

The Era of Multi-Cloud Services Has Arrived

VMware executive outlook for
IT leaders and decision-makers

Executive summary

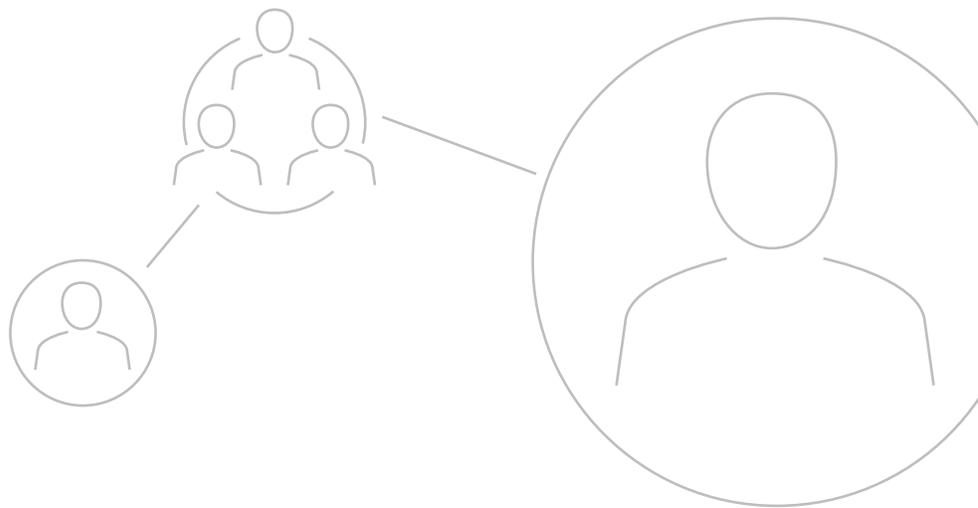
Today, companies are choosing to deploy apps and consume services on the clouds that best suit their business objectives. This model, known as multi-cloud, uses two or more public clouds, two or more private clouds, or a combination of public, private and edge clouds to distribute applications and services.

Multi-cloud strategies offer several advantages, from increased resiliency and agility to data sovereignty. A key benefit is that organizations can leverage a specific cloud provider's bespoke capabilities and services to further transform their application stacks.

The challenge, however, is that each cloud runs as a separate, isolated silo with its own development and operating model, taxonomy, services, APIs and management tools that do not extend functionality to other cloud platforms. This lack of interconnection increases security and financial risk while forcing companies to manage their multi-cloud environments through a patchwork of off-the-shelf, custom-built and native cloud service provider tools.

Rather than managing cloud environments in a piecemeal or customized manner, how can organizations implement a multi-cloud strategy that simplifies and streamlines development, operations, networking and security across clouds?

This white paper explores an emerging category of IT services that addresses the increasing complexity of enterprise cloud environments. It also details the capabilities required to successfully transform your multi-cloud estates.



Advantages and challenges of multi-cloud

Enterprises are accelerating multi-cloud strategies to

- Achieve regulatory compliance or availability in different regions and markets
- Meet the diverse needs and requests of users, teams and business units
- Distribute applications and services to the edge to be closer to physical devices and end users
- Improve application and workload performance, scalability and security
- Avoid vendor lock-in and ensure data sovereignty

As applications continue to serve as the primary channels to deliver customer services and optimize business processes, organizations are looking for ways to support the needs of applications both new and old. For example, some workloads or apps perform better on a particular platform while others work best with a service that is uniquely offered by a specific cloud vendor. As organizations continue to mature cloud native development, a multi-cloud strategy provides access to the broadest array of services when building new apps.



A multi-cloud strategy also supports the needs of your entire application portfolio while overcoming the challenges of legacy infrastructure. Organizations can accelerate digital transformation by using a multi-cloud environment to methodically migrate workloads and modernize application portfolios with cloud-specific services best suited for each application.

