



THE STATE OF  
**MULTI-CLOUD ARCHITECTURE**

**TURBONOMIC** SURVEY REPORT

---

# THE STATE OF MULTI-CLOUD ARCHITECTURE

## PART TWO

---



## CONTENTS

|   |           |
|---|-----------|
| <b>Executive Summary</b>                | <b>3</b>  |
| <b>Demographics</b>                     | <b>6</b>  |
| <b>On The Multi-Cloud Landscape</b>     | <b>12</b> |
| <b>On Business Requirements</b>         | <b>15</b> |
| <b>On Technology Requirements</b>       | <b>18</b> |
| <b>On Vendor Selection</b>              | <b>20</b> |
| <b>On Implementation and Management</b> | <b>22</b> |

# THE STATE OF MULTI-CLOUD ARCHITECTURE

## PART TWO

---

### Executive Summary

In the first report in this series, The State of Multi-Cloud Architecture, Part One, we examined the existing and anticipated adoption of various cloud approaches; the business needs driving adoption; and the concerns preventing organizations from proceeding with a cloud strategy. In this report, The State of Multi-Cloud Architecture, Part Two, we address the question what can go wrong, and what does go wrong when implementing Multi-Cloud? It will investigate implementation challenges, vendor selection, and technical and business requirements facing organizations adopting multi-cloud; as well as how organizations could overcome the challenges they face through an autonomic approach to multi-cloud.

This survey report will look at why organizations choose multi-cloud; what technological capabilities they make their priority; who they choose to provide those capabilities; and what challenges they will face both when implementing and managing their multi-cloud architecture.

The key findings of the report are that organizations' stated business drivers for a multi-cloud strategy are not supported by the business and technical requirements they have for choosing multi-cloud vendors. For instance, while reducing capital and operational expenditure are priorities for businesses, the vast majority focus on pricing when choosing a vendor, which will not on its own

guarantee the best economic results from a multi-cloud strategy.

Possibly due to this narrow focus, organizations are also encountering multiple challenges when implementing and managing multi-cloud, with the majority lacking the skills to solve those challenges in-house. Ultimately, organizations cannot satisfy their key business drivers if these challenges are not met, and will not realize the economic benefits of multi-cloud. As the number and complexity of multi-cloud environments increases, it is no longer enough to manage the environment using traditional methods. Instead, organizations need a self-organizing, self-managing autonomic system to ensure their multi-cloud environment is always supporting business and economic demands.

### Definitions

To delineate between virtualization models and ensure consistent interpretation of the data herein, a set of definitions for the terms used throughout this report is in order. These are:

**Virtualization** is a construct that practices installing a hypervisor on traditional x86 servers such that multiple virtual machines running diverse operating systems may run on those servers.

**Private Cloud** is a construct that practices using virtualization plus automated provisioning and orchestration to deliver a cloud service model – Infrastructure-as-a-Service (IaaS) or Platform-as-a-

# EXECUTIVE SUMMARY

Service (PaaS) – on infrastructure assets owned and maintained by the organization delivering these services.

**Public Cloud** is a construct that offers IaaS and/or PaaS cloud service models as a utility, on infrastructure assets that can be accessed by any paying customer with appropriate credentials. Although the best-known examples include Amazon Web Services, Google Cloud and Microsoft Azure, there are many smaller public cloud providers offering similar services.

**Infrastructure-as-a-Service (IaaS)** is a cloud service model that enables end-users to provision virtual computing resources on-demand through a self-service portal. The initial and ongoing maintenance, including anti-virus, monitoring and patching is the responsibility of the end-user.

**Platform-as-a-Service (PaaS)** is a cloud service model that enables end-users, usually application developers, to input or upload source code into a command-line interface (CLI) or destination folder, and the platform automates the deployment, capacity provisioning and orchestration of the application instance designed by the developer.

**Multi-Cloud** is an architecture whereby an organization delivers application services out of multiple virtualized, private cloud and public cloud availability zones without actively porting workloads between these zones.

**Hybrid Cloud** is an architecture whereby an organization delivers application services out of multiple virtualized, private cloud and public cloud availability zones and actively ports workloads between these zones for reasons including cost, performance and availability.

