


WHY YOU NEED MULTI-CLOUD AND HOW TO MANAGE IT



Mikhail Gloukhovtsev

Sr. Solutions Architect

Digital Solutions, Cloud & IoT

Orange Business Services

Mikhail.Gloukhovtsev@orange.com

Table of Contents

1. Introduction	4
1.1 Multi-cloud Concept in Cloud Computing Evolution: the Next Generation Cloud.....	4
1.2 What Is Multi-cloud?	5
2. Multi-cloud Architecture.....	8
2.1 Conceptual Multi-cloud Architecture	8
2.1.1 Requirements for Multi-cloud Architecture	9
2.1.2 Multi-cloud Architecture Layered Model	10
2.2 Network Connectivity in Multi-cloud Architecture	11
2.3 Security in Multi-cloud Architecture.....	15
3. Pros and Cons of Multi-cloud	16
3.1 Multi-cloud Benefits	17
3.2 Multi-cloud Challenges	18
4. Multi-cloud Strategies.....	19
5. Use Cases for Multi-cloud	21
5.1 Storage in Multi-cloud	21
5.2 Archiving in Multi-cloud.....	23
5.3 DR in Multi-cloud	23
5.4 DevOps in Multi-cloud.....	24
5.5 Big Data in Multi-cloud: Global Data Fabric	24
5.6 Low-code Development in Multi-cloud	24
6. How to Implement Multi-cloud.....	25
6.1 General Principles for Implementing Multi-cloud.....	25
6.2 Methodologies for Multi-cloud Migration.....	26
6.2.1 Rehosting	26
6.2.2 Multi-cloud Refactor.....	27

6.2.3 Multi-cloud Rebinding	27
6.2.4 Multi-cloud Rebinding with Cloud Services Brokerage	27
7. Multi-cloud Services Brokerage.....	28
7.1 Roles of Cloud Services Brokerages in Multi-cloud Ecosystem	28
7.2 Open Service Broker	30
8. Multi-cloud Integration and Inter-Cloud Services. Integration Platform as a Service (iPaaS)	30
9. Multi-cloud Management.....	31
9.1 Architecture and Core Capabilities of Multi-cloud Management.....	31
9.2 Automation and Orchestration in Multi-cloud	33
9.3 Workload Mobility in Multi-cloud	34
9.3.1 Interoperability in Multi-cloud	34
9.3.2 Portability in Multi-cloud.....	35
9.4 Overview of Multi-cloud Management Systems	37
10. Conclusion	38
11. References.....	38

Disclaimer: The views, processes or methodologies published in this article are those of the author. They do not necessarily reflect Dell Technologies' views, processes or methodologies.

1. Introduction

1.1 Multi-cloud Concept in Cloud Computing Evolution: the Next Generation Cloud

Many companies moving to cloud find later that they need to use both private and public clouds and different cloud providers for these cloud types. The wide variety of business requirements results in a need for various cloud offerings, each focused on a particular service set such as enterprise applications in Azure, fast expanding spectrum of services in Amazon Web Services (AWS), enterprise resource planning (ERP) in Oracle Cloud, and machine learning in Google Cloud. Regulatory requirements – for example, data locality in a given country – makes local cloud providers preferable because the applications can run locally where their data are stored. Cost optimization, availability and performance requirements are other factors contributing to selection of multiple cloud offerings.

At the same time, companies experience that separate non-integrated cloud services are not an optimal cloud strategy for their business as they create service/data silos, result in “cloud sprawl,” lead to new security challenges and increase the complexity of cloud service management.

The multi-cloud concept has been developed to address these issues. In contrast to disparate multiple clouds, multi-cloud enables unified management of different cloud services as an integrated pool for each type of resources regardless of the resource location: the next generation cloud.

This article explores the emerging multi-cloud concept, benefits of multi-cloud computing, architecture of multi-cloud solutions, implementation of multi-cloud strategy, cloud services brokerage, unified management of multi-cloud services, application and data portability across different clouds, and other aspects of multi-cloud services.

1.2 What Is Multi-cloud?

Multi-cloud is an evolutionary spiral development of cloud-enabled IT service architecture beyond hybrid cloud. Multi-cloud development is driven by a need to provide highly scalable and reliable applications to meet business goals that are difficult to achieve using private-only or hybrid cloud architectures.

There are various definitions of multi-cloud. Most of them are based on differences between hybrid cloud and multi-cloud.¹ Both terms describe an IT architecture using more than a single cloud, and some IT experts consider the differences between these terms merely semantic. However, there are some key distinctions between these two terms (Table 1). While hybrid cloud is always a combination of both a public cloud and a private cloud, multi-cloud can be any mixture of public and/or private clouds. Hence, a hybrid cloud can be considered as a subset of multi-cloud with applications deployed on both on-premises as well as cloud platforms. In contrast, multi-cloud makes no distinction between the types of cloud deployments it includes. A multi-cloud may not have a private cloud environment at all. For example, all the IT services used by a company can be delivered by AWS and Microsoft Azure clouds without any use of a private cloud. Table 1 provides a comparison of hybrid cloud and multi-cloud.

Hybrid Cloud	Multi-cloud
It is based on different deployment models and always includes both on-premises private cloud and public cloud.	A combination that includes multiple public clouds and, in many cases, private clouds as well and focuses on cloud services rather than on specific deployment models, for example, a multi-cloud consisting of several only public PaaS clouds.
One vendor may provide both cloud platforms and managing only a single provider is required.	There are usually multiple providers and managing and overseeing these providers are required.
The same type cloud service models but different deployment models, for example, using IaaS with resources from both on-premises private cloud and public cloud.	Multi-cloud virtualization: virtualization of cloud deployment models where separate infrastructure domains present a single

	virtualized set of IT services abstracted from platforms that service delivery is based on.
Modified legacy applications can be used.	A need to use cloud-native applications.
Both private on-premises cloud and public cloud serve a single purpose and the components of a hybrid cloud typically work together in the same service model: IaaS, PaaS, SaaS, etc.	Different clouds are typically used by different applications. Resource usage in the multi-cloud environment can be fragmented, as some applications need resources only from one public cloud whereas other applications utilize resources only in other public clouds.
Workload mobility between two clouds in a hybrid framework. Since compute resources perform the same function in both a private cloud and a public cloud with the same service delivery model, it is possible to manage the usage of compute resources in either cloud based on performance and cost.	Multi-cloud may or may not need workload mobility between the clouds.
Data centers are used on-premises or on private cloud (co-location).	Applications run in public clouds and the users work in a data center-less environment as they do not own/rent data center services.
Relatively simple integration of cloud management.	Heterogeneous architecture of multi-cloud makes development of unified multi-cloud management tools difficult.

Table 1: Multi-cloud vs. hybrid cloud

In a hybrid cloud environment an application can use load balancing, web and application services provided by a public cloud whereas the database and storage services run in a private cloud. Since compute resources perform the same function in both a private cloud and a public cloud, it is possible to manage the use of compute resources in either cloud based on performance and cost. In a hybrid cloud, the private cloud platform can be integrated with the public cloud to provide workload mobility. For example, disaster recovery is a typical use case for hybrid clouds. Integration of the public cloud environment with the in-house private cloud makes it possible for

operations to continue as normal after failover if a disaster occurs. Other hybrid cloud use cases include archiving, cloud bursting or application development in the public cloud with production run in the private cloud.

In contrast, resource usage in the multi-cloud environment can be fragmented, as some applications need resources only from one public cloud while other applications utilize resources only in other public clouds. In other scenarios, some applications may use resources only in the public cloud while other applications leverage resources only in the private cloud.

A hybrid cloud deployment model is often used when critical data needs to be stored in-house in corporate data centers because of regulatory or business requirements. Multi-cloud takes it to the next level when an application uses data stored in various countries to meet their data location laws. Organizations can use both hybrid cloud and multi-cloud simultaneously since they are not mutually exclusive.

Multi-cloud can be seen as virtualization of cloud deployment models where separate infrastructure domains present a single virtualized set of IT services abstracted from the service delivery platforms. In other words, multi-cloud services are presented to the users as services delivered by a single cloud platform amalgamated from different underlying platforms. To provide coherent services, multi-cloud needs to use cloud-native applications as multi-cloud is a fabric that binds together application components.

Some workloads in non-hybrid multi-cloud environments run in one cloud and others run in another cloud because some services may be available only in a given cloud; or because cost considerations, security and compliance needs limit the cloud selection; or for other reasons. While some tools are already available to manage and monitor multi-cloud architecture (Section 9), most multi-cloud environments are not unified in the same manner as hybrid clouds.

Keep in mind the key distinction between “multiple clouds” and multi-cloud. The IT architecture consisting of cloud-specific domains, each having its own siloed infrastructure and operations, cannot deliver multi-cloud benefits. What differentiates a true multi-cloud is its ability to function as a single entity across all the clouds in the multi-cloud environment.

2. Multi-cloud Architecture

2.1 Conceptual Multi-cloud Architecture

A single heterogeneous architecture is the key feature of multi-cloud. As multi-cloud environments use a broad spectrum of various deployment models, it is difficult to pinpoint a single reference architecture that can serve as a blueprint for all conceivable multi-cloud implementations. Multi-cloud architecture should be able to quickly evolve in order to support new digital initiatives, such as Big Data analytics, IoT platforms, and AI that will impose new service requirements.

While requirements for multi-cloud architecture are determined by the application portfolio that is unique for every company, some common patterns can be identified. These patterns fall into two main categories:

Patterns based on a distributed deployment of applications. The goal of these patterns is to provide the computing environment that is the best for a given application. *The partitioned multi-cloud pattern* is an example of this architecture. It enables application deployment in the optimal environment by combining multiple public cloud environments provided by different vendors. In a typical partitioned multi-cloud pattern, application A runs in one public cloud – for example, AWS – whereas application B including the SQL Server database uses Azure. The partitioned multi-cloud can also be considered as a federated multi-cloud, in which components forming an application stack can reside on different clouds operated by various providers. A microservices scale-out application environment is an example. Workload portability becomes a key requirement if there is a need to move workloads between the public cloud environments. In this case, we must abstract away the differences between the environments. I review workload portability in more detail in Section 9.3. Key advantages of partitioned multi-cloud architecture include avoidance of vendor lock-in and ability to optimize operation by shifting workloads between different computing environments.

