Research Article OPEN ACCESS

Evaluating the Efficiency and Limitations of Configuration Strategies in Hybrid Cloud Environments

Raghava Satya SaiKrishna Dittakavi
Independent Researcher

Abstract

The increasing complexity and scale of modern computing needs have led to the development and adoption of cloud computing as a ubiquitous paradigm for data storage and processing. The hybrid cloud model, which combines both public and private cloud infrastructures, has been particularly appealing to organizations that require both the scalability offered by public clouds and the security features of private clouds. Various strategies for configuring and managing resources have been developed to optimize the hybrid cloud environment. These strategies aim to balance conflicting objectives such as cost-efficiency, performance optimization, security, and compliance with regulatory standards. This exploratory research focused on evaluating the efficiency and limitations of different configuration strategies in hybrid cloud environments. Findings indicate that each approach presents distinct advantages. Improving resource utilization and automating governance processes are significant advantages of Policy-based Resource Management, which leads to cost-effectiveness. Intelligent routing of traffic is a feature of Cross-cloud Load Balancing, resulting in optimized performance and higher service availability. By centralizing control, the Hybrid Cloud Service Mesh allows for secure and streamlined cross-service communication. A notable feature of Cross-cloud Container Orchestration is its ability to simplify the migration of applications across diverse cloud environments. For immediate threat detection and regulatory compliance, real-time monitoring is facilitated by Log Management and Analytics. However, Policy-based Resource Management can be complex and inflexible. Extra costs for data transfer between different cloud providers are a drawback of Cross-cloud Load Balancing. Additional network hops create latency issues in Hybrid Cloud Service Mesh configurations. If configured incorrectly, Cross-cloud Container Orchestration could expose the system to security risks. Finally, Log Management and Analytics require both ample storage and advanced analytical capabilities.

Keywords: Cloud computing, Hybrid cloud, Resource management, Strategies, Limitations

Declarations

Competing interests:

The author declares no competing interests.

© The Author(s). **Open Access** 2022 This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution, and reproduction in any medium or format, as long as appropriate credit is given to the original author(s) and source, a link to the Creative Comsmons license is provided, and changes are indicated. Unless otherwise stated in a credit line to the source, the photos or other third-party material in this article are covered by the Creative Commons license. If your intended use is not permitted by statutory law or exceeds the permitted usage, you must acquire permission directly from the copyright holder if the material is not included in the article's Creative Commons lice

Introduction

Cloud computing has fundamentally altered the way computing resources are utilized, allocated, and delivered. In traditional models, companies would have to invest heavily in physical hardware and software licenses, leading to significant capital expenditure and ongoing maintenance costs. Cloud computing shifts this model by using virtualization and Internet technologies to provide resources as a service [1], [2]. This means that rather than owning physical servers or software, users can lease or rent these resources as needed. The immediate advantages of this approach include lower upfront costs, the ability to scale resources dynamically based on demand, and simplified management and maintenance. Furthermore, cloud computing's centralized nature enables high levels of automation, which in turn results in operational efficiencies and cost savings for organizations.

Among the services offered in the cloud computing model, Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS) are the most prevalent [3]. IaaS provides virtual machines, networks, and storage over the Internet. This allows businesses to run their own applications on rented servers, without the need to invest in physical hardware. PaaS takes this concept a step further by providing a complete environment where developers can build, deploy, and manage applications without worrying about the underlying infrastructure. SaaS delivers software applications over the Internet, eliminating the need for end-users to install and maintain software on their own machines [4]. All these services are generally offered under a subscription-based or pay-per-use model, making it cost-effective for businesses and individual users alike [5].

Table 1. Features of contemporary Hybrid cloud computing

Feature	Description
Flexibility	Facilitates integration of legacy infrastructure with public cloud services; supports workload migration.
Cost Management	Allows for strategic allocation of capital and operational expenses through environment selection.
Agility and Scalability	Enhances resource provisioning and deployment; accommodates demand spikes via public cloud bursting.
Resiliency and Interoperability	Provides workload redundancy and supports component interoperability between private and public environments.
Compliance	Ensures data locality in accordance with regulatory guidelines; permits selective public cloud utilization.

