



**Point of View**

---

# **Ten Trending Topics** in the Hybrid/ Multi-Cloud Space

Author: Sakthivel Sabanayagam

# Introduction

**Hybrid and multi-cloud computing is becoming the norm in the industry and is here to stay for the foreseeable future.**

**With rapid transitions into hybrid and multi-cloud environments, enterprises and organizations of all sizes are at somewhat of "a cloud computing" crossroads.**

Most companies already use more than one cloud, including their private data centers. A recent study of 750 odd organizations confirmed that they use around 2.7 clouds per organization. As such, it is safe to say that multi-cloud environments are becoming widespread. It has become the default model of cloud computing across the globe.

With it, as new service providers and services become available, your ability to understand, optimize your infrastructure, and align it to your business goals has become increasingly important.

---



# Table of Content

<b>1.</b> Enterprise infrastructure defaulting to hybrid/multi-cloud	Page <b>04</b>
<b>2.</b> Increased adoption of hybrid/multi-cloud management tools	Page <b>09</b>
<b>3.</b> Rise in demand for workload portability in cloud platforms	Page <b>13</b>
<b>4.</b> Increased usage of Kubernetes Cluster Federation	Page <b>17</b>
<b>5.</b> The rise of multi-cloud data integration & exchange patterns	Page <b>21</b>
<b>6.</b> Hybrid cloud DevOps: Shift from VM and containerized DevOps	Page <b>25</b>
<b>7.</b> Federated architectures within single pane observability	Page <b>29</b>
<b>8.</b> Central policy enforcement need across cloud storage systems	Page <b>32</b>
<b>9.</b> Rise of FinOps from multi-cloud management platforms	Page <b>35</b>
<b>10.</b> Architect for hybrid/multi-cloud	Page <b>38</b>
<b>11.</b> Conclusion	Page <b>41</b>
<b>12.</b> What can LTIMindtree do for your company?	Page <b>42</b>

01

## Enterprise infrastructure defaulting to hybrid/multi-cloud



In 2019, only half of enterprises utilized hybrid cloud strategies. As we reach the middle of a new decade, that number is much higher at almost 90%. Furthermore, businesses see their hybrid/multi-cloud usage expanding in the coming years. Because this transition happened so quickly, many companies lack the expertise and tools to manage their hybrid/multi-cloud environments. Poor hybrid/multi-cloud management weighs down IT, DevOps, and security teams.

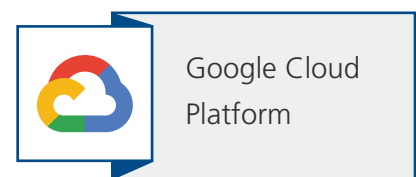
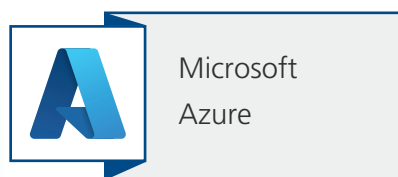
As recently as a decade ago, these same enterprises debated about adopting multi-cloud infrastructures; however, it is the default today. Because almost all companies have adopted multiple clouds in different parts of their organizations, we have moved past hybrid clouds as an option and have new pillars to manage.

**As such, today,  
IT enterprise teams know that we currently aim to understand that:**

- ✓ You are already in a multi-cloud world. It is no longer hypothetical. **How** do you take full advantage of this ecosystem that you're most likely already paying for?
- ✓ **How** do you convert your company structure into a hybrid where you can have synergy and synchronization between these multiple cloud platforms?
- ✓ **How** do you drive that synergy?
- ✓ **How** do you rebalance your workloads between these multiple platforms?

## Hybrid cloud vs. multi-cloud defined

Hybrid cloud combines computational storage IT services from on-premises data centers, private cloud services, and public cloud services like:



These hybrid clouds support orchestration and management of your various services as a single cohesive entity, with a focus on achieving optimal workload placement, portability, and flexibility.



On the other hand, multi-cloud is simply a mix of resources and services from more than one public cloud service. It may not necessarily be orchestrated and managed as one cohesive entity. A hybrid cloud can be classified as a multi-cloud when more than one public cloud service is combined with private cloud/on-premise resources.

## Why is hybrid/multi-cloud becoming the default model?



### Specialization of services

One of the reasons the industry has welcomed hybrid/multi-cloud management is that many specialized services are available on specific cloud platforms.

For example, if we're talking about analytics, Azure has Synapsis and so many other services which strengthen their analytics as a service. Likewise, if you look at GCP and AWS, their data and AI solutions are competitive. Of course, on-prem data centers that companies run by themselves specialize in certain things tailored to their respective organizations.

These, among others, always have a certain efficient model. One may have cheaper hardware, and another may have a low latency setup, while another may just be faster. Their differences can go on and evolve as quickly as each fiscal quarter.

As each of these has its own strengths, this hybrid cloud model helps your IT team leverage multiple cloud service providers, including your on-prem. Armed with this infrastructure, your team will efficiently achieve your business goals and thereby enhance its business value.