**Autonomic** describes a system that reacts automatically, with no manual intervention required, in response to external stimulus or other changes, in order to regulate lower-order processes (like heart rate in the human autonomic system) and ensure that the system is performing as appropriate for its environment. Such a system enables focus on higher-order functions.

## The Multi-Cloud Survey Series

This installment is the second in a series of four survey reports, listed below, being published over the course of 2016. Collectively, they illustrate the reality facing IT organizations today. All reports will be available for download on [turbonomic.com](http://turbonomic.com); Volume one is available [here](#).

1. *The State of Multi-Cloud Part One*
2. *The State of Multi-Cloud Part Two*
3. *The State of Open Source*
4. *The State of Hybrid Cloud*

# METHOD

## Purpose

Analysts and experts predict that the end state of the current IT evolution is an architecture called hybrid cloud, whereby an organization delivers application services out of multiple virtualized, private cloud, and public cloud availability zones and actively ports workloads between these zones for reasons including cost, performance, and availability. Despite this projection, few if any organizations have achieved this end state.

Verizon and Turbonomic have collaborated on a survey series, which will collectively establish a baseline for where organizations currently reside along the journey to hybrid cloud, identify the challenges they face, and establish a framework for proceeding on the path to hybrid cloud over the next three years.

Our goal is that the results will instigate a data-driven conversation across the broader virtual and cloud community.

## Sample

The data in this report were collected through an online survey conducted from April 22, 2016 to July 5, 2016. The 1,821 survey respondents came from across the Enterprise IT and data center landscape. Respondents are of 18 years of age and older. In order to reveal the range in characteristics, respondents were identified demographically by their business and environment characteristics, such as role, business type, hosts in production and virtual machines in production.

This sample represents organizations spanning SMB to large enterprise, with various roles and responsibilities in those organizations.

## Procedure

This survey recruited participants from an internal email database and on the social media platforms Facebook, LinkedIn and Twitter. Participants were given an

opportunity to win a \$100 gift card (USD) by entering their email address and contact information at the completion of the survey. Additionally, participants were given the option to participate in a follow-up conversation centered on Verizon's Intelligent Cloud Control (ICC) service subsequent to completing the survey.

While the survey successfully recruited a significant sample size, the distribution of the sample weighs highly in System Administrator as a role and was well distributed across business types. The survey itself was designed by product management, product marketing and subject matter experts.

## Survey Flow

Progression through the twenty-three survey questions depended on respondents' adoption or planned adoption of multi-cloud. All respondents were asked the same demographic questions, but proceeded differentially based on responses. The survey took between 5 and 15 minutes to complete, depending on responses. Not all questions were mandatory, and participants could exit the survey at any time.

## Margin of Error

The margin of error on the initial participant sample is  $\pm 2.7\%$ . Due to question branching and optionality, the initially robust population of 1,821 did not respond to all questions. Additionally, responses are segmented by company size in this analysis (1-200 Employees; 201-1,000 Employees; 1,001+ Employees). Therefore, smaller sample sizes, particularly far into the question sequence, introduce a wider margin of error. Data should be interpreted with this in mind.


# METHOD

## **Citing this Survey**


We welcome your use of the results in this survey as you share insights with members of the broader IT community. Please reference Turbonomic and include our homepage URL, [turbonomic.com](https://turbonomic.com) as you do so. A downloadable version of the complete dataset is available at [github.com/turbonomic/turbonomicsurvey](https://github.com/turbonomic/turbonomicsurvey).

# DEMOGRAPHICS

Which of the following best describes your role?




Which of the following best describes your industry?




# DEMOGRAPHICS

What is your approximate company size?




What is your approximate annual Information Technology budget (USD)?






# DEMOGRAPHICS

How many physical servers are in your IT environment (All known on-premises and off-premises assets)?




How many developers does your organization employ?



# DEMOGRAPHICS


Does your company currently leverage multiple clouds to deliver IT services, or plan to in the next 3 years?