The user primarily interacts with the system through a 3rd Party API, enabling the processing of their requests. The central hub for these interactions is the private cloud, housing both integrated server capabilities and a specific application which has ties to a Kubernetes cluster [6]. These elements collaboratively serve to manage and streamline user requests, making the overall experience intuitive and efficient. In addition, there's a defined frontend cloud architecture system. This system showcases the more visible aspects of the cloud environment, including an enterprise interface that facilitates access to various applications. To ensure security and regulate access, a firewall mechanism is in place. This firewall interfaces with various devices, like laptops and mobile phones, permitting authorized users to connect and utilize the services provided [7], [8].

The most commonly used deployment models are Public Cloud and Private Cloud.

In a Public Cloud model, computing resources are provided by a third-party cloud service provider and are made available to the general public. These resources are typically owned and operated by the cloud service provider and are delivered over the Internet. Users of a Public Cloud can take advantage of its scalability and cost-effectiveness since resources are shared among multiple tenants [9], [10]. This multi-tenancy can raise concerns about data security, compliance, and performance, as users have less control over the infrastructure.

Figure 1. Hybrid cloud user interaction and frontend components

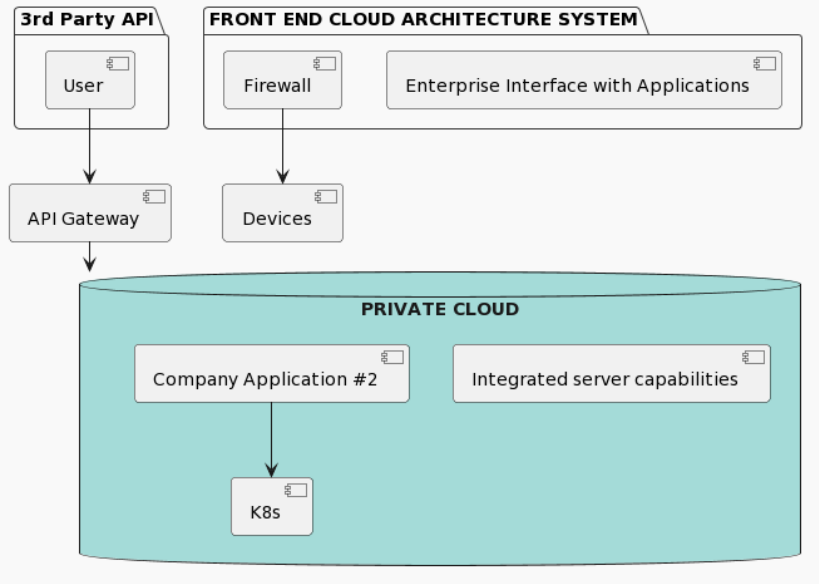

Private Cloud model offers a more controlled environment, as the computing resources are used exclusively by a single organization. A Private Cloud can either be hosted on-premises or externally by a third-party provider. In either case, the organization has greater control over its data, more customization capabilities, and higher levels of security and compliance. Nevertheless, the Private Cloud model usually comes with higher costs, both in terms of initial setup and ongoing maintenance. It also may lack the kind of elasticity and scalability that a Public Cloud can offer, particularly if the Private Cloud is hosted on-premises [11], [12].

The hybrid cloud model serves as a versatile computing environment that seeks to harmonize the advantages of both private and public cloud deployments. In a hybrid cloud configuration, an organization utilizes a private cloud for specific, sensitive tasks that require a high level of security and control, and a public cloud for tasks that can benefit from greater scalability and cost-efficiency. Importantly, the private and public cloud components in a hybrid cloud remain distinct but are interconnected through a set of technologies that allow for seamless data and application portability. This interoperability enables organizations to distribute their workloads more

strategically, choosing the most appropriate environment for each task [13], [14].

Hybrid cloud deployments offer significant flexibility in how organizations manage and allocate their computing resources. One common approach is for an organization to operate its own private cloud while also leveraging services from a public cloud provider. Alternatively, some organizations may choose to partner with a vendor that specializes in private cloud solutions and also has a relationship with a public cloud provider. The overarching objective is to enable an organization to tailor its cloud computing environment to its specific needs. For instance, an organization might opt to use the public cloud for handling large-scale data analytics where the cost and scalability advantages are most pronounced, while keeping sensitive customer data in a private cloud where it can exercise more stringent security controls.