Patterns based on redundant deployments of applications. In these patterns, the same applications are deployed in multiple computing environments to increase scalability or resilience. The business continuity hybrid pattern uses a public cloud–based computing environment for failover purposes. The business continuity multi-cloud pattern is less common than the hybrid cloud DR. In the multi-cloud DR the production environment uses one cloud provider and the DR environment uses a different cloud provider. By deploying copies of workloads across multiple cloud providers, availability can be increased beyond what single-provider multi-region deployment offers.² I discuss multi-cloud DR solutions in Section 5.3.

2.1.1 Requirements for Multi-cloud Architecture

There are various multi-cloud architecture requirements.³ The most essential of them fall into categories of security, network, and management.

1. Multi-cloud security effectiveness.

A trust-nothing security model applied to multi-cloud traffic flows should be implemented with policy-based border control, packet inspection and real-time security analytics. Intercloud data transfers should be encrypted.

2. Minimization of intercloud data transfers.

Moving large data sets between clouds or placing sensitive data in some public clouds can be avoided by co-locating private data in a co-location hub (traffic exchange point). This will provide secure, low-latency access from multiple cloud platforms. While intracloud data transfer costs are usually very low with a single cloud provider, data transfer between two clouds operated by different providers can be costly, as both providers charge for outbound traffic.

3. Ability to implement unified multi-cloud management (see Section 9).

4. Scaling geographical multi-cloud demand.

Growth in multi-cloud traffic, data volumes and processing can be managed by distributing workloads across geographically placed co-location hubs in proximity to users and cloud availability regions.⁴ Load balancing of resources allows for scaling and business continuity.

5. Simplification of multi-cloud connectivity and operations.

Multi-cloud application workflows can be optimized by placing co-location hubs close to users and clouds. Traffic should be localized with direct, secure, low-latency interconnection.

6. External business requirements.

Deploying a multi-cloud architecture may be a requirement of customers rather than an internal business decision. Indeed, some customers have a preference for one cloud provider or another because of business competition, so you may have to run their applications in the cloud of their choice.

2.1.2 Multi-cloud Architecture Layered Model

The primary building blocks for multi-cloud can be categorized into three layers (see Figure 1):

- **Foundational Resources.** The underlying compute, storage, network, and security resources provide the foundation for any workload infrastructure. The underlying network spans the data center, cloud, and anywhere applications and users reside. As a result, it should include a mix of physical and virtual devices deployed both on-premises and in public clouds. This layer also includes major management constructs.

- **Workload Management.** The workload layer includes constructs such as bare metal servers, hypervisors, virtual machines, containers, and serverless architectures as well as workload/application management platforms such as OpenStack and OpenShift and services provided by various public clouds.

- **Service Consumption.** This layer abstracts the foundational and workload management layers into a set of services for each application by decoupling infrastructure components and services. This workload abstraction is the initial step in automating workload management in multi-cloud. The multi-cloud orchestration tools must be agnostic to the underlying infrastructure.

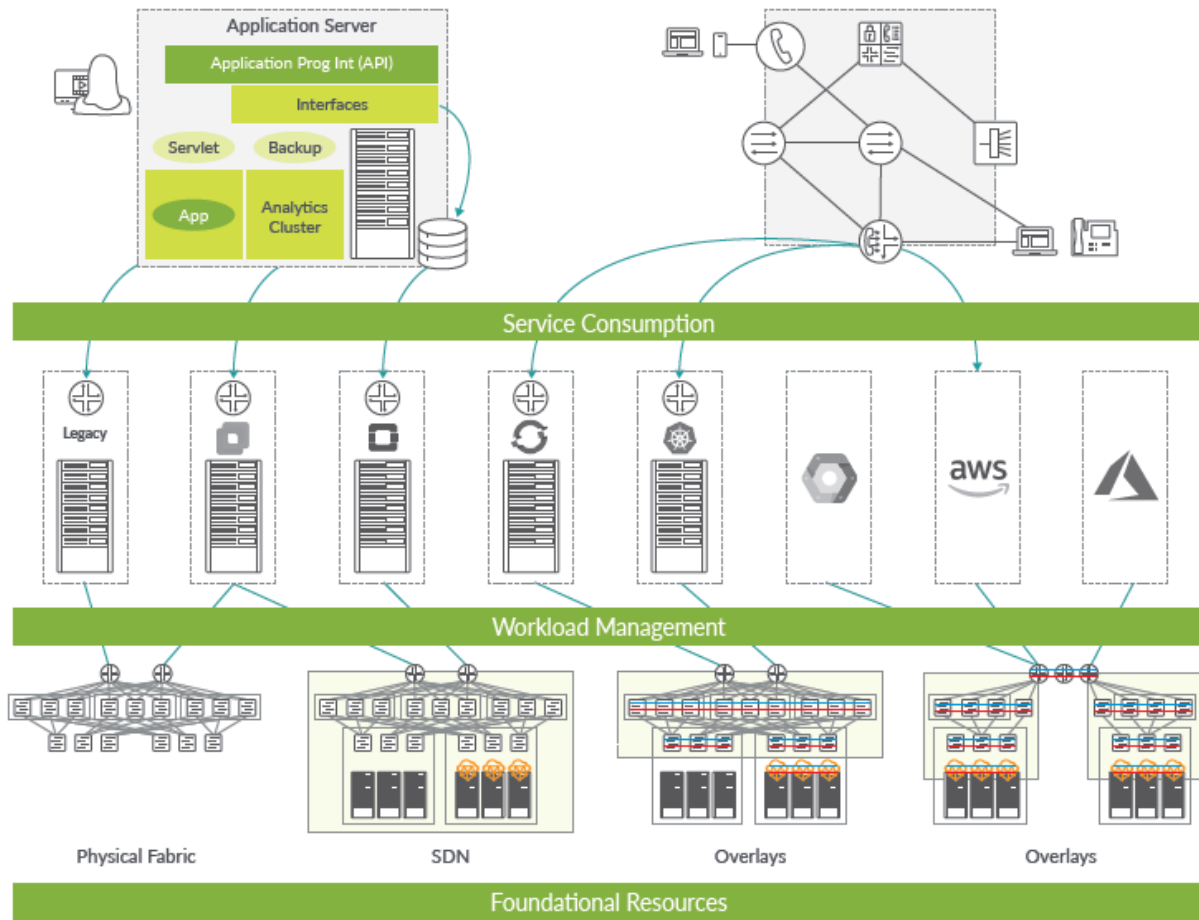


Figure 1: Multi-cloud Architecture (Ref.5)

2.2 Network Connectivity in Multi-cloud Architecture

Implementing a multi-cloud means that the last mile connection is no longer mainly from branch offices to the corporate data center. A new traffic type – intercloud traffic – goes between various public clouds and private clouds. Therefore, multi-cloud architectures require WAN connectivity. Telecommunication providers offer private network services such as Orange Business VPN Galerie⁶ and AT&T NetBond for Cloud. Orange Galerie extends the corporate VPN to the cloud via fully secured gateways with end-to-end high performance and reliability. As shown in Figure 2, cloud providers such as AWS, Azure, and Google Cloud Platform connect to the Orange Business Services network via Business VPN Galerie to make their services seamlessly available to any Business VPN customer.

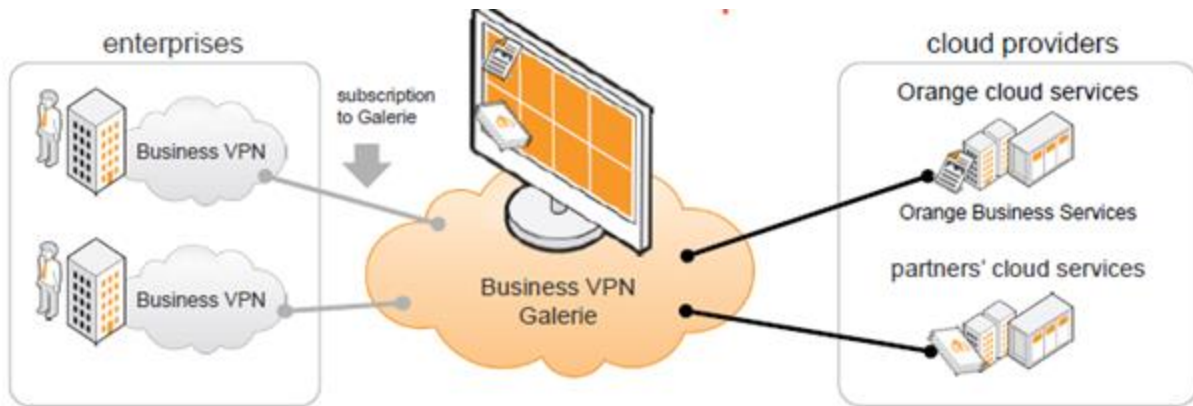


Figure 2: Orange Business VPN Galerie (Ref.6).

Orange Business VPN Galerie provides a foundation for the Managed Multi-Cloud Orchestration Platform that offers customers choices of public and private cloud services and access to various cloud services with the same user experience control. As a result, it presents a secure way to migrate applications and application data between cloud provider services.

Multi-cloud Fabrics. From a networking perspective, networking in multi-cloud should provide: (1) multi-domain connectivity; (2) multi-vendor network orchestration; (3) end-to-end visibility; (4) pervasive security; (5) reduced complexity.