1-200 Employees




N = 531

201-1,000 Employees



N = 445

1,001+ Employees




N = 407

● Yes


● No

# DEMOGRAPHICS


Which public cloud does your company leverage or plan to leverage? (1001+ employees) n=495



Which public cloud does your company leverage or plan to leverage? (200-1000 employees) n=248

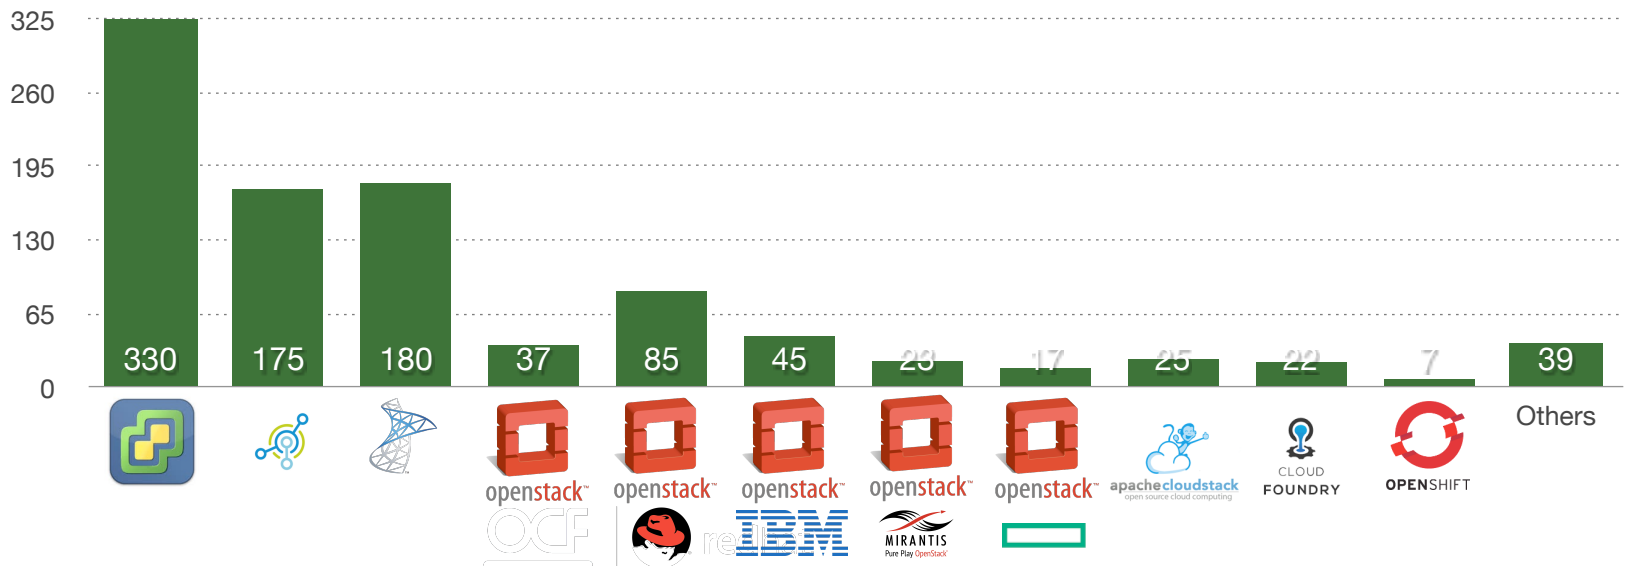


Which public cloud does your company leverage or plan to leverage? (1-200 employees) n=292




# DEMOGRAPHICS


Which private cloud does your company leverage or plan to leverage? (1001+ employees) n=495



Which public cloud does your company leverage or plan to leverage? (200-1000 employees) n=248



Which private cloud does your company leverage or plan to leverage? (1-200 employees) n=292



**ON THE  
MULTI-CLOUD  
LANDSCAPE**

---


**ANALYSIS**

# ON THE MULTI-CLOUD LANDSCAPE

## ANALYSIS


### The Future

According to the research, most organizations do not believe that public clouds will replace the need for private data centers this decade, if ever. 45.4% of survey participants say that public clouds will never fully replace private data centers. 23.1% believe that it will happen by 2025; 7.5% by 2030; and 3.2% later than 2030. Only 20.8% believe that public clouds will replace the need for private data centers by 2020. This points towards a multi-cloud future, where organizations use a mixture of public and private clouds to provide their IT services in the most economic manner.




### Multi-Cloud Adoption

Survey participants were clear in their intentions to use multi-cloud. **60.9%** of participants' companies either leverage multiple clouds to deliver IT services, or plan to in the next three years.