Companies that embrace multi-cloud architecture experience tangible results. In a global study of more than 1,800 IT leaders, IT decision-makers and developers, VMware found that organizations that leverage a multi-cloud environment to modernize IT infrastructure and apps, automate operations, and provide secure access to apps and data from any device and location experience the following:²

42%

faster time to release applications



41%

fewer IT hours spent on infrastructure



35%

improvement in business productivity



1. VMware Inc. "VMware FY22 Q4 Executive Pulse." January 2022.

2. VMware, Inc. "FY22 H2 Benchmark, Digital Momentum." August 2021.

The challenge of multi-cloud

Top multi-cloud challenges

- Differing infrastructure, APIs, database, network and security constructs
- Cost of refactoring applications for a new public cloud environment
- Risks related to security, data and privacy issues associated with regional data regulations
- Increased complexity from specialized tools
- Need to train staff or hire for specialized skills to support public clouds

Despite the benefits, operating in a multi-cloud environment is complex. Each cloud environment has its own infrastructure and operational requirements, requiring different skillsets. There are also new skills to learn to leverage the innovative capabilities of each cloud environment. Because of these differences, application and development teams struggle to quickly troubleshoot performance issues when they arise, and operators find it hard to consistently apply policies to ensure that apps are secure and compliant wherever they are deployed.

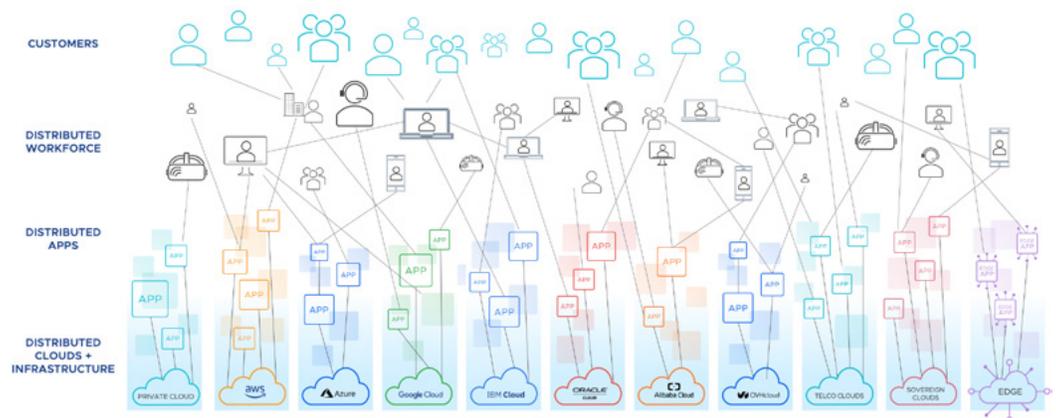
The ways in which companies arrive at multi-cloud can also add complexity. Many organizations become multi-cloud by accident through ad hoc cloud adoption from lines of business or through acquisitions. The result can be a hodgepodge of cloud environments and operating models. Achieving the business outcomes that multi-cloud can deliver requires eliminating complexity across multi-cloud operations and architecting for a modern, multi-cloud world.

Multi-cloud demands a modern approach

The issue of inconsistency—and the complexity that it causes—is a familiar one. IT teams have seen variations of this problem before and tackled it by adding a layer of abstraction to the environment. Virtualization made it possible to abstract finite physical infrastructure assets, storage arrays and networking devices that all operated independently in the data center, and create unified and scalable compute, storage and networking resources based entirely in software.

As virtualization advanced by adding automation and on-demand self-service access to infrastructure resources, the door opened to creating infrastructure as a service (IaaS) and the era of hyperscale cloud providers, such as Amazon Web Services (AWS), Google Cloud and Microsoft Azure.