To meet these requirements, multi-cloud fabrics are created. A single, common federated fabric that spans the entire multi-cloud environment uses a unified operational framework. Typically two types of fabrics are built: Ethernet fabrics that operate at Layer 2, and IP fabrics that operate at Layer 3. As more and more applications are designed to operate at L3, many organizations are choosing to operate their underlay network as an IP fabric.⁷

An IP fabric provides direct L3 connectivity to applications. If L2 is required, the standard for Ethernet connectivity over IP fabrics is Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN), which allows for building fabrics with standard protocols without using proprietary protocols. This makes the deployment of multivendor networks possible.

SD-WAN in Multi-clouds. SD-WAN (software-defined wide area network) can much better manage the complexity of connecting on-premises private clouds to various public clouds than traditional WANs.^{8,9} SD-WAN is a significant enhancement to a multi-cloud strategy¹⁰ with benefits including:

1. *Simplified routing.* With SD-WAN, it is easier to route traffic from branches and to establish local Internet breakouts. Flexible policies, defined in a single cloud management console, determine the best path every time, regardless of the underlying transport, whether it is commodity Internet, MPLS, LTE or satellite.
2. *Cost savings on bandwidth.* SD-WANs can use low-cost Internet links, in addition to augmenting MPLS circuits. As a result, organizations can avoid costly and inefficient carrier-specific lock.
3. *Reliability.* SD-WAN technology can create reliable connections to multiple clouds directly from the on-premises private cloud and branch offices.
4. *Efficiency and Security.* Network efficiency is one of the main goals of SD-WAN design. By simplifying the network, SD-WAN significantly reduces the cost of multi-cloud connections. To reduce security risks, SD-WAN segregates data based on application or source and can also restrict access to SaaS applications for certain groups of users.

Dedicated Network Connectivity to Clouds. For bandwidth-heavy applications that are sensitive to network latency, cloud providers offer dedicated network connections, which are more expensive. For building multi-cloud connectivity, Equinix offers Equinix Cloud Exchange, which is an application programming interface (API)-driven platform, with Equinix performing the “meet me” function between the customer and the cloud provider.

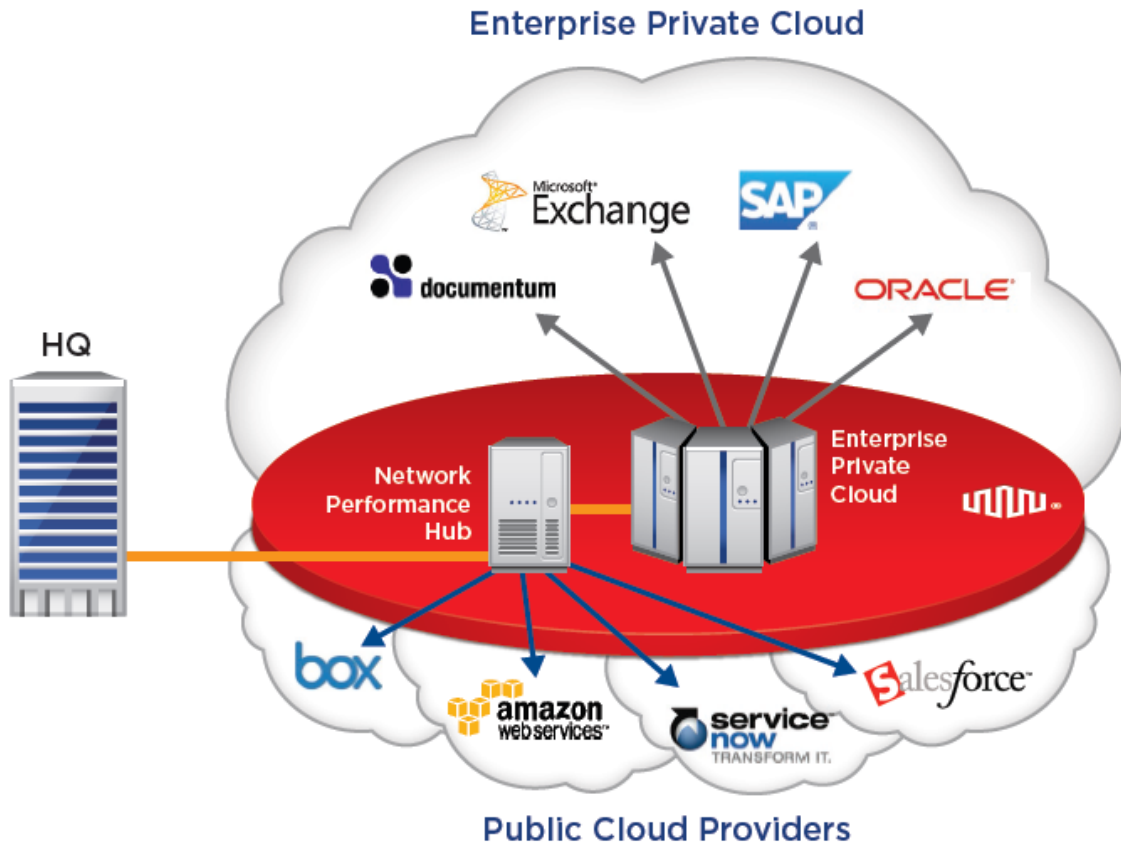


Figure 3: Equinix Cloud Exchange (Ref.4)

The Cloud Exchange platform (Figure 3) is well-suited to customers who want to take a hybrid, multi-cloud approach, simultaneously connecting to multiple clouds and programmatically adding or removing new cloud services. The Equinix Cloud Exchange allows users to quickly connect to new clouds or switch clouds at any time to support different workloads.⁴ For example, secure and reliable connection of a branch office to the Cloud Exchange provides branch office users and applications low-latency and cost-efficient virtualized connections to multiple cloud providers such as AWS, Google Cloud and Microsoft Azure, as well as optimized connections to SaaS applications such as Microsoft Office 365.

2.3 Security in Multi-cloud Architecture

Using a multi-cloud platform opens the floodgates to various security issues. The more clouds a multi-cloud environment contains, the bigger the attack surface. Indeed, you need to ensure security of distributed applications and data that now reside across multiple clouds. The best way to do this is to integrate automated security policies into the multi-cloud management platform.

If you use multiple security tools from various vendors for specific use cases, your security environment becomes fragmented and your security teams will need to manually correlate data to implement actionable security protections. This level of human intervention increases the likelihood for human error, leaving organizations exposed to threats and data breaches.

Cross-cloud data transfers lead to serious security risks. As a result, various encryption technologies should be used for data in flight and at rest and this will increase the complexity and costs of multi-cloud solutions.¹¹ Solutions for multi-cloud encryption are available. For example, IBM Multi-cloud Data Encryption (MDE) is a software-defined encryption solution that can be run in any public or private cloud configuration that offers IaaS to protect data residing in a multi-cloud environment.¹²

Other security challenges include managing trust between clouds and establishing a common identity between multi-cloud segments so that systems can authenticate securely across environment boundaries. This requires development of federated and consistent identity management and authentication processes. Securing API traffic exchanges can also be required.

Consistent security policies should apply across the entire multi-cloud environment. Security policy and control must exist at both the tenant and the application level, and it must be managed in a unified way. This requires a single point of management and a single point of end-to-end validation to ensure regulation compliance.

To address these challenges in multi-cloud security, a security framework – Multi-Cloud Security Architecture (MCSA) – has been suggested.¹³ It includes key design components required to provide security in multi-cloud environments:

- *Distributed security*: Security controls are placed close to the protected assets.
- *API-driven extensibility*: Ease of integration and extensibility from various vendors and across multi-cloud.

- *Unified security architecture*: If separate security architectures and isolated management systems are used, it often results in creation of blind spots.
- *Decoupling from infrastructure* is required to protect workloads capable of moving throughout the clouds. For example, VMware NSX Cloud provides networking and security for applications running natively across multiple public clouds.
- *Visibility* of the entire multi-cloud environment.
- *Correlation between security tools*: Essential threat intelligence needs to be immediately shared and correlated.
- *Automation*: Effective automation required to close the gap between detection and response is particularly important in a multi-cloud environment.
- *Deep and actionable analytics* to enable detection of threats and attack remediation.

3. Pros and Cons of Multi-cloud

Multi-cloud allows companies to realize the main benefits cloud services offer such as agility, scalability, redundancy, and cost-effectiveness. As a result, multi-cloud adoption is gaining momentum rapidly. Companies turn to multi-cloud to achieve business objectives which are difficult to accomplish by using private-only and/or hybrid cloud architectures. However, multi-cloud is not a panacea and comes with its own challenges. Let us briefly discuss pros and cons of multi-cloud, which are listed in Table 2.

Pros	Cons
Lock-in risk reduction: reducing reliance on any single vendor	Complexity of management of multiple providers
Services that best suit demand (functional criteria or SLA)	Broader need for specialized skills
Greater service availability by leveraging multiple clouds	Integration costs with the development of customized software
OpEx reduction by leveraging public cloud cost benefits	Security risks associated with cross-cloud data transfers
"Shadow IT" reduction	Complexity and costs associated with potential encrypting

Flexibility: ability to select superior-level services in a "best of breed" approach	Network latency and/or bandwidth
Data location and compliance satisfaction	Increase in attack surface due to system complexity. Managing trust between clouds.
Ability to compare internal IT services functioning costs and external cloud services	Management overhead
Agility in addressing new business needs. Time-to-market reduction.	Conflicting IT policies

Table 2: Pros & Cons of Multi-cloud

3.1 Multi-cloud Benefits

Cost-Performance Optimization. Cost optimization allows organizations to get the most appropriate cloud pricing for each application across providers. Using cloud solutions at different pricing tiers with multiple providers, companies can meet diverse business, security and performance needs. However, you need to consider that many cloud providers offer discounts depending on how many of their services you use. Of course, it is easy to get volume discounts if you use a single cloud provider. While the service cost is important, note that focusing primarily on cost overlooks some other main benefits of multi-cloud.