## Increased demand for dynamic workloads

Additionally, since several business interactions occur online, enterprises have expanded their cloud infrastructures to meet new demands. The hybrid/multi-cloud model's flexibility, functionality, availability, scalability, and cost-saving potential allow enterprises to adapt it for their needs quickly.

For example, different cloud providers have varying strengths. A company might use Azure Synapse Analytics for analysis and AWS for cost savings. Other hybrid cloud models have become the best way to get the most bang for your buck from any one of the cloud platforms.



## Scalability flexibility

Private clouds are necessary for many enterprises, but scaling them is costly and takes additional hardware, maintenance, and monitoring. Companies with such heavy private cloud and on-premise infrastructure look at the public cloud as an option, to speed up the transformation and save IT costs on workloads—making hybrid cloud structures attractive.

Ultimately, hybrid/multi-cloud models increase enterprise agility, independence, and capability to meet consumer demands while simultaneously separating private and public data types. It helps them avoid being vendor locked, i.e., locked into one cloud service. For this reason, it is clear why the hybrid/multi-cloud model has become the default for enterprises.

Lastly, when poorly planned, hybrid/multi-cloud models can be complex, cause operational issues and incur unforeseen costs. When thinking about adopting a hybrid model, it's essential that the planning for the hybrid cloud is customized to the organization's business objective and technology vision.



## High availability, failover, and Disaster Recovery (DR)

Enterprises must also consider multi-cloud disaster recovery strategies for critical workloads instead of cross-region DR.

For example, if the functioning "on-prem" data centers are disrupted, then applications could fail over to the cloud and continue to run without any disruption to the users. A single region failure might cause a huge load on another region and trigger a domino effect.



## Edge workloads are increasingly becoming a critical part of a multi-cloud/ hybrid infrastructure

Cloud-managed edge computing devices bring the computing power of the public cloud to the private cloud, closer to where the IoT devices reside, including data residing in applications, connected devices, and mobile consumer services.





02

## Increased adoption of hybrid/ multi-cloud management tools



There are some already clear advantages to using a hybrid or multi-cloud infrastructure. The question now is how best to manage these services. With this, organizations are finding the need to take a step back and ease out optimal configurations and management schemes for their infrastructure.

## The Problem

When you have multiple cloud solutions (AWS, Azure, GCP), your team will need specialists or Subject Matter Experts (SMEs) for each one of these cloud platforms. These people are certified in that particular cloud domain to help build multiple smaller teams which specifically work on each one of these clouds.

While these teams are great, they tend to become siloed. For example, there can be no synergy between an Azure cloud operator and an AWS cloud operator. This setup can become challenging to manage.

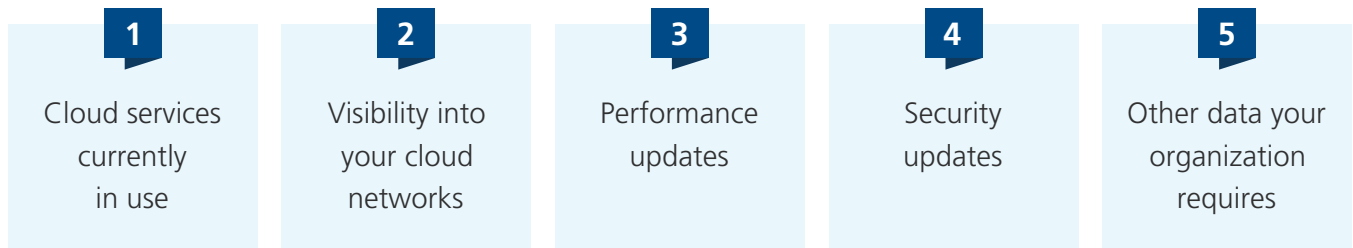
What you need is a tool that will help your teams with Multi-Cloud Management. This is called a hybrid/ Multi-Cloud Management tool. It is also called MCM, Multi-Cloud Management.

## The Solution

Multi-Cloud Management tools enable organizations to experience the full benefits of multi-cloud adoption and remove the need for additional specialists to manage resources across cloud platforms. A tool-driven management approach streamlines processes, freeing developers time to work on business-focused projects and enhance business value.

## Multi-Cloud Management (MCM): What is It?

Multi-Cloud Management tools integrate the cloud services being used by a company and serve as high-level organizational tools that optimize:



Multi-Cloud Management can be done in-house, through a third party, or with a combination of the two. The obvious drawbacks of creating and maintaining Multi-Cloud Management tools are hardware breakdown, constant developer education, maintenance, and above all, costs. For these reasons, many organizations opt to work with flexible third-party services that offer pre-made multi-cloud management tools.

**Some companies that have developed MCM tools include:**



IBM's MCM



VMware's vRealize



Google's Anthos

It is not a fully functional cloud management tool, but it does a little bit of coordination across clouds.



Red Hat's Advanced Cluster Management (ACM)



Appvia



Azure Arc

It is not a fully functional cloud management tool, but it does a little bit of coordination across clouds.