With most organizations now using multiple cloud providers for their apps and infrastructure, IT teams are revisiting a similar pattern of operational complexity. For developers, each cloud provider has unique infrastructure, interfaces and APIs that add work and slow the pace of their releases. For operators, each additional cloud increases the complexity of their architecture, fragmenting security, performance optimization and cost management.



As hypervisors did previously, tackling multi-cloud complexity starts with a layer of abstraction. This abstraction layer must span across heterogeneous cloud platforms or environments to enable horizontal development and operations capabilities without hindering access to each cloud provider's unique portfolio of services.

In practice, this abstraction means that developers can write apps using their preferred framework without worrying about the infrastructure or the cloud on which it runs. Operators can deploy, manage and monitor apps and container infrastructure in the same way for every cloud. Security teams can apply policies consistently to every cloud and app. And ultimately, the business realizes quicker time to market and quantifiable improvements in app performance, efficiency and security. Those in regulated industries can also meet their unique sovereign cloud requirements and maintain jurisdictional control while achieving cutting-edge transformation at scale.



The arrival of multi-cloud services

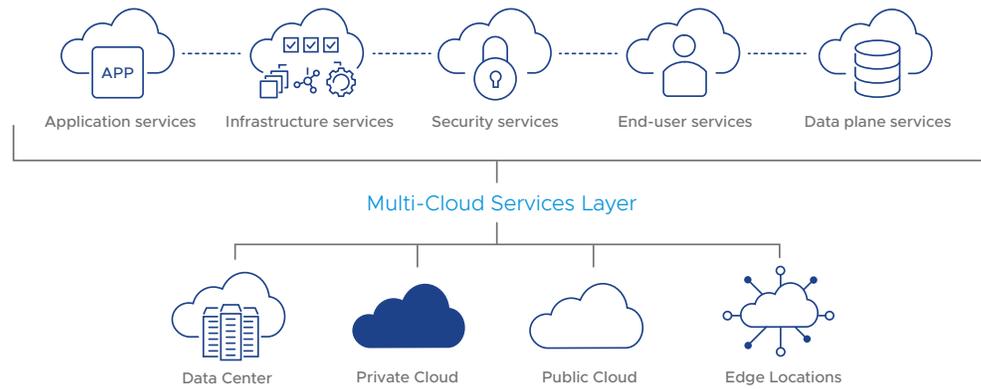
Abstraction by itself, however, does not solve the challenge of how best to build, manage, govern and optimize apps and workloads across a multi-cloud environment. This is where multi-cloud software services come in.

A multi-cloud service provides a consistent API, object model, identity management and other core functions across clouds. The service has one or more of the following characteristics.

- **Runs on a single cloud but supports interactions with at least two different clouds** – For example, a Kubernetes cluster management service control plane exists only on AWS but can deploy and manage Kubernetes clusters on a private cloud, edge, AWS, Azure or Google Cloud.
- **Runs on multiple clouds and supports interactions with at least two different clouds** – For example, a Kubernetes cluster management service control plane is deployed on a private cloud, AWS, Azure and Google Cloud and supports deploying and managing Kubernetes clusters on those clouds.
- **Runs on any cloud or edge, even in disconnected mode, and basic operations are fully automated** – For example, the control plane is delivered as software with embedded operational logic, giving it some aspects of a live cloud-delivered service.

Organizations can use multi-cloud services to abstract and standardize cloud infrastructure and operations, development and security capabilities into one platform to reduce or eliminate the complexity of individually building or consuming the equivalent native services from multiple clouds.

In this model, public clouds, data centers, private clouds and edge locations are all verticals, and multi-cloud services are horizontals, providing functionality across these locations. These horizontal capabilities integrate with and complement the native services of each cloud while providing the consistency and standardization that development, operations and security teams need.