Ability to Select the Best in the Class Services. As mentioned in Section 1.1, cloud vendors offer different service capabilities and constantly change them. These capabilities vary significantly in feature richness, performance, and price. No single cloud provider is the best in all service categories. Multi-cloud gives companies freedom of choice among the best cloud services that meet their business needs. For example, AWS can be used for new cloud native applications, whereas services with large “Microsoft affinity” such as SQL Server DBs can run on Azure. Despite the additional delays for intercloud communication and the additional costs for intercloud data transfer, research using real cloud performance and cost data has shown that multi-cloud allocation outperforms single-cloud allocations in a variety of realistic scenarios.¹⁴

Availability. Concerns about cloud service availability initiated the organizations’ interest in multi-cloud. Taking the proverb “do not put all of your eggs in one basket” as a kind of guidance, they looked for protection against data loss and the ability to meet requirements on business continuity. Indeed, while providers of a single or hybrid clouds use data centers across multiple geographical regions to ensure a high level of service redundancy, there is still the possibility of an event affecting the data and services globally. By distributing an applications’ workload

across multiple clouds, each offering high availability in the SLA, organizations can significantly lower the risk of concurrent and simultaneous downtime across all clouds.

Avoiding Vendor Lock-in. This is a very important advantage of multi-cloud, as businesses can easily outgrow the cloud they use or a better deal becomes available at other providers. When a company limits itself to IT service compatibility with a single cloud vendor, it is both time-consuming and expensive to move applications anywhere else.

Data Privacy and Compliance. Data privacy and protection laws vary by region or country. Organizations may be required to know where their data is located at all times. This favors local cloud providers over global hyper-scale operators so that workloads can be co-located with the data they service. Multi-cloud provides flexible data storage options for companies facing data sovereignty and privacy compliance issues.

Proximity. Clouds may differ by the performance you get for a given application. In a multi-cloud environment, organizations may choose to run some applications using local clouds offering shorter round-trip times for traffic. Furthermore, it has been shown that an application deployment that spans multiple cloud services can offer lower latencies to its clients compared with using a single cloud service.¹⁵ This is because a multi-cloud application deployment has a larger set of data centers to choose from when serving its users. As a result, an application can take advantage of the fact that 1) one cloud provider may have a data center in a particular region while another may not, or 2) even if various cloud providers in a multi-cloud environment have several data centers in a region, one of these providers may offer lower latencies to the application users in that region due to less circuitous routing.

3.2 Multi-cloud Challenges

Multi-cloud Financial Management Strategy. Complexity of multi-cloud environments is also reflected in understanding multi-cloud costs and billing that is a new challenge for many customers.

As the concept of multi-cloud is relatively new, standard approaches to multi-cloud billing across various industry verticals are still in development. It does not help that individual cloud vendors have little interest in multi-cloud billing. Neither AWS, Microsoft Azure nor Google offers multi-cloud billing that covers billing from other providers. It is not likely that these competitive cloud providers will team up to offer unified multi-cloud billing. However, a growing number of cloud

management platforms such as those from Scalr, BMC, Solarwinds, Cloudyn, and others have begun to offer various cloud cost monitoring tools.

Complexity of Management of Multiple Providers. Challenges in managing multi-clouds are leading to development of new cloud management technologies (Section 9). The main goal of these technologies is a unified management platform with minimal management overhead and the capability to implement consistent IT policies across multiple clouds.

Complex Monitoring. While multi-cloud solutions address many data compliance challenges as discussed above, new regulations like General Data Protection Regulation (GDPR) can impose relatively strict data auditing requirements. It is possible that some existing monitoring tools are not capable of providing logging and aggregation features necessary to create adequate logging records across multiple clouds.

Increase in Attack Surface due to Multi-cloud Complexity. There are several barriers that still slow down multi-cloud adoption, and security is among the top concerns (see Section 2.3 for detail). Multi-cloud environments may involve cross-cloud data transfers resulting in security risks. Use of data encryption to mitigate these risks increases the multi-cloud service cost.

Development Platform Integration for Multi-cloud. Development toolchain needs to be unified for agile development across disparate platforms in various clouds.

4. Multi-cloud Strategies

Organic cloud adoption by different groups within the same organization commonly creates a shadow IT and results in using various cloud providers in a multi-cloud environment. Hence, the organization has become a multi-cloud user but it happened without any IT governance. This brings up a question about developing a multi-cloud strategy.¹⁶ Multi-cloud strategy is particularly effective in cases like the following:

- Users are widely distributed geographically across the globe. For example, an organization can use Google Cloud to serve the users in North America and Microsoft Azure for the customers in Europe.
- Countries having regulatory requirements for local data storage within the country, e.g. EU.
- Cloud-based application is not resilient and its deployment in multi-cloud environments for increasing capacity or resiliency is required.

There are several considerations for developing a multi-cloud strategy:¹⁷⁻¹⁹

1. **Short- and Long-Term Business Objectives.** You need to consider business growth plans (both organic growth and M&A), new product launches, moving into new markets in different geographical areas or industry verticals, introduction of new technologies and many more.

2. **Existing Infrastructure.** What issues in performance, availability, scalability and cost have recently arisen? Do you have solutions to address them? What platforms are on the premises and what are in the cloud? Is your network multi-cloud-ready? What is your technology refresh cycle?

3. **Current Cloud Providers.** How many cloud providers do you currently have? Do you have a handle on shadow IT? Do some business units have favorite providers and will they protect their choice of these providers?

4. **IT Governance for the Use of Multi-cloud Services.** The IT team should become a broker of cloud services by having management control and visibility and provide the on-demand self-service capabilities that developers and application users expect. How will you monitor the cost of the multi-cloud services? How will you be able to avoid multi-cloud sprawl?

5. **Current Application Portfolio.** Review your application taxonomy. What is your strategy for legacy applications? How many applications will use partitioned multi-cloud pattern of distributed application deployment (Section 2.1)? Assess the strategic advantages of a partitioned multi-cloud environment against the additional complexity it leads to. Will you need to implement redundant deployments for some applications? Do some applications need workload portability between clouds? If so, what solutions can be used for achieving it? Keep in mind that achieving workload portability and consistent tooling in multi-cloud environments will increase costs for development, testing, and operations.

6. **Security and Compliance Requirements.** Critically review your current security policies, toolsets, security controls, etc. How will moving to multi-cloud services affect compliance requirements your company has?

7. **Management Systems.** Are your existing management systems capable of managing multi-cloud environments? If not, what management platforms do you plan to evaluate?

8. User Experience. What is the current level of user satisfaction? What applications do your users complain about and why? How can you address these complaints in your plans for migrating to a multi-cloud environment?

9. External Customer Experience. Will migration to multi-cloud help you deliver more value to your company's customers? What customer feedback can be helpful in designing your multi-cloud strategy? Do your customers have special security requirements and if so, how will you address them by moving to multi-cloud? Will your multi-cloud strategy allow you to react quickly to rapidly changing customer needs?

10. Competence with Multi-cloud. Does your IT team have knowledge and operational experience in the use of one or multiple public clouds — AWS, Azure and/or others? What additional training does the IT team need to successfully support multi-cloud services?

5. Use Cases for Multi-cloud

The use cases for multi-cloud represent capitalizing on multi-cloud advantages we discussed in Section 3.1. Let us briefly review some of the most typical multi-cloud use cases.

5.1 Storage in Multi-cloud

In Section 3, I outlined various reasons companies decide to use multi-cloud storage technology by storing data on various clouds. While on-premises storage still provides the best security, control and performance, public cloud providers offer compelling services for off-site compliance and long-term archiving, such as AWS Glacier and Microsoft's Azure Blob. As discussed (Table 2), new government regulations on data privacy and locality, i.e. GDPR can also be a reason to move to multi-cloud storage.

To meet requests from customers for multi-cloud storage solutions, storage vendors are adding features that enable multi-cloud storage. Modern multi-cloud storage²⁰ offerings can deliver:

- **Enterprise-Grade Features:** Compared to traditional storage, multi-cloud storage offerings can provide far better data durability and greater copy data management capabilities.
- **Global Visibility:** The goal is to provide advanced analytics to track and monitor multi-cloud storage globally so that users have global visibility into storage services provided by private and public clouds. The ability to track current usage and costs makes it possible to optimize data placement in the clouds.

A challenge in achieving global visibility for storage resources is the fact that various cloud service providers store and manage data in different ways.

The discussion of multi-cloud storage solutions in detail deserves a separate article. I will just briefly list some of them and refer readers to the review of Taneja Group on multi-cloud storage vendors.²¹

Dell Technologies. Dell EMC UnityVSA Cloud Edition is the recent addition to Dell EMC Cloud Data Services and its capabilities enable users to easily deploy Dell EMC Unity block and file storage with VMware Cloud on AWS. The next generation ECS platform – ECS EX-Series – is designed for cloud-native workloads. Large-scale data migrations between cloud environments and the users' on-premises Dell EMC Unity environments are made possible by using the VMware Hybrid Cloud Extension.

Hedvig. Hedvig offers a multi-protocol storage solution that provides block, file and object services simultaneously across AWS, Google and Azure. Distinctive features of the Hedvig solution are end-to-end security, data locality control, advanced capacity optimization and performance tiering.

HPE. HPE Nimble Cloud Volumes offer an enterprise-grade storage service with multi-cloud mobility. Predictive-analytics capabilities of HPE Nimble Cloud Volumes give users control of their data across multiple clouds.

IBM. The IBM Spectrum Virtualize and IBM Spectrum Virtualize for Public Cloud leverage a multi-cloud architecture for enhancing IT service availability and business resilience. Together they support mirroring between on-premises and cloud data centers or between clouds.