At present, Multi-Cloud Management tools are not heavily used. So approximately 40 to 48% of the companies have some kind of a tool in their hand. It can be homegrown, or it can be a commercial tool. Either way, the market is ripe for it because of its benefits.

## Benefits of Multi-Cloud Management MCM

Overall, Multi-Cloud Management tools help organizations:

1

Reduce resource provisioning

2

Serve their needs in an al-a-carte manner, saving them money

3

Analyze multi-clouds as one system

4

Enable better workflow automation between services

5

Be free from the stress of managing data

6

Boost innovation creativity in new ways previously unseen

7

Maintain flexibility because they will not have to be locked into one cloud vendor

Once configured, this model can be adjusted to operate between multiple clouds to automate the compatibility process and, ultimately, free developer time. In addition, the universal error reporting and analysis functions mean it's easier to pick up a cloud service for a particular project or change a major one when needed.

It's also easier to detect errors by using this top-down approach. Additionally, MCM works in the customer's favor because multiple clouds create redundancy that allows circumvention of disasters.

In short, these tools reduce your risk of falling behind the competition, downtime, costs resulting from disaster recovery and increase flexibility to scale at crucial moments. They also offer a holistic view of the usage, budget, and balance of applications. For example, logs from each cloud platform can be aggregated and analyzed to identify issues.

03

## Rise in demand for workload portability in cloud platforms



Effectively, workload portability realizes the ability to migrate elements of a business service like applications, operating systems, and other data to the appropriate infrastructure, and to serve the changing needs of business and application owners.

Even though 90% of services are commoditized to Azure or AWS, they all give equivalent services. However, each cloud platform has a specialty. For example, when it comes to low code/ no code, Power BI from Microsoft is well renowned. Analytics on Microsoft is also popular.

## The Problem

For each one of these, customers place-specific applications in certain clouds. However, managing these specializations in a single pane gets challenging and is not dynamic. For example, imagine your company is looking for a facility to move applications from one cloud to the other cloud. The customer will have to move applications again to a different cloud platform if, **for example:**

You have your analytics running on Azure, but AWS has better discounts

The company has grown and now wants to move closer to GCP strategically

Furthermore, suppose this company now brings their on-premise data centers or on-premise private cloud into the mix. In that case, they'll quickly see that they may have a reason to move away from your private cloud and into one of the public clouds because of the redundancies.

Sometimes, your company may want to do the reverse. You might want to move applications running in your public cloud, back to the private cloud. With a hybrid cloud, this is all just a few clicks away.

## The Solution

By combining both workload portability with hybrid/multi-cloud platforms' the ability to store and manage data and configurations online increases. This flexibility increases the business value and reaps many benefits.



On a well laid out Hybrid Cloud, engineers can often, with few clicks, move a workload from one cloud or on-premises data center to another, without extended disruptions to the application's availability.

However, achieving workload portability isn't as simple as writing it down once and running it anywhere a company wants--but it is doable. The complication often comes in data, environment variables, secrets, and other configurations.

## **Why workload portability and flexibility are important?**

As cloud platforms add functionality, an organization may want to switch from one platform to another. If the organization has coded its system to fit the cloud they are currently using, flexibility and portability are vital for taking advantage of other offerings and moving workloads and data from one place to another with ease.

Additionally, if your company has an application under development, it may need to be ported between development and live environments for testing.

