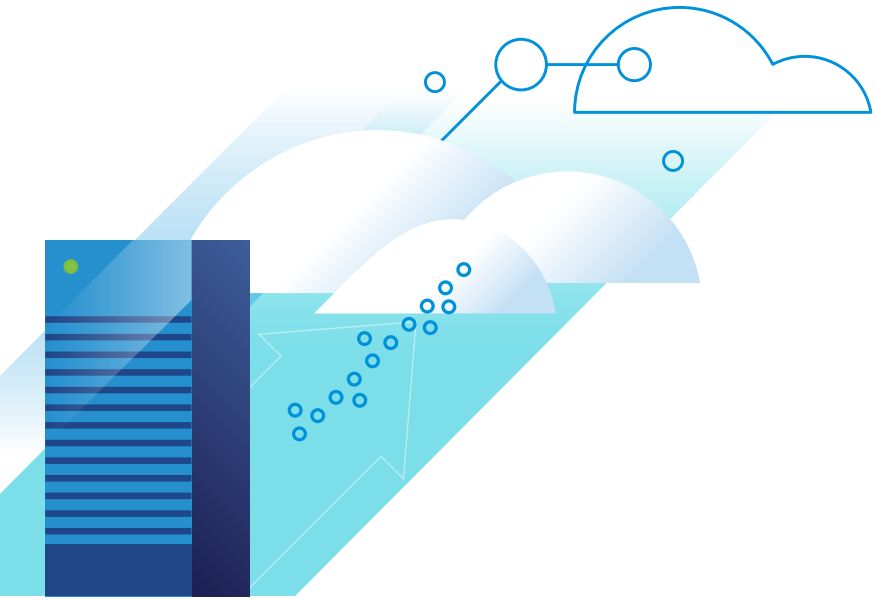

1. On-demand self-service. Users can request and access services as needed without requiring manual setup or configuration.
2. Network access. Services and resources are made available over a network (most commonly via the internet), providing access any time and from any location.
3. Shared resources. Resources are pooled to serve multiple consumers rather than limiting one resource to one user.
4. Rapid elasticity. Services are elastically provisioned and released, enabling users to scale usage up or down quickly and easily as changes are required.
5. Measured service. Resource usage can be monitored, controlled, and reported.

In practice, each of these characteristics of cloud computing can be deployed in different ways, called cloud deployment models. The definitions of the different types of cloud deployment models can be causes for debate among the cloud community, with seemingly no one definition to rule them all. This is especially the case when defining hybrid and multi-cloud deployment models. To help provide a complete picture, we outline the most well-known and commonly accepted definitions, along with how we are defining them for the purposes of this guide.

It's important to note that an organization's cloud deployment model is not set in stone—you can switch from one model to another as your business and technology needs evolve.

Private cloud

In a private cloud deployment model, cloud computing services and infrastructure are exclusive to a single organization or consumer, and resources are not shared with other organizations.

You might find some discrepancies in this definition regarding *where* the infrastructure is hosted and *how* it is managed. For example, some limit the definition of private cloud to when the infrastructure is hosted within a company's own servers or data center, with the company overseeing the infrastructure management, maintenance, and operations, and users gaining access through the company's virtual private network (VPN).

However, others broaden the definition of private cloud to include infrastructure that exists off-premises and is owned, managed, and operated by a third party or a combination of a third party and the company.

In this guide, if an environment is dedicated to a particular consumer, even if it is hosted by a third party, we refer to it as a private cloud, because what makes it private is that the underlying hardware layer is segregated from any other client's infrastructure. When the private cloud lives on the organization's own servers or data center, we refer to this as a "hosted private cloud."

Public cloud

In a public cloud, a third-party provider owns and manages the cloud services and infrastructure, and it is available on demand to multiple tenants (organizations) via the internet.

Like the private cloud, public cloud providers often leverage virtualization in their data centers to maximize the usage of their servers. But what most distinguishes a public cloud from a private cloud is that, in a public cloud, virtualized resources are in a centralized pool that different organizations can access on demand.

Public cloud service providers—think AWS, Azure, and Google Cloud Platform—offer various types of cloud-based services, such as infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS), to consumers. Each type of service encompasses varying degrees of functionality and responsibility. We'll go into these different types of cloud services in more detail in the next section.

Hybrid cloud

In a hybrid cloud deployment model, an organization uses a combination of at least one private cloud and at least one public cloud, with consistent and coordinated infrastructure and operations working between them.

To be considered a hybrid cloud, some level of interoperability is needed between the private and public cloud environments.

Multi-cloud

A multi-cloud deployment model uses a combination of clouds, which can be two or more public clouds, two or more private clouds, or a combination of public and private clouds.

To summarize the key differences between multi-cloud and hybrid cloud:

- Unlike hybrid cloud, multi-cloud does not require use of a private cloud. An organization can be multi-cloud if it is using more than one public cloud.
- Unlike hybrid cloud, multi-cloud does not require unified or coordinated operations to work between the different cloud environments, although seamless visibility and operations between environments is an ideal scenario to achieve the most value from a multi-cloud strategy.

With these definitions in mind, it is also possible—and common—to have a multi-cloud environment that is also a hybrid cloud. Here is a simplified example: An organization runs a private cloud built on its own data center powered by VMware vSphere® and VMware vSAN™ solutions. It also uses AWS' public cloud offering for compute and storage and Google Workspace for collaboration and productivity. If we stopped here, this model would be considered a multi-cloud deployment: one private cloud environment and two public cloud environments or providers.

However, the company also leverages VMware Cloud Foundation™ to unify its on-premises private cloud and public cloud services to deliver consistent infrastructure and operations across the cloud environments. We can now consider this both a multi-cloud and hybrid cloud deployment model.

Why confusion exists

As a cloud consumer or someone thinking about adopting cloud, you will find that the various deployment models are not as cookie-cutter as one might hope. Confusion often exists because of the way the private and public cloud environments were traditionally thought of and how they have evolved to encompass many more things today.

In the past, public clouds were run only off-premises. Now, public cloud providers are running cloud services on their clients' on-premises infrastructure, or inversely, organizations are building private clouds on rented, vendor-owned data centers located off-premises. For this reason, Accenture predicts that we will see deviation from using the terms "private" and "public," and instead use "dedicated" and "shared."[1]

Given the blurring of private and public clouds, it follows that what we think of as hybrid and multi-cloud have also branched off into varying viewpoints and definitions. All perspectives and definitions are valid and depend on which perspective you take. What is important is that when engaging in conversation on these topics, there is a clear understanding of which definitions and perspectives are being used in that particular discussion to ensure that all parties involved are aligned and speaking the same cloud language.