Qumulo. Qumulo provides scale-out file services operating in AWS and private clouds. Other public clouds can be added later. Qumulo Core can store many times the number of files compared to legacy scale-out NAS products.

Scality. Scality delivers object-based storage software for on-premises and public clouds. Scality takes advantage of AWS S3's popularity by making a compatible API available for all public clouds.

SoftNAS. SoftNAS offers software-defined network attached storage (NAS) operating across popular public, private and hybrid cloud computing platforms, including AWS, Microsoft Azure and VMware vSphere.

SwiftStack. SwiftStack offers multi-cloud object storage services and provides a data control plane across the clouds. It maintains the public cloud's native object storage API and provides OpenStack Swift API support.

5.2 Archiving in Multi-cloud

There are several vendors who offer multi-cloud archiving. For example, Aparavi provides the active archive solution with on-premises and multi-cloud mobility.²² This, along with the use of open data format, avoids vendor lock-in. The Nasuni Cloud File Services offering eliminates the distinction between primary NAS data and secondary archive data by using the combined solution.²³ It can cost-effectively store any amount of file data for any length of time, align costs with access requirements, and give users and applications fast access. Rubrik offers Alta 4.1 as a multi-cloud archival solution that allows customers to orchestrate data storage across any public or private cloud for long-term retention.²⁴

5.3 DR in Multi-cloud

Using a multi-cloud strategy for disaster (DR) significantly enhances your DR options. Indeed, while basing your DR solutions on services from a single cloud provider with data copies stored across multiple geographical regions is better than no data redundancy, a single provider can experience a temporary service disruption that globally impacts the provider services. This results in downtime that affects business operations. If your DR strategy includes a failover solution across multi-cloud platforms, it makes your IT services even more resilient, regardless of what types of service disruptions happen. While multi-cloud DR solutions are typically more complex than doing everything on one cloud platform, they may be cost-efficient, as the backup target is less expensive storage in a different cloud.²⁵

There are several vendors of multi-cloud DR solutions. For example, Dell offers Cloud Data Protection for data stored on various cloud platforms. Dell EMC Data Domain Cloud DR provides application-consistent cloud DR in AWS and recovery to VMware Cloud on AWS.²⁶ Zerto Virtual Replication (ZVR) is the foundation for Zerto IT Resilience Platform and provides DR solutions for multi-cloud environments (Fig. 4). Zerto Analytics offers multi-cloud visibility across the entire protected environment.²⁷




Figure 4: Multi-cloud ZVR (Ref.27)

If the recovery time objective (RTO) is a few hours or even days, DR solutions are typically based on using backups. Analogous to the 3-2-1 backup strategy, storing backup data in a multi-cloud environment has a clear advantage for reliability. To be effective, multi-cloud backup requires establishing a centralized process for managing backups in various clouds. Veritas, Veeam and other vendors offer cross-cloud backups and data replications.^{28,29}

5.4 DevOps in Multi-cloud

Various surveys have shown that the foundational platform for DevOps processes shifts to multi-cloud environments.³⁰ Multi-cloud environments pose particular challenges to DevOps teams, and companies have to develop a vision for DevOps in multi-cloud.^{31,32}

Moving DevOps from a single cloud to a multi-cloud environment requires new platforms and tools. An example of such tools is MODAClouds Multi-cloud DevOps Toolbox,³³ which is a set of tools and best practice methods created for multi-cloud scenarios.

5.5 Big Data in Multi-cloud: Global Data Fabric

Benefits of multi-cloud we discussed in Section 3.1 open new opportunities for Big Data Analytics solutions resulting in developing the Global Data Fabric concept. For example, MapR offers Converge-X Data Fabric that operates as a globally distributed data store for managing files, objects, and containers across multiple on-premises and public cloud environments.³⁴ Another example is BlueData EPIC software, which provides a unified solution for Big Data workloads across multiple cloud providers.³⁵

5.6 Low-code Development in Multi-cloud

A low-code/no-code programming tool enables software development with minimal custom programming by using drag-and-drop for application components. This type of development platform allows specialists like citizen integrators who are not software developers to build and test applications quickly. One of the challenges of implementing low-code development is a risk

of lock-in if the low-code platform is bound to a particular cloud service provider. As adoption of multi-cloud environments grows, flexibility given by multi-cloud architectures becomes more and more appealing for low-code/no-code development.³⁶

6. How to Implement Multi-cloud

6.1 General Principles for Implementing Multi-cloud

I have discussed various benefits provided by multi-cloud services in Section 3.1. But how should you build a multi-cloud environment to realize these benefits to full extent? Multi-cloud is not just another new technology, and its deployment requires transformation in IT governance, architecture, technology, people, processes, etc.³⁷ Comprehensive analysis and planning are needed to make these transformations successful. The list of best practice recommendations you find below is by no means complete but it illustrates the complexity of the transformation to a multi-cloud environment.

1. Establish processes and policies in your organization to implement best practices for multi-cloud management.
2. Developing strategy and adequate controls for implementing and enforcing cloud provider-agnostic standards becomes an imperative. For multi-cloud deployment, applications need to be platform-agnostic. Instead of relying on any particular provider's proprietary service or technology, they should use open-standard services that are available with any of the large cloud providers. Using containers and Kubernetes will allow you to abstract the differences between provider environments.
3. Make sure that your data are always in an open, cloud-native format so that they can be accessed regardless of where they are stored and can be easily moved if needed.
4. Standardize data access, control and security across the multi-cloud environment by using standard object (the Amazon S3 API) and file (NFS and SMB) interfaces.
5. Develop a unified, software-defined operating model for all application workloads.
6. Enable transparent data brokering so that data placement becomes dynamic and automatically follows predefined business policies (see Section 7).

7. Use PaaS multi-cloud platforms such as Pivotal Cloud Foundry and Pivotal Container Service to implement a “cloud-first” strategy so that all the new applications for your business will be built as cloud-native software capable of running in every major private and public cloud.

8. Develop searching capabilities across the multi-cloud environment to ensure that data can be found regardless of where they reside.

6.2 Methodologies for Multi-cloud Migration

Analogous to the 5R approach to cloud migration suggested by Gartner³⁸: 1) Rehost on infrastructure as a service (IaaS); 2) Refactor for platform as a service (PaaS); 3) Revise for IaaS or PaaS; 4) Rebuild on PaaS; 5) Replace with software as a service (SaaS), similar migration methodologies have been conceptualized for migrations to a multi-cloud environment.³⁹ Due to the article size limitations, we will only briefly review some of these methodologies.

6.2.1 Rehosting

Rehosting or “Cloudification” in the case of multi-cloud represents an application migration from on-premises hosting to different cloud platforms.

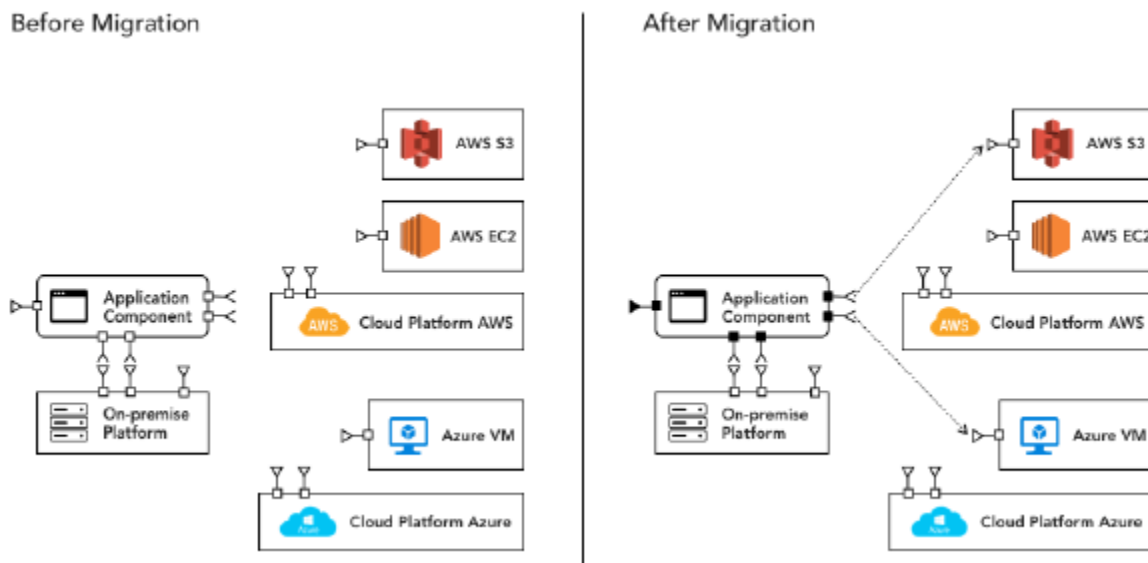


Figure 5: Rehosting (Ref.39)

For example, an application component migrated to a multi-cloud environment starts to use AWS S3 storage service and Azure compute service (Figure 5). Some application components can still run on-premises or the entire application can migrate to clouds.

6.2.2 Multi-cloud Refactor

Refactoring of an on-premises application enables deployment on a multi-cloud platform so that application components can be optimized independently. Refactoring improves application scalability and performance by making various multi-cloud deployment options possible.

6.2.3 Multi-cloud Rebinding

A re-architected application has a component deployed as two redundant components on different clouds for achieving high availability. For example, as shown in Figure 6, AC1 component remains on-premises whereas two AC2 components are deployed on AWS and Azure. AC1 and two AC2 components are connected via an integration message broker like Azure Service Bus. AC2 failover between AWS and Azure will still result in some downtime.

