

Connecting Clouds in a Multicloud Strategy

Research on a Multicloud Approach and Practical Implementation

Lorenzo ELIAS
Student Bachelor Electronics-ICT – Cloud & Cyber Security

Preface

I am Lorenzo Elias, a student at Thomas More in Geel, studying a bachelor Electronics ICT with a specialization in Cloud & Cyber Security. With sincere gratitude for all the knowledge and experiences gained throughout my studies, I am proud to present this research.

To conclude my academic journey, I completed a 18-week internship at Cegeka in Hasselt, within the Network & Security team. Cegeka is a leading European IT solutions provider with a strong focus on innovation and digital information. This research not only marks the end of my student life, but also the beginning of a new chapter full of opportunities and challenges. During my internship, I worked on the technical and strategical aspects of multicloud strategies. With full commitment and enthusiasm, I researched the differences between major cloud providers and developed a proof of concept (PoC) for a fully integrated multicloud environment. The aim was not only to understand how such architectures can be built, but also to evaluate when and why they are worth implementing in real business contexts.

At Cegeka, I had the privilege of being surrounded by a dedicated and inspiring team. The support and guidance gave me space to grow professionally. Despite the pressure that sometimes comes with such a project, I always seek challenges that allow me to push my own boundaries. My curiosity, passion, and perseverance kept me focused and motivated to deliver the best result possible.

I would like to express my sincere thanks to my mentors, especially Toon Vanoeteren, Dirk Stegen, and Jan Weyens. Toon guided me from the beginning and was always available for feedback and advice. Jan and Dirk, my internship supervisors, encouraged me to work independently and to think critically. I also want to thank the entire Network & Security team for making me feel welcome, and for creating an environment in which I could be myself. This team made Cegeka feel more than just a workplace, it felt like a place where I truly belonged.

Looking back at my time at Thomas More and Cegeka, I feel a deep sense of appreciation. This experience has confirmed my passion for the IT world and has laid the foundation for the professional I want to become.

Lorenzo Elias

Abstract

Cloud computing has become an essential component of modern IT infrastructure, with organizations increasingly adopting public cloud platforms such as Microsoft Azure, Amazon Web Services (AWS), and Oracle Cloud Infrastructure (OCI). This thesis explores how a multicloud environment can be implemented, and under what circumstances such a strategy is justified. Beyond a theoretical comparison of the technical and strategic differences between platforms, a proof of concept (PoC) was developed to demonstrate the technical feasibility of integrating multiple cloud providers, such as Azure and OCI, and securely connecting them together with an on-premises datacenter using site-to-site VPN tunnels. The setup uses hub-and-spoke networks, firewall clusters, and custom routing. While a multicloud setup offers benefits such as flexibility, and workload optimization, this research also shows that multicloud should never be the default, but it must be a deliberate strategy driven by clear business or technical needs. Results demonstrate increased complexity and hidden costs. In many cases, these factors make multicloud more expensive than single cloud setups. However, in specific scenarios, such as compliance requirements, workload specialization, or client driven demands, a multicloud strategy can be a valid and strategic choice. This project contributes to Cegeka's internal expertise and offers practical insights for organizations considering a multicloud approach. For such organizations, developing multicloud expertise offers a competitive advantage, enabling support for a broader client base and unlocking new commercial opportunities.

Table of contents

INTRODUCTION	7
1. CLOUD ENVIRONMENTS	9
1.1. History and Market Share of Cloud Providers	10
1.2. Definition of a Multicloud Environment	11
2. DIFFERENT CLOUD PROVIDERS	12
2.1. Global Infrastructure and Terminology	12
2.2. Services	14
2.3. Strengths and Specialized Services	14
2.4. Identity and Access Management (IAM)	15
2.4.1. IAM in AWS	15
2.4.2. IAM in Azure	16
2.4.3. IAM in OCI	16
3. MULTICLOUD STRATEGIES	17
3.1. Benefits and Challenges of a Multicloud Strategy	17
3.2. Justified Use Cases	18
3.2.1. Use Case 1: Workload Specialization Across Providers	19
3.2.2. Use Case 2: Business Mergers and Customer Constraints	19
3.2.3. Use Case 3: Customer or Compliance-Driven Platform Separation	19
3.3. When Multicloud is Not Justified	19
3.4. Cost Considerations	20
3.4.1. Scenario 1: Multicloud Deployment	20
3.4.2. Scenario 2: Single Cloud Deployment	21
3.4.3. Cost Interpretation and Broader Considerations	22
4. PROOF OF CONCEPT	25
4.1. Infrastructure Overview	25
4.2. Infrastructure in Azure	27
4.2.1. VNets and Subnets	27
4.2.2. FortiGate Cluster Configuration	28
4.2.3. Load Balancer	28
4.2.4. Azure VPN Gateway and VPN Connectivity	29
4.2.5. Routing and Peering	29
4.3. Infrastructure in OCI	30
4.3.1. VCNs, Subnets, and FortiGate Cluster	31
4.3.2. VPN Connectivity and DRG	31
4.3.3. Routing and VCN Peering	31
4.4. Proof of Implementation	31
4.4.1. Connectivity in Azure	32
4.4.2. Connectivity in Azure to the Internet	32

4.4.3.Connectivity between Azure and the Datacenter	33
4.4.4.Connectivity between OCI and Azure	34
4.4.5.Connectivity in OCI	35
5. CONCLUSION	36
REFERENCE LIST	38
ATTACHMENTS	41

Introduction

Cloud computing differs from earlier computing models by not only facilitating resource sharing, but also providing elasticity, scalability, self-service provisioning, and centralized management. These capabilities allow organizations to have more efficient operations, reduce costs, and increase flexibility. Instead of having large upfront capital expenditures, businesses can operate on the pay-as-you-go model, where they only pay for what they use. This reduces financial risks and offers more possibilities for innovation and scalability as companies have the freedom to innovate without limitations to their infrastructure (Susnjara & Smalley, 2025). Beyond cost savings, cloud adoption improves collaboration, as employees can access data and applications from virtually anywhere. Given the growing need for flexibility, speed, and resilience in today's IT, and other, industries, the cloud has become a foundational element in organization's strategies. In today's business landscape, it's no longer a question whether to move to the cloud, but what to move and how it is best to do it (Google Cloud, s.a.).

Figure 1 illustrates the estimated growth of the global cloud computing market, highlighting its expected expansion from \$626.4 billion in 2023 to \$1,266.4 billion in 2028, which would result in a compound annual growth rate (CAGR) of 15.1%. The surge is driven by increasing enterprise adoption, as businesses seek to decrease their costs while increasing all-round efficiency and gain a competitive advantage. Additionally, macro-economic factors such as supply chain disruptions and evolving data privacy regulations are shaping investment strategies, further accelerating cloud adoption worldwide (Brophy, 2025).




Figure 1. Expected CAGR of the Global Cloud Computing Market

Note. Adapted from "Cloud Computing Market," by MarketsandMarkets, 2023, MarketsandMarkets. <https://www.marketsandmarkets.com/Market-Reports/cloud-computing-market-234.html>

Multicloud is a cloud computing strategy that can offer several potential benefits. In this so-called multicloud strategy, organizations can meet their needs by combining services from different cloud providers. The strategy wants to base itself on the optimization of performance on workloads, increased flexibility and limiting the risk of relying on one supplier. However, the use of a multicloud environment should be a well-considered choice. In addition to the many benefits that this strategy offers, a multicloud environment can also have its disadvantages.

This thesis addresses how a multicloud environment can be implemented and under which circumstances such an approach is justified to use. The research examines and compares key services offered by leading cloud providers, specifically Amazon Web Services (AWS), Microsoft Azure and Oracle Cloud Infrastructure (OCI). Furthermore, the focus lies on identifying which of these services are essential to implement a multicloud architecture. Each cloud provider is also assessed individually to highlight its specific strengths. In addition to the theoretical analysis, a proof of concept (PoC) is developed, emphasizing the focus on interconnectivity between the different cloud platforms. Although AWS is part of the technical research, it is not included in the PoC due to technical and time constraints. The cloud environments in the PoC will have their own hub-and-spokes architecture. In addition, the setup includes firewall clusters, private and public connectivity, and site-to-site VPN connectivity. This research aims to contribute to Cegeka's internal knowledge base by comparing different cloud providers, exploring multicloud use cases, and identifying the technical components required for such environments.

1. Cloud Environments

Before diving into specific cloud providers or multicloud environments, it is important to understand the building blocks of cloud computing. Two essential concepts are used when an organization approaches the cloud: the way environments are deployed, called deployment models, and the level of services offered in the cloud, called service models. Together, these terms define how cloud infrastructure is structured and managed.

Depending on the level of control, security, and flexibility that's required, organizations can choose between different deployment models. Generally speaking, these are subdivided into three models. First, a public cloud, where IT resources such as compute and storage are provided by third party vendors like AWS, Azure or OCI. The services are accessible over the internet, with multiple customers sharing the same underlying infrastructure in isolated environments. In contrast to that, there is a private cloud, where an environment is dedicated to one organization. It can be hosted on-premises or by a third party provider and offers more control over networking, security, and compliance, often important for businesses dealing with sensitive data. Between the previous models is the hybrid cloud, which combines public and private environments. It allows workloads to move between the two, offering more flexibility and scalability (Microsoft, 2025).

In addition to deployment models, cloud computing is defined by its service models: Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). As shown in Figure 2, IaaS provides infrastructure such as servers and storage, where the customer remains responsible for the operating systems, applications, and data running on them. PaaS often offers a managed environment for developers to build and deploy applications without the need to handle underlying infrastructure. While SaaS delivers fully managed applications, such as Microsoft 365 or Google Workspace, where the provider handles everything, from hosting to updates and security (IBM, 2021).