Figure 2. Hybrid cloud backend components and data management

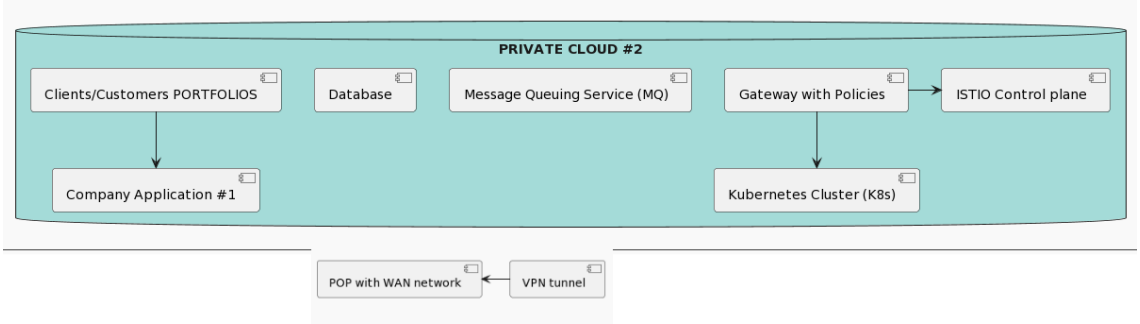

The decision to employ a hybrid cloud is often driven by a combination of operational requirements, strategic objectives, and regulatory considerations. Operational efficiency can be improved by dynamically allocating resources across public and private environments based on current needs. This allows organizations to optimize costs while maintaining performance and security standards. Regulatory compliance is another major factor; certain data may be subject to laws that require it to be stored and processed in a specific way or within a specific jurisdiction. A hybrid cloud offers the flexibility to comply with such regulations while still enjoying the benefits of cloud computing.

The hybrid cloud model aims to offer organizations the best of both worlds by combining the scalable, cost-effective aspects of public clouds with the greater control and performance of private clouds. This allows businesses to seamlessly integrate cloud-based services without the need for substantial modifications to their existing IT infrastructure. This is particularly advantageous for organizations that have made significant investments in on-premises systems but still want to capitalize on the benefits of cloud computing for specific workloads or business processes. By enabling a more flexible resource allocation, hybrid clouds can help organizations optimize costs and improve operational efficiencies.

Hybrid cloud management

Hybrid Cloud Management represents an advanced approach in managing computing resources by combining on-premises infrastructure and cloud-based services, including public and private clouds as well as services from multiple cloud providers. The objective is to offer an integrated platform that streamlines various organizational needs [15], [16]. Resource provisioning allows for the effortless creation and scaling of IT resources across both local and cloud-based environments. This is particularly useful for organizations that require rapid resource allocation for fluctuating workloads. Monitoring and management capabilities allow for the constant observation of resource performance, system health, and security metrics across all organizational environments. Automation features contribute to the streamlining of resource provisioning and routine tasks, reducing the potential for human errors and increasing overall efficiency.

The unified platform often includes tools for chargeback and showback, providing a transparent account of resource consumption and enabling better budget planning. Security and compliance features play an integral role in hybrid cloud management. The platform can enforce uniform security measures such as access controls and encryption across multiple environments. This ensures that security policies and compliance requirements, such as GDPR or HIPAA, are consistently applied, regardless of where the data resides or how resources are being utilized.

The key elements of Hybrid Cloud Management include scalability, flexibility, resource optimization, and multi-cloud management. Scalability and flexibility allow organizations to scale their applications and services up or down based on demand. These features also permit the choosing of the most suitable environment for specific workloads, whether it is on-premises or in the cloud. Resource optimization aims to minimize waste by allocating resources effectively, based on real-time demand and performance metrics.

Policy-based resource management

Policy-based resource management is a structured approach that involves the application of predefined rules or policies to manage various resources within a hybrid cloud environment. The term 'resources' here includes computational power, storage, bandwidth, and other components crucial for the functioning of cloud-based applications and services. The policies can serve diverse objectives, including, but not limited to, controlling access, specifying computational resource allocation, or adhering to external legal and regulatory mandates. By delineating clear policies, organizations can establish standards for how resources should be allocated and used, thereby eliminating ambiguity and potential for misuse [17], [18].