Respondents also show clear preferences for the vendors they will use to deliver cloud services. Of those respondents either leveraging, or planning to leverage, public clouds, 57.3% use or plan to use Microsoft Azure; 56.6% Amazon Web Services; 23.7% VMware vCloud Air; and 22.0% Google Cloud Platform.




# ON THE MULTI-CLOUD LANDSCAPE

## ANALYSIS

Similarly, there are clear preferences for respondents' private cloud platforms. A majority of respondents either leverage or plan to leverage VMware vSphere (62.1%). Other platforms are far less popular choices; the next most-leveraged are Microsoft System Center (33.0%), VMware vCloud (30.7%) and OpenStack (Red Hat) (14.2%).

As evidenced in these statistics, there are organizations using or planning to use multiple public and private clouds. As we will discuss later, organizations will need to be sure that the clouds they use meet their business and technical requirements if their strategy is to succeed; ideally acting as a balanced, autonomic system that can react to economic and other pressures.



**MULTI-CLOUD**  
**ON BUSINESS**  
**REQUIREMENTS**

---

**ANALYSIS**




# ON BUSINESS REQUIREMENTS ANALYSIS

## Key Business Drivers

The first stage of any multi-cloud strategy should be identifying the key drivers behind the move. Without knowing and understanding those drivers, an organization will struggle to choose the right technology requirements needed to support them – or vendors that can meet those requirements.

Survey participants were asked to rate the importance of business drivers as they pertained to their organizations' implementation of multi-cloud. The highest-rated business driver was the need to ensure business continuity across multiple sites, which 77.4% of respondents said was "important" or "very important". Others that were rated highly included the need to increase resilience without additional capital (74.0%), the need to reduce operational and capital expenditure (70.4% and 68.7%); the need for consistent performance of services in all geographies at all times (67.0%); and the need to support a wide range of complex services (66.8%).



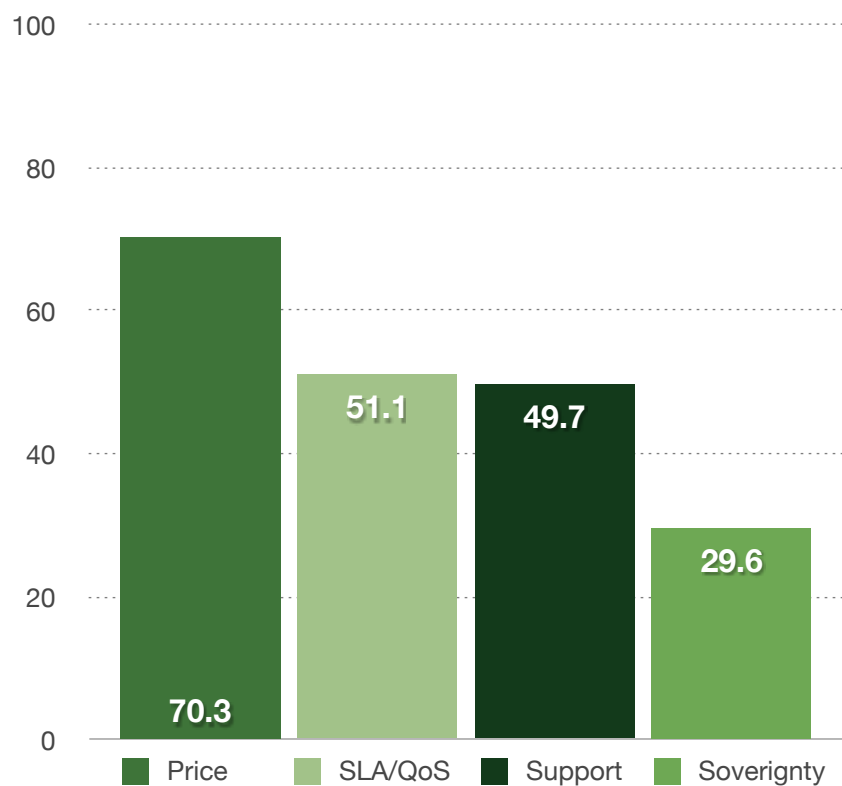
While we can see that reducing costs, whether CapEx or OpEx, is important to organizations, it is not the number one or even number two choice. It is also notable that supporting a wide range of services, and ensuring consistent performance, come very close behind reducing costs. As a result, we wouldn't expect to see organizations' choice of vendor dominated by economic requirements.