Figure 2. Cloud Service Models: SaaS, PaaS, IaaS Compared

Note. Adapted from "IaaS vs PaaS vs SaaS: Explaining the Key Differences," by A. Dua, 2024, Turing. <https://www.turing.com/blog/iaas-vs-paas-vs-saas-key-differences>

1.1. History and Market Share of Cloud Providers

AWS was founded in 2006, when Amazon stumbled on the fact that it was, and still is, hard to provision and manage IT infrastructure, which distracted their talented teams from actually innovating. Amazon wanted to rethink IT infrastructure completely, so every individual could access the same power of cloud computing as the world's largest companies. It started with foundational services such as Amazon Simple Storage Service (S3 Buckets), which allowed individuals and organizations to store and manage their data in a cost-effective way. They also offered the Elastic Computing Instance (EC2), scalable virtual servers in the cloud, allowing users to run applications by launching instances with the chosen Operating System (OS), CPU, memory, storage, and networking capacity. AWS quickly gained popularity for their scalability and pay-per-use pricing models and grew into a world leading cloud platform. As of today, it is estimated that AWS holds a market share of 30%, as shown in Figure 3, making it the largest provider in the global cloud computing industry (AWS, 2025).




Figure 3. Market Share of Different Cloud Providers

Note. Adapted from "Cloud Infrastructure Market: Amazon and Microsoft Stay Ahead in Global Cloud Market," by F. Richter, 2025, Statista. <https://www.statista.com/chart/18819/worldwide-market-share-of-leading-cloud-infrastructure-service-providers/>

Following AWS's early success, other technology giants began building their own cloud platforms. Microsoft launched Azure in 2010, entering the market with a strong focus on enterprise clients. Thanks to smooth integration with existing Microsoft products, such as Active Directory, Office 365, and Windows Server, Azure quickly became a natural choice for organizations already embedded in the Microsoft ecosystem. Its support for hybrid cloud environments, made it especially attractive to businesses transitioning from traditional on-premises infrastructure to cloud environments. As a result, Azure captured the second-largest position in the cloud market. Holding an estimated 21% market share in the end of 2024 (Abandy, 2022).

Meanwhile, OCI followed a more specialized trajectory. Launched several years later, in 2016, OCI did not aim to immediately compete with AWS or Azure in terms of service breadth. Instead, Oracle focused on performance, cost efficiency, and secure infrastructure, positioning itself as the go-to platform for high performance workloads and large scale enterprise databases, particularly Oracle's own database products. OCI is especially appealing in industries where performance per core, predictable pricing, or database licensing optimization are critical. While its global market share remains modest, estimated 3%, OCI has gained ground through strategic partnerships, such as its direct interconnect with Azure, and is increasingly considered in multicloud strategies for specific, high demand workloads (Gartner et al., 2024).

As shown throughout this chapter, each cloud provider has taken a different path in reaching its current position in the market. AWS leads with an early-mover advantage and vast service portfolio, Azure capitalizes on its strong enterprise integration, and OCI emphasizes performance and cost optimization. However, market presence alone does not determine which provider is the best fit for particular use cases. In the next chapter, a deeper comparison is made between the providers, giving better insights that will form the bases for the multicloud PoC that follows later in this thesis.

1.2. Definition of a Multicloud Environment

A multicloud environment refers to the use of cloud services from two or more different cloud providers within a single organization. Rather than depending on one vendor for all infrastructure, database, or software services, organizations can adopt multiple cloud platforms, such as AWS, Azure, and OCI in parallel. This approach can offer great advantages, which will be explained in Chapter 3, Multicloud Strategies. Multicloud environments can take different forms, depending on the organizational needs and on how the cloud platforms are used and whether they are connected to each other (Chen, 2025).

One option is an isolated multicloud environment, where each cloud platform is used for specific workloads, but the environments are not technically connected. For example, an organization might use Azure for its Microsoft based applications and OCI for hosting Oracle databases. In this case, the platforms operate independently, and data or services are not exchanged between them (Bell, 2019).

Multicloud environments can also use an integrated multicloud architecture, where multiple cloud environments are technically linked to each other, typically through the use of VPN tunnels, allowing secure communication and the exchange of data between the environments. This form of multicloud makes it possible to distribute workloads or other services across multiple environments. In practice, this means that multiple service models, such as IaaS or PaaS may coexist within the same environment. While this can bring added value, it also introduces complexity (F5, s.a.).

2. Different Cloud Providers

In a single-cloud scenario, understanding a provider's terminology is relatively straightforward. However, in multicloud setups, like the one developed in this thesis, it is essential to understand how similar concepts are implemented differently. The foundational elements will be further examined in the following sections, where each cloud provider is analyzed more in detail based on their unique characteristics and how they were applied in the PoC developed for this thesis.

2.1. Global Infrastructure and Terminology

In the rapidly evolving landscape of cloud computing, a handful of cloud providers have emerged as global leaders in the multi-billion dollar industry. In this thesis, there will be a deeper dive into three of them: Amazon Web Services, Microsoft Azure, and Oracle Cloud Infrastructure. As the providers strive to gain the most market share, they need to be competitive. Because of this, each platform offers a wide range of services across infrastructure, platform, and software layers. While all three serve the same fundamental goal "delivering scalable and on-demand computing resources", they differ in terms of focus, market positioning and ecosystem integration. To strengthen their presence in the market and ensure low-latency access to services, cloud providers invest heavily in expanding their geographical footprint. Figure 4 illustrates the presence of regions across the globe from AWS (yellow), Azure (blue), and OCI (brown). Important to note is that this only points out the regions, and not the actual datacenters within it, as one provider may have multiple datacenters in one region where another may only have one.




Figure 4. Presence of Regions Across the Globe: AWS, Azure and OCI

Note. Adapted from "Microsoft Datacenters," by Microsoft, 2025, Microsoft. <https://datacenters.microsoft.com/globe/explore?view=map>; and "AWS Global Infrastructure," by AWS, 2025, AWS Amazon. <https://aws.amazon.com/about-aws/global-infrastructure/>; and "Public Cloud Regions," by Oracle, 2025, Oracle. <https://www.oracle.com/be/cloud/public-cloud-regions/>

Although each cloud provider uses its own terminology and internal structure, the foundational architecture is largely the same. Despite the naming differences, the core concept remains: cloud infrastructure is globally distributed and logically segmented to

provide scalability, fault tolerance, and low latency. To illustrate this structure, Figure 5 shows how Azure organizes a region. Each Azure region consists of at least one but often multiple availability zones. An availability zone in Azure represents a geographical location where one or more datacenters can reside, each with its own power, cooling, and networking. The datacenters are interconnected through high-speed fiber networks. While the exact implementation may differ, all cloud providers enable the possibility of critical services to be deployed across multiple regions or zones, ensuring continuity even if one zone experiences a failure (Microsoft Learn, 2025).




Figure 5. Cloud Region Architecture with Redundancy

Note. Adapted from "What are Azure Regions?," by Microsoft Learn, 2025, Microsoft. <https://learn.microsoft.com/en-us/azure/reliability/regions-overview>

Although Figure 5 visualizes how Azure structures its regions and availability zones, it is important to understand that the same architectural principles exist across cloud providers, only the naming is different. For instance, AWS also organizes its infrastructure into regions and availability zones, using near identical definitions. Each AWS region contains multiple availability zones, which are isolated datacenters with independent power, cooling, networking, connected to each other with low-latency links (AWS, 2025).

OCI on the other hand, structures its regions using availability domains instead of availability zones. Each availability domain typically consists of a single datacenter, unlike Azure's or AWS's zones, which often span multiple datacenters. OCI's availability domains still offer fault isolation and redundancy, but the physical separation is usually less extensive than that of its competitors. However, this design allows for ultra-low latency

and high-bandwidth connections between availability domains in the same region, which is beneficial for high-performance and database-heavy workloads (Oracle, 2025).

2.2. Services

Cloud providers may vary greatly in branding, ecosystem design, and platform experience, but at their core, they offer comparable service categories designed to meet the same foundational needs. Whether deploying virtual machines, storing data, managing identity, or deploying databases, AWS, Azure, and OCI all provide a similar range of services that cover the fundamental parts of the different cloud computing models.

While the names, interfaces, and some implementation details differ, the core functionality remains consistent. This allows organizations to design cloud architectures based on service type, rather than provider specific-features, at the foundation level. It is only when diving into specialized offerings, pricing models, performance optimizations, and ecosystem integrations that meaningful differences begin to emerge. These differences are explained in the next chapter (Oracle, 2024).

Table 1 offers an overview of the core cloud service categories and the corresponding naming used by each provider.

Table 1. Key Cloud Services per Provider

Functional Area	AWS	Azure	OCI
Compute	EC2	Virtual Machines	Compute Instances
Object Storage	S3	Blob Storage	Object Storage
Relational Database	RDS	Azure SQL DB	Autonomous Database
Container Services	ECS/EKS	Azure Kubernetes Service	OCI Kubernetes Service
Identity & Access	IAM	Microsoft Entra ID	IAM

2.3. Strengths and Specialized Services

While the core functionalities of AWS, Azure, and OCI often overlap in multiple areas, each provider also offers specialized services that differentiate them in the market. These unique strengths can be a decisive factor in multicloud strategies, where organizations selectively use the best capabilities of each platform to optimize their workloads.

