



Architecting Resilient Hybrid Multi-Cloud Platforms across on-Prem, AWS, and Azure

Sridhar Nuti

Solutions Architect, Nutanix, USA

ABSTRACT: Multi-cloud strategies are now taking the form of hybrid ones, in order to be resilient, cost effective, and vendor neutral. In this paper, the researcher is suggesting a reference architecture that brings together on-premise infrastructure with AWS and Azure. The framework focuses on the unified networking, identity federation, workload portability, infrastructure automation and observability. The study compares the speed of the deployment, resource consumption, operational controls, and cost effectiveness using the Infrastructure as Code (IaC) tools and multi-cloud orchestration. Findings indicate increased deployment speed, equalized resource consumption, greater success in automation and a maximum of 34 cost reduction. The results offer a real-life advice on how to build resilient, scalable, as well as efficient hybrid multi-clouds.

KEYWORDS: Hybrid Multi-Cloud, AWS, On-Premises Infrastructure, Azure, Workload Portability, Infrastructure as Code (IaC)

I. INTRODUCTION

Companies are also shifting to hybrid multi-cloud systems to achieve greater resilience, regulatory compliance, and flexibility of their operation. Conventional single-cloud or on-prem infrastructures are normally limited in terms of scalability, performance, and reliance on the vendor. Hybrid multi-cloud is a type of infrastructure merging on-premise as well as the public clouds including AWS and Azure that offer work load portability, bursting clouds, and automation. Modern tools such as Terraform, Ansible and Kubernetes are helpful to manage Infrastructure as Code (IaC) and coordinate the work across different environments. This study provides a reference architecture and assesses this architecture on quantitative measurements, such as deployment velocity, resource consumption, operational governance and cost optimization and gives actionable advice to enterprise cloud adoption.

II. RELATED WORKS

Evolution and Need for Hybrid and Multi-Cloud Architectures

Cloud computing has also developed very fast by having single-provider platform up to multi-cloud strategy to hybrid cloud strategy. Cloud services initially were mostly confined to data centers of only one provider which was very inflexible and limited in scaling [5]. Gradually, businesses started engaging more than one provider of cloud computing services in order to achieve vendor indifference, performance and resilience to disruption [3][9]. Hybrid cloud represents a hybrid of both a private and a public cloud solution and became one of the effective solutions applicable to companies that aim to maximize cost without losing the control over the data of high sensitivity [2][10]. Focusing on combining on-premises infrastructure with the public cloud platform e.g. AWS and Azure, hybrid cloud allows issuing the horizontal scaling of applications and workloads and this idea is commonly known as cloud bursting [2].

Studies indicate that hybrid deployments enable enterprises to comply with regulatory and compliance demands and enable them to have operational flexibility [3][7]. The multi-cloud strategies enable elimination of the threat of vendor lock-in by distributing workloads among the numerous cloud vendors [3][9]. The architectures enable the enterprises to also take advantage of niche services provided by various vendors e.g. utilizing high-performance computer power on one cloud but storing confidential data safely on-premises [5]. The trend to the hybrid multi-cloud developments is indicative of an increasing desire to have scalable, flexible, and resilient computing platforms.

Elasticity and Virtual Cluster Approaches

Dynamic provisioning of resources in environments is one of the major advantages of hybrid and multi-cloud environments. Elastic computing is offered by virtual clusters that can be scaled with a greater or smaller number of



nodes and can easily support each of the High-Performance Computing (HPC) and High Throughput Computing (HTC) computational loads [1]. Conventionally, this elasticity applied only to one cloud provider, though according to recent research, architectures exist that may extend virtual clusters to multiple cloud sites. These hybrid virtual elastic clusters are made using open-source, and Infrastructure as Code (IaC). The clusters are automatically deployed and configured to consist of on-premises and cloud resources [1]. Such clusters also adopt the progressive network like automated VPN tunnelling and therefore are appropriate in massive distributed calculations [1].

In the same way, the application platforms, such as Aneka are used to contemplate multi-cloud deployment, by dynamically provisioning resources provided by various cloud vendors [4]. The Aneka allows distributed apps to use resources provided by Azure and AWS EC2, and provides scripting management with scripting languages such as power shell [4]. This will enhance the performance of the distributed applications, and show the performance benefits of multi-cloud deployments. These solutions bring out the possibility of hybrid clouds not only to the conventional enterprise workloads but also to research, simulation, and AI-based computation [6].

Integration, Networking, and Security in Hybrid Clouds

Although hybrid multi-cloud environments provide flexibility and scaling, the integration of on premises and cloud technologies brings a lot of complexity, in particular, in network and security [3][8]. The connectivity between the private and the public clouds will be vital, and VPNs or software-defined networking tools will normally be used to implement such connectivity [2][8]. Multi-tenant hybrid environments make identity and access management (IAM) more difficult, where the identities on the on-premises have to be recreated on the public cloud platforms and still have secure authentication and authorization [8].