## Business Requirements When Choosing Multi-Cloud Vendors

Based on organizations' key business drivers, it might be expected that their choice of technology requirements and vendor would follow this. For instance, focusing on vendors' ability to treat the multi-cloud as an autonomic system so that it always responds appropriately to economic and other conditions to ensure consistent performance and availability.

Survey participants were asked which business requirements are primary considerations in their organization's selection of multi-cloud vendors.

# ON BUSINESS REQUIREMENTS ANALYSIS



70.3% say that pricing is a primary consideration, while other requirements that would support the other business drivers respondents identified are a much lower priority. For instance, Service Level Agreements / Quality of Service are only a priority for 51.1%, cloud provider service and support for 49.7%, and data sovereignty for 29.6%.

These are all considerations that would be important in supporting key business drivers, such as ensuring business continuity across multiple sites that may be located in multiple regions; increasing resilience; and helping ensure consistent performance of services across geographies. However, it is clear that all of these considerations are second to pricing.

## Implications Of Focus On Pricing

A focus on pricing could be explained as organizations feeling it is an important part of ensuring an overall economic strategy. However, reducing capital and operational expenditure will be reliant on many more factors than pricing, which may in fact have little impact. For instance, insufficient Service Level Agreements / Quality of Service; poor service and support leading to increased downtime; and the risk of regulatory action due to a lack of data sovereignty and compliance can all produce far higher costs than choosing a higher-priced cloud vendor.

As a result, the overwhelming focus on pricing may be causing organizations to incur greater economic costs overall than if they focused on other business requirements.

**MULTI-CLOUD**  
**ON TECHNOLOGY**  
**REQUIREMENTS**

---

**ANALYSIS**

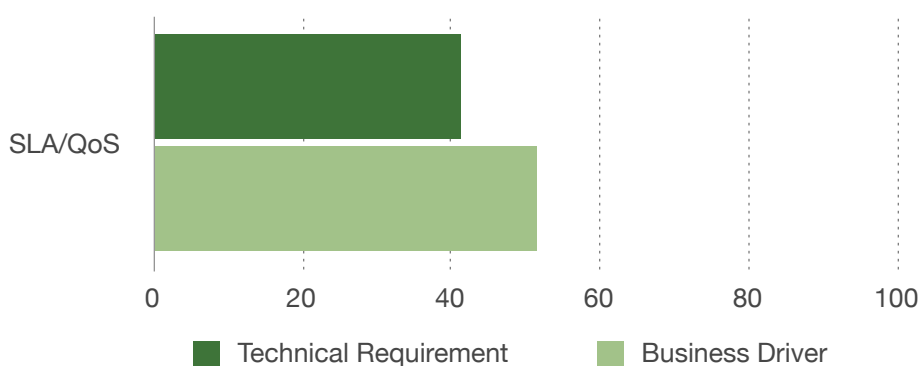
# ON TECHNOLOGY REQUIREMENTS

## ANALYSIS

### Minimum Technology Requirements

When an organization identifies the business drivers behind its multi-cloud strategy, and the business requirements of any solution, it can identify the technology requirements needed to meet these demands. For instance, an organization whose priority was to reduce OpEx might require that their multi-cloud architecture include automated orchestration, so that they are always getting the best performance for the best price.

Survey participants selected the top five minimum requirements for their organizations' selection of multi-cloud vendors. Security was the most-selected, and a top five requirement for 75.1% of respondents. Performance came second, identified by 67.7%. Despite them being critical to the majority of organizations' key business drivers, **Service Level Agreements / Quality of Service were a distant third**, chosen by 41.3% of respondents.