AWS has long been recognized as a pioneer in serverless computing. It offers a vast array of services across nearly every domain, compute, Artificial Intelligence (AI), databases, storage, analytics, and more. This can make AWS a perfect choice for startups but also for large organizations like product-driven companies aiming to build and scale software quickly and cost efficiently. From a technical standpoint, AWS is widely recognized as a pioneer in serverless computing, with AWS Lambda enabling highly scalable, event-driven architectures (Barr, 2024). AWS’s pricing flexibility and global reach can make it a good choice for organizations that value cost control and breadth of features (AWS, 2025).

Azure is strongly positioned as the natural choice for enterprises already embedded in the Microsoft ecosystem. Its integration with Active Directory, Office 365, and Windows

Server enables organizations to extend their existing infrastructure into the cloud with minimal disruptions. Microsoft also performs well in hybrid cloud scenario, with services like Azure Arc allowing on-premises and multicloud environments to be managed as one unified platform. This makes Azure especially appealing within existing organizations who already have a large environment running in the Microsoft ecosystem and seek to shift their infrastructure to the cloud (Krishnakumar, 2025).

OCI targets a more specific segment of the cloud market: large enterprises with demanding back-office and data workloads, especially those already using Oracle software. OCI's strength lies in performance-optimized, cost efficient infrastructure, which can be a great solution for critical applications like ERP, finance, and database intensive systems or applications. OCI also distinguishes itself with its capabilities in bare-metal infrastructure, often at lower cost than traditional hyperscalers. This makes OCI particularly attractive to organizations who seek predictable performance for internal workloads with tight licensing (Oracle, 2024).

2.4. Identity and Access Management (IAM)

IAM is a critical pillar in securing any cloud platform. It defines how users and systems authenticate and gain access to cloud resources, and how those permissions are enforced. Although AWS, Azure, and OCI all provide IAM solutions, they differ significantly in structure, terminology, and integrated options. Each cloud provider enforces permissions using its own native model, making it practically necessary to use their platform specific IAM model to implement granular access control. While basic authentication and authorization concepts are consistent, such as roles, policies, and group-based access, the scope, structure and enforcement mechanisms vary, especially when dynamic permissions are involved (Mishra, 2024).

In multicloud environments, this becomes more complex. Managing identities across providers is not natively integrated, and federation between IAM systems requires third party tools or manual configuration. Maintaining user identities in multiple locations is both inefficient and increases the chances on configuration errors. The key takeaway is that granular and secure access control can only be achieved by adopting the native IAM model of each cloud platform, which increases the operational complexity when more than one provider is involved (Mishra, 2024).

2.4.1. IAM in AWS

AWS IAM provides flexible access control through users, groups, and roles. Granular permissions are defined using JSON-based policies, but secure and scalable use requires adopting AWS's native structure fully. Federated identity is supported, but integration with non-AWS systems depends on additional configuration and external tools (AWS, s.a.).

2.4.2. IAM in Azure

Azure uses Entra ID for identity management, with strong hybrid capabilities due to its integration with on-premises Active Directory. Access control is enforced using role assignments at multiple scopes. While federation is supported, fine-grained access requires the full use of Microsoft's Entra framework and IAM model (Microsoft Learn, 2025).

2.4.3. IAM in OCI

OCI organizes identities using tenancies, compartments, and group-based policies in a readable syntax. Dynamic groups enable cloud-native access control, but granular access control still depends on the use of OCI-native tools. Federated identity is possible but requires a manual setup (Oracle, 2023).

3. Multicloud Strategies

As mentioned before, a multicloud strategy refers to the use of cloud services from multiple cloud providers within a single organization, based on strategic, technical, or operational considerations. It goes beyond the incidental use of different platforms and instead represent a deliberate architectural approach. The concept is visually represented in Figure 6.



Figure 6. AWS, Azure, and OCI: In a Multicloud Perspective

Note. Figure created by the author.

When an organization considers moving its infrastructure or workloads to the cloud, the central question is not merely how to migrate, but where to migrate. With the growing maturity of multiple cloud providers, the possibility of a multicloud strategy becomes increasingly relevant. However, adopting a multicloud strategy is not a default recommendation. Instead, it must be the result of a deliberate and well-justified decision, keeping in mind the organization's technical needs, business priorities, and operational capacity (Google Cloud, s.a.).

For IT service providers such as Cegeka, the question "*Which cloud platform is most suitable for this client?*" cannot be answered in isolation. It depends on a combination of factors such as workload requirements, positioning of data, licensing conditions, integration needs, and long-term flexibility. Only when the advantages of a multicloud environment clearly outweigh the disadvantages like the added complexity and cost, a multicloud architecture is a responsible choice.

3.1. Benefits and Challenges of a Multicloud Strategy

A well-considered multicloud strategy offers organizations the opportunity to take full advantage of the strengths of multiple cloud providers. By leveraging platforms, such as AWS, Azure, and OCI in parallel, organizations can gain more flexibility in how workloads are deployed and managed. This increased freedom results in multiple benefits.

One of the primary advantages is avoiding vendor-lock in. By distributing workloads across multiple platforms, organizations are less dependent on a single provider's

ecosystem, reducing the risk of contractual or technical constraints. Additionally, multicloud enables optimal workload placement: each cloud platform has its own specializations. Azure's enterprise integration, AWS's broad service portfolio, or OCI's performance-focused compute and database services, allowing organizations to match workloads with the most appropriate infrastructure. Another important benefit is increased resilience. The ability to run services across different cloud environments can significantly improve availability and fault tolerance. In case of a provider outage or service degradation, workloads could failover to another provider. The geographic diversification can also help organizations with compliance and data privacy requirements, like GDPR in Europe. Cost optimization is another important benefit. By comparing pricing models, which will be done later in this thesis, organizations can control expenses more effectively and avoid being locked into premium services unnecessarily. Furthermore, innovation is accelerated by multicloud strategies, as organizations get access to a broader range of services, ranging from machine learning and analytics to specialized storage offers, without having to wait for a single provider to catch up (Teske, 2024).

Despite these multiple advantages, adopting a multicloud strategy introduces several challenges. From a technical perspective, managing identity, networking, routing and security policies across multiple different platforms can become complex. Each provider uses its own standards, tools, and terminology, which complicates unified usage. Without careful planning, this can lead to inefficiencies, security gaps, and unforeseen costs. Visibility is also a common concern, monitoring and logging must often be unified by using third-party solutions, as native tools typically lack cross-provider integration. This increases the overhead of ensuring consistent observability, cost tracking, and incident response. Furthermore, maintaining expertise across multiple platforms requires much broader skillsets within IT teams, potentially increasing training costs. Security and compliance management also become more nuanced. Ensuring consistent security policies across multiple providers requires extra attention, especially when dealing with access control, multiple firewalls in different cloud environments, encryption policies, and auditing. Similarly, deploying Infrastructure-as-Code (IaC), where organizations can setup cloud environments automated with code, can become more challenging as different platforms can require specific templates and workflows (Sasovets & Teres, 2025).

3.2. Justified Use Cases

As previously mentioned, the decision to adopt a multicloud strategy depends on a variety of factors. Not every scenario will profit from a multicloud approach, and it should never be seen as a default. In this section, a selection of representative use cases is provided to illustrate when a multicloud strategy is applicable. These examples highlight practical motivations and serve to demonstrate the strategic value of multicloud strategies beyond the technical implementations.

The use cases discussed do not cover every possible scenario, but rather aim to provide a relevant set that aligns with the objectives of this thesis. They form the conceptual foundation for the PoC designed in this thesis, and help answer the question of when a multicloud strategy is justified.

3.2.1. Use Case 1: Workload Specialization Across Providers

One of the most common strategic justifications for multicloud is functional specialization. As mentioned in previous chapters, each cloud provider offers unique services or cost advantages. For example, Oracle's database services are more performant and license-efficient in OCI, while Azure excels in application hosting and enterprise integration with services like Microsoft Entra ID. In such cases, splitting workloads across OCI and Azure can result in higher performance and lower costs, provided that the connectivity between the platforms is secured and monitored. The PoC built in this thesis demonstrates that such communication between cloud environments is technically feasible using encrypted VPN tunnels and centralized inspection through firewalls, although such a setup can also occur in a non-technical integrated multicloud (Chen, 2025).

3.2.2. Use Case 2: Business Mergers and Customer Constraints

In the context of mergers, acquisitions, or IT standardization, multicloud can be a necessity. Different business units may have already invested heavily in different cloud platforms. Rather than forcing a costly and time consuming migration from one platform to another, organizations can profit from a multicloud strategy by connecting their environments and allow for phased transitions or cross cloud communication. One of the merged companies might require Azure for sovereignty reasons, while the other demands Oracle-based infrastructure. In such cases, service providers profit from offering multicloud support to remain compliant and competitive (Chen, 2025).

3.2.3. Use Case 3: Customer or Compliance-Driven Platform Separation

In some cases, client requirements dictate the use of specific cloud providers. For example, an IT service provider may serve a government agency that mandates the use of Azure due to data residency and compliance standards, while other clients prefer AWS for flexibility and cost. The service provider can support both by deploying applications and services in the foreseen clouds, without connecting them at infrastructure level. This setup supports a multi-tenant architecture across clouds, while keeping environments isolated for security, legal, or contractual reasons. Yet, it still qualifies as a multicloud strategy due to the intentional provider split and coordinated management (Appvia, 2024).