So, while other definitions are valid, this guide uses the definitions of private, public, hybrid, and multi-cloud as we have described.

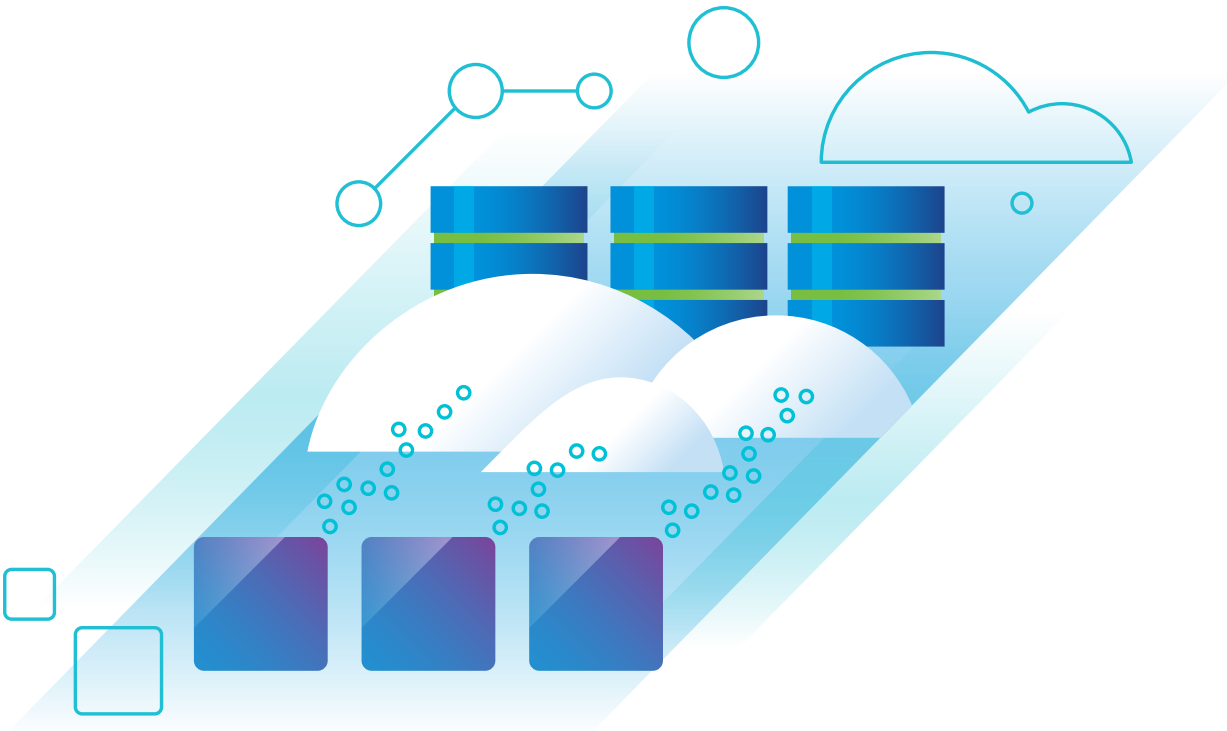

Other cloud deployment models

In addition to the four primary cloud deployment models (private, public, hybrid, and multi-cloud), you will also hear these different deployment models.

The edge

Edge computing, or "the edge," takes place at or near the source of the user or the source of the data. With edge computing, you are essentially extending your cloud environment and services out to other locations where data is consumed and processed in real time without having to go back to a central server.

With edge computing, organizations benefit from reduced latency, faster service times, and increased reliability, depending on the organization's use case. Edge computing can take on many different approaches.

Some organizations have a large infrastructure footprint at the edge—essentially hundreds of mini data centers—because they need more processing power and independent edge locations. Examples of this can include cruise ships, distribution centers, and service centers. Then there is the smaller infrastructure footprint at the edge—thousands of edge sites with resource-constrained, two- to three-server deployments. Examples of this can include remote or branch offices, of which banks and storefronts are common examples.

Local cloud

A local cloud is a dedicated cloud service that runs on-premises and includes provisioning of physical infrastructure, software, and full lifecycle management by the cloud service provider.

Also known as local cloud as a service (LCaaS), it offers a flexible cloud-like model of private cloud resources that are deployed on-premises and dedicated to an individual organization. The cloud service provider takes responsibility for and maintains control of software maintenance, such as patches and upgrades, for all pre-bundled software in the cloud platform, regardless of its physical location.

LCaaS solutions currently available in the industry include VMware Cloud on Dell EMC, AWS Outposts, Azure Stack, and Google Anthos.

Distributed cloud

In a distributed cloud, the public cloud services are distributed to different physical locations while being centrally managed from a public cloud provider. Put another way, a distributed cloud is where all a customer's cloud services and environments, including public clouds, private clouds, on-premises data centers, and edge locations, are managed from a central control plane by the originating public cloud provider.

Gartner predicts that by 2024, most cloud service platforms will provide at least some distributed cloud services, with the key value propositions being increased productivity and innovation for the consumer, while support, security, and reliability remains intact.[2]

Cloud service models

Now that we have outlined the primary cloud deployment models, we can dive deeper into the different types of cloud service models. Generally, a cloud deployment model describes how resources are provided or hosted in the cloud, and a cloud service model describes to what extent the resources are managed by the consumer or by a cloud service provider.

In cloud computing, "levels of abstraction" distinguish the different types of cloud service models. Levels of abstraction are layers of encapsulated functionality, with each level encompassing varying services and degrees of functionality available to the consumer.

To help explain, think about when businesses started renting physical servers from internet data centers. The data center provided the hardware, so the business was relieved of the need for a secure physical hosting environment—in other words, "abstracted" away from the business.

Today, cloud service providers offer numerous cloud services at varying degrees of abstraction that offload responsibilities from the consumer. There are three primary levels of abstraction, otherwise known as the three main categories of cloud computing services: IaaS, PaaS, and SaaS.