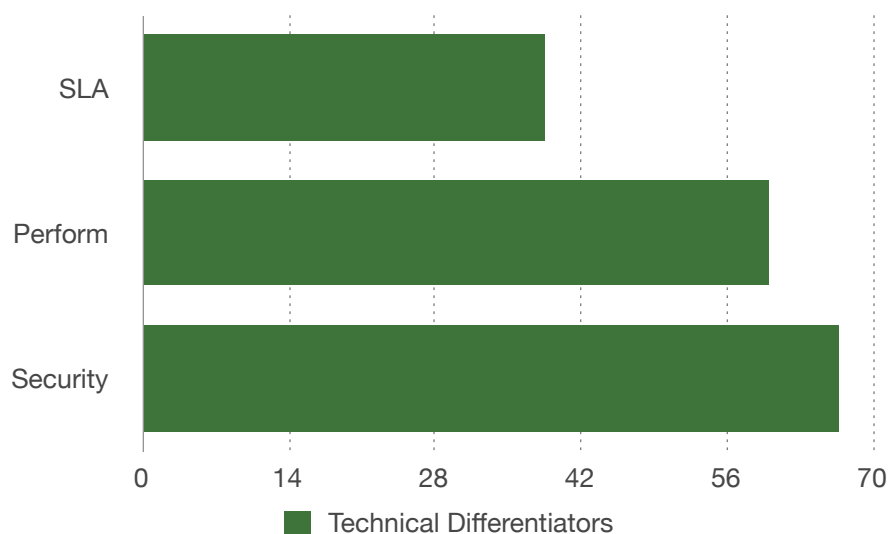


### Technology Differentiators Between Vendors

When selecting vendors based on technical requirements, there will often be vendors that provide a near-identical service or product, at near-identical cost. Survey respondents were asked to select the top five differentiators in their organizations' selection of multi-cloud

vendors. In other words, those factors that would cause them to choose one vendor over an otherwise identical competitor.

The preference for differentiators was much the same as for minimum technology requirements. The most commonly chosen differentiator was still security (66.7%), followed by performance (60.0%). Service Level Agreements and Quality of Service were again the third choice, and again were significantly less popular than the top two (38.5%).



### Attitude Towards Technology Requirements

Notably, respondents are not focusing on technology requirements that would give them greater control over and flexibility in using their multi-cloud architecture, such as: self-service provisioning; automated orchestration; change management and logging; application and OS support; virtual instance selection and variety; and native management and monitoring. This in turn will create a lack of control over and visibility of multi-cloud environments, ultimately making it less likely organizations can meet the requirements needed for their business drivers.

**MULTI-CLOUD**  
**ON VENDOR**  
**SELECTION**

---

**ANALYSIS**

# ON VENDOR SELECTION


## ANALYSIS

### Vendor Selection On Business Requirements:

When an organization knows the business and technology requirements it needs from multi-cloud vendors, and has ensured that those requirements support its key business drivers for adopting multi-cloud, it can select those vendors that best meet its requirements.

Survey respondents identified which clouds had been excluded from their organizations' multi-cloud implementation for failing to satisfy business requirements. The most common reason given for excluding a vendor was pricing. Other concerns, such as compliance, Service Level Agreements / Quality of Service or data sovereignty, were more rarely identified, and with little overall consistency.

Every cloud was excluded by at least one quarter of the respondents which considered it. The least-excluded cloud was Microsoft Azure, which was excluded by 25.9% of respondents who examined it. The most-excluded was Oracle Cloud, which was rejected by 38.3% of respondents.



The other clouds examined were IBM Softlayer (excluded by 37.4%), Rackspace (32.7%), Verizon (32.2%), Virtustream (30.9%), CenturyLink (30.3%), Google Cloud Platform (30.0%), 1&1 (29.8%), Interoute (29.4%), VMware vCloud Air (29.2%), Dimension Data (27.9%) and Amazon Web Services (27.5%).

### Vendor Selection On Technical Requirements:

Survey participants also identified which clouds had been excluded from their organizations' multi-cloud implementation for failing to satisfy technical requirements. Here there was less consensus than when rejecting vendors for failing to meet business requirements. There was no technical requirement that stood out above an others as a reason to reject vendors, and no vendor that is rejected to an unusual degree.

This suggests that pricing is still the primary reason for many organizations to select a multi-cloud vendor, above any other business or technical requirements. While this is understandable, it also increases the risk that organizations will find themselves at a disadvantage after implementation, when they realize their cloud provider cannot provide a truly autonomic system that would provide the performance and economic benefits needed to meet business drivers. This is especially true with a multi-cloud approach since, once a cloud is chosen, that is where the workload resides; meaning **the initial decision on where to place workloads is critical.**

**MULTI-CLOUD  
ON IMPLEMENTATION  
AND MANAGEMENT**

---

**ANALYSIS**


# ON IMPLEMENTATION AND MANAGEMENT

## ANALYSIS

### Challenges Implementing Multi-Cloud


So far we have examined what can go wrong during the planning stages of a multi-cloud adoption. For instance, if an organization cannot reconcile its business and technical requirements with the key business drivers behind multi-cloud, or select a vendor that will meet both its requirements and business drivers, there is a risk that its multi-cloud adoption will not provide the benefits that the business needs. However, there are further challenges that organizations can experience when implementing and managing multi-cloud.