3.3. When Multicloud is Not Justified

While multicloud strategies can offer flexibility and resilience, they also introduce architectural and operational overhead. There are numerous situations in which multicloud is not recommended, either because the risks outweigh the benefits, or because simpler alternatives are more appropriate. Small to medium large organizations often lack the internal knowledge and expertise to properly manage and secure multiple cloud environments. Each platform comes with its own tools, services, configurations, and best practices. Managing all that requires specialized personnel and constant training. So if an organization's workloads can be efficiently hosted on a single cloud without performance, compliance, or service limitations, then introducing a second provider adds unnecessary complexity.

3.4. Cost Considerations

While the technical feasibility and strategic goals of a multicloud architecture may be clear, cost remains one of the most decisive factors in determining whether such a strategy is justified or not. The adoption of the strategy is not just a technical choice, but a fundamental business decision. To illustrate how cost can influence a cloud strategy, this section presents a comparison of a multicloud deployment versus a single cloud deployment. The scenario is based on the practical context explored in this thesis, but it is important to note that this represents only one of the countless possible configurations. Cloud pricing is highly variable and context dependent. Influencing factors on the price are workload types, traffic patterns, service tier selection, the tier of support that is chosen, and not to forget, the negotiated enterprise agreements. The analysis made below does not aim to generalize pricing scenarios, but rather to highlight the importance of including cost as a core decision criteria in any multicloud evaluation (Tata Communications, 2025).

The following cost comparisons are based on theoretical configurations using the official pricing calculators provided by AWS, Azure, and OCI. While the PoC developed for this thesis involved deploying virtual machines merely to answer the question *"What is technically needed to deploy a multicloud environment?"* and to prove that interconnectivity is possible. Therefore, the cost estimates should be interpreted as indicative rather than actual data from an operational environment. The comparison is based on a simplified SaaS application that allows end users to submit data through a web interface and view personalized analytics in a dashboard. It includes a static frontend, serverless backend for processing user input, persistent storage, and a reporting layer.

3.4.1. Scenario 1: Multicloud Deployment

Table 2 provides a detailed overview of the cloud components used per provider, along with their estimated monthly cost (see Attachment 4 to Attachment 6). The multicloud setup distributes services across Azure, AWS, and OCI to leverage specific strengths of each provider. A web frontend is hosted in AWS using S3, providing static content deliver. Backend logic is executed using AWS Lambda and API gateway, handling user input. Persistent data is stored in OCI's Autonomous Database, which is optimized for transactional workloads. Analytical dashboards are served through the use of Azure's Power BI embedded, maintaining a consistent reporting experience. Connection between the environments is done through site-to-site VPN's, and bandwidth usage is accounted for.

Table 2. Multicloud Cost Breakdown by Component per Provider

Component	Cloud	Service	Monthly Cost (USD)
Frontend	AWS	S3	\$1.28
Backend/API	AWS	Lambda + API Gateway	\$32,30
VPN + Egress	AWS	VPN	\$165,16
Database	OCI	Autonomous DB	\$118,41
Storage	OCI	Object Storage	\$1,25
Dashboarding	Azure	Power BI Embedded	\$353,43
VPN Gateway	Azure	VPN Gateway	\$138,70
Internet Egress	Azure	Bandwidth	\$72,00
Blob Storage	Azure	Blob	\$5,88
Total monthly Cost			\$888,41

3.4.2. Scenario 2: Single Cloud Deployment

In the analysis of the single cloud deployment, there is a comparison made between Azure, AWS, and OCI. To better understand the price differences between each cloud in this scenario, the same SaaS solution was priced across the cloud providers using the official pricing calculators, which can be seen in Attachment 1 to Attachment 3.

Each platform was configured using equivalent services where possible. For instance, AWS Lambda was compared with Azure Functions and Oracle Functions, while storage and database options were selected based on default configurations and with the bring your own license pricing where possible.

Figure 7 below illustrates the cost per component per provider for running the solution within a single environment. Such comparisons should always be made within a company when selecting a cloud provider, although they may not only focus purely on pricing of services, as explained in Chapter 3.4.3, Cost Interpretation and Broader Considerations.

Some services, like the VPN/egress, dashboard, and database show strong variations. These differences are primarily because of how each provider structures its pricing model and the specific services used to fulfill the same functions. For example, Azure prices its VPN gateway component, but a fixed amount of money is paid for 0 to 10 site-to-site VPN tunnels. AWS on the other hand, charges per site-to-site connection. OCI strongly focuses on cost optimization and charges nothing for the first 10TB egress data.

The comparison helps illustrate that even when functionally almost identical, the same cloud-native applications can vary greatly in monthly costs depending on the provider used. However, it is important to remind that these values are theoretical and based solely on public pricing calculators, and no billing data was used from active usage.




Figure 7. Cost Comparison of Components in a Single Cloud Setup

3.4.3. Cost Interpretation and Broader Considerations

This section aims to go beyond a simple comparison of pricing, and is made to highlight the underlying variables that need to be accounted for when evaluating multicloud versus single cloud deployments. By presenting the theoretical costs alongside the factors that influence them, this analysis gives a more informed and realistic conclusion on which cost are often unforeseen and may not be immediately visible during initial planning.

Figure 8 below summarizes the total monthly cost per provider and a multicloud scenario, based on a simple SaaS application in different architectures. It offers a simplified but informative comparison of how pricing differs depending on the selected cloud strategy.




Figure 8. Comparison of Monthly Cost in a Multicloud vs. Single Cloud Setups

While the results appear to provide clarity on paper, the reality is more complicated. Cloud cost is not a fixed number, it is formed by a wide variety of technical, operational, and business factors. The pricing calculators used in the analysis only give a part of the picture and are based on the public pricing tiers, these don't reflect actual enterprise agreements or consumption patterns.

A large factor that is excluded, and often overlooked in calculator based consumptions, is personnel and operational overhead. Managing multiple cloud environments greatly increases complexity across the whole environment. Each platform has its own services, terminology, and tooling, requiring IT engineers to have knowledge in multiple environments. This means that organizations must invest heavily in training, certification, and possibly hiring staff familiar with different platforms. Organizations that already have broad experience within their IT teams may experience less overhead and training cost, while the smaller or less mature IT organizations will experience a steep learning curve with associated costs.

Licensing conditions further complicate the comparison. Most providers allow customers to bring their own license when it comes to databases, virtual machines, and other services. While others include licensing in the service price, depending on the service. So customers using Microsoft technologies may benefit from Azure Hybrid benefit or discounted rates on reserved instances. OCI offers advantages for running Oracle workloads natively, and AWS provides cost flexibility without various saving plans. Configurations in this comparison used pay-as-you-go pricing, which offers maximum flexibility but comes at a higher price. In practice, large organizations can often negotiate custom pricing agreements or get committed use contracts that offer better pricing. Enterprises consuming large volumes of cloud services frequently receive better benefits not reflected in the standard pricing calculators. These elements are critical when determining the right strategy, and they can cause a theoretical gap in cost in comparison to real world deployments.

Another crucial variable is data volume and data transfer behavior. The scenario used for the comparison assumes that there is minimal inter-service traffic and relatively small data volumes, as only outbound internet traffic is estimated at 50GB/month. In real deployments, workloads may produce way more data, require more bandwidth, or may need more frequent interaction between components in different clouds. This is especially relevant in multicloud environments where intern communication of workloads may have to traverse site-to-site VPNs, often generating outbound traffic charges in multiple environments. Even intern network traffic in clouds can lead to unexpected costs, such as egress between regions or availability zones. These data movement patterns are highly scenario specific and can scale costs drastically beyond what is initially estimated.

Each provider also includes or excludes certain fees in its default pricing model, adding a layer of unpredictability. AWS, for example, charges for NAT gateway usage, load balancer processing, or traffic between availability zones. Azure can generate extra costs through monitoring and diagnostics services like Log Analytics. OCI includes relatively large outbound data allowances, but when these thresholds are exceeded, costs can rise quickly. These types of charges are often hard to calculate during the planning phase, especially for organizations with dynamic or scaling workloads.

Ultimately, while cost is an important factor, the operational goals and business context of an organization also help to determine which strategy is the right choice. A provider that is more expensive may still be the better fit if it offers tighter integration, licensing advantages, or reduced time when it comes to bringing new services to the market. Cloud selection should therefore never be based on cost alone, but also in context of strategic alignment, service specialization, and organizational capability.

4. Proof of Concept

The goal of this Proof was to explore the practical implementation of a technically integrated multicloud architecture, connecting Azure and OCI, while ensuring secure access to the on-premises datacenter from one cloud provider. This PoC was designed in the context of Cegeka's internal innovation initiatives, aiming to increase its knowledge based around multicloud environments.

A central requirement was to ensure secure and reliable network connection between the cloud environments and a connection to the on-premises datacenter. This involved configuring VPN tunnels, firewall clusters, and designing custom routing paths to direct traffic appropriately. The PoC focuses on infrastructure-level connectivity and control in a multicloud setup.

The implementation of the PoC was carried out in several stages, each focusing on a specific cloud environment or connectivity layer. The goal was to incrementally build the underlying infrastructure, ensuring that each cloud environment was functional in isolation before establishing cross-cloud communication.

4.1. Infrastructure Overview

The multicloud environment designed for this PoC consists of two cloud platforms, Azure and OCI, each configured with a distinct, yet comparable architecture. Both cloud environments follow a hub-and-spoke topology, where the hub virtual network or cloud network acts as the central point of connectivity, and the spokes host isolated workloads.

In Azure, the hub VNet includes a FortiGate firewall cluster in an active-active setup and a VPN gateway, while the spokes contain VM's that could simulate workloads within isolated subnets. Similarly, in OCI, a VCN was created for the hub, containing a FortiGate firewall cluster with spokes deployed in different VCNs and subnets.