Types of multi-cloud services

The common horizontal functions that organizations seek to standardize across clouds and that can be provided by a multi-cloud service include the following:

- **Application services** – Enable teams to use core services, such as observability, replication, backup and restore, in a standardized manner across clouds to speed innovation. Application services include developer tooling, modular development platforms, databases, messaging, AI and machine learning, serverless and CI/CD capabilities.
- **Infrastructure services** – Deliver consistent infrastructure as code across public, private and edge cloud environments. Multi-cloud infrastructure services interact with individual cloud providers' IaaS APIs to create, configure and release infrastructure resources. Infrastructure services can improve application and service resiliency, security, interoperability, performance and portability. Examples of multi-cloud infrastructure services include core compute, storage and network services presented through virtual machines or containers, infrastructure automation and Kubernetes solutions, observability, and services for monitoring and managing performance, costs and compliance across clouds.
- **Security services** – Deliver security and networking as a built-in distributed service across users, apps, devices and workloads in any cloud. Multi-cloud security services include network detection and response (NDR), endpoint detection and response (EDR), next-gen antivirus (NGAV) and secure access service edge (SASE).
- **End-user services** – Enable organizations to provide the correct level of role-based access to applications and data across clouds. End-user services include virtual desktops, mobile device management, and end-user application delivery.
- **Data plane services** – Enable the workloads, analytics and data that make up applications to operate across containers, services and clouds.

How multi-cloud services reduce complexity

Multi-cloud services enable organizations to architect multi-cloud environments that best match their applications, with the flexibility to build, deploy and manage from the data center to the cloud to the edge. Let's take a closer look at what the different types of multi-cloud services offer.

Multi-cloud application services



Native cloud services, such as Amazon RDS, Azure SQL and GCP Big Query, give enterprises the opportunity to innovate quickly by increasing developer velocity and leveraging pre-built offerings for data analytics.

However, organizations run into trouble when they extend development to application services on multiple cloud providers.

Multi-cloud application services provide the abstraction layer at which most standardized DevSecOps activities occur. By abstracting disparate APIs and core services, such as observability, replication, backup and restore across clouds, multi-cloud application services simplify CI/CD pipelines and streamline provisioning and managing multi-cloud workloads.

Multi-cloud application services can route requests for a particular service, such as a database, and deploy the service on the correct individual cloud using the most appropriate technology for the use case, such as Amazon RDS, Azure SQL or Oracle DB, or multiple individual endpoints, such as MongoDB or Snowflake. When an enterprise has the ability to control and consume these services in a standardized manner through multi-cloud application services, the speed of innovation increases and technologies are easier to consume.

Multi-cloud application services can also be consumed via a platform as a service (PaaS). PaaS solutions provide developers flexibility while enabling organizations to create guardrails to ensure that security and compliance requirements have been met.

You can implement multi-cloud application services provided by PaaS solutions on most cloud providers, including private and edge clouds, for full application portability.

“The best part about this is that my development team doesn't know any difference between going to OpenStack and AWS. The developer experience is almost exactly the same.”

Greg Meyer, Director and distinguished engineer
Cerner

Multi-cloud services



[Multi-cloud application services](#)



[Multi-cloud infrastructure services](#)



[Multi-cloud security services](#)



[Multi-cloud end-user services](#)



[Multi-cloud data plane services](#)

Multi-cloud infrastructure services



Most enterprise applications, whether off-the-shelf or home grown, are traditional applications that run best on traditional infrastructure. This infrastructure is most commonly implemented using virtualization. As an enterprise ventures into public clouds, teams must decide whether to retire, replatform, or refactor an application. Some applications are modernized by dividing the functionality into microservices and running them in containers. Kubernetes is the most used orchestration engine for container runtimes. Other applications are modernized by converting the functionality into components and workloads run on application services offered by public cloud providers.

Whether an enterprise wants to continue to run its applications as virtual machines, containers or both, using a multi-cloud infrastructure service to standardize management and operations across clouds can reduce the complexity associated with running workloads in multiple clouds. Multi-cloud infrastructure services interact with the IaaS APIs of individual cloud providers to create, configure, manage and release infrastructure resources. Examples of these resources are virtual compute, storage, networking and other IaaS platforms like Kubernetes. The abstraction of multiple cloud infrastructure services is also an automation point that can be leveraged by DevSecOps and site reliability engineers.