Research lays stress on best practices in inter-cloud networking, workload orchestration and automated security enforcement [3]. As an example, hybrid clouds need to have collective identity fusion, role-based access control, and zero-trust security systems, to effectively regulate distributed assets [8]. These to ensure compliance with operational governance and regulatory compliance especially on sensitive data is a must [10]. Monitoring, logging and cost optimization, which makes the deployment resilient, secure and efficient can also be provided through the multi-cloud management platforms [3][9].

Application Deployment and Performance Considerations

Applications distributed in hybrid clouds need to be planned architecturally to be performance, reliable and cost-efficient [7][10]. Cloud bursting that enables applications to be spread to the public clouds on top of the existing private infrastructure, enhances the throughput and latency of some of the workloads [2]. The experiments on using distributed database systems in the hybrid clouds have demonstrated that distributed database systems such as MongoDB and MySQL Cluster can take advantage of cloud bursting, whereas other systems such as Cassandra and Redis could suffer due to the absence of cloud bursting [2].

Furthermore, orchestration systems, such as Kubernetes, and platforms, such as ECHO, can also be used to schedule tasks in a dynamic fashion on edge, on-prem, and cloud assets [6]. These models can process IoT and smart cities data in real time, which is evidence of how widely hybrid multi-cloud applications can be deployed in various settings, as opposed to standard enterprise computing. Through automation, containerization, and IaC, organizations will be able to lower operational overheads, time to deployment, as well as increase reliability of the application [1][4][6].

According to the available literature, the hybrid multi-cloud systems are becoming urgent to the enterprises, which require resilience, performance, and regulatory compliance. Aneka and other virtual clusters and platforms, as well as the orchestration frameworks, show how one can dynamically provision large workloads in both on-premises and cloud environments [1][4][6]. Security, networking and identity management is a big issue, which has been handled using VPNs, IAM architecture and models of zero-trust [2][8] and [10]. Hybrid cloud applications make it possible to achieve cloud bursting, portability of work load and optimize costs and minimize risks of vendor lock-in [2][3][7]. The literature is informative on architectural work, patterns of implementation, and best practices that underlie the design of resilient, scaleable and secure hybrid multi-cloud platforms based on the use of on-premises infrastructure along with AWS, Azure, and other cloud services.



III. METHODOLOGY

The research methodology applied in this study is quantitative research in order to determine the adherence of a hybrid multi-cloud architecture to integrate on-premises infrastructure with AWS and Azure. The main goal is to quantify the effects of the suggested reference architecture concerning velocity of deployment, cost management, and operational management. The approach is an integration of experimental application, metrics-based assessment, and statistical reporting to offer practical solutions to the enterprises aiming at gaining a robust hybrid multi-cloud system.

Reference Architecture Design

It consists of the initial design of reference architecture of hybrid multi-cloud deployment. The architecture incorporates on-premises infrastructure together with Aws and Azure cloud infrastructure assuring unified networking, verification federation, workload mobility, infrastructure automation, and observability. SDN and VPN software-defined networking has become standardized in networking across environments to facilitate secure communications between on-premises and cloud environments. The identity federation is applied by using centralized authentication and single sign-on which maintains the consistency of the user access to all platforms. The containerization of applications with Docker and the management of the applications with Kubernetes in a way that is independent of environment, whether on-premises or a public cloud is a challenge that is taken care of. Infrastructure as Code (IaC) tools, (including Terraform and Ansible) are used to provision, configure, and manage the resources in a programmatic way, which is an infrastructure automation tool. The centralized logging, monitoring, and alerting tools are used to implement the observability to monitor system performance and identify anomalies in real time.

Implementation and Experimental Setup

It creates a hybrid multi-cloud testbed that is an on-premise data center and virtualized environments within AWS and Azure. Typical enterprise workloads (compute-intensive applications, distributed databases, and containerized microservices) are taken to provide the selection of workloads. All workloads are deployed based on the reference architecture in IaC scripts, and the deployment of workloads is automated to create consistency. The experiment carries out various deployment cases, such as the different allocation of resources, workload allocation, and scaling pattern to test the system performance in different conditions.

Data Collection and Metrics

The deployment and runtime quantitative measurements are used to measure the success of the architecture. Key metrics include:

- **Deployment Velocity:** It is time required to deliver and implement workloads in the hybrid cloud.
- **Resource Utilization:** CPU, memory utilisation and network utilisation efficiency of on-premise and cloud resources.
- **Operational Governance:** Adherence to set policies, configuration management automated and a minimization of manual interventions.
- **Cost Efficiency:** Cost comparisons of deployment between having hybrid and single-clouds.