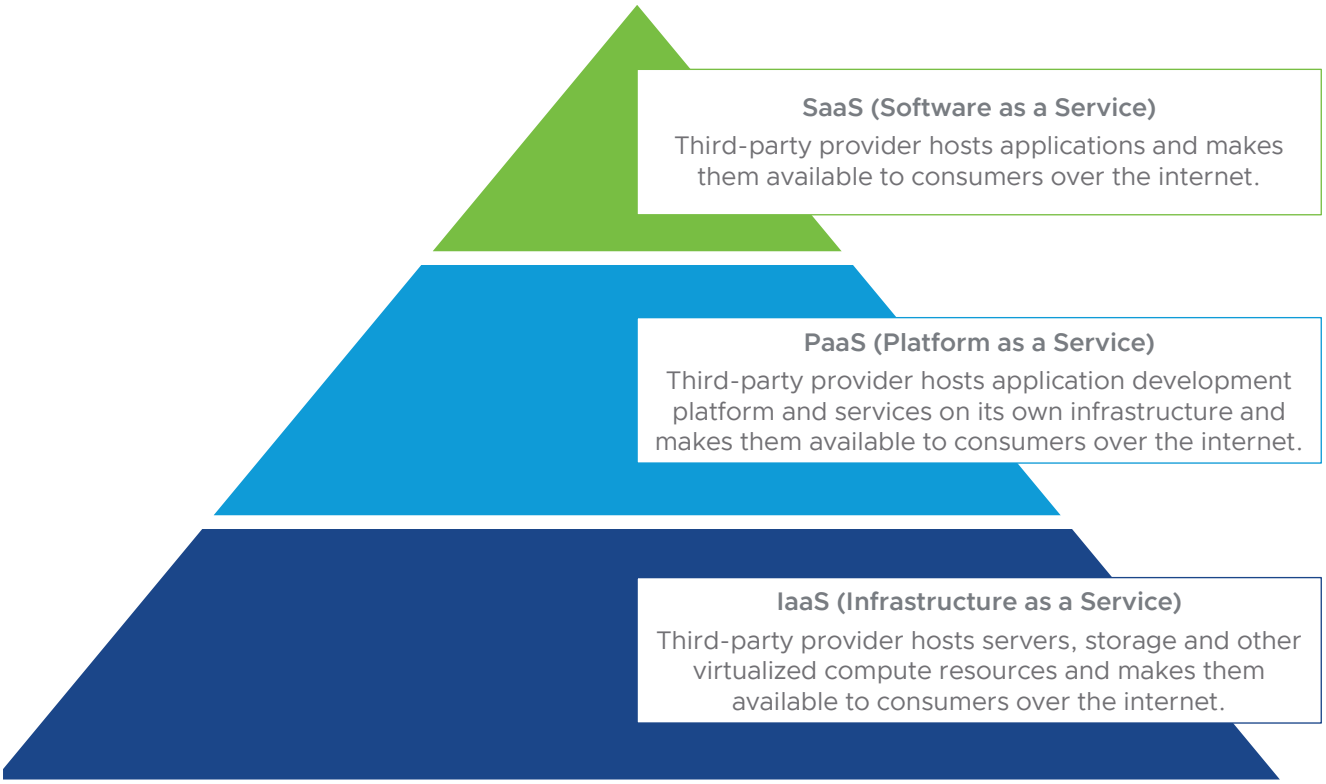

IaaS

With IaaS, the cloud provider manages the infrastructure that is typically included in an on-premises data center, such as the servers, storage, and networking hardware, as well as the virtualization or hypervisor layer. This infrastructure is provided to the consumer via virtual machines accessible through the internet. Essentially, it is a virtual data center in the cloud that provides resources to the consumer on demand.

IaaS is the lowest level of abstraction because the consumer has a greater degree of control and also greater responsibility for resource utilization, security, and operations.

PaaS

With PaaS, the level of abstraction is taken one step further than with IaaS. The cloud provider not only provides the underlying infrastructure (as with IaaS) but also a platform for customers to build, run, and manage applications. Because the provider handles hosting and maintaining the infrastructure and development platform, developers have more freedom to focus on building and running applications.

Some of the most popular services at this level of abstraction are AWS Elastic Beanstalk, Google App Engine, OpenShift, and VMware Tanzu.