To connect both cloud environments, a site-to-site VPN between Azure and OCI is used. Using IPsec tunnels, this connection ensures encrypted, reliable communication between the clouds. Additionally, a second VPN tunnel connects Azure to the on-premises datacenter, allowing connectivity from on-premises to the cloud.

Each firewall cluster acts as the primary ingress and egress point in its respective environment, and custom routing tables were configured to direct all traffic, including intern traffic, through the firewalls. Figure 9 illustrates the high-level architecture.




Figure 9. Multicloud Architecture Between Azure and OCI with a Connection to an On-Premises Datacenter

4.2. Infrastructure in Azure

This section describes the design and deployment of the Azure environment within the multicloud PoC. As Figure 10 shows, the implementation follows a hub-and-spoke network model, equipped with a firewall cluster to control and inspect the traffic.




Figure 10. Azure Cloud Architecture with a Hub-and-Spokes Model

4.2.1. VNets and Subnets

The Azure implementation started with the creation of the hub-and-spoke network topology. Central in the topology lies the HUB VNet, which serves as the central communication point in the environment. Within this VNet three subnets were provisioned:

- I. WAN subnet: Designed to host interfaces that require external connectivity, such as the public facing network interfaces of the FortiGate firewalls.
- II. LAN subnet: Dedicated to internal traffic flow between Azure spokes and the firewalls, ensuring protected communication.
- III. Gateway subnet: A required subnet that hosts the VPN Gateway used for site-to-site connectivity with OCI and the on-premises datacenter.

Additionally, a separate Bastion VNet was deployed to support secure administrative access. This VNet includes the Azure Bastion service which allows secure connection to the Windows VM that is also deployed in the same VNet.

4.2.2. FortiGate Cluster Configuration

A high availability FortiGate firewall cluster was deployed within the HUB VNet to handle traffic inspection. Each FortiGate instance was configured with two network interfaces, one interface in the WAN subnet for external traffic, and one interface in the LAN subnet for internal traffic. This setup makes centralized inspection of traffic moving between spokes, the internet, and external clouds possible. The cluster makes sure that only permitted traffic, based policies, reaches the resources in the cloud. The active-active setup supports automatic failover, which brings high availability and redundancy with it.

4.2.3. Load Balancer

To make distribution of traffic and session handling possible, an Azure Internal Load Balancer (ILB) was configured in front of the cluster. The ILB receives outbound traffic from internal resources and distributes it across the firewall nodes. Load balancing rules were based on source IP and protocol to ensure session persistence, which is particularly useful for stateful inspection, where traffic from a session must be handled by the same firewall instance. The ILB plays a critical role in providing high availability and redundancy, when one firewall becomes unavailable, the load balancer automatically reroutes traffic to the healthy instance.

A conscious architectural decision was made to not use Azure's external load balancer for inbound traffic. This decision was based on two limitations: the lack of dynamic port allocation, as shown in Figure 11. And the absence of support of Internet Control Message Protocol (ICMP), which was required for diagnostics and reachability tests in this PoC.




Figure 11. SNAT Exhaustion Risk with Azure Load Balancer vs NAT Gateway

Note. Adapted from "Optimize your Azure deployments for outbound connectivity," by Microsoft Learn, 2025, Microsoft. <https://learn.microsoft.com/en-us/azure/load-balancer/troubleshoot-outbound-connection?>

4.2.4. Azure VPN Gateway and VPN Connectivity

To enable secure communication with external networks, an Azure VPN Gateway was deployed in the dedicated subnet. The gateway supports multiple VPN connections, which have been used. The gateway forms one point of the connection of the site-to-site VPN tunnels with OCI and the on-premises datacenter. For this PoC, site-to-site VPN tunnels were chosen due to the fact that they are designed for permanent, secure connections between networks. This aligns with the PoC's objective to simulate an integrated multicloud setup where interconnection comes first.

Alternatives such as Azure ExpressRoute were considered, but ultimately not used for this setup. ExpressRoute offers private, high-throughput connectivity to other networks, but it is more expensive and thus not applicable for this setup. In contrast, a site-to-site VPN offers a more accessible solution, especially in the context of a test environment.

The gateway forms one endpoint of two IPsec site-to-site VPN tunnels, of which one goes to OCI, and the other to the on-premises datacenter. These tunnels enable secure, encrypted communication between the Azure environment and both external networks, being a critical component in a multicloud setup.

During implementation, an additional technical challenge emerged due to the use of overlapping IP address ranges in the on-premises datacenter and the OCI environment. Both environments use identical network prefixes, which would cause routing conflicts and incorrect policy matching with Azure. To resolve this, NAT rules were applied at the VPN gateway, translating incoming traffic from the on-premises environment into a dedicated, non-overlapping IP range, as shown in Figure 12.

Name	Type	IP Configuration ID	Mode	Internal Mappings	External Mappings
ONPREM-TO-AZURE	Static		IngressSnat	172.22.20.64/28	192.168.0.64/28
AZURE-TO-ONPREM	Static		EgressSnat	192.168.0.64/28	172.22.20.64/28

Figure 12. IP Address Translation Rules Between Azure and On-Premises Environment

4.2.5. Routing and Peering

To enable traffic between the HUB VNet and the spoke VNets, VNet peering was configured with bidirectional access. Each spoke is peered with the HUB VNet, and no peering exists between the spokes. This ensures that all traffic is routed to the HUB VNet, where inspection can occur. To achieve this, User Defined Routes (UDRs) were configured on the routing tables of all the corresponding subnets. The UDRs override the default routes of Azure when overlapped. So in the case of this PoC, the default 0.0.0.0/0 routes from Azure were overwritten so that all traffic from the spokes is directed to the IP address of the ILB, as seen in Figure 13.

Associated route table: ①

RT-SPOKE-LE-POC

Effective routes

Source	↑↓	State	↑↓	Address Prefixes	↑↓	Next Hop Type	↑↓	Next Hop IP Address	↑↓	User Defined Route Name	↑↓
Default		Active		10.1.0.0/16		Virtual network		-		-	
Default		Active		10.0.0.0/16		VNet peering		-		-	
Default		Invalid		0.0.0.0/0		Internet		-		-	
User		Active		0.0.0.0/0		Virtual appliance		10.0.1.132		allTraffic	

Figure 13. User-Defined Route Overriding Default Route in Azure

4.3. Infrastructure in OCI

This section outlines the design and implementation of the OCI environment. As shown in Figure 14, it followed the same principles as in Azure, with a hub-and-spokes model being deployed. Other key elements include the Virtual Cloud Networks, firewall cluster, and VPN connectivity.




Figure 14. OCI Cloud Architecture with a Hub-and-Spokes Model

4.3.1. VCNs, Subnets, and FortiGate Cluster

In OCI, the network design followed a similar hub-and-spoke approach as in Azure. A dedicated HUB VCN was created to act as a central point for connectivity, routing, and traffic inspection, also with a WAN and a LAN subnet assigned for external and internal traffic. However, unlike Azure, the Dynamic Routing Gateway (DRG) in OCI is not associated with a specific subnet. Instead, it is a region level resource and multiple VCNs, IPsec tunnels, remote peerings, and more services can be attached to it. It can act as the central point for connectivity between OCI regions, other cloud networks, and on-premises networks. Therefore, no dedicated subnet was needed for VPN connectivity.

Separate spoke VCNs were created to host private workloads. These VCNs include a single subnet containing one test VM used to validate internal routing and connectivity to other networks.

Also similar to Azure's environment, a pair of FortiGate firewalls were deployed in active-active setup that are responsible for all traffic between the spoke VCNs, the internet, and the external cloud environments. To manage load distribution across the cluster, a Network Load Balancer (NLB) was configured within the hub VCN. The setup closely mirrors the one in Azure, providing uniform behavior across platforms.

4.3.2. VPN Connectivity and DRG

In OCI, connection with external environments is enabled via DRG, which had the HUB VCN attached to it. An IPSec VPN tunnel was established between the DRG and the Azure VPN gateway to enable secure, encrypted communication.

Static routes were configured to direct traffic from the OCI spoke VCN through the FortiGate, towards the Network Load Balancer (NLB), then through the DRG and finally through the VPN tunnel to Azure. Similarly, traffic from Azure is routed back to OCI.

The choice for the site-to-site VPN in OCI was made for the same reasons as in Azure: flexible, cost-effective connectivity and full compatibility with multicloud environments. While OCI also supports FastConnect for dedicated, high-speed, private links, this was not a requirement for the PoC.

4.3.3. Routing and VCN Peering

To control traffic flow, route tables were configured at both the subnet and the DRG attachment levels. All traffic from the spoke VCNs was routed to the FortiGate firewall cluster via the internal IP address of the NLB, making sure that traffic is inspected before reaching its destination. Local VCN peering was used to connect the spoke VCNs to the HUB VCN, enabling connectivity with the HUB as central point, similar to Azure.

4.4. Proof of Implementation

To validate the design and ensure that the multicloud architecture operates as intended, several connectivity tests were performed. The test focused on internal routing, VPN connectivity, firewall inspection, and NAT behavior across Azure, OCI, and the on-premises datacenter. Each scenario was tested using real network traffic, with firewall session logs to confirm proper traffic flow.