The automation enabled by policy-based resource management is a significant benefit, particularly for large-scale organizations that manage vast arrays of resources across different cloud environments. Automation helps in reducing manual intervention, thereby minimizing human errors and streamlining administrative workflows. It allocates resources where they are most needed, according to rules set forth by the organization. For example, a policy could automatically allocate additional server capacity for a retail website during peak

shopping seasons to handle increased traffic, while another policy could restrict access to sensitive data, ensuring that only authorized personnel can view it. Whether it is adhering to data sovereignty laws or following industry-specific compliance standards like the Health Insurance Portability and Accountability Act (HIPAA) or the Payment Card Industry Data Security Standard (PCI DSS), effective policy enforcement ensures that an organization can meet its legal and ethical obligations. This has the added benefit of enhancing an organization's reputation and could potentially mitigate legal and financial repercussions associated with non-compliance.

Various platforms exist to facilitate the implementation of policy-based resource management in hybrid cloud environments. VMware vRealize Suite, Microsoft Azure Policy, and AWS Organizations are among the cloud management platforms that offer capabilities to define, implement, and enforce policies. These platforms usually feature a centralized dashboard that allows administrators to manage policies across different cloud environments and services. Through this centralized control, organizations can ensure consistency in the application and enforcement of policies, thereby achieving a more coherent and efficient management structure.

Specifying conditions and criteria for the allocation and deployment of cloud resources enables organizations to use these assets more efficiently. For instance, a policy could stipulate that virtual machines should be consolidated onto fewer physical servers during periods of low demand, thereby saving energy and reducing costs. Such optimal utilization not only maximizes the value derived from the resources but also reduces waste, which is particularly important in large-scale operations where even minor inefficiencies can lead to significant cumulative losses.

Another prospect offered by policy-based resource management is automated governance. The method allows for automatic enforcement of compliance and security protocols, making it easier for organizations to meet regulatory requirements and maintain security standards. Automated governance eliminates the need for manual monitoring to a great extent, thereby reducing the chances of human error, which can be both costly and risky. By translating governance requirements into executable policies, organizations can more reliably adhere to established guidelines, ensuring that compliance is consistently maintained across different cloud environments.

The automated allocation of resources based on pre-established policies and usage patterns can lead to substantial operational cost savings. For example, a policy might dictate that less critical workloads be moved to cheaper, lower-performance storage during off-peak hours, and then moved back when performance is more critical. Such dynamic reallocation based on real-time needs ensures that organizations only pay for the resources they actually need, minimizing unnecessary expenditures.

There is inherent complexity in defining and maintaining the policies. Crafting policies that adequately meet an organization's complex needs and regulatory requirements often necessitates specialized expertise. This complexity can act as a barrier to entry for smaller organizations or departments that lack the technical skills to implement and manage policy-based

systems. Even in larger organizations, the initial setup and ongoing management can be time-consuming and complicated.

While policies can be effective for routine operations, they may not easily accommodate ad-hoc or exceptional circumstances. For example, if an unforeseen surge in demand occurs, rigid policies may not allocate resources quickly enough to meet this new requirement, thereby impacting performance or availability. Additionally, administrative overhead can become a concern, as frequent updates to policies may be necessary to adapt to changing conditions or regulatory frameworks, requiring continual oversight and management efforts.

Cross-cloud load balancing

Cross-cloud load balancing is a strategy for distributing incoming network traffic across multiple cloud environments, and in some cases, extending it to on-premises infrastructures. This technique aims to optimize the performance of applications by ensuring that the computational workloads and data traffic are evenly distributed among available resources. Such an approach prevents any single cloud environment or server from becoming a bottleneck, thereby improving application response times and user experience. Furthermore, cross-cloud load balancing enables more efficient use of resources, as it can direct traffic to servers that are less busy or closer to the end-users geographically.

The load balancer itself is specially designed to be platform-agnostic, software-based, and globally functional. This load balancer manages traffic across different cloud deployments, referred to as multi-cloud or hybrid cloud environments. A multi-cloud environment could comprise multiple public cloud providers, while a hybrid cloud environment consists of at least one public cloud and one on-premise data center. The load balancer directs traffic to servers located in these different environments. The servers in individual clouds or data centers handle the incoming traffic as directed by the load balancer. The interaction between the DevOps team, the load balancer, and the cloud or data center servers ensures that traffic is efficiently distributed regardless of where the servers are located. This allows for a highly available and resilient system, capable of serving global user traffic.