**Flexibility and portability are solved through better architectures that:**

1

Have an emphasis on portability

2

Drive an organization to adopt flexible and modular architectures

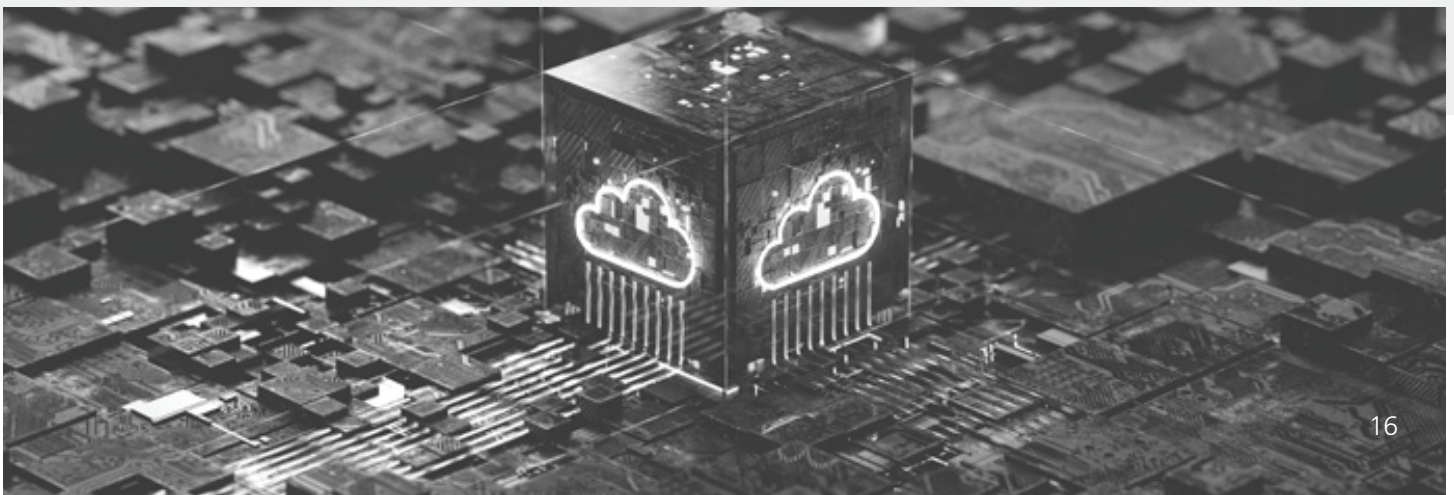
3

Prevent vendor tie-in, keeping the customer-vendor agnostic

Architectures that have an emphasis on portability allow for efficiency and continuous optimization. Performance, efficiency, and capacity become essential when an organization needs to restack and rebalance applications between clouds. In these situations, optimal porting proximity to data and latency ratios must be considered.

Furthermore, companies can consider cost savings and capacity when using higher density applications to reduce fragmentation over time. These choices will ensure that your company uses its resources efficiently.

Effectively, building architecture that supports these functionalities gives organizations a strategic advantage because they can easily use multiple cloud service providers that provide exactly what they need at competitive prices.



04

## Increased usage of Kubernetes Cluster Federation



### What is Kubernetes?

Kubernetes is an open-source container management tool that brings agility, automation, and optimization to containers in a hybrid environment. Arguably the most popular open-source after Linux, Kubernetes and its enterprise offshoots like RedHat OpenShift, VMWare Tanzu, Rancher is the most commonly used container orchestration platform.

The rise of hybrid/multi-cloud environments has increased the need for combining, orchestrating, and operating with multiple Kubernetes clusters across public and private cloud platforms. Kubernetes Federation or KubeFed has become a popular way to coordinate the configuration of numerous clusters from a single set of applications in a hosting cluster for single-pane management. When used in conjunction with multi-cloud management tools, teams can manage multi-cluster deployments, look across cloud platforms and deploy their virtual machines.

## How does Kubernetes work?

With Kubernetes, you don't operate on the individual nodes of these virtual machines but, instead, at the cluster level. With Kubernetes, you can deploy your applications on the cluster. A layer within the cluster will determine the free node and place your application in that node.

It does it automatically.

If that node goes down, it can recreate that application in some other node and bring it up for you within a few seconds. As an end-user on your IT team, you can add more nodes and increase the size of the cluster.

Kubernetes works as an abstraction on top of the lower-level nodes, VMs, or workstations.

With it, you can have Kubernetes clusters in Azure, AWS, or GCP. Each one of them might be hosting a different set of applications.

## What is the Kubernetes cluster federation?

Kubernetes cluster federation is all about bringing in these multiple clusters in different cloud platforms, including your on-prem private data center, into a single pane so that you have a much larger cluster.

Since KubeFed is an abstraction of lower-level nodes, including virtual machines, workstations, and others, its reach can then be telescoped out to the entire system allowing companies to view every piece as a cluster. So, KubeFed helps replicate clusters from one place to another more efficiently. It coordinates the configuration of multiple clusters at multiple levels.

**With this larger cluster, you can:**

**Deploy your applications**

**Observe the clusters from a single place**

**Extend the functionality of your multi-cloud management tool**

**Free up operations engineers and bandwidth by providing a single interface for managing multiple clusters across the globe**

**Improve your ability to manage multiple Kubernetes clusters and reduce the required workforce and bandwidth**

With Kubernetes cluster federation, your Multi-Cloud Management tools look at the Kubernetes cluster and the virtual machines you provision onto the cloud as a resource.

Kubernetes cluster federation offers an outstanding level of abstraction that an organization needs. Because this technology brings all your nodes together using cluster federation, you don't deal with anything in the cloud, like VMs and other SaaS services. KubeFed is another way of tying multiple cloud platforms together.

## KubeFed and containers

In recent years, its use has increased as IT has prioritized containers for more applications. Containers are convenient for workloads and data that must be ported from one platform to another. KubeFed allows end-users to target clusters rather than individual containers. This feature is also beneficial if a container fails, the cluster will reproduce it within minutes across platforms reducing time spent manually porting containers between clouds or on-prem data sources.

Ultimately, KubeFed is more efficient because your team can spread workloads between clusters. It minimizes latency by contacting the closest clusters, and it offers a good level of abstraction that an organization needs, to solve certain problems.





05

## The rise of multi-cloud data integration & exchange patterns



Organizations have had to find ways to seamlessly integrate and exchange data between environments now that hybrid and multi-cloud computing has exploded to increase agility development speed and meet new customer demands.