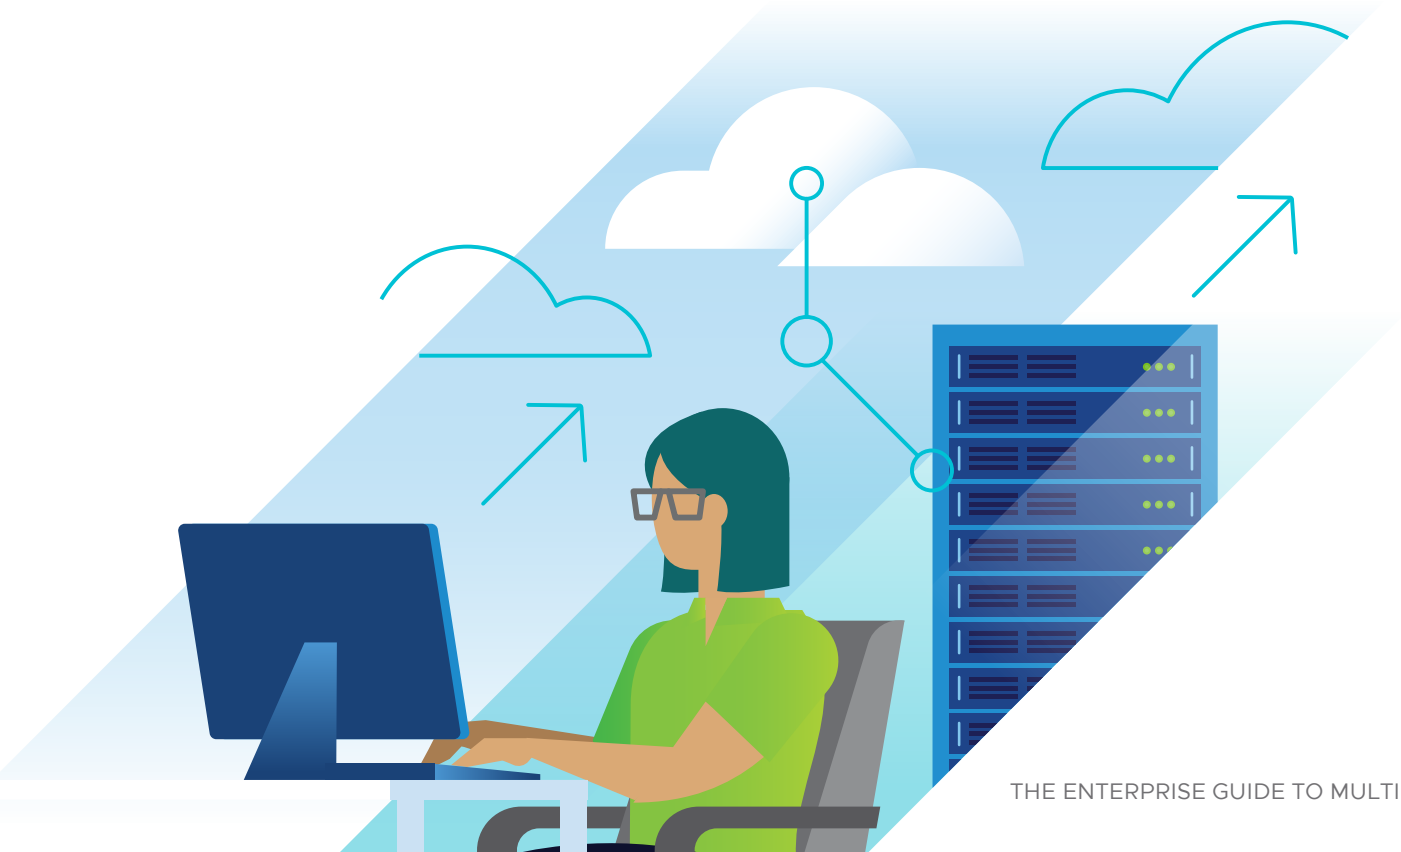

SaaS

At the highest level of abstraction is SaaS. With SaaS, the provider hosts applications and makes them available for customers to consume over the internet. SaaS is most widely known and understood by the general population, given that most people interact with SaaS applications on a daily basis. Think Netflix, Salesforce, or Slack.

SaaS removes the need for organizations to install and run apps in their own computers or data centers and usually offers flexible payment structures, scalable usage, and automatic updates.

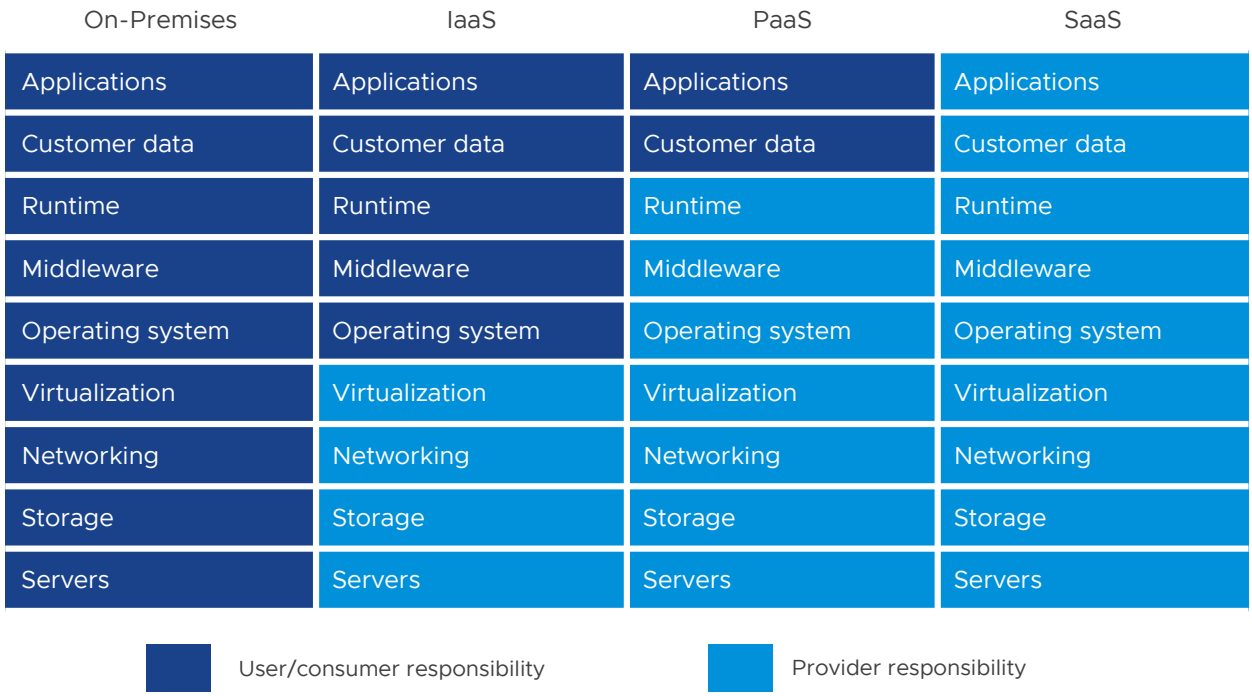

IaaS, PaaS, and SaaS are the three primary cloud service models, with most organizations using a combination of services across each model. However, there are several other ways to consume cloud services.

Bare metal services

Bare metal services let organizations deploy virtual machines (VMs) directly onto their cloud provider's hardware (bare metal) rather than in a virtualized environment. As described by AWS, bare metal services can be valuable for customers who want "access to the physical resources for applications that take advantage of low-level hardware features that are not always available or fully supported in virtualized environments, and also for applications intended to run directly on the hardware."[3]

Bare metal services are foundational to the VMware Cloud on AWS service, which brings VMware's Software-Defined Data Center (SDDC) to AWS Cloud with optimized access to AWS. IT teams can seamlessly migrate and run business-critical vSphere workloads in a familiar environment while modernizing with AWS cloud services.

Containers

Simply put, a container is a unit of software that bundles up code and all its dependencies into a single package so that an application can run quickly and reliably from one computing environment to another, such as moving from staging to production or from a physical machine in a data center to a VM in a private or public cloud.

Whereas VMs virtualize at the hardware level (carry their own app, operating system, and dependencies), containers virtualize at the operating system level (run on one shared operating system). For this reason, containers are faster to spin up and much more portable across multiple clouds and platforms.

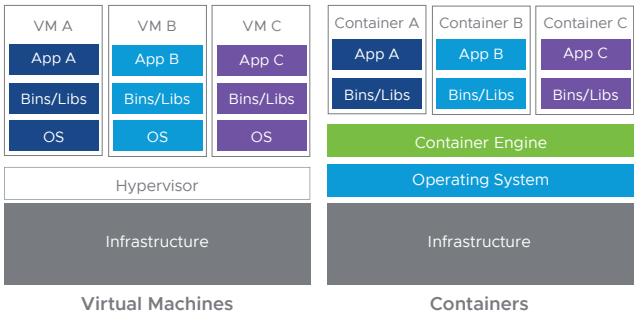

Modern applications are increasingly being built using cloud containers because of the fast deployment speed, workload portability, and the ability to simplify resource-provisioning for time-pressed developers, largely because developers can trust that the application environment will be consistent no matter where it is eventually deployed. This approach means less time diagnosing and debugging and more time for shipping new functionality to users.

Kubernetes

As organizations expand their use of containers across multiple services and environments, managing, securing, and operating them becomes increasingly complex. To help alleviate this issue, organizations turn to container orchestration tools, such as Kubernetes. Kubernetes, also known as K8s, is the industry's leading solution designed to simplify and automate processes involved with deploying, scaling, and operating containerized workloads.