4.4.1. Connectivity in Azure

To verify internal routing and firewall integration within Azure, the VM in spoke VNet 1 attempted to reach another VM in spoke VNet 2. Figure 15 shows the successful ICMP traffic traversal (ping) results and demonstrate that VNet peering is operational. Figure 16 then shows that traffic correctly traverses the FortiGate cluster for inspection and forwarding. This confirms that internal communication between spoke networks is functional via the hub-and-spoke model. Additionally, the FortiGate session table confirmed that all traffic passed through the firewall, validating that routing and inspection are consistent across the environment.

```
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state
link/ether 00:22:48:5a:fa:f8 brd ff:ff:ff:ff:ff:ff
inet 10.1.0.4/24 metric 100 brd 10.1.0.255 scope global eth0
    valid_lft forever preferred_lft forever
inet6 fe80::222:48ff:fe5a:faf8/64 scope link
    valid_lft forever preferred_lft forever
lepoc@UBUNTUS-P1-LE-POC:~$ ping -c 3 10.2.0.4
PING 10.2.0.4 (10.2.0.4) 56(84) bytes of data.
64 bytes from 10.2.0.4: icmp_seq=1 ttl=63 time=7.87 ms
64 bytes from 10.2.0.4: icmp_seq=2 ttl=63 time=5.71 ms
64 bytes from 10.2.0.4: icmp_seq=3 ttl=63 time=2.47 ms

--- 10.2.0.4 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2002ms
rtt min/avg/max/mdev = 2.472/5.351/7.871/2.218 ms
lepoc@UBUNTUS-P1-LE-POC:~$
```

Figure 15. Connectivity Test Between Two VM's in Azure

Time	@	Source	Device	Destination	Application Name	Result	Policy ID
i:52		10.1.0.4		10.2.0.4	PING	✓ Accept (252 B / 252 B)	AzureIntern(4)

Figure 16. FortiGate Firewall Session Log

4.4.2. Connectivity in Azure to the Internet

To validate outbound internet access from within Azure's spoke VNets, a test was performed by sending ICMP traffic from a VM in a spoke VNet to google.com, as shown in Figure 17. The response confirmed that outbound internet connectivity was functional. It also demonstrates that no public IP is required on the virtual machine itself, as all traffic is routed centrally through the FortiGate firewall.

```
lepoc@UBUNTUS-P1-LE-POC:~$ ping -c 3 google.com
PING google.com (142.250.181.238) 56(84) bytes of data.
64 bytes from fra16s56-in-f14.1e100.net (142.250.181.238): icmp_seq=1 ttl=112 time=4.06 ms
64 bytes from fra16s56-in-f14.1e100.net (142.250.181.238): icmp_seq=2 ttl=112 time=3.14 ms
64 bytes from fra16s56-in-f14.1e100.net (142.250.181.238): icmp_seq=3 ttl=112 time=3.85 ms

--- google.com ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2003ms
rtt min/avg/max/mdev = 3.135/3.681/4.060/0.395 ms
lepoc@UBUNTUS-P1-LE-POC:~$
```

Figure 17. Connectivity Test From a VM in Azure to the Internet

In Addition to the ping result, the FortiGate session table was again required to verify that the traffic passed through the firewall, as seen in Figure 18. Figure 19 shows the output of the *diagnose sys session list | grep 10.1.0.4* command on the FortiGate firewall and shows that the internal source IP of the VM was translated (SNAT) to the FortiGate's private WAN interface IP, before reaching the public destination of the IP address from Google. The return traffic was also handled properly with DNAT, mapping the response back to the original source. This confirmed that FortiGate actively processes outbound traffic, NAT is correctly applied, and routing and session handling behaves as intended.

Date...		Source	Device	Destination	Application Name	Result	P...
2025/0...		10.1.0.4		142.250.185.78	PING	✓ Accept (2.44 kB / 2.44 kB)	ping (5

Figure 18. FortiGate Firewall Session Log

```

le poc-fgt-1 # diagnose sys session list | grep 10.1.0.4
hook=post dir=org act=snat 10.1.0.4:2->142.250.186.46:8(10.0.1.69:5118)
hook=pre dir=reply act=dnat 142.250.186.46:5118->10.0.1.69:0(10.1.0.4:2)

```

Figure 19. FortiGate CLI Output: Translated Session from 10.1.0.4 to the Internet

4.4.3. Connectivity between Azure and the Datacenter

Connectivity with the on-premises datacenter was tested using a ping from the Azure VM to a known reachable host in the datacenter, in Figure 20 is shown that the source ping from the host in the datacenter is 172.22.20.65, and it arrives in Azure as 192.168.0.65, this is because of the NAT rules that are applied at the VPN gateway, as explained in a previous chapter. The successful response, in Figure 22, verified that the IPsec site-to-site VPN tunnel was active, traffic is properly routed through the FortiGate and Azure VPN Gateway, and that NAT applied at the gateway does not interfere with reachability or return paths. Figure 21 and Figure 23 again show that the traffic is effectively traversing the needed firewalls.

```

CLI Console (1)
649.697860 192.168.0.65 -> 10.1.0.4: icmp: echo request
649.697860 192.168.0.65 -> 10.1.0.4: icmp: echo request
649.697864 192.168.0.65 -> 10.1.0.4: icmp: echo request
649.697864 192.168.0.65 -> 10.1.0.4: icmp: echo request
649.699806 10.1.0.4 -> 192.168.0.65: icmp: echo reply
649.699806 10.1.0.4 -> 192.168.0.65: icmp: echo reply
649.699826 10.1.0.4 -> 192.168.0.65: icmp: echo reply
649.699826 10.1.0.4 -> 192.168.0.65: icmp: echo reply
649.699829 10.1.0.4 -> 192.168.0.65: icmp: echo reply
649.699829 10.1.0.4 -> 192.168.0.65: icmp: echo reply

^C
--- 10.1.0.4 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 13.6/17.6/23.0 ms

FGT-CGK-FORTI # execute ping-options source 172.22.20.65
FGT-CGK-FORTI # execute ping 10.1.0.4
PING 10.1.0.4 (10.1.0.4): 56 data bytes
64 bytes from 10.1.0.4: icmp_seq=0 ttl=63 time=14.5 ms
64 bytes from 10.1.0.4: icmp_seq=1 ttl=63 time=13.8 ms
64 bytes from 10.1.0.4: icmp_seq=2 ttl=63 time=14.0 ms
^C
--- 10.1.0.4 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 13.8/14.1/14.5 ms

```

Figure 20. Connectivity Test from On-Premises Datacenter and Azure

Da...		Source	Device	Destination	Application Name	Result	Policy ID
2025...		192.168.0.65		10.1.0.4	PING	✓ Accept (252 B / 252 B)	testvpn(6)

Figure 21. FortiGate Firewall Session Log

```

lepoc@UBUNTUS-P1-LE-POC:~$ ping -c 3 192.168.0.65
PING 192.168.0.65 (192.168.0.65) 56(84) bytes of data.
64 bytes from 192.168.0.65: icmp_seq=1 ttl=254 time=15.6 ms
64 bytes from 192.168.0.65: icmp_seq=2 ttl=254 time=14.7 ms
64 bytes from 192.168.0.65: icmp_seq=3 ttl=254 time=14.3 ms

--- 192.168.0.65 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2004ms
rtt min/avg/max/mdev = 14.293/14.869/15.611/0.550 ms
lepoc@UBUNTUS-P1-LE-POC:~$

```

Figure 22. Connectivity Test from Azure to On-premises Datacenter

	Source	Device	Destination	Application Name	Result	Policy ID
202...	10.1.0.4		192.168.0.65	PING	✓ Accept (252 B / 252 B)	testvpn (6)

Figure 23. FortiGate Firewall Session Log

4.4.4. Connectivity between OCI and Azure

Figure 24 shows that to confirm connectivity between the cloud environments, a ping test was initiated from an Azure VM towards a VM hosted in OCI and back. The response confirmed that the IPsec VPN tunnel between the Azure VPN gateway and OCI DRG was active, and that routes and firewall rules are correctly configured on both sides. The test validated the multicloud aspect of the PoC and confirmed that traffic between the clouds is both encrypted and inspected. Figure 25 shows that the traffic traverses through the FortiGate firewalls.

```

lepoc@UBUNTUS-P1-LE-POC:~$ ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq s
    link/ether 00:22:48:5a:fa:f8 brd ff:ff:ff:ff:ff:ff
    inet 10.1.0.4/24 metric 100 brd 10.1.0.255 scope global et
        valid_lft forever preferred_lft forever
    inet6 fe80::222:48ff:fe5a:faf8/64 scope link
        valid_lft forever preferred_lft forever
lepoc@UBUNTUS-P1-LE-POC:~$ ping 172.22.20.68
PING 172.22.20.68 (172.22.20.68) 56(84) bytes of data.
64 bytes from 172.22.20.68: icmp_seq=1 ttl=59 time=14.2 ms
64 bytes from 172.22.20.68: icmp_seq=2 ttl=59 time=12.0 ms
64 bytes from 172.22.20.68: icmp_seq=3 ttl=59 time=11.9 ms
^C
--- 172.22.20.68 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2003ms
rtt min/avg/max/mdev = 11.904/12.725/14.241/1.072 ms
lepoc@UBUNTUS-P1-LE-POC:~$

ubuntu@spoke3-vnic:~$ ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOW
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: ens3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 9000 qdisc pfifo_fast st
    link/ether 02:00:17:03:3b:a5 brd ff:ff:ff:ff:ff:ff
    altname enp0s3
    inet 172.22.20.68/28 metric 100 brd 172.22.20.79 scope global ens3
        valid_lft forever preferred_lft forever
    inet6 fe80::17ff:fe03:3ba5/64 scope link
        valid_lft forever preferred_lft forever
ubuntu@spoke3-vnic:~$ ping 10.1.0.4
PING 10.1.0.4 (10.1.0.4) 56(84) bytes of data.
64 bytes from 10.1.0.4: icmp_seq=1 ttl=60 time=15.1 ms
64 bytes from 10.1.0.4: icmp_seq=2 ttl=60 time=12.0 ms
64 bytes from 10.1.0.4: icmp_seq=3 ttl=60 time=14.6 ms
^C
--- 10.1.0.4 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2002ms
rtt min/avg/max/mdev = 12.017/13.917/15.130/1.360 ms
ubuntu@spoke3-vnic:~$

```

Figure 24. Connectivity Test between OCI and Azure

...	Source	D...	Destination	Application Name	Result	Policy ID
/0...	10.1.0.4		172.22.20.68	PING	✓ Accept (252 B / 252 B)	testvpn (6)

Figure 25. FortiGate Firewall Session Log

4.4.5. Connectivity in OCI

To validate internal within the OCI environment, a ping test was performed between two virtual machines in two different subnets across peered VCNs, which can be seen in Figure 26. The successful response together with the session table of the FortiGate confirmed that traffic was successfully routed through the firewall cluster, ensuring that intern traffic also passes the HUB VCN.