There are accuracy and repeatability in data collection by means of monitoring tools, IaC logs and cloud provider APIs. Every test is repeated several times in attempts to cover the variability and enhance the statistical reliability.

Data Analysis

The obtained data is managed through the descriptive and inferential statistics. The improvement percentages of the hybrid multi-cloud reference architecture are measured by means, standard deviation and performance improvement. On-premises only, single-cloud, and multi-cloud hybrid deployments are compared to each other. Resource allocation, deployment configurations and the performance metrics are assessed using regression analysis. The outcomes are then utilized in the validation of the architecture proposed, in pointing out the best deployment patterns as well as in making decision matrices in order to guide the enterprise adoption.



IV. RESULTS

Deployment Velocity

The speed of deployment is the speed of workloads that can be brought up and delivered to the hybrid multi-cloud environment in a manner that is as fast as possible. Three deployment cases have been tested in the experiments, all on-prem, single-cloud (AWS or Azure), and hybrid multi-cloud (on-prem + AWS + Azure) deployment. Application computing time, distributed database, and containerized microservice deployment time was measured through IaC automation.

The findings show that hybrid multi-cloud environments are deployed faster with regards to overall deployments than an on-premises only environment. This has been enhanced to a large extent because it offers parallel provisioning across several environments and automated configuration with the help of Terraform and Ansible scripts. Single-cloud was also quicker than on-premises and slower than hybrid multi-cloud because of a small level of parallelism.

Table 1: Deployment Time Across Different Environments (in minutes)

Workload Type	On-Premises	Single-Cloud	Hybrid multi-cloud	Improvement (%) Hybrid vs On-Prem
Compute-Intensive App	45	32	28	38%
Distributed Database	52	37	30	42%
Containerized Microservices	48	35	29	40%

The hybrid multi-cloud model exhibits a steady decrease in the time of deployment in all workloads. This implies that automated provisioning and orchestrating are useful in enhancing the operational efficiency.




Figure 1: Deployment velocity comparison across On-Prem, Single-Cloud, and Hybrid Multi-Cloud



Resource Utilization and Workload Portability

Resource utilization is useful in determining the level of efficiency of CPU, memory, and network bandwidth when workloads are executed. The experiments were recorded with regards to the average resources used on on-premises servers and on the cloud instances. The portability of workloads was evaluated through the result of the testing of performance consistency during a workload migration to different environments.

The findings indicate that hybrid multi-cloud systems keep resource use even, preventing over-provisioning of servers on on-premises/servers and under-utilization in the instances with the cloud system. The performance of containerized workloads deployed using Kubernetes was also similar across platforms, which is good evidence of the efficiency of workload portability.

Table 2: Average Resource Utilization (%)

Resource Type	On-Premises	Single-Cloud	Hybrid multi-cloud
CPU	72	65	68
Memory	70	63	67
Network Bandwidth	60	55	62

Hybrid deployments indicated some improvement in the efficiency of the network with the optimization of SDN routing, as well as the VPN tunnels linking on-premises and cloud nodes. Additional latency variations had a slight variation as the distributed databases burst to public clouds, which contributed to the fact that workload portability has been effective [2].




Figure 2: CPU and memory utilization trends across deployment scenarios.

Operational Governance

Operation governance measures the system in accordance with the policies set, automation rules, and standards of compliance. The metrics that have been gathered are frequency of manual process, the instances of policy violation, and the success rate of an automated configuration.



The hybrid multi-cloud environment decreased the number of manual interventions as construction of IaC created provisioning and orchestration between AWS, Azure, and on-premises. The technology of automated logging and monitoring also enhanced the level of governance and deviation to the policies of operations was easier to detect.

Table 3: Operational Governance Metrics

Metric	On-Premises	Single-Cloud	Hybrid multi-cloud
Manual Interventions per Week	8	5	2
Policy Violations per Month	6	4	1
Automation Success Rate (%)	78	85	94

These findings demonstrate that hybrid multi-cloud configurations yield better compliance, less error tolerance and a lower operational overhead, allowing IT individuals to invest in higher-level workflows as opposed to manual applications configuration and monitoring.




Figure 3: Comparison of operational governance metrics across deployment types.

Cost Optimization

A major factor that is important among businesses that implement hybrid multi-clouds is its cost-efficiency. The number of compute hours, storage use and consumption of network have been measured in terms of total operational costs depending on all the three deployment scenarios. The impact of cloud bursting of the cloud to public clouds to address the peak workload demands was also taken into account in the study.