Figure 3. load balancing in hybrid and multi-cloud

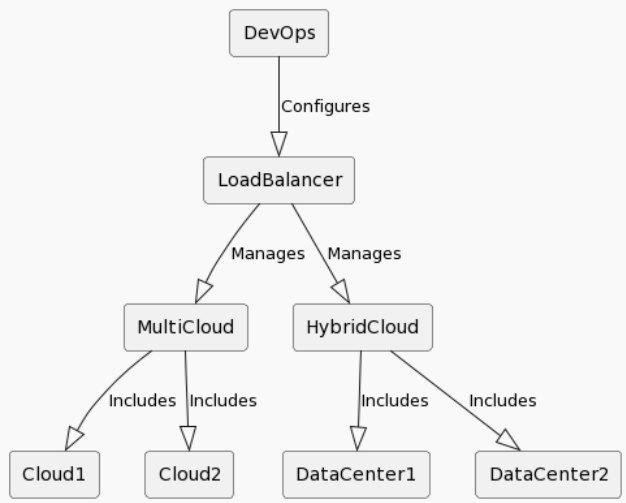

The focus on high availability is another aspect of cross-cloud load balancing. Dispersing traffic across multiple environments minimizes the risk associated with the failure of a single cloud service provider or data center. In case one environment faces an outage, the load balancer can automatically reroute the traffic to other operational environments, thereby ensuring that the application remains accessible. This fault tolerance is

particularly crucial for mission-critical applications where downtime can have severe financial or operational repercussions.

Automatic routing of traffic to the most suitable resources is another capability of cross-cloud load balancing solutions. Advanced load balancers use algorithms to assess various parameters such as server health, current load, and latency to determine where to direct incoming requests. This intelligent routing allows organizations to adapt to fluctuating demand dynamically and ensures optimal performance regardless of the volume of traffic or specific conditions in any given cloud environment. As a result, not only does the application perform better, but the resources are also used more efficiently.

Distributing network traffic across multiple cloud environments approach minimizes the impact of a single point of failure. If one cloud provider experiences downtime or other operational issues, the load balancing mechanism can automatically reroute the traffic to other available and functioning cloud services. This ensures that applications and services remain accessible, meeting the high availability requirements that are often crucial for business continuity and customer satisfaction.

As the demand for an application changes, whether increasing or decreasing, the load balancing mechanism can adapt swiftly. This is particularly useful for applications with varying workloads, such as e-commerce platforms that experience seasonal traffic spikes. The ability to scale resources across multiple cloud providers offers flexibility and ensures that the application can handle increased traffic without suffering performance degradation.

Cross-cloud load balancing solutions often employ algorithms that can route traffic to the closest or most performant server location. This reduces latency and enhances the overall user experience. It can also contribute to more efficient utilization of resources, as workloads can be directed to servers based on real-time performance metrics, ensuring that no single environment is overwhelmed while others are underutilized [19], [20].

Managing multiple cloud providers, each with their own set of tools, APIs, and billing structures, can add a layer of operational intricacy that might be challenging to navigate. This complexity necessitates specialized expertise and may require additional management tools to maintain coherency across different cloud environments.

Data transfer costs can be significant and may offset some of the advantages gained through improved performance and availability. Additionally, potential inconsistency is a concern. Maintaining consistent configurations, security protocols, and policies across multiple providers can be challenging, increasing the risk of configuration errors or security vulnerabilities. Therefore, while cross-cloud load balancing offers various advantages, these limitations should be carefully considered during implementation.

Hybrid cloud service mesh

Service mesh in hybrid cloud environment is found to be useful for modern organizations that rely on complex, distributed systems for their operations. A service mesh essentially serves as a communication control plane that stands between microservices, facilitating their intercommunication, managing data flow, and implementing policies. In a hybrid cloud

service mesh, this layer is extended to not just manage services within a single cloud environment, but across multiple clouds as well as on-premises infrastructure. This offers a unified approach to managing inter-service communication, regardless of where those services reside. With handling tasks like load balancing, traffic routing, and service discovery, the mesh allows developers and IT teams to focus on application logic rather than networking intricacies.

Hybrid cloud environment faces the challenges of security of communication channels between different microservices, which may reside on different cloud platforms or on on-premises servers. A hybrid cloud service mesh addresses this concern by providing built-in security features, such as identity-based authorization, encryption, and mutual TLS authentication between services. This ensures that data transmission across services is secure, irrespective of the underlying infrastructure, which is crucial for compliance with industry regulations and internal security policies. It essentially establishes a zero-trust security model across the hybrid cloud, ensuring that every service is authenticated and authorized before it can communicate with another.