Survey respondents ranked challenges as they pertained to their organizations' implementation of multi-cloud. A minimum of 88.6% of respondents identified one or more challenges that they ranked as at least 'somewhat challenging'. These include controlling unsanctioned cloud adoption, i.e. "Shadow IT" (79.7%), choosing which workloads should reside in which cloud (81.1%), ensuring workloads are correctly balanced between performance and cost (86.9%) and adhering to budget (87.0%).



### Challenges Managing Multi-Cloud

An organizations' challenges do not stop when it implements multi-cloud. Survey respondents also ranked challenges as they pertained to their organizations' ongoing management of multi-cloud. At least 89.6% of respondents specified management challenges that they ranked as somewhat challenging or worse. These included ensuring a consistent balance between performance and cost (89.6%), adapting to changes in demand (83.5%) and changes to business strategy (83.6%), ensuring that application performance is consistent (86.3%), ensuring that IT services are delivered on budget (88.8%), and ensuring workloads always reside in the most appropriate cloud (84.2%).





# ON IMPLEMENTATION AND MGMT ANALYSIS

## Implementation & Management Challenges

These implementation and management challenges are affecting a large majority of organizations. An organization cannot expect to solve every single one of these challenges by itself. Yet if these challenges are not addressed, the organization cannot be certain that any multi-cloud adoption will meet the needs of the key business drivers. For instance, without consistent performance a multi-cloud cannot guarantee improved resilience or business continuity.

Overcoming these challenges can prove a significant advantage to organizations looking to fulfill their business drivers for multi-cloud. For example, choosing which workloads should reside in what cloud when implementing multi-cloud is a complex decision, which also contains significant consequences for the organization. Placing the wrong workload in the wrong cloud could result in increased costs, reduced performance and resiliency, and potential compliance issues.

Similarly, ensuring a consistent balance between performance and cost is an important concern for any organization that is looking to both improve its application performance and resilience, and to reduce CapEx and OpEx. Even if an organization has the necessary skills in-house, attempting to manually balance performance and cost at all times, or to choose the location of every workload, will be a complex, time-consuming, and ultimately uneconomic process. This is especially true as the factors affecting an organization's decision will change by the day, if not faster.

An autonomic system could make these decisions based on business rules and economic demands, ensuring that workloads are placed in the best possible cloud for the organization's needs and that

performance and cost are always balanced, based on business rules, economic demands and any stimuli on the environment.

Ultimately, by seeing their multi-cloud as an autonomic system that can meet all key business drivers - such as business continuity, resilience, cost and performance - at all times, organizations can have confidence that their multi-cloud is being driven by the best strategic and economic decisions at any individual point in time.

---

# THE STATE OF MULTI-CLOUD ARCHITECTURE

## PART TWO

---



**The State of Multi-Cloud Architecture Part Two** is the second installation of this four-part survey series and investigates the implementation challenges, vendor selection, and technical and business requirements facing organizations adopting multi-cloud. It addresses the question *What can go wrong, and what does go wrong when implementing Multi-Cloud?*

### About Turbonomic



Turbonomic's autonomic cloud platform enables heterogeneous environments to self-manage to assure performance of any application in any cloud. Turbonomic's patented autonomic decision engine dynamically analyzes application demand and allocates shared resources to all applications in real time, to maintain environments in a healthy state.

Launched in 2010, Turbonomic is one of the fastest growing technology companies in the virtualization and cloud space. Turbonomic's autonomic platform is trusted by thousands of enterprises to accelerate their adoption of virtual, cloud, and container deployments for all mission critical applications. To learn more, visit [turbonomic.com](https://turbonomic.com).

### About Verizon Intelligent Cloud Control (ICC)



Verizon Intelligent Cloud Control generates intelligent decisions on which workloads to run on which public cloud services, based on your specific performance, price and resource needs – all of which can be automated to reduce your personal workload. Helping you rest at night, knowing Intelligent Cloud Control is has your apps exactly where they need to be to deliver the productivity your business needs.