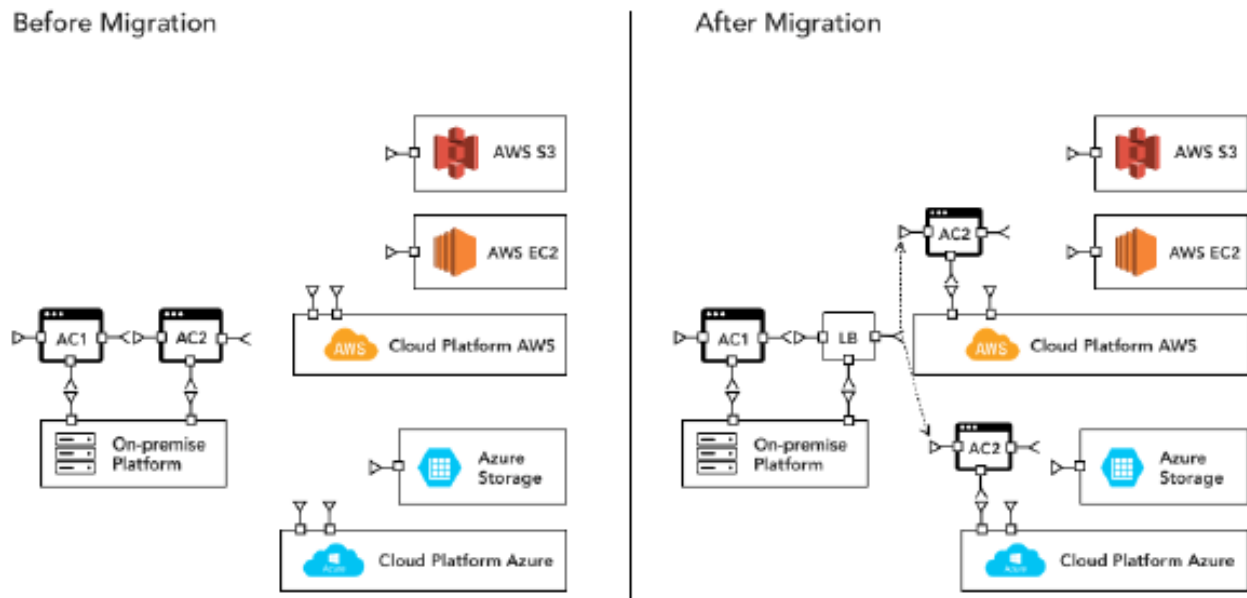



Figure 6: Multi-cloud Rebinding (Ref.39)

6.2.4 Multi-cloud Rebinding with Cloud Services Brokerage

In this case, services provided by the cloud broker integrate application components and offer flexibility to choose services from multiple cloud providers. An example is shown in Figure 7: AC1 component is deployed on-premises and two re-architected AC2 components are deployed on two different cloud platforms, AWS and Azure. Cloud services broker integrates all three components and provides flexibility in choosing services from multiple cloud providers.

Before Migration



After Migration




Figure 7: Multi-cloud Rebinding with Cloud Brokerage (Ref.39)

The importance of cloud services brokerage in multi-cloud ecosystems warrants a closer look at their functions.

7. Multi-cloud Services Brokerage

7.1 Roles of Cloud Services Brokerages in Multi-cloud Ecosystem

Multi-cloud environment is complex (Section 3.2) and if not managed properly, IT services quality may become unsatisfactory, cloud services sprawl takes place and cloud services costs can spiral out of control. These challenges have significantly increased the importance of cloud services brokers (CSB) who have come on the scene as experts in managing the complexity of multi-cloud ecosystems.

The primary roles of a CSB can be defined as adding value to cloud services by integrating cloud services with on-premises applications and with each other, and ensuring security of cloud services.⁴⁰

A CSB is mostly used for multi-cloud environments to provide intercloud support. CSB activities typically include service life cycle management, and aggregating, comparing, negotiating, purchasing, customizing, managing, orchestrating or administering of cloud services. CSBs can be categorized into three types:

1. Cloud Aggregators bring together multiple services such as cloud-scale provisioning, single sign-on (SSO), unified billing, and unified management by integrating multiple service catalogs from different service providers in a multi-cloud ecosystem into a single user interface. While some CSBs function primarily as re-sellers, most CSBs offer additional services such as security and governance on top of the services provided by individual clouds in the multi-cloud environment. As CSBs typically have pre-existing relationships with a number of cloud providers, they may be asked by the clients to negotiate contracts with cloud service providers on behalf of the client for service cost reduction.

2. Cloud Integrators add value by bringing multi-cloud services together and making them work through a single orchestration.

3. Cloud Customizers modify capabilities of the existing cloud services to meet specific functional requirements of the client by providing bespoke services.

Sometimes a line is drawn between cloud brokers and cloud integrators (Table 3) but in many cases the line is blurring as cloud integration capabilities become a part of the CSB portfolio.

Cloud Brokerage Platforms & Tools	Delivery of cloud brokerage capabilities through a platform, software license or subscription that is used by the end customer for the purpose of aggregating, comparing, negotiating, purchasing, customizing, managing or administering cloud services
Cloud Brokerage Services	Delivery of cloud brokerage capabilities through a services engagement, under which the third-party cloud brokerage service provider aggregates, compares, negotiates, purchases, customizes, manages or administers cloud services
Cloud Integration Platforms & Tools	Delivery of cloud integration capabilities through a platform, software license or subscription that is used by the end customer to achieve benefits including data sharing, analytics, capacity sharing, or simplified management

Cloud Integration Services	Delivery of cloud integration capabilities through a services engagement that provides the end customer benefits including data sharing, analytics, capacity sharing, or simplified management
----------------------------	--

Table 3: Cloud Brokerage and Integration Platforms (Ref.41)

7.2 Open Service Broker

The Cloud Foundry Foundation has initiated the Open Service Broker (OSB) project in an effort to bridge the gap between proprietary cloud services such as databases, messaging and object storage and containerized applications deployed in Kubernetes. OSB consolidates CaaS and PaaS platforms through an open API specification. With OSB, it is possible for an application running on Google Kubernetes Engine to connect to the SQL Database Services provisioned by Azure. The combination of Kubernetes and OSB is a way to deliver portability in multi-cloud.

8. Multi-cloud Integration and Inter-Cloud Services. Integration Platform as a Service (iPaaS)

Multi-cloud architectures extend the cloud concept to creation of cloud of clouds operating by using intercloud services. The level of integration between clouds in your multi-cloud environment depends on how you use these cloud services. In some cases, you do not necessarily need to integrate them. Deploying of components of a distributed application in multiple clouds with some redundancy for failure protection is an example. However, more commonly organizations want to integrate clouds composing a multi-cloud environment to maintain a high level of efficiency.

Standardization efforts are critically important for cloud integration. They aim to create a unified interface for managing cloud resources and services across cloud vendors. Such efforts are exemplified by the Topology and Orchestration Specification for Cloud Applications (TOSCA),⁴² the Cloud Infrastructure Management Interface (CIMI), the Open Cloud Computing Interface (OCCI), and the Cloud Data Management Interface (CDMI). The TOSCA standard, which provides support for applications whose components are deployed on different clouds, is used to define a model describing the topology of cloud applications and the required resources in a provider-agnostic and resources-independent way. At the same time, some vendors work to develop minimalistic APIs to support the most common interactions.

Multi-cloud strategies lead to growing interest in integration Platform as a Service (iPaaS), a suite of cloud services supporting cloud-to-cloud and cloud-to-on-premises integration scenarios.⁴³ For example, Dell Boomi iPaaS,⁴⁴ which is a 100% native cloud and low-code, enables users to quickly and easily connect cloud and on-premises applications. Another example of iPaaS is the Informatica Intelligent Cloud Services, which is an end-to-end solution capable of managing data in multi-cloud environments.⁴⁵

9. Multi-cloud Management

9.1 Architecture and Core Capabilities of Multi-cloud Management

In the previous sections of this article, I have discussed various benefits of multi-cloud and methodologies for implementing multi-cloud. However, you can realize these benefits only if you can successfully manage your multi-cloud environment. As we have already seen (Sections 3.2), multi-clouds present their own management challenges. Addressing these challenges requires the development of new technologies for unified cloud management. In addition to having to manage resource utilization, performance and costs of various public and private cloud services, multi-cloud management platforms must be capable of managing applications running on-premises and various clouds and interoperating with the iPaaS connecting them.

As shown in Figure 8, multi-cloud management and integration architectures can be presented in terms of three systems: systems of engagement that serve for user interaction, systems of execution for work orchestration, and systems of record that are transactional and/or data/content stores.