Apart from security, traffic management is another vital functionality provided by a hybrid cloud service mesh. Organizations can define complex routing rules, implement canary deployments, and conduct A/B testing with ease. The mesh enables the monitoring and control of traffic flow between services, allowing for intelligent routing based on a variety of parameters, such as latency, service health, and other custom metrics. In case of service failure or slowdown, the mesh can automatically reroute traffic to healthy instances, thus enhancing the overall resilience and availability of applications [21], [22].

In distributed architectures, services are often dynamically scaled and may move between various nodes, making it challenging to keep track of service endpoints. The service mesh alleviates this problem by maintaining a real-time registry of services and their locations. When a service needs to communicate with another, the mesh provides the current endpoint information, ensuring that the communication is efficient and without unnecessary latency. This dynamic service discovery is particularly useful in hybrid cloud environments, where services may be scattered across multiple cloud platforms and on-premises infrastructure.

The implementation of a hybrid cloud service mesh can offer organizations improved operational efficiency and flexibility. Unifying communication control across various environments simplifies the network architecture and makes it more manageable. Teams can apply consistent policies and rules across the board, reducing the likelihood of errors and inconsistencies. Furthermore, the centralized monitoring and logging capabilities provided by the mesh offer invaluable insights into system performance and security, facilitating quicker troubleshooting and better decision-making. These benefits make a compelling case for the adoption of a hybrid cloud service mesh in organizations that aim for agility, security, and operational excellence.

The prospects of employing a hybrid cloud service mesh in an organization are numerous and carry substantial benefits.

The architecture offers centralized control over multiple environments, enabling easier observability of all service-to-service communications. This centralization facilitates quick troubleshooting, advanced analytics, and performance monitoring, thus allowing for informed decision-making and problem resolution. It provides an overarching view of metrics and logs, reducing the time and effort needed to aggregate data from different sources.

The built-in features often include strong encryption methods and identity-based access controls. This ensures secure communication channels between microservices, irrespective of whether they are deployed on-premises or across multiple cloud environments. The identity-based access controls further validate that only authenticated and authorized services can access certain resources or data. By essentially implementing a zero-trust model across the hybrid cloud, the service mesh ensures compliance with stringent industry regulations and helps in fulfilling internal security mandates.

Organizations typically utilize a variety of technologies, and a service mesh ensures that these diverse services and infrastructures can communicate with one another seamlessly. Given that services may be written in different programming languages and may be running on different cloud providers or on-premises servers, the ability to standardize communication across these heterogeneous environments is vital. The service mesh handles this by abstracting the communication layer, allowing services to interact without knowing the details of their respective environments.

The service mesh introduces an additional layer in the network, and as messages pass through this layer for routing, security checks, and other tasks, they may incur extra network hops that introduce latency. While the latency is often minimal, in real-time or high-throughput applications, even slight delays can be critical. Lastly, there is the risk of vendor lock-in. Some service mesh technologies are designed to work best with specific cloud services or platforms, which may limit an organization's flexibility to switch or integrate other vendors in the future. This makes it crucial for organizations to carefully evaluate the long-term implications of adopting a particular service mesh technology. The architecture is inherently complex and necessitates specialized knowledge for its configuration and maintenance. Organizations often require skilled engineers who understand both cloud computing and service mesh technologies, which can be a resource-intensive requirement. Incorrect configurations can lead to issues like service disruptions or security vulnerabilities, making it imperative to have an expert handle these setups.

Cross-cloud container orchestration

Cross-cloud container orchestration is fundamentally concerned with the automated configuration, coordination, and management of containerized software applications across various cloud service platforms as well as on-premises data centers. At the core of this orchestration is a centralized orchestration engine, typically managed by orchestration software such as Kubernetes. This engine communicates with each cloud provider's API to initiate tasks such as container deployment, scaling, and load balancing. It translates higher-level directives into API calls specific to each cloud provider, allowing for consistent application deployment and

management across diverse infrastructures.

The orchestration engine is responsible for determining where to place each container based on a set of predefined policies and current system metrics. It takes into consideration factors such as CPU and memory availability, data locality, and network latency when making these decisions. Once the optimal location has been determined, the orchestration engine will deploy the container and dynamically adjust resources as needed. This involves scaling containers vertically (adjusting CPU and memory allocation) or horizontally (adding or removing container instances) based on real-time demand and pre-set rules.