An integration architecture for multi-cloud allows companies to implement multi-cloud infrastructures in any way needed. Without data exchange between clouds, they cannot be used for redundancy purposes, disaster recovery, or sequential operations between clouds. This makes multi-cloud integration and exchange necessary for companies that want to optimize their multi-cloud architectures for best performance and cost savings.

## Why has multi-cloud data integration and exchange become so important?

When you have your applications running on multiple cloud platforms, they need to talk to each other. For example, you might have a presentation layer running in AWS, your back-end services running in your on-prem data center, some other service running in your Azure or GCP. For whatever enterprise you work in, they need to talk to each other for you to have the meaningful business function they'll deliver.

They usually talk to each other directly, through peer-to-peer communication, or they may have to go through some hub. It is traditional to use enterprise service buses and API gateways to tie multiple services together. With this route, they multiplex, i.e., the communication will hit the hub and reach out to anybody afterward.

This system is multifaceted and delicate. With the rise of multi-cloud, there are different ways that you need to solve the problem now while the clouds continue to grow.

There is no quick fix because the problem is not simple.

### **In the past when you integrated within a single cloud platform or a single data center:**

- The communication was faster.
- The data exchanges were more reliable.
- You were only using a particular cloud platform.
- It was secured, and there were no charges usually for the data movement.

**However, with the growing need for multi-cloud spaces:**

- ✓ Data movement between AWS, Azure, and multiple clouds has to go through the internet.
- ✓ You are most likely going to be charged a fee.
- ✓ There is increased latency.
- ✓ Increased chance for cybersecurity infringement.

Your team will need more robust data protection with encryption at each step because your data might be at risk. With the rise of multi-cloud, integration patterns for multi-cloud and hybrid cloud are a critical aspect to look at.

## The main advantage of using multi-cloud integration and exchange



### Decreased latency

The primary advantage of multi-cloud integration and exchange is decreased latency. When an organization expands its cloud services and needs to port workloads and data between clouds as redundant functions, the operations will occur much more slowly than when everything was on a single cloud. One way to solve this problem is by using data distribution applications.

Data distribution applications operating as management services over multiple or hybrid cloud architectures regulating streaming, eventing, and caching solve the latency problem when working between clouds.



## Distributed data and event streaming

Distributed streaming platforms enable efficient communication and data exchange across cloud platforms in a Hybrid/Multi-cloud setup. They help organize data from different sources and streams to various other subscribing applications and services across cloud platforms.

There are many distributed streaming platforms on the market, including Apache Kafka. Apache Kafka is widely used due to its flexibility and fault tolerance.



## Distributed caching

Distributed caching is a distribution application that helps with crosstalk and data storage between clouds. These applications serve to connect and store cache information across clouds. As one would imagine, distributed caching cuts down on cloud crosstalk and time spent waiting for data and workflows to move from one cloud to another. Many other distributed data and event streaming applications can help with latency issues between clouds.



## Multi-cloud integration, exchange, and data security

Different cloud service providers will offer varying levels of security when using their platforms. For example, Azure provides the Azure Security Center, which integrates and protects customers on ASC, AWS, and GCP.

There are many security solutions in the market; it is best to do your due diligence when looking into them.

06

## Hybrid cloud DevOps: Shift from VM and containerized DevOps



DevOps have traditionally focused on the provisioning and development of virtual machines. This configuration makes collaboration and the use of multiple clouds or environments very difficult. Hybrid cloud computing, management tools, and the invention of containers and container orchestration applications like Kubernetes have allowed DevOps to take place on a much higher plane to address the underlying complexity. The goal of DevOps is all about building the application and creating a pipeline so that the application can go and sit into a specific physical infrastructure where it can actually run. As many organizations have, the process is multifaceted when you have multiple cloud platforms. You will need to move the workloads between these cloud platforms. Efficiency is the key.

## DevOps in hybrid cloud environments

Hybrid clouds introduce issues for DevOps when they do not have easy ways to manage changes between platforms. In a hybrid cloud environment where teams use containers and VMs to deploy workloads between clouds and hybrid cloud management tools, assessing data between environments is easier. Many hybrid cloud management tools are available, including IBM Watson AIOps, Azure DevOps, and more. This situation is more flexible for organizations because they do not have to become vendor-dependent.

This design, however, also creates more work for the DevOps team than using a PaaS environment, but a PaaS environment leaves DevOps tied into a single cloud provider.

Imagine this scenario. Your team wants to deploy your application and use Azure's ARO, but you have to provide a VM for the scripting language you use. However, additionally, AWS uses Cloud Formation to create that infrastructure. That's a lot of workloads to move.

VM and containerized DevOps are fading away because to change from AWS to Azure, you have to change your entire DevOps process again. Alternatively, you'll have to rewrite the DevOps process so that it can target Azure from AWS. Likewise, if tomorrow GCP comes out with a new update, you may have to rewrite your program. That's hardly an efficient system, isn't it? These events would cause DevOps to re-code their application to fit the new parameters, costing valuable time.