In VMware's State of Kubernetes report, **98 percent of software development and IT professionals reported seeing clear benefits from Kubernetes**. Respondents cited these primary benefits:[4]

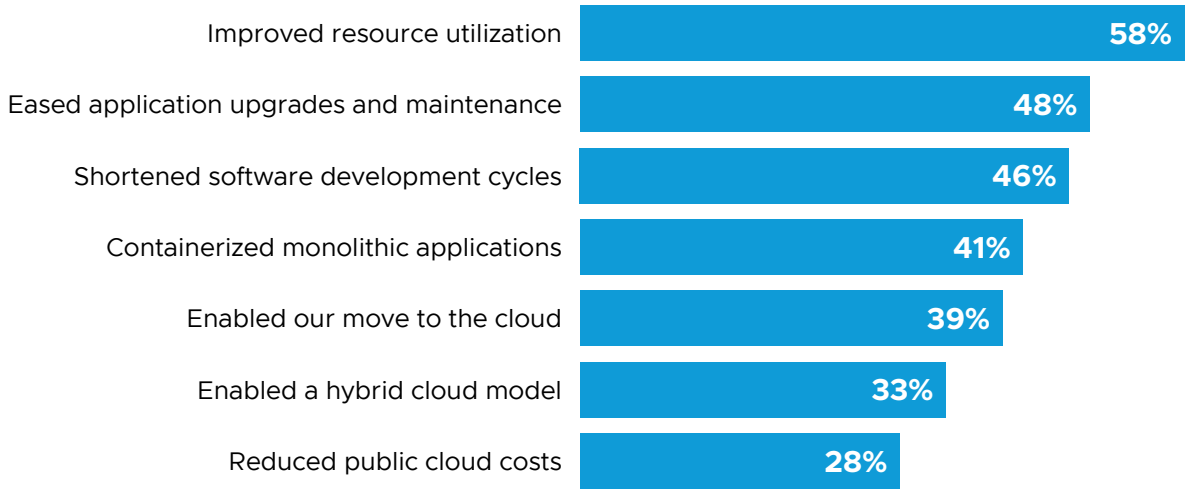

See all the results and lessons learned in The State of Kubernetes 2021 Report.

Microservices

Microservices refers to an architecture approach in which a single application is composed of many loosely coupled and independently deployable smaller services that communicate via APIs.

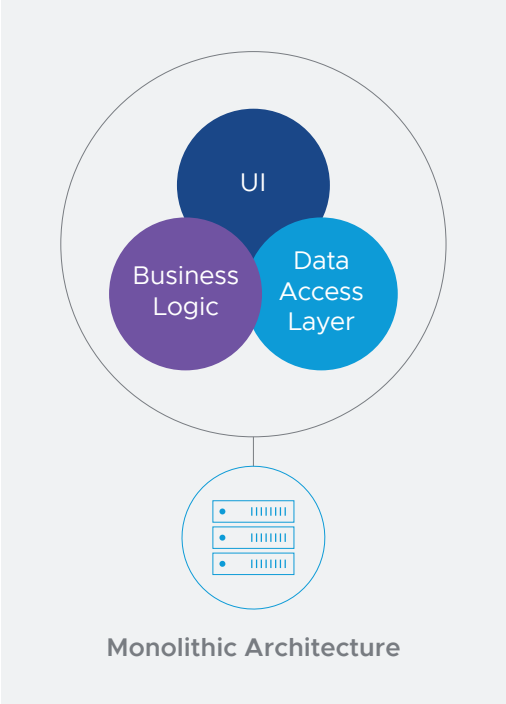

Monolithic Architecture

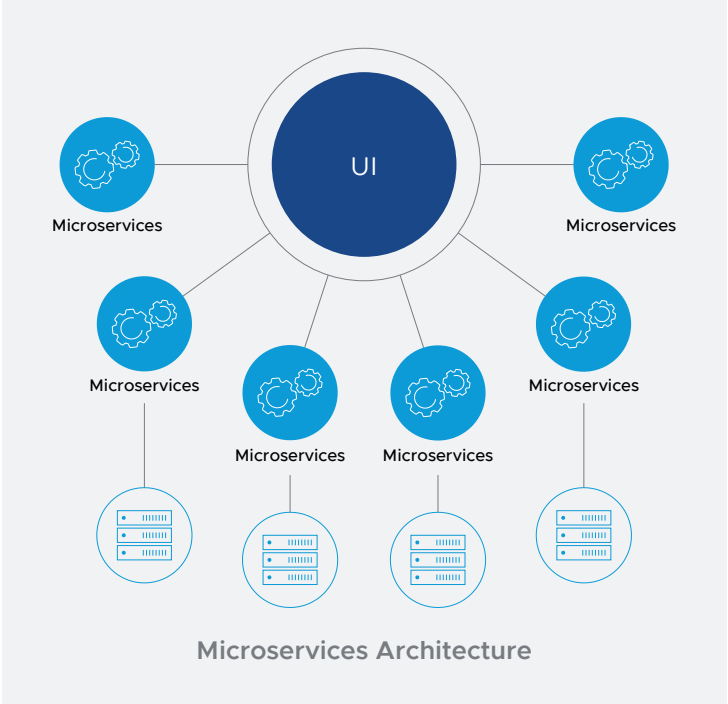

Microservices Architecture

Traditionally, organizations operate from a monolithic architecture in which all processes are tightly coupled and run as a single service. So if one process of the application changes, the entire architecture needs to change. But with a microservices architecture, each service performs a single function and can be updated, deployed, or scaled without affecting the functioning of other services.

Why Multi-Cloud?

Given the number of cloud deployment and service models we've just covered, it's clear that the options available for cloud consumers today are numerous. Research shows that most organizations are opting for a multi-cloud deployment model to leverage the differentiated cloud services that are best suited to meet the specific needs and goals of the business.

In fact, a recent a survey indicates that 80 percent of organizations have a distributed model for cloud apps across public, private and edge environments, and expect to support this combination in three years' time, with 46 percent of organizations expecting to add more cloud environments by the end of 2021.[5]

Whether technical in nature or business-driven, considerations for multi-cloud can vary by organization as well as individual departments and business units. But at a foundational level, organizations opt for multi-cloud to increase business agility, enable scalability, and optimize developer productivity.

Even organizations that might not outright choose a multi-cloud architecture will likely result in a multi-cloud reality in their search to accelerate innovation and meet business demands—for example, when individual users or departments procure their own cloud services without IT knowledge, or when a merger or acquisition brings a new provider into the enterprise cloud fold.