Table 4: Operational Cost Comparison (USD per Month)

Workload Type	On-Premises	Single-Cloud	Hybrid multi-cloud	Cost Savings Hybrid vs On-Prem
Compute-Intensive App	2,400	1,800	1,650	31%
Distributed Database	3,200	2,400	2,150	33%
Containerized Microservices	2,800	2,000	1,850	34%



The concept of hybrid multi-cloud deployment can lead to cost reductions with the use of on-demand load on the cloud infrastructure during the peak load and on-premises infrastructure supporting baseline loads. This elasticity of use minimizes the requirement of over-provisioning of the local servers, as well as streamlines the use of cloud resources [2] [4].




Figure 4: Cost comparison across deployment scenarios for all workloads.

Summary of Findings

The quantitative analysis proves the hypothesis that multi-cloud hybrid strategies promise substantial advantage in the deployment speed, and realizing the use of resources, as well as governance of the operations and optimizing the costs. The deployment speed is enhanced by a maximum of 42 percent to the use of hybrid models over the on-premises only systems. There is also a balanced utilization of resources in the environment and the practicality of portability of workload is supported using containerization and orchestration software. Measures of operational governance show the decrease of manual interventions and the increase of the success rates of automation. Last, hybrid deployments allow saving more than 30 percent of operational expenses monthly on most workloads. The discoveries confirm the suggested reference architecture of the hybrid multi-cloud framework in which on-premises infrastructure is involved with AWS and Azure. Infrastructure as Code, integrated networking, identity federation, and automated orchestration together is an operation of a scalable, resilient, and economical way of enterprises to adopt hybrid cloud computing. The findings are also used to create a basis of decision matrices and architectural blueprint that would lead organizations in planning and implementing hybrid multi-clouds deployments.

V. CONCLUSION

The paper shows that hybrid multi-cloud deployment has important advantages compared to traditional or single-cloud one. The deployment speed increased up to 42%, the utilization of resources was done efficiently, the governance of operations was strengthened and the costs were lessened with an average of 34%. Containerization and orchestration led to the success of workload portability whereas IaC automation minimized manual interactions. The proposed reference architecture which combines on-premise infrastructure with AWS and Azure offers a model that offers scalability, resilience, and low costs to the enterprises. Such findings underscore the fact that hybrid multi-cloud



strategies have the potential of streamlining performance, security, and operational efficiency and can help companies design and deploy powerful multi-cloud.

REFERENCES

- [1] Caballer, M., Antonacci, M., Šustr, Z., Perniola, M., & Moltó, G. (2021). Deployment of elastic virtual hybrid clusters across cloud sites. *Journal of Grid Computing*, 19(1). <https://doi.org/10.1007/s10723-021-09543-5>
- [2] Mansouri, Y., Prokhorenko, V., & Babar, M. A. (2020, June 4). *An automated implementation of hybrid cloud for performance evaluation of distributed databases*. arXiv.org. <https://arxiv.org/abs/2006.02833>
- [3] Kumar, R. (2021). Multi-Cloud and hybrid cloud strategies – balancing flexibility, cost, and security. *International Journal for Multidisciplinary Research*, 3(2). <https://doi.org/10.36948/ijfmr.2021.v03i02.39459>
- [4] Buyya, R., & Barreto, D. (2015, November 28). *Multi-Cloud Resource Provisioning with Aneka: A Unified and Integrated Utilisation of Microsoft Azure and Amazon EC2 Instances*. arXiv.org. <https://arxiv.org/abs/1511.08857>
- [5] Varghese, B., & Buyya, R. (2017). Next generation cloud computing: new trends and research directions. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.1707.07452>
- [6] Ravindra, P., Khochare, A., Reddy, S. P., Sharma, S., Varshney, P., & Simmhan, Y. (2017). \mathbb{E} {ECHO}: An Adaptive Orchestration Platform for Hybrid Dataflows across Cloud and Edge. In *Lecture notes in computer science* (pp. 395–410). https://doi.org/10.1007/978-3-319-69035-3_28
- [7] Venkateswaran, S., & Sarkar, S. (2017). Architectural partitioning and deployment modeling on hybrid clouds. *Software Practice and Experience*, 48(2), 345–365. <https://doi.org/10.1002/spe.2496>
- [8] Deochake, S., & Channapattan, V. (2022). Identity and Access Management Framework for multi-tenant resources in hybrid cloud computing. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.2203.11463>
- [9] Saxena, D., Gupta, R., & Singh, A. K. (2021). A survey and comparative study on Multi-Cloud Architectures: Emerging issues and Challenges for Cloud Federation. *arXiv (Cornell University)*. <https://doi.org/10.48550/arxiv.2108.12831>
- [10] Deb, M., & Choudhury, A. (2021). Hybrid Cloud: A New Paradigm in Cloud Computing. *Hybrid Cloud: A New Paradigm in Cloud Computing*, 1–23. <https://doi.org/10.1002/9781119764113.ch1>