To solve that, you need to look at hybrid cloud DevOps, wherein you do not worry about the specifics of each cloud platform. Instead, you can use technologies that are typically offered by your multi-cloud management (MCM) tools. These can include IBM CloudPak for multi-cloud management, RedHat advanced cluster management, or vRealize. With these, you target your DevOps process to hit that MCM tool and let the MCM tool abstract all the complexities of placing your workload in any of the clouds.

In this route, it'll translate your script to adapt to Azure or AWS platforms so that you achieve the DevOps without having too many rewrites whenever you want to rebalance your workloads between clouds.



## Multi-Cloud Management Tools (MCM) support efficient DevOps

Your MCM tool is your best bet to manage these changes, encourage abstraction and streamline DevOps support.

### Benefits of using Multi-Cloud Management Tools in hybrid cloud DevOps include:

- ✓ The DevOps process is abstracted from the underlying diversity and complexity of the Cloud infrastructure (multiple public cloud services, private cloud). The CI/CD pipeline is simplified and avoids using multiple technology stacks from cloud services.
- ✓ Policy-driven workload placement further abstracts the CI/CD process from the complexity of scripting the placement.
- ✓ Ability to deploy parts of applications across multiple cloud platforms and their replicas for high availability and scalability.
- ✓ Single pane operations and monitoring of DevOps processes.

## MCM makes containerization an auxiliary technology

With MCM tools, containerization acts as a technology that helps the process instead of hindering it. This means that your team doesn't have to necessarily package anything differently if your application is containerized as Docker images or Rocket images. It is portable by default. It can go and run on any cloud platform. You only need to make sure about the way, the application reaches the cloud platform.

Using multi-cloud management tools for hybrid cloud DevOps is the preferred way for many teams because of the benefits they can afford by using a hybrid cloud environment. Multi-cloud management tools remove the need for re-coding to fit the specifications of using another cloud. The tools translate the script for the DevOps team, so there is little rewriting when teams need to switch or rebalance workloads between clouds.



07

## Federated architectures within single pane observability



Whether viewing multiple individual components through a multi-cloud management tool or a Kubernetes Federation, any technique that builds abstraction will need to be viewed and managed through a single pane. Observing what's happening in each one of these platforms is called single pane observability.

## What are the current tools available for Federating Architectures?

There is much to gain for an IT Team with the ability to view multi-cloud systems and the applications they contain from a single pane of glass or through one monitor.

**Kubernetes-** One of the most commonly referenced 'types of single pane observability' in workplace project management comes in Kubernetes clusters. Using the Kubernetes dashboard, a company can coordinate the configuration of multiple clusters.

However, issues in your Kubernetes infrastructure can arise from virtual machines and data storage that the Kubernetes dashboard does not monitor. So, the Kubernetes dashboard is limited in what it can do for system operators, and operators need a way to view the entire system from one frame.

For example, you could spin a new cluster only to find that it is not as stateless as you thought, the data is improperly stored, or the container image fails. In all these instances, your Kubernetes Dashboard shows no errors. Situations like this are frequent in the hybrid cloud world and, without the proper tools, take a lot of time to fix.

There are numerous ways to solve this problem. Still, a common one for hybrid cloud users is to employ management tools, thereby instrumenting your application to give in-depth monitoring information. Of course, you can always code observability into your application and services– this is not usually the most popular choice because it requires a lot of extra work.

There are many management tools available on the market. You can even hire developers specializing in Kubernetes observation in conjunction with your system. As you build out your cloud infrastructure, installing some type of single pane observation application that includes cluster visibility for the best monitoring possible is highly recommended.

## Advantages for Federated Architectures for single pane observability

With it, you aggregate:


- Logs, traces, and events from multiple sources**  
To run analytics on top of these aggregated, time-sequenced logs, traces, and events.
- Cost and time usage aspects**
- Liveness of services**  
To determine which service is down and why globally or locally.
- The entire network**  
So, it can be protected, organized, and managed together

This way, analytics can be run on them for several reasons, and the sequence of events across cloud platforms can be seen. For example, you will see where an event started in one cloud, then ported to another, and kicked off there.

You can also monitor proprietary components. With single pane observability, you have a broader look at it all. With this federation, when you have challenges in your systems, you simply need to look at observability and localization. Single pane observability is coupled with your hybrid cloud, and it combines logs, trays, and events from multiple sources. This is invaluable because you'll only look through part of the information if you don't aggregate the logs.

08

## Central policy enforcement need across cloud storage systems

A photograph showing a woman with glasses and a man in a light-colored shirt sitting at a desk, looking at a laptop. The man is pointing at the screen. There is a white coffee cup on the desk. The background shows a modern office with large windows.

With the continued rise of remote and hybrid workplace roles, central policy enforcement and security central policy enforcement have become imperative within the cloud platforms where the workloads are maintained. Central policy enforcement allows companies to create rules for performing actions and automation across managed systems. Security-focused central policy enforcement helps businesses with challenges such as cyber security breaches, ransomware, and malware have solidified the need for tight security.