Further considerations, such as market competition, licensing agreements, data privacy, regulatory compliance, and existing partnerships often introduce competing needs and bespoke views, which will heavily influence cloud adoption decisions.

Organizations also have a need to ensure that regulatory data-resiliency requirements can be achieved. Depending on geographical location and data requirements, these requirements could create a situation where an organization's cloud provider of choice does not have sufficient regional or geographic presence to meet this need.

Get more insights into how to modernize with multi-cloud in the 2021 VMware Market Insights Report, Driving Digital Business with App and Cloud Transformation.

Ultimately, the ideal infrastructure for the majority of organizations is one that enables a combination of cloud environments: private, public, data center, and the edge. Here are some critical reasons why an organization pursues a multi-cloud strategy:

- **Optimization and operational efficiency.** Not all clouds are created equal, and not all applications, services, workloads, data, or business requirements are either. A multi-cloud strategy enables organizations to leverage the environment that is best suited for its intended purpose rather than trying to fit square pegs into round holes.
- **Regulatory compliance alignment.** Compliance requirements, such as the European Union's GDPR, often require customer data to be held in specific regions. The implementation of this requirement frequently leads to a multi-cloud scenario in which data is hosted in locations that satisfy regulatory and compliance requirements, and application code is hosted elsewhere.
- **Application modernization and native services integration.** Different preferences exist among business units for higher-level cloud services that integrate with existing applications and accomplish specific tasks in areas such as machine learning, artificial intelligence, Internet of Things, and containers. This in turn, results in an organization needing to use multiple cloud providers to meet the combined needs of all business units.
- **Cost, consumption, and licensing models.** It's often the case that one cloud provider's pricing or licensing model for a particular service does not align with how the customer's business operates, which can be due to any number of factors, including financial, technical, or operational. It's also the case that one provider's prices are more or less advantageous for a particular service, enabling the customer to optimize costs depending on the services it uses.
- **Reduced reliance on a single provider.** Also referred to as "vendor lock-in," organizations do not want to be tied down to the specific standards and services of one vendor. Businesses that leverage multiple cloud providers have greater agility and adaptability when faced with market disruptions and industry changes because they are not dependent on the performance of a single provider.
- **Enhanced service availability and resiliency.** Application resilience and the protection from outages offered by not having everything running in one location, or even one public cloud, is a common reason for operating in more than one cloud.

The advantages of multi-cloud are evident, and the stats are revealing, but the reality is not so simple. Countless roadblocks can prevent organizations from realizing the benefits of successful multi-cloud transformation.

Multi-Cloud Challenges

With a multi-cloud environment, customers contend with applications in different environments and different stages of modernization. Incompatible architectures between clouds, diverse management and operations, and inconsistent security policies combine to bring a level of complexity and risk that challenge the inherent benefits of multi-cloud. In this next section, we cover some of the most common challenges an organization can face along its multi-cloud journey.

Where to start

When looking at a portfolio of hundreds or thousands of applications and facing a myriad of cloud and container choices with only minor differences, it can be difficult to choose which path to take. In fact, from recent market research, 80 percent of senior business and IT leaders feel that matching the best platform to each app is extremely challenging.[6]

While some enterprise apps can continue to run on-premises, you might choose to use containers to accelerate development of new modern apps. You might also want to integrate capabilities available through public cloud providers, such as artificial intelligence, machine learning, Kubernetes, or database management and analytics.

Before embarking on a cloud transformation journey that involves migrating or modifying workloads, an organization needs to strategically consider which scenario is best for a particular resource, depending on the expected cost, complexity, and the benefits that will be achieved.

To help facilitate this process, these five primary strategies—commonly known as the 5 Rs—can help customers choose which approach to take for which workload.

- **Retain** – Leave workloads in a private cloud environment.
- **Rehost** – Move applications from one virtualized environment to another, such as from a private cloud environment hosted on-premises to a public cloud environment.
- **Replatform** – Retool an application to take advantage of cloud services and technology without changing its core architecture. Replatforming usually means leveraging container technology or Kubernetes for app development and modernization.
- **Refactor/Build** – Rewrite an application at the source code level to better support its environment. Here, you're usually taking advantage of modern application design, microservices, or cloud native principles.
- **Retire** – Decommission an application or replace it with a SaaS alternative.

Cloud migration

In an effort to maintain growth and innovation at an exponential rate, many organizations choose to rehost, or migrate, applications from an on-premises environment to the public cloud. However, as a result, many have rushed into cloud migration projects without fully understanding or planning for the challenges they will face.

In a recent survey, 70 percent of companies reported being actively engaged in public cloud migration and plan to migrate more than half of their applications. However, only 33 percent of companies are more than halfway to their migration goal.[7] Additionally, a survey by Accenture shows that nearly two-thirds of companies have not achieved the benefits they expected from cloud migration.[8] Why is this the case?

- **Lack of strategic planning**. Gartner estimates that less than one-third of organizations have a documented cloud strategy.[9] Without a strategic and agreed-on strategy, many organizations' cloud migration attempts backfire, causing projects to stall or requiring apps to be moved back to their original on-premises environments.[9]
- **Skills shortages**. As many as 86 percent of organizations report shortages in IT staffing, with many acute skills gaps related to cloud platforms.[10] Attempting to manage the complexities of cloud migration with the right teams or personnel in place can open organizations up to additional delays, costs, and risks.[11]
- **Incompatible infrastructure between environments**. One of the most prevalent migration challenges is incompatible and legacy infrastructure, which according to an Accenture study, is the No. 1 challenge for 43 percent of organizations.[8] Different environments have different operating systems, APIs, and integration requirements, which can cause performance issues, or in worst-case scenarios, the application does not work at all in the new environment.
- **Lack of support for existing security policies**. Each public cloud provider has its own unique security policies, which might not align with the policies that an organization has established on-premises. This difference can lead to increased security risks while moving data from one environment to another, as well as increased costs and complexity to ensure seamless and adequate protection.
- **Cost of refactoring and rewriting applications**. Many enterprises approach cloud migration with a lift-and-shift strategy, inevitably embarking on rounds of refactoring or rewriting their applications to suit their public cloud host configurations. This process is costly, time consuming, and resource intensive, often delaying project timelines and taking time away from more strategic initiatives.

App modernization

As we have previously discussed, organizations are choosing cloud strategies that enable them to keep their applications in the most suitable environment, which often means supporting a mix of existing applications while also deploying new cloud-native application architectures—both on-premises and in the public cloud.

The 5 Rs outline the primary approaches that organizations can take when choosing what to do with each application. While cloud migration (rehosting) is a common approach, research shows that app modernization projects (including replatforming or refactoring and rebuilding) are among the most popular plans for existing applications.[7] In this scenario, on-premises applications take on cloud-like capabilities and are integrated with public cloud services.