Figure 26. FortiGate Session Log and Connectivity Test in OCI

5. Conclusion

This thesis investigated the question of how a multicloud environment can be implemented and when such a strategy is justified. Through a combination of theoretical research and a technical PoC, both the strategic motivations and the practical challenges of multicloud adoption were analyzed and demonstrated in a realistic setup.

From a theoretical standpoint, it became clear that all cloud providers, AWS, Azure, and OCI offer a broadly similar set of core services, such as compute, storage, networking, and IAM. However, the way in which these services are implemented, managed, and priced differs significantly. Azure excels in hybrid integrations and enterprise ecosystems, AWS leads in breadth of services and global scale, while OCI positions itself as a cost-effective and high performance platform for data intensive workloads. These differences create a strategic opportunity to leverage each provider's strengths in a multicloud architecture.

The PoC validated that such integration is technically possible. By designing and implementing a hub-and-spoke architecture in both Azure and OCI, and connecting them via a VPN connection, the PoC demonstrated that cloud environments from different providers can be interconnected securely and reliably. Routing was centralized through firewall clusters to make sure that there is full traffic inspection and control. Furthermore, the PoC showed that integrating a third environment, such as the on-premises datacenter, is possible using the same principles and technologies. Although it introduces additional challenges, such as overlapping VPN IP ranges, which were mitigated through NAT strategies.

Throughout the thesis, the benefits of multicloud were confirmed: flexibility, reduced vendor lock-in, workload optimization, possibility of increased resilience, and faster innovation. At the same time, a lot of challenges emerged. Managing IAM across different cloud platforms requires multiple configuration and federation with third party tools, monitoring lacks cross provider integration without external tools, and cost visibility can be limited. These challenges make it clear that multicloud is not a default path, but it must be a calculated strategy that aligns with the organization's needs and capabilities.

In terms of cost, the analysis and the cost considerations showed that multicloud deployments often carry a higher operational and financial burden compared to single cloud setups. While theoretical cost savings are possible through service specialization, they are easily outweighed by indirect costs such as personnel, training, certification, and governance overhead. These findings confirm that multicloud is most suitable for organizations that already possess cloud maturity and have clearly defined business or technical motivations, such as workload specialization, compliance constraints, or customer driven requirements.

It's safe to assume that a multicloud environment can be effectively implemented, provided that the organization has the technical expertise, the right capacity, and a valid strategic reason for doing so. The decision to adopt a multicloud strategy should never be taken lightly or purely based on provider comparison, it must be driven by actual business needs, workload characteristics, and long term flexibility goals. While theoretical cost advantages may exist, the hidden operational costs mean that multicloud is, in practice, almost always more expensive than a single cloud setup. However, there are

scenarios where this added cost is justified: for example, when platform specific strengths need to be combined, compliance require it, or client expectations demand support across providers.

For a service provider like Cegeka, these complexities present an opportunity rather than a barrier. By building expertise in multiple cloud environments, Cegeka can support a broader range of client environments, respond to increasingly diversifying customer demands, and differentiate itself in a competitive IT market. Ultimately, this capability contributes not only to technical innovation but also to commercial growth, expanding the client base and increasing potential revenue coming from provided cloud services. The thesis demonstrates that both the opportunity and the risks are there, and finding a balance between them asks more than just technical knowledge, it requires a well thought out, well informed architectural vision.

REFERENCE LIST

Abandy, R. (2022, August 24). *The History of Microsoft Azure*. Microsoft. Retrieved April 08, 2025 from <https://techcommunity.microsoft.com/blog/educatordeveloperblog/the-history-of-microsoft-azure/3574204>

Appvia. (2024, February 19). *How a Multi-cloud Strategy Can Simplify Compliance*. Appvia. Retrieved March 12, 2025 from <https://www.appvia.io/blog/how-a-multi-cloud-strategy-can-simplify-compliance>

AWS. (2025). *AWS Global Infrastructure*. AWS Amazon. Retrieved March 19, 2025 from <https://aws.amazon.com/about-aws/global-infrastructure/>

AWS. (2025). *How do you pay for AWS?*. AWS Amazon. Retrieved March 28, 2025 from https://aws.amazon.com/pricing/?nc2=h_ql_pr_ln&aws-products-pricing.sort-by=item.additionalFields.productNameLowercase&aws-products-pricing.sort-order=asc&awsf.Free%20Tier%20Type=*all&awsf.tech-category=*all

AWS. (s.a.). *Identity providers and federation*. AWS Amazon. Retrieved May 13, 2025 from https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_providers.html

AWS. (2025). *Our origins*. AWS Amazon. Retrieved April 08, 2025 from <https://aws.amazon.com/about-aws/our-origins/>

AWS. (2025). *Regions and Availability Zones*. AWS Amazon. Retrieved April 15, 2025 from https://aws.amazon.com/about-aws/global-infrastructure/regions_az/?p=ngi&loc=2

Barr, J. (2024, November 18). *AWS Lambda turns ten – looking back and looking ahead*. AWS. Retrieved February 27, 2025 from <https://aws.amazon.com/blogs/aws/aws-lambda-turns-ten-the-first-decade-of-serverless-innovation/>

Bell, H. (2019, January 31). *What is the Difference Between Multicloud and Hybrid Cloud?*. Nutanix. Retrieved March 04, 2025 from https://www.nutanix.com/en_gb/blog/hybrid-cloud-vs-multi-cloud-whats-difference

Brophy, M. (2025, May 19). *Cloud Computing Market Trends to Watch in 2025*. AlphaSense. Retrieved May 21, 2025 from <https://www.alpha-sense.com/blog/trends/cloud-computing-market-trends/>

Chen, M. (2025, February 20). *What Is Multicloud? Uses and Benefits*. Oracle. Retrieved March 11, 2025 from <https://www.oracle.com/cloud/multicloud/what-is-multicloud/>

Dua, A. (2024, February 23). *IaaS vs PaaS vs SaaS: Explaining the Key Differences*. Turing. Retrieved March 18, 2025 from <https://www.turing.com/blog/iaas-vs-paas-vs-saas-key-differences>

F5. (s.a.). *What Is Multi-Cloud Network Architecture?*. F5. Retrieved February 19, 2025 from <https://www.f5.com/glossary/multi-cloud-network-architecture>

Gartner, Magic Quadrant for Integration Platform as a Service, Guttridge, K., Comes, A., Pasricha, S., van den Berk, M. & Humphreys, A. (2024, February 19). *Oracle Restricted*

IT Executive's Guide to Oracle Cloud Infrastructure. Oracle. Retrieved April 08, 2025 from <https://www.oracle.com/a/ocom/docs/cloud/oracle-cloud-infrastructure-platform-overview-wp.pdf>

Google Cloud. (s.a.). *Advantages and Disadvantages of Cloud Computing*. Google Cloud. Retrieved March 05, 2025 from <https://cloud.google.com/learn/advantages-of-cloud-computing?hl=en>

Google Cloud. (s.a.). *What is multicloud?*. Google. Retrieved March 13, 2025 from <https://cloud.google.com/learn/what-is-multicloud>

IBM. (2021, October 20). *What are Iaas, Paas and Saas?*. IBM. Retrieved March 18, 2025 from <https://www.ibm.com/think/topics/iaas-paas-saas>

Krishnakumar, V. (2025, April 11). *Azure Arc: Simplifying Hybrid and Multi-Cloud Management*. Cloudoptimo. Retrieved April 24, 2025 from <https://www.cloudoptimo.com/blog/azure-arc-simplifying-hybrid-and-multi-cloud-management>

MarketsandMarkets. (2023, December). *Cloud Computing Market*. MarketsandMarkets. Retrieved March 05, 2025 from <https://www.marketsandmarkets.com/Market-Reports/cloud-computing-market-234.html>

Microsoft. (2025). *Microsoft Datacenters*. Microsoft. Retrieved March 19, 2025 from <https://datacenters.microsoft.com/globe/explore?view=map>

Microsoft. (2025). *What are public, private, and hybrid clouds?*. Azure Microsoft. Retrieved May 13, 2025 from <https://azure.microsoft.com/en-us/resources/cloud-computing-dictionary/what-are-private-public-hybrid-clouds>

Microsoft Learn. (2025, May 28). *Azure Identity Management and access control security best practices*. Microsoft. Retrieved May 13, 2025 from <https://learn.microsoft.com/en-us/azure/security/fundamentals/identity-management-best-practices>

Microsoft Learn. (2025, March 21). *What are Azure regions?*. Microsoft. Retrieved February 26, 2025 from <https://learn.microsoft.com/en-us/azure/reliability/regions-overview>