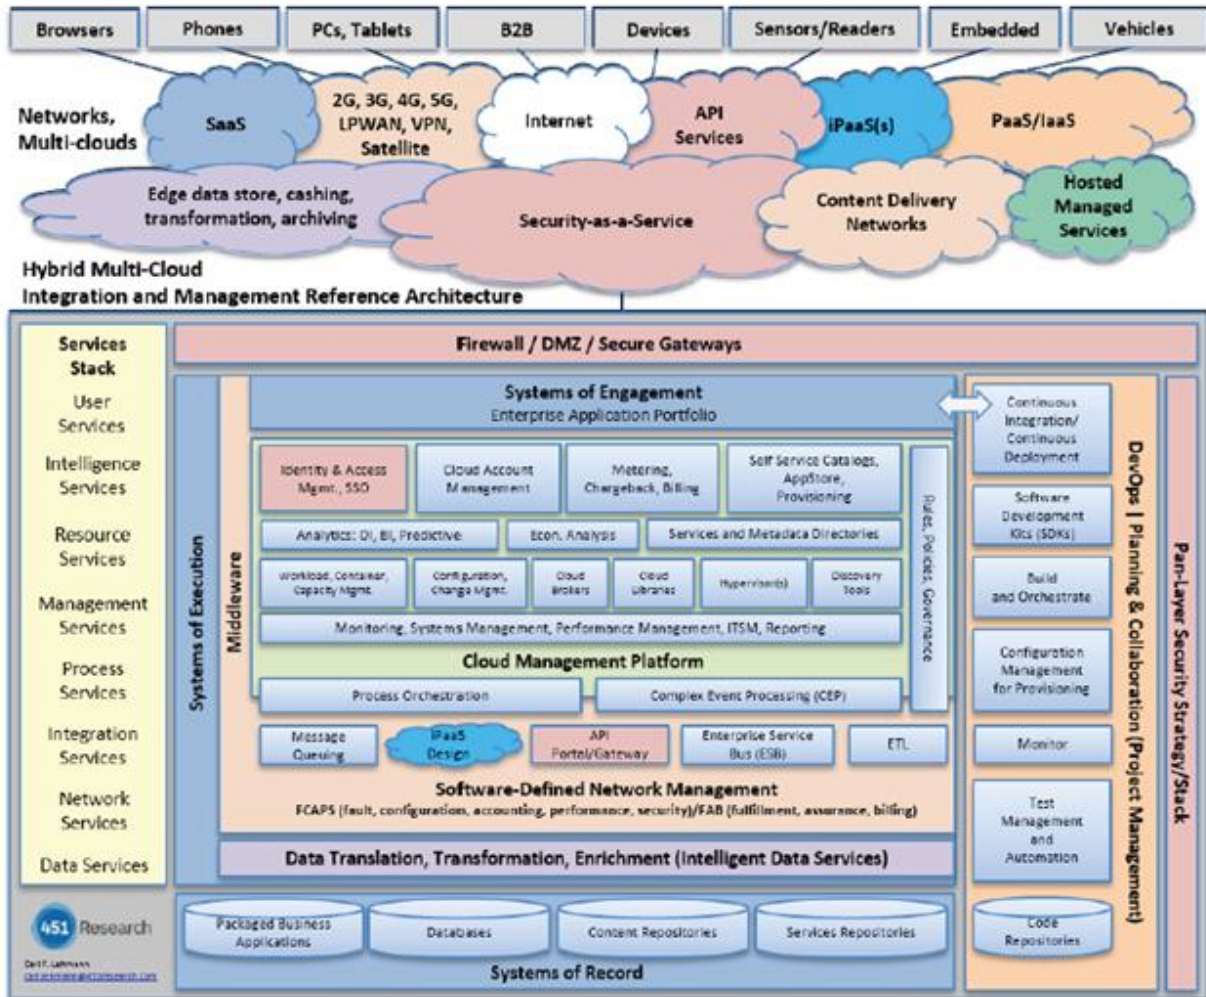


Figure 8: A Multi-cloud Integration and Management Reference Architecture (Ref.46)

While not all of the components presented in the reference architecture (Figure 8) are required and most enterprises will implement only a select subset, a primary requirement is a unified self-service portal enabling management of services across multiple clouds. Some cloud management vendors extend their management platform functionality by coupling cloud account management software with monitoring and performance management tools.

The core multi-cloud management capabilities required to conquer multi-cloud complexity and limit cloud sprawl in multi-cloud ecosystems are presented in the following list:

- 1) Automate the provisioning of compute, storage, network, security, application stacks and data across clouds.

- 2) Intelligently deploy workloads and services based on economic analysis and policies across on-premises infrastructure and private and public clouds.
- 3) Manage identity authentication and access control, ideally using SSO.
- 4) Manage runtime execution and performance of all the clouds composing the multi-cloud environment and enact policies to automate scaling, bursting, high availability and disaster recovery.
- 5) Maintain a service component library, including operating system images, databases, etc.
- 6) Deploy a unified self-services application catalog based on access control and governance policies.

9.2 Automation and Orchestration in Multi-cloud

Since a broad range of capabilities is required in multi-cloud deployments, they are typically multivendor and create a challenge for unified management. Constructing common automation and orchestration layers on top of a heterogeneous multivendor underlay has historically proven difficult; the track record of development of multivendor element management systems can serve as evidence.

As mentioned, a key attribute of a successful multi-cloud implementation is effective IT governance (Section 4). Instead of relying on application-specific or cloud vendor-specific automation tools, organizations have to look for automation tools capable of operating in a multi-cloud environment and apply unified IT governance across multi-cloud domains by using multi-cloud policy automation. The Dell Multi-cloud Manager, BMC Control M, RightScale multi-cloud policy engine (part of the RightScale Multi-Cloud Platform), Cisco CloudCenter, and Morpheus are just a few examples of such tools providing automation and orchestration of data, applications, and workloads across different cloud environments.

Orchestration tools for multi-cloud should be able to determine deployment targets based on costs, security, traffic, affinity/anti-affinity rules for clouds composing the multi-cloud environment and other criteria. Multi-cloud orchestration platforms can use various standard protocols such as network configuration protocol (NETCONF),⁴⁷ open source libraries like NAPALM (network automation and programmability abstraction layer with multi-vendor support, a Python library used for network automation)⁴⁸ or vendor-neutral data models like OpenConfig.⁴⁹ The key requirement is that orchestration be built on an open platform.

Infrastructure as code (IaC) tools enabling automation of deployment and configuration of compute resources – servers VMs, and containers – can simplify multi-cloud management. IaC is considered a virtual hosting model for applications. Using IaC makes changes in the infrastructure transparent to applications and, as the IaC model applies across multiple clouds, addition of a new cloud provider requires just defining it in IaC.

Cloud providers offer native IaC tools such as AWS CloudFormation and Azure Resource Manager templates. However, they may be not the best choice for multi-cloud deployments since cloud providers' native IaC tools mainly focus on single-cloud or simple hybrid cloud deployments. Terraform, an open source product from HashiCorp, seems to be the most popular third-party IaC tool. It is multi-cloud compatible and supports a range of cloud vendors. Reusing Terraform templates allows IT teams to standardize infrastructure configurations across various cloud providers.

9.3 Workload Mobility in Multi-cloud

9.3.1 Interoperability in Multi-cloud

Interoperability in multi-cloud computing can be defined as the capability of public and private clouds to support each other's application and service interfaces associated with different aspects of cloud services, authentication and authorization processes, data formats, etc. Different aspects of interoperability which can be positioned as separate layers are presented in the EU 4-level model (Table 4). In this model, exchange of information between two systems takes place on four levels and the higher layers utilize the lower layers:

#	Layer	Aim	Objects	Solutions
1	Technical	Technically secure data transfer	Signals	Protocols of data transfer
2	Syntactic	Processing of received data	Data	Standardized data exchange formats, e.g. XML
3	Semantic	Processing and interpretation of received data	Information	Common directories, data keys, ontologies
4	Organizational	Automatic linkage of processes among different systems	Processes (workflow)	Architectural models, standardized process elements

Table 4: The EU 4-level Interoperability Model for Cloud Computing [50]

Technical interoperability deals with the protocol used for information exchange. Syntactic interoperability is related to the format of the exchanged data, for example, XML data structures or JSON data streams. Semantic interoperability concerns the data structure.




Figure 9: Elements of Interoperability and Portability in Cloud Services (Ref.50)

Figure 9 shows three main interfaces between customer roles and the cloud service: the Functional interfaces, the Admin interfaces and the Business interfaces. The main functional capabilities offered by the cloud service rely on the Functional interfaces. The Admin interfaces are associated with managing the cloud service and include capabilities such as monitoring the services and managing security features such as user identification, authentication and authorizations. Finally, functions of the Business interfaces relate to the commercial aspects of the cloud service including subscription information, billing and invoicing. It is important to mention that the interoperability of one interface does not guarantee interoperability of the others.

SaaS applications present the greatest interoperability challenge in multi-cloud environments. Indeed, there are very few standard APIs for SaaS applications. Furthermore, even switching from one SaaS application to another SaaS application having comparable functionality typically requires a change in interface.

9.3.2 Portability in Multi-cloud

While some multi-cloud implementations may not have intelligent workload portability as the main goal, optimal placement of workload is one of the key advantages of a multi-cloud environment (Section 3.1). Data and workload portability requires the ability to move not only VMs or containers

and application data between clouds but the environment metadata as well. Environment metadata is often very cloud provider-specific. As a result, account structures, users, permissions, policies, etc. vary between clouds forming a multi-cloud system.

Portability can be categorized into two separate areas: data portability and application portability:

- *Cloud data portability* is the ability to easily move data from one cloud service to another cloud service without needing to re-enter the data. Data is typically the difficult component of a workload to move. The data inertia is characterized by the so called “data gravity”.⁵¹ Data gravity increases in a multi-cloud environment because data transfer takes place between clouds in multi-cloud data fabric (Section 2.2). If you keep redundant data sets, you face one more challenge – how will you keep all the copies in sync?

Data portability requires data export from one cloud and import into the target cloud. This can be achieved by using an API associated with the cloud service. If the API used for the source cloud service is different from the API used by the target cloud service, some data migration tools may be required. Consideration of the next higher level in the interoperability/portability model – Syntactic level (Table 3) – shows that if there is no match between the data syntax of the source cloud service and the syntax of the target service (for example, the source uses JSON syntax but the target uses XML), it may still be possible to map the data using mapping tools such as Dell Boomi AtomSphere or Informatica Cloud Data Integration. Moving up to the semantic level in the interoperability/portability model leads to more challenges in data portability. Indeed, mismatch in source and target service data semantics makes data portability very difficult or may be even impossible to attain.

- *Application portability* can be defined as the ability to move applications between cloud vendors so that running the applications in the target cloud service will require a minimal level of integration. While application recompiling or relinking for the target cloud service may be needed in some cases, making significant changes to the application code should not be necessary if the application is portable. Containerizing applications and using Kubernetes enable application portability. Using containers means that programmers may not need to rewrite the code for each new operating system and cloud platform. However, containerizing applications with monolithic architecture can be a significant challenge.

The goal of the TOSCA⁴² standard that we briefly discussed in Section 7 is to improve the portability and manageability of applications by composing a service once and running it on any cloud.