Multi-cloud infrastructure services could also include services to deploy, migrate or operate applications and infrastructure; observe, monitor or optimize performance, costs and compliance; and otherwise manage clouds in a consistent manner.

Application services from public cloud providers might be available only in their public cloud. Therefore, teams best achieve application portability at the infrastructure layer or by using a PaaS to develop and deploy applications. **Enterprises that run traditional or modern applications and want workload portability or the use of additional cloud providers for capabilities, like disaster recovery or burst capacity, should prioritize implementing multi-cloud infrastructure services.**

Multi-cloud services



[Multi-cloud application services](#)



[Multi-cloud infrastructure services](#)



[Multi-cloud security services](#)



[Multi-cloud end-user services](#)



[Multi-cloud data plane services](#)

Multi-cloud security services



Multi-cloud security services centralize security operations and secure software supply chains. One of the most complex aspects of running workloads in multiple clouds is maintaining a robust security posture against ever-increasing cyberthreats. Each cloud provider has its own security tools and approaches, which makes it difficult to establish a consistent security posture across clouds.

In addition to implementing security controls in individual clouds, enterprises must also secure communication between clouds and their respective workloads, applications and end users. For example, a cross-cloud application delivery controller can centralize how traffic flow is balanced and secured across various apps and services. Another example is the introduction of a global namespace through a network service mesh that centralizes global policy enforcement across clouds.

SASE provides a platform where the core components of a Zero Trust architecture (ZTA) are available in multiple clouds, ensuring the edge and corporate devices connecting to services and apps are secure. In addition to ZTA, organizations should monitor endpoints, workloads and containers with NDR, EDR, NGAV and runtime protection.

Multi-cloud security services offer cohesive visibility, context and control by providing a single interface. Multi-cloud security services also offer capabilities not available natively by a public cloud provider to perform security operations. **With an increasing number of government entities rolling out strict software supply chain compliance mandates, multi-cloud security services are a necessity.**

Multi-cloud services



[Multi-cloud application services](#)



[Multi-cloud infrastructure services](#)



[Multi-cloud security services](#)



[Multi-cloud end-user services](#)



[Multi-cloud data plane services](#)

Multi-cloud end-user services



Regardless of the number of employees, applications and external users you have, providing the correct level of access to applications and data by role is difficult when running in multiple clouds.

One example of end-user services is unified endpoint management. All enterprise devices must be able to be managed easily and from a single interface. Device management should include technology that allows employees to perform their job functions anywhere with any device that is aligned with enterprise compliance and security policies. Doing so gives employees flexibility but also ensures that the enterprise can quickly grant and remove access to individual cloud platforms or services.

Another example of end-user services are virtual desktops. These provide a uniform way for enterprise employees, especially administrators and developers, to have quick and secure access to the systems that they work with most routinely. Virtual desktop solutions not only provide access to systems in multiple clouds, but they can also be implemented in multiple clouds and managed from a single location via multi-cloud end-user services to meet compliance or performance requirements.

It is possible to leverage separate identity and access management solutions based on workload location or cloud platform, but in doing so an enterprise loses the flexibility and velocity provided by multi-cloud end-user services.

Multi-cloud data plane services



Traditionally, the availability of applications or data analytic services has been restricted to a single data plane contained within a single public cloud provider, private cloud or edge location. As organizations adopt multi-cloud services, they have the opportunity to run workloads in the multi-cloud data plane.