Microsoft Learn. (2025, February 25). *What are availability zones?*. Microsoft. Retrieved February 28, 2025 from <https://learn.microsoft.com/en-us/azure/reliability/availability-zones-overview?tabs=azure-cli>

Mishra, S. (2024, September 09). *Multicloud Masterclass: OCI, Azure and AWS Compared*. Rainfocus. Retrieved March 13, 2025 from https://static.rainfocus.com/oracle/ocw24/sess/1721138693474001PhzP/finalsessionfile/09.03.24%20OCW%20Multicloud%20Masterclass%20Oracle%20Cloud%20Infrastructure%2C%20Azure%2C%20and%20AWS%20Compared%20V1_1725356172709001XH2m.pdf

Oracle. (2023, May). *Best Practices for Identity and Access Management (IAM) in Oracle Cloud Infrastructure*. Oracle. Retrieved May 13, 2025 from <https://docs.oracle.com/en-us/iaas/Content/Resources/Assets/whitepapers/best-practices-for-iam-on-oci.pdf>

Oracle. (2024, October 31). *Compare OCI with AWS, Azure, and Google Cloud*. Oracle. Retrieved March 06, 2025 from <https://www.oracle.com/cloud/service-comparison/>

Oracle. (2024, October 23). *Oracle Again Recognized as a Leader in the 2024 Gartner Magic Quadrant™ reports for Strategic Cloud Platform Services and Distributed Hybrid Infrastructure*. Oracle. Retrieved March 27, 2025 from <https://www.oracle.com/news/announcement/oracle-again-recognized-as-a-leader-in-the-2024-gartner-magic-quadrant-reports-for-strategic-cloud-platform-services-and-distributed-hybrid-infrastructure-2024-10-23/>

Oracle. (2025). *Public Cloud Regions*. Oracle. Retrieved March 19, 2025 from <https://www.oracle.com/be/cloud/public-cloud-regions/>

Oracle. (2025, May 21). *Regions and Availability Domains*. Oracle. Retrieved April 17, 2025 from <https://docs.oracle.com/en-us/iaas/Content/General/Concepts/regions.htm?>

Richter, F. (2025, February 27). *Cloud Infrastructure Market: Amazon and Microsoft Stay Ahead in Global Cloud Market*. Statista. Retrieved April 08, 2025 from <https://www.statista.com/chart/18819/worldwide-market-share-of-leading-cloud-infrastructure-service-providers/>

Sasovets, I., & Teres, K. (2025, February 11). *Multi-Cloud Security: Key Challenges and Effective Practices from Experts*. Techmagic. Retrieved May 08, 2025 from <https://www.techmagic.co/blog/multi-cloud-security>

Susnjara, S., & Smalley, I. (2025, February 10). *What is cloud computing?*. IBM. Retrieved February 19, 2025 from <https://www.ibm.com/think/topics/cloud-computing>

Tata Communications. (2025, April 11). *Multi-cloud vs. single-cloud: Which is right for your business?*. Tatacommunications. Retrieved April 02, 2025 from <https://www.tatacommunications.com/knowledge-base/multi-cloud-vs-single-cloud/>

Teske, K. (2024, June 17). *Multi-Cloud Strategy: Secure Data & Compliance*. Veeam. Retrieved May 14, 2025 from <https://www.veeam.com/blog/multi-cloud-strategy.html>

ATTACHMENTS

[AWS Pricing Calculator](#) > My Estimate

My Estimate [Edit](#) [Info](#)

[Export](#) [Share](#)

Estimate summary [Info](#)

Upfront cost	Monthly cost	Total 12 months cost
0.00 USD	754.96 USD	9,059.52 USD
		Includes upfront cost

Getting Started with AWS

[Get started for free](#)

[Contact Sales](#)

My Estimate [Duplicate](#) [Delete](#) [Move to](#) [Create group](#) [Add support](#) [Add service](#)

< 1 > [Settings](#)

☐	Service Name		Status	Upfron...	MonthL...	Descrip...	Region	Con
☐	Amazon Simple Storage Service (S3)	Edit	-	0.00 USD	1.28 USD	-	US East (...)	S3 S
☐	AWS Lambda	Edit	-	0.00 USD	0.20 USD	-	US East (...)	Inve
☐	Amazon API Gateway	Edit	-	0.00 USD	32.10 USD	-	US East (...)	HTT
☐	Amazon QuickSight	Edit	-	0.00 USD	510.00 USD	-	US East (...)	SPIC
☐	Amazon Virtual Private Cloud (VPC)	Edit	-	0.00 USD	66.50 USD	-	US East (...)	Wor
☐	Amazon Simple Storage Service (S3)	Edit	-	0.00 USD	0.92 USD	-	US East (...)	S3 S
☐	Amazon RDS for SQL server	Edit	-	0.00 USD	143.96 USD	-	US East (...)	Stor

Attachment 1. Calculation of Price in a Single Cloud Setup in AWS

OCI
About
Services
Solutions
Pricing
Partners
Resources
Search
Flag
Language
Sign in to Oracle Cloud

My Estimate
Configure and estimate costs for OCI services (Learn more)
Send to Sales
Start for Free
EUR - Euro
Estimated Monthly Cost
€628.80

Add configuration

Object Storage
Estimated Monthly Cost
€2.09

Virtual Cloud Networks
Estimated Monthly Cost
€0.00

Oracle Analytics Cloud - Enterprise
Click to edit title
Estimated Monthly Cost
€502.34

Autonomous Database
Estimated Monthly Cost
€122.84

OCI Functions
Estimated Monthly Cost
€0.03

Object Storage
Estimated Monthly Cost
€1.50

Attachment 2. Calculation of Price in a Single Cloud Setup in OCI

single cloud

Static Web Apps	Standard tier, 1 app, 100 GB of Bandwidth...	Upfront: \$0.00	Monthly: \$9.00
Azure Functions	Consumption tier, Pay as you go, 512 MB...	Upfront: \$0.00	Monthly: \$0.20
Storage Accounts	Block Blob Storage, General Purpose V2, ...	Upfront: \$0.00	Monthly: \$4.84
Power BI Embedded	1 node(s) x 176 Hours, Node type: A2, 2 ...	Upfront: \$0.00	Monthly: \$353.43
Bandwidth	Internet egress, 300 GB outbound data tr...	Upfront: \$0.00	Monthly: \$16.00
VPN Gateway	VPN Gateways, VpnGw1 tier, 730 gateway...	Upfront: \$0.00	Monthly: \$138.70
Azure SQL Database	Single Database, vCore, Hyperscale, Provi...	Upfront: \$0.00	Monthly: \$213.43

Support
SUPPORT:
Basic (Included)
\$0.00

Select your program/offer
LICENSING PROGRAM:
Microsoft Customer Agreement (MCA)
Log in to see your Azure agreement pricing.
Show Dev/Test Pricing

Estimated upfront cost
\$0.00
Estimated monthly cost
\$735.60

Attachment 3. Calculation of Price in a Single Cloud Setup in Azure

[AWS Pricing Calculator](#) > My Estimate

My Estimate [Edit](#)

[Export](#) [Share](#)

Estimate summary [Info](#)

Upfront cost	Monthly cost	Total 12 months cost
0.00 USD	198.74 USD	2,384.88 USD
		Includes upfront cost

Getting Started with AWS

[Get started for free](#)

[Contact Sales](#)

My Estimate

[Duplicate](#) [Delete](#) [Move to](#) [Create group](#) [Add support](#) [Add service](#)

☐	Service Name	Status	Upfront...	Monthl...	Descrip...	Region	Confi
☐	Amazon Simple Storage Service (S3)	-	0.00 USD	1.28 USD	-	US East (...)	S3 St
☐	AWS Lambda	-	0.00 USD	0.20 USD	-	US East (...)	Invok
☐	Amazon API Gateway	-	0.00 USD	32.10 USD	-	US East (...)	HTTP
☐	Amazon Virtual Private Cloud (VPC)	-	0.00 USD	165.16 USD	-	US East (...)	Work

Attachment 4. Calculation of Price in a Multicloud Setup in AWS

multicloud

☐
☐
☐
☐

✓ Power BI Embedded	1 node(s) x 176 Hours, Node type: AZ, 2 Vi...	Upfront: \$0.00	Monthly: \$353.43
✓ Storage Accounts	Block Blob Storage, General Purpose V2, F...	Upfront: \$0.00	Monthly: \$5.88
✓ VPN Gateway	VPN Gateways, VpnGw1 tier, 730 gateway ...	Upfront: \$0.00	Monthly: \$138.70
✓ Bandwidth	Internet egress, 1000 GB outbound data tr...	Upfront: \$0.00	Monthly: \$72.00

Support

SUPPORT:

Basic (Included) [Info](#) **\$0.00**

Select your program/offer

LICENSING PROGRAM:

Microsoft Customer Agreement (MCA) [Info](#) [Log in](#) to see your Azure agreement pricing.

☐ Show Dev/Test Pricing [Info](#)

Estimated upfront cost **\$0.00**

Estimated monthly cost **\$570.01**

Attachment 5. Calculation of Price in a Multicloud Setup in Azure

My Estimate

Configure and estimate costs for OCI services ([Learn more](#))

Send to Sales

Start for Free

EUR - Euro

Estimated Monthly Cost

€119.66

Add configuration

Object Storage

Estimated Monthly Cost €1.25

Autonomous Database

Estimated Monthly Cost €118.41

Attachment 6. Calculation of Price in a Multicloud Setup in OCI

Internship – Connecting Clouds – Bachelor Thesis

44