9.4 Overview of Multi-cloud Management Systems

The requirements for multi-cloud management platforms and their reference architecture were discussed in Section 9.1. The scope and size limitations of this article do not allow a detailed review of multi-cloud management systems. Readers can find very good overviews in the recent Gartner review and publication by CSCC (OMG).^{52,53}

I will mention just a few popular multi-cloud management tools.

Dell Multi-Cloud Manager is a cloud-agnostic standalone cloud management solution that includes automated provisioning, a self-service portal, application and infrastructure templates, SSO, scaling, governance, security, monitoring, and integrations with a broad range of private and public cloud platforms. The Multi-Cloud Manager enables deployment and management of applications across private, public and hybrid clouds.⁵⁴ It is available as a SaaS subscription or for on-premises deployment.

BMC Multi-cloud Management automates provisioning of IT services across cloud platforms.⁵⁵ By integration with IT service management processes such as change management, configuration management, compliance and patching, it enables governance and compliance controls for cloud workloads. BMC Discovery for Multi-cloud platform provides a view of both on-premises and distributed cloud resources and their relationships.

Embotics vCommander supports various cloud environments by providing provisioning automation and self-service capabilities.⁵⁶

Flexera (RightScale) Multi-cloud Platform includes a self-service portal that enables developers to access public and private cloud infrastructures and helps automate cloud application deployment.⁵⁷ It can also monitor, forecast and optimize costs across a multi-cloud platform.

Nutanix Enterprise Cloud Platform has a technology-agnostic architecture, allowing IT teams to manage applications across multi-cloud environments.⁵⁸ The Nutanix Enterprise Cloud OS includes Beam, a multi-cloud cost optimization service. Beam provides analytics detailing cloud consumption patterns that enable cost optimization in multi-cloud environments.

Red Hat CloudForms (based on open source ManageIQ) is a multi-cloud management platform that enables setup of policy-controlled, self-service cloud environments.⁵⁹ CloudForms offers unified management for hybrid environments by providing a consistent cross-platform experience.

Scalr Cloud Management Platform is an open source cloud-agnostic management solution that offers a full range of cloud management functionality in key areas such as governance, security, and compliance; business agility; cost optimization and visibility.⁶⁰ It enables organizations to manage, automate and control multi-cloud environments by using an administrative console with a single UI and API standardizing multiple clouds to orchestrate, automate workload deployments and enforce policy across multi-cloud environment.

VMware vRealize Suite enables organizations to extend their cloud environment by including multiple public clouds. Its recently added application life-cycle management helps DevOps in building multi-cloud applications. vRealize can manage infrastructure and applications across private and public clouds and provision IT resources in an optimal way.⁶¹

10. Conclusion

Rapid adoption of emerging multi-cloud computing brings a lot of opportunities and offers various benefits. However, implementation of a multi-cloud strategy has many challenges. Every organization designs and deploys multi-cloud differently to meet its needs. I hope my article provided an overview of key aspects of multi-cloud computing that will help you successfully implement multi-cloud for your organization.

11. References

1. <http://www.bmc.com/blogs/hybrid-cloud-vs-multi-cloud-whats-the-difference>
2. <http://cloud.google.com/solutions/hybrid-and-multi-cloud-architecture-patterns>
3. <http://www.stratoscale.com/blog/cloud/7-considerations-building-multi-cloud-solution>
4. <http://blog.equinix.com/blog/2018/08/21/how-do-colocation-and-interconnection-enable-hybrid-multi-clouds>
5. Multi-cloud Technical Guide for Network and Cloud Architects. Juniper Networks, White Paper. 2018.
6. D. Loi. Aligning Applications and Connectivity to Enable Fast And Safe Cloud Computing. Orange Business Services, 2015.

7. http://www.juniper.net/documentation/en_US/release-independent/solutions/topics/task/configuration/ip-fabric-underlay-cloud-dc-configuring.html
8. <http://www.citrix.com/blogs/2018/03/07/from-sd-wan-to-secure-multi-cloud>
9. <http://www.talari.com/blog/3-reasons-you-need-sd-wan-for-your-multi-cloud-strategy>
10. <http://www.orange-business.com/en/products/flexible-sd-wan>
11. <http://searchcloudcomputing.techtarget.com/ehandbook/How-to-secure-a-multi-cloud-architecture>
12. <http://www.ibm.com/cloud/garage/architectures/securityArchitecture/security-for-data>
13. <http://www.varmour.com/resources/blog/entry/introducing-the-multi-cloud-security-architecture>
14. S. S. Woo, J. Mirkovic. Optimal application allocation on multiple public clouds. *Computer Networks*. v. 68, p. 138, 2014.
15. Z. Wu and H. V. Madhyastha. Understanding the Latency Benefits of Multi-Cloud Webservice Deployments, *ACM SIGCOMM Computer Communication Review*, v.43, p.14, 2013.
16. <http://www.delltechnologies.com/en-us/perspectives/seeing-through-the-multi-clouds-navigating-your-public-and-private-cloud-strategy>
17. B. Felter. <http://www.vxchnge.com/blog/determine-needs-multi-cloud>
18. M. Roy. <http://searchcio.techtarget.com/blog/TotalCIO/Multi-cloud-strategy-Determine-the-right-cloud-for-your-workloads>
19. J. Edwards. <http://searchstorage.techtarget.com/tip/8-multi-cloud-storage-strategies-that-target-performance>
20. Multi-cloud Storage: Planning, Deployment and Management. TechTarget e-Guide, 2018.
21. Emerging Market Report on Multi-Cloud Primary Storage. Taneja Group, 2017.
22. Multi-Cloud Active Archive. White Paper, Aparavi. 2018.
23. <http://www.prnewswire.com/news-releases/nasuni-releases-first-multi-cloud-solution-that-combines-primary-and-archive-file-storage-and-automatically-reduces-costs-as-files-age-300717393.html>
24. <http://www.rubrik.com/blog/rubrik-enhances-multi-cloud-management>
25. J. Edwards. <http://searchdisasterrecovery.techtarget.com/tip/5-crucial-multi-cloud-disaster-recovery-principles>
26. <http://www.emc.com/about/news/press/2018/20180827-01.htm>

27. <http://www.zerto.com/products/disaster-recovery-to-the-cloud/multi-cloud-agility/>
28. <http://www.veritas.com/solution/cloud>
29. <http://www.veeam.com/multi-cloud-enterprise.html>
30. <http://www.redhat.com/en/resources/automation-devops-multi-cloud-world-financial-idc-infobrief>
31. T. Steinborn. <http://opensource.com/article/18/1/future-devops>
32. <http://www.infoworld.com/article/3315077/cloud-computing/devops-is-mandatory-for-multi-cloud-deployments.html>
33. <http://multi-clouddevops.com/technologies.html>
34. <http://mapr.com/datasheets/perspective-series-orbit-cloud-suite>
35. <http://www.bluedata.com/blog/2018/09/hybrid-and-multi-cloud-playbook-for-ai-and-big-data-workloads>
36. C. Tozzi. <http://searchcloudcomputing.techtarget.com/tip/Cloud-native-low-code-platforms-rival-third-party-options>
37. http://www.dellemc.com/en-us/collaterals/unauth/white-papers/solutions/whitepaper_utilizing_a_multi_cloud_framework_to_enable_collaborative_care.pdf
38. <http://www.gartner.com/newsroom/id/1684114>
39. J. Solanki. 6 Multi-Cloud Architecture Designs for an Effective Cloud Strategy. <http://www.simform.com/multi-cloud-architecture>.
40. A CIO Primer on Cloud Services Brokerage. Gartner. 2012.
41. A. Krans and C. Mooshian. The Rise of Cloud Brokers and Integrators Amid Hybrid Integrations. TBR, 2015.
42. S.Tummalapalli, R. Kanth .P.Yuvaraj, and S.Velagapudi. TOSCA Enabling Cloud Portability. International Journal of Advanced Research in Computer Engineering & Technology. vol. 2, p.974, 2013.
43. <http://www.gartner.com/it-glossary/information-platform-as-a-service-ipaas>
44. <https://www.dell.com/learn/us/en/04/shared-content~data-sheets~en/documents~dell-multi-cloud-manager-datasheet-25278.pdf> <http://boomi.com/ja/integration-topics/cloud-data-integration>
45. Solving Multi-cloud and Hybrid Data Management Challenges with iPaaS. Informatica Solution Brief. 2018.
46. C. Lehmann. Hybrid multi-cloud architecture and the vendors aiming to enable and manage it. 451 Research, 2016.

47. <http://www.netconfcentral.org>
48. <http://napalm-automation.net>
49. <http://www.openconfig.net>
50. Interoperability and Portability for Cloud Computing: A Guide. Cloud Standards Customer Council, 2014.
51. A. McDonald, M. Tidwell, and E. Lakin. Cloud Mobility and Data Movement. SNIA, 2018.
52. Magic Quadrant for Cloud Management Platforms. Gartner, 2019
53. <http://www.omg.org/cloud/deliverables/practical-guide-to-cloud-management-platforms.htm>
54. <http://www.dell.com/learn/us/en/04/shared-content~data-sheets~en/documents~dell-multi-cloud-manager-datasheet-25278.pdf>
55. <http://www.bmc.com/it-solutions/multi-cloud-management.html>
56. <http://www.embotics.com/hybrid-cloud-management-platform>
57. <http://www.rightscale.com/products-and-services/multi-cloud-platform>
58. <http://www.nutanix.com/go/enterprise-apps-in-a-multi-cloud-world.php>
59. <http://www.redhat.com/en/resources/red-hat-cloudforms-unified-management-for-hybrid-environments>
60. <http://www.scalr.com/>
61. <http://www.vmware.com/it-priorities/integrate-public-clouds/manage-multi-clouds.html>

Dell Technologies believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” DELL TECHNOLOGIES MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying and distribution of any Dell Technologies software described in this publication requires an applicable software license.

Copyright © 2019 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.