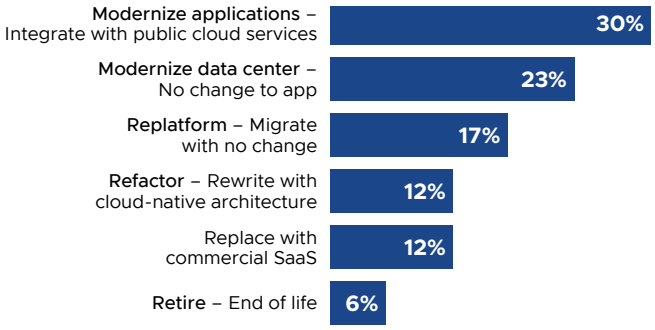

To help facilitate app modernization, many organizations are turning to containerization and Kubernetes, with research showing that today, an average 19 percent of all applications are containerized. This number is expected to grow to 25 percent over the next three years as a way to improve availability, increase developer productivity, accelerate development efforts, and ease feature updates.[12]

However, this is easier said than done. Common challenges that organizations experience in their app modernization journey include:

- **Increased risk** while managing infrastructure and the applications with different operations teams, tools, and processes in the data center and cloud
- **Slower time to market** as different IT processes for different environments negatively impact responsiveness and agility
- **Higher cost** due to less flexibility and the inability to quickly migrate and scale to meet infrastructure demands
- **Limited visibility** across all environments, causing higher security risk and service-level agreement (SLA) violations due to multiple disparate management tools and processes
- **Skills shortages** because organizations need to build competence with multiple environment-specific tools and processes

Management and optimization

After an organization has settled on a cloud infrastructure and app modernization approach (with the understanding that changes and adjustments occur on an ongoing basis), the next challenge is operationalizing and managing the multi-cloud environment. With hundreds or even thousands of users across multiple clouds consuming numerous services every day, cloud operations, security, and costs can quickly spiral out of control.

Key challenges that organizations face when it comes to operationalizing and managing a multi-cloud environment include:

- **Skills shortages**. From a Gartner survey of infrastructure and operations leaders, the majority of respondents indicated insufficient skills as a top priority for the years ahead.[12] Many organizations that are either new to the cloud or rapidly scaling cloud usage, quickly realize that the technical skills and operational frameworks of the past do not always apply to the nuances of managing a cloud environment—and especially a multi-cloud environment. If each cloud is run as a separate, isolated silo with its own unique development and operating model, taxonomy, and set of APIs, new skills are needed to leverage the innovation capabilities of each cloud environment.
- **Change management**. As the saying goes, "the only constant is change." And this is certainly true when it comes to multi-cloud management. Whether migrating workloads previously run on-premises to a public cloud, introducing modern apps and container technology, or integrating a new cloud architecture from an acquisition, managing the sheer number of changes that occur in a multi-cloud environment can be overwhelming, to say the least. What differentiates organizations with successful change management depends on their people, processes, and technology and how each of these components work together in an iterative and agile way.
- **Lack of visibility**. It is impossible to control what you cannot see, and this is exacerbated by a decentralized multi-cloud environment. Without visibility across all environments, organizations struggle to identify opportunities to optimize infrastructure performance and operations, predict and forecast costs, and prevent security and compliance risks. Visibility is foundational to success with a multi-cloud strategy, not only with ongoing management and operations but also when choosing whether to migrate or modernize an application (back to the 5 Rs).

- **Integration incompatibilities**. When introducing a new environment, application, or service, organizations usually find that it does not integrate seamlessly into their existing infrastructure—or at least as seamlessly as they hoped. Different environments have different operating systems, APIs, and integration requirements, which can lead to performance issues and require additional time, money, and resources to integrate successfully.
- **Security and compliance**. The traditional means of security and compliance are not enough to keep pace with the dynamic nature of the cloud, combined with the complexity of a multi-cloud environment and the pace at which attackers are able to take action. Security parameters and controls vary depending on the environment and the cloud provider, which can cause confusion around who is responsible for which aspects of cloud security. The possibility of misconfigurations is multiplied, and without the right tools in place, notifications or automated actions around misconfigurations and vulnerabilities can lag, leaving your environment open to potential risk.
- **Cloud costs**. The potential to save money in the cloud is huge, but when up and running, many organizations find that they are not saving as much as they anticipated or that they are even spending more than they were before. This does not mean that moving to the cloud is a mistake. Overspending often stems from a few reasons, each of which are exacerbated when adding multiple, disparate cloud environments: incomplete or limited visibility into resources and activity, the complexity of cloud pricing and billing, lack of governance and policies to drive accountability and keep costs in control, or insufficient tools to accurately track cloud usage and costs by dynamic business groupings (application, team, department, individual, and so on).
- **Operations and automation**. Managing fluctuations in cloud usage, juggling governance policies, and keeping track of reservations and discounts across several different cloud environments, departments, cost centers, locations, and needs is challenging, especially when administrators are more accustomed to the traditional way of managing data centers and physical servers. Many organizations find that as their cloud grows, staff cannot keep up with the volume of information and the complexity of running workloads across multiple clouds. Implementing automation is a best practice for organizations who want to operate fast and at scale. But an automated, hands-off approach does not always work with every scenario due to app or environment exceptions and constantly changing technical or business requirements.

The Ideal Multi-Cloud Environment

As you may have noticed, there are common themes across the multi-cloud challenges we've discussed that span migrating, modernizing, and operationalizing. These include a failure to plan and align, a lack of skills and expertise, and the incompatibility or complexity of technology and operations.

So, how can organizations break through these barriers for effective and successful multi-cloud use?

In a recent survey of about 1,200 IT and Line of Business decision-makers and practitioners, respondents were asked about their ideal multi-cloud environment and ranked the following as the most important characteristics:[7]

1. No matter where an application is deployed—it's secure and protected.
2. IT can consistently manage applications regardless of where they are deployed.
3. Applications are portable from cloud to cloud without refactoring.
4. A single set of tools is used to manage the entire application portfolio wherever applications are deployed.
5. Developers and operations teams can collaborate easily.
6. Developers can build and deploy applications to any public cloud.

From the results, we can gather that respondents favor a **single operating model that delivers consistent operations wherever applications are deployed—from the data center to the edge and across cloud environments.**

A New Approach: The Cloud Operating Model

What is needed is an approach to cloud that is designed for multi-cloud and for every app. Because you shouldn't have to choose between your past and your future, between private and public, between containers and VMs, or between speed and security.