As the need for security has increased, central policies ensure data security and applications on your cloud platforms to authorized users. With security inside an enterprise's cloud, reliable access will always exist for the team's files and workloads when they need to enter secure databases and other content. With all data under central policy enforcement, the risk of the unauthorized dissemination of files on the cloud platform significantly decreases.

Specifically, cloud security control is a set of security controls that protects cloud environments against vulnerabilities, reducing the effects of malicious attacks. It includes the best practices, procedures, and guidelines implemented to secure cloud environments.

## How does a business increase cloud security with Flexibility?

Flexibility in the cloud security setup allows companies to scale their capacity up or down while protecting their servers from crashing. Once established, you can control cloud security by implementing secure cloud environments. There are various pillars of secure cloud environments that are safe and keep company data safe.

 <p>Secure access controls</p>  <p>Data protection</p>  <p>Network protections</p>	 <p>Change management</p>  <p>Continuous monitoring</p>  <p>Visibility and compliance</p>	 <p>Web application firewalls</p>  <p>Identity security</p> <ul style="list-style-type: none"> <li> Enforcement of the principle of least privilege</li> <li> Encryption &amp; key management</li> <li> Curation of secrets and vaults</li> <li> Tenants isolated in a multi-tenant environment</li> </ul>
---	--	---

These pillars are foundational requirements for comprehensive cloud security that companies can utilize to provide a secure background for all team members operating on the cloud platform. Working together, the implementation of cloud security in hybrid and remote workplaces available in today's cloud platforms is not only a possibility but a reality. Finally, companies should optimize policy management for quality governance by visualizing the actualization of their engineering culture and finding the right balance for developers' flexibility within it. With an appropriate level of flexibility, governance ensures that enterprises can meet customer needs and deploy critical services uniformly and consistently.



09

## Rise of FinOps from multi-cloud management platforms



With the rise of multi/hybrid cloud computing, FinOps (cross-department cloud financial operations) has become necessary. A FinOps team usually involves the IT, production, executive, and finance departments in an enterprise. Like an enterprise's DevOps teams, FinOps works as one unit to optimize product delivery times and the market cost.

## What does a FinOps team do?

Traditionally, developers or IT approached procurement or the finance department with their needs. Needs were signed-off with little or no knowledge about the underlying operations, sometimes leading to unnecessary expenditures. A FinOps team works together to instantiate best practices for choosing and purchasing technology solutions that make all departments aware of the benefits and trade-offs of each new deal.

Most enterprises had plans to move many digital on-premise operations to the cloud and multi-cloud architectures; however, the COVID-19 pandemic pushed many organizations to implement strategies ahead of schedule. This inadvertently resulted in cloud spending influxes that affected the bottom line. Because of the recent aggressive move to cloud vendor use, FinOps most commonly refers to a company's cloud spending for their planning.

FinOps teams control spending and get all capable hands on deck when making technical financial decisions. Using this team brings transparency and predictability to the cost of producing new products or implementing new processes in virtually any area of the company.

The team's functionality does not stop there; best practices also help companies see better services to drive revenue, where cloud computing dollars have the best bang for their buck. So, FinOps teams are all about defining and finding value in the marketplace.

Additionally, FinOps teams also negotiate rates and service terms with vendors.

## Looking for the best way to form a FinOps team?

FinOps is becoming a necessity for companies of all sizes. Some of the more prominent cloud vendors offer FinOps services, including IBM MCM and VMware vRealize, to analyze company data usage and future spending. You can also hire specialists and seek out training.

LTIMindtree offers a range of solutions to help your company get cloud spending under control and better understand what future spending should and will look like. They can help reimagine entire organizations or help you find a multi-cloud architecture optimized for your needs.



10

## Architect for hybrid/ multi-cloud



Minus the buzzwords and appeal of newly minted methods, what options exist, and what decision-making process will result in optimal data storage system configurations? With the rise of hybrid and multi-cloud computing, architecture options and decisions have become important. It's clear that enterprises have spent the last year or two busily migrating applications and workloads to various cloud services. Innovation gives rise to further progress, and it can be tough to keep up with the rapid evolutions of cloud software regularly.



## What's the difference between a hybrid and multi-cloud architecture?

To have a hybrid cloud, all an organization needs is an on-premise network and associate that with a public or private cloud. A multi-cloud architecture is realized when more than one public cloud is used to run workloads. So, using these definitions, many companies are already at the hybrid-multi-cloud stage of general cloud architectures.

Your cloud architecture strategy depends on what data can leave the confines of company boundaries and the amount of work you need to offload for smooth DevOps and other operations.

## Multi-cloud architecture options

Because multi-cloud is a new term and encompasses diversified structures, it can mean different things to different companies. As such, multi-cloud configurations are entirely customizable and can be structured in different ways, including arbitrary, segmented, choice, parallel, and portable.