Multi-cloud services require a set of APIs at the infrastructure layer or cloud layer to function, regardless if the cloud is private, public or edge. When a multi-cloud service calls to the infrastructure or cloud API, it invokes a process that deploys an object to the multi-cloud data plane. While some workloads in the multi-cloud data plane might be distributed applications that are architected for the speed, latency, resiliency and other factors necessary to run in multiple clouds, an object in the multi-cloud data plane does not need to be a member of a distributed application to take advantage of the multi-cloud data plane. **Native cloud data planes no longer exist when an organization transitions to using multi-cloud services that support workloads in the clouds that they are consuming.**

An example of workloads running in the multi-cloud data plane is a hybrid application that runs in multiple cloud provider environments. Part of the application might run in the private cloud, and the other services are deployed in a public cloud. When this application is connected through a service mesh, the multi-cloud data plane is where the containers, services and data exist and where user interaction occurs.

Multi-cloud services



[Multi-cloud application services](#)



[Multi-cloud infrastructure services](#)



[Multi-cloud security services](#)



[Multi-cloud end-user services](#)



[Multi-cloud data plane services](#)

Benefits of multi-cloud services



Reduce operational overhead by managing applications and infrastructure with the same toolsets across clouds



Create “skill portability,” enabling developers and operators to use the same skills across multiple cloud platforms via consistent services and APIs

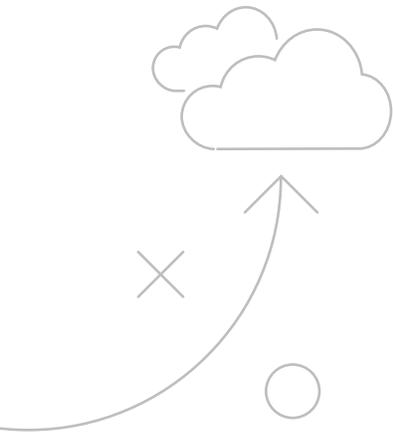


Improve observability at all layers consistently across clouds, which in turn improves application performance and security



Enhance security posture by leveraging ZTA and secure software supply chains

Strategy for adopting multi-cloud services



The velocity and consistency at which developers can build and deploy applications in multiple clouds correlates to their skillsets and the tools that they use. Multi-cloud services can provide flexibility to the organization and create the foundation for an excellent developer experience. Understanding developer needs regarding training, processes and tools is critical when creating a multi-cloud strategy. It is important to determine how much of the underlying cloud provider technology and implementation details the developer must be aware of to perform tasks. This information might guide the organization to focus on a specific multi-cloud service first.

For example, if it is best for the developer not to focus on the underlying cloud or edge capabilities, implementing a multi-cloud infrastructure service with Kubernetes might be the first action to take. Another solution for the same use case is to leverage a PaaS for standardization and infrastructure abstraction. Other developers might need to consume public cloud services directly, so providing self-service developer tools with the ability to create and consume flexible CI/CD pipelines could be the best solution in these situations.

DevSecOps is another area that needs tight integration. DevSecOps methodology requires that management and response to an issue or outage is automated through a set of tools. DevSecOps tools should either exist in a multi-cloud service or integrate with the multi-cloud service or single cloud platform APIs. It is possible to leverage individual cloud provider DevSecOps tools, but this choice should be intentional because it limits the ability to easily expand to another cloud platform.

Conclusion

Choosing one cloud provider or another no longer needs to introduce added complexity for IT. Instead, organizations can focus on the innovations that improve customer experience and drive growth, built on any cloud they choose.

Developers do not want to think about underlying infrastructure, either at the container level or cloud level. And operators want consistent visibility over deployments, spend and security configurations for every cloud. Meeting everyone's needs requires software interfaces that operate across clouds, abstracting away the complexity of the underlying multi-cloud infrastructure.

VMware believes multi-cloud complexity is best addressed through a rich layer of multi-cloud services that equip enterprises with a broad set of capabilities to build, run, manage and secure apps consistently across clouds. By abstracting the complexity of multi-cloud, businesses can reach new levels of agility and growth without compromising sovereignty or security.