Enter the cloud operating model. The cloud operating model is a framework that binds your business, application, and cloud strategies together to accelerate agility, optimize performance, and control your multi-cloud environment. It brings people, processes, and technology together for consistent service delivery, operations and governance wherever workloads reside, for today and the future.

A cloud operating model is made possible by a cloud management platform that abstracts away the differences between environments to unify cross-cloud operations, while at the same time enables access to the unique and innovative portfolio of cloud services delivered by each cloud environment and provider.

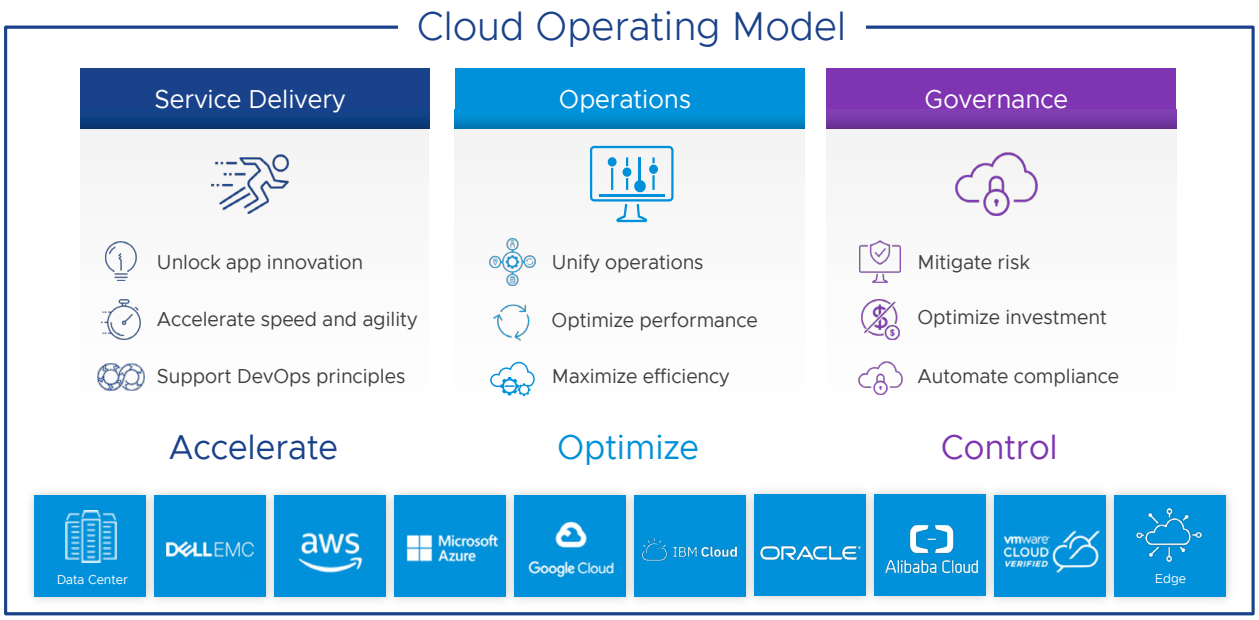

How You Can Achieve a True Multi-Cloud Operating Model

A cloud operating model sounds great, but how do you make it a reality? At VMware, we make a cloud operating model possible for our customers with VMware Cloud™, a robust, flexible solution that enables the innovation and freedom of multi-cloud with the simplicity of one.

VMware Cloud delivers multi-cloud services that span the data center, edge, and any cloud, including native cloud services. Optimized for both traditional and modern apps, VMware Cloud provides a single platform for all apps in any environment, providing the foundation for a true multi-cloud operating model.

Designed to preserve your choice in cloud and application architectures, it unifies all environments with consistent infrastructure and operations, delivering the agility, reliability, and security you need while reducing total cost of ownership.

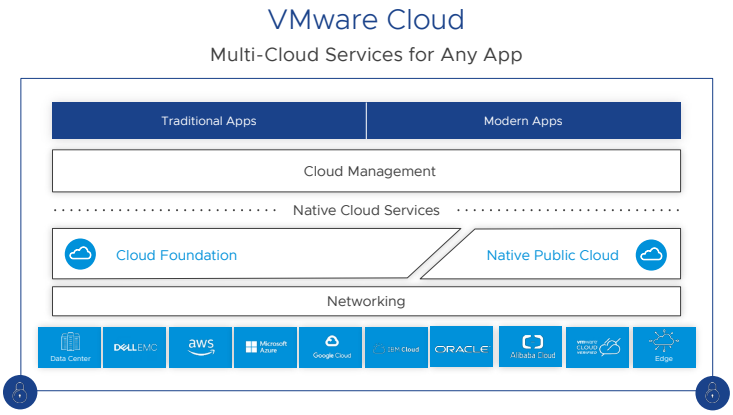

Choice: The flexibility of any cloud, infrastructure, and application architecture without risk.

Speed: The fastest and simplest path to modernization and migration for any cloud at unparalleled performance.

Control: Unified and simplified management across clouds and apps to reduce risk and lower total costs.

Learn more about operationalizing your multi-cloud environment with VMware at www.vmware.com/cloud-solutions/multi-cloud, or feel free to get in touch with our team directly.

1. Accenture, Inc. “Cloud computing: Understanding what cloud is and what it can do for you.” 2021.
2. Gartner, Inc. “The CIOs Guide to Distributed Cloud.” Katie Costello. August 12, 2020.
3. Amazon Web Services. “Compute Abstractions on AWS: A Visual Story.” Massimo Re Ferre. September 6, 2018.
4. VMware, Inc. “The State of Kubernetes 2021.”
5. VMware, Inc. “Driving Digital Business with App and Cloud Transformation: 2021 VMware Market Insights Report.” March 2021.
6. VMware, Inc. “Architecting Your Multi-Cloud Environment.” May 2021.
7. VMware, Inc. “The State of Application Modernization and Hybrid Cloud Computing.” February 2020.
8. Accenture. “Cloud outcomes survey: Expectation vs reality.” June 6, 2019.
9. Gartner, Inc. “Cloud Strategy Leadership.” 2017.
10. 451 Research: Cloud Trends in 2020: The Year of Complexity and its Management, January 2020.
11. VMware, Inc. “Top 5 Barriers to Cloud Migration and Modernization.” May 2020.
12. Gartner, Inc. “The Cloud Infrastructure and Platform Services Skills I&O Teams Require for the Future.” Raj Bala, Ross Winser. September 2, 2020.

Join us online:

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273
Fax 650-427-5001 www.vmware.com Item No: FY22-6442-VMW-ENT-GUIDE-MULTI-CLOUD-EBK-USLET-WEB-20210813 8/21