### **Arbitrary Multi-Cloud Model**

Generally speaking, most businesses do not aim to have an arbitrary architecture but end up with it anyway. In an arbitrary multi-cloud architecture, the distinction between clouds and their functions is unclear. It is most often the result of "use due to external factors," including adding cloud vendors to take advantage of a well-marketed sale.

### **Choice Multi-Cloud Model**

The choice architecture allows workload deployment across chosen clouds. It reduces the chances of vendor lock-in and enables DevOps to transition between proprietary and other services easily.

### **Segmented Multi-Cloud Model**

There is more of a method to the madness in the other architecture types— a segmented multi-cloud architecture results from employing specific clouds to run specific workloads. Segments are commonly categorized by work, type of data, or product type. With this cloud model, decision-making teams need to remain focused on their computing needs rather than being coerced into purchasing products they don't need.

### **Parallel Multi-Cloud Model**

In a parallel architecture, businesses deploy workloads and applications to multiple clouds. The structure yields redundancy and computing power above what a single vendor might offer.

### **Portable Multi-Cloud Model**

Lastly, the portable architecture gives you all the benefits of the previously mentioned architectures. They provide the customer with the highest range of flexibility and optimization when carefully planned.

## **What multi-cloud architectures are right for you?**

If your cloud architecture is growing increasingly complex, LTIMindtree is the partner for you. LTIMindtree offers a range of solutions with all the top cloud vendors and has access to niche platforms and space providers.

## Conclusion

To help you stay ahead of the curve and increase efficiency while managing your multi-cloud space, you must:

- ✓ Formulate a team of people throughout the company that can knowledgeably represent each department's cloud needs.
- ✓ Coordinate these automation changes to your company's business and technology strategy.
- ✓ Operationalize the new systems.
- ✓ Maintain them.

### **How can you do it?**

Staying on top of the latest developments without dedicating each task to employees is the future of enterprise efficiency. Seeking help from B2B businesses such as LTIMindtree, which specializes in cross-cloud architecture and solutions, is a great way to maximize your hybrid and multi-cloud efficiencies.

# What can LTIMindtree do for your company?

- ☑ LTIMindtree can help with your search for the optimal multi-cloud configuration.
- ☑ We help enterprises find and tailor cloud infrastructures to give them the performance they need to unlock their full potential.
- ☑ We work with companies to help organize and put an architecture in place that allows them to take full advantage of multiple cloud service providers. Find out how much your company can save with a streamlined architecture that enables you to use the services you need without hiring specialists for each service.
- ☑ Federating your architecture is not difficult for most but implementing KubeFed most efficiently can be difficult. LTIMindtree helps organizations achieve efficient architectures and aids with installing and making decisions about the right processes for your clusters.
- ☑ LTIMindtree also offers full integration and exchange services. It will work with your company to ensure your data, events, and cache are being managed in the best way possible.

- ☑ **Multi-Cloud Management tools help remove production bottlenecks organizations might encounter as they use additional cloud providers. LTIMindtree tailors hybrid clouds and development environments that fit your company's needs.**
- ☑ **We help companies structure their architecture for single pane observability and overall efficiency.**
- ☑ **LTIMindtree offers a range of solutions to help your company get cloud spending under control and better understand what future spending should and will look like.**
- ☑ **We can help reimagine entire organizations or help you find a multi-cloud architecture optimized for your needs.**

LTIMindtree provides many digital transformation services to enterprises and businesses such as banking and finance, insurance, energy, utility, CPG and retail, hi-tech, life science, healthcare, media and entertainment, and manufacturing. Our services include:

- ☑ Consulting (reimagining enterprises, experience design, privacy, OCM, blockchain, and next-generation workplace)
- ☑ Cloud and infrastructure consulting, agility, assurance, and quality engineering
- ☑ Cyber defense resiliency
- ☑ Application management (applications development and maintenance and system integration)

If your enterprise is running into inefficiency issues with your current cloud architecture or believes there is a better way to compute, consider setting up a consultation with LTIMindtree. You can find several examples of work we have done through the case studies available on our website.

## Author Profile



### **Sakthivel Sabanayagam**

Head of Cloud Native  
Engineering & Consulting

Sakthi is a principal cloud strategist and a hands-on solution architect at LTIMindtree, with in-depth knowledge and advisory of cloud advisory services with cloud-native engineering across technologies and business domains.



**LTIMindtree** is a global technology consulting and digital solutions company that enables enterprises across industries to reimagine business models, accelerate innovation, and maximize growth by harnessing digital technologies. As a digital transformation partner to more than 700+ clients, LTIMindtree brings extensive domain and technology expertise to help drive superior competitive differentiation, customer experiences, and business outcomes in a converging world. Powered by nearly 90,000 talented and entrepreneurial professionals across more than 30 countries, LTIMindtree — a Larsen & Toubro Group company — combines the industry-acclaimed strengths of erstwhile Larsen and Toubro Infotech and Mindtree in solving the most complex business challenges and delivering transformation at scale. For more information, please visit [www.ltimindtree.com](http://www.ltimindtree.com).

LTIMindtree Limited is a subsidiary of Larsen & Toubro Limited