When new instances of a service are deployed or existing ones are terminated, the orchestration engine updates a central registry. This registry helps in routing incoming requests to the appropriate instances. Load balancers distribute incoming application or network traffic across multiple servers, thereby ensuring that no single server is overwhelmed. The orchestration engine may implement load balancing algorithms and update the load balancer dynamically, ensuring optimal distribution of network traffic.

Monitoring and logging mechanisms are also integrated into cross-cloud container orchestration frameworks. These mechanisms collect performance metrics, logs, and other key indicators from each container instance, regardless of where it resides. The aggregated data is used for real-time monitoring, debugging, and performance tuning. Monitoring tools can be configured to send alerts or trigger automatic actions such as scaling operations in response to specific conditions, thus contributing to a self-healing infrastructure.

Cross-cloud container orchestration enables organizations to deploy and scale applications quickly, responding to market demands or computational requirements in real-time. The centralized orchestration engine automates the distribution and scaling of containers, thereby significantly reducing manual intervention and accelerating time-to-market for new features or services.

This orchestration approach simplifies the migration of applications across different cloud providers or between cloud and on-premises environments. Containers encapsulate all the dependencies required for an application to run, making it possible to move applications seamlessly. Coupled with a cross-cloud orchestration engine, this containerization allows businesses to avoid vendor lock-in and leverage best-of-breed services from multiple cloud providers [23].

The orchestration engine ensures that policies, monitoring, and management practices are uniform across all cloud platforms and on-premises data centers. This consistent approach simplifies governance and compliance, as organizations can implement policies or make adjustments in a single place, confident that changes will propagate across all environments.

Managing multiple orchestrators, or even a single orchestrator with various configurations for different cloud providers, can become challenging. The intricacy of configuration files, networking setups, and policy definitions can result in a steep learning curve and ongoing management overhead. While the orchestration engine

incorporates various security measures, such as role-based access controls and encryption, an improperly configured engine can expose the system to risks. In a multi-cloud environment, security must be rigorously maintained across all platforms, requiring a robust understanding of each cloud provider's specific security features. There is also a concern of resource consumption. The orchestration services themselves require computational power and memory, which may necessitate additional infrastructure costs.

Log management and analytics

Log management and analytics in a hybrid cloud environment involve a set of activities and technologies that interact with each other to provide a holistic view of the system. The process starts with the collection of logs from various sources. These logs are typically generated by different components such as applications running on virtual machines, databases hosted on dedicated servers, and networking devices like switches and routers [24]. Agents or collectors are often deployed on these sources to capture the logs and forward them to a centralized log management system. Some log management systems also provide agentless options, using protocols like Syslog or APIs to collect data. The collected logs are then stored in a centralized database that could be on-premises or cloud-based, depending on the organization's infrastructure strategy.

Once collected, the logs are usually subjected to a normalization process to convert them into a standard format. This involves extracting useful information from the raw logs and mapping them into predefined fields. The normalization is essential for analytics because it ensures that logs from different sources can be analyzed coherently. After normalization, the data often undergoes enrichment, which means adding additional contextual information such as geographical location or threat intelligence data. This enrichment process helps in better understanding the log entries and facilitates more effective analytics [25].

The analytics component commonly applies various algorithms and statistical models to detect patterns, anomalies, or specific events that are of interest. For instance, machine learning algorithms might be used to identify anomalous behaviors that could signify a security breach. These analytics engines can generate alerts based on predefined rules or dynamic baselines. The alerts could then trigger automated actions like blocking an IP address or could be sent to the administrators for manual intervention.

In a hybrid cloud setting, the log management system must also deal with the challenges posed by the diversity of cloud environments. Data privacy laws, latency issues, and network costs can influence where logs are stored and how they are transferred between on-premises and cloud systems. Some log management solutions offer features to automatically segregate sensitive data and store it in compliance with regulations like GDPR or HIPAA. Additionally, they can provide capabilities for logs to be transferred securely over the internet using encryption technologies.

Container orchestration, on the other hand, is concerned with automating the deployment, scaling, and management of containerized applications across multiple cloud environments. Kubernetes, a leading container orchestration platform, achieves

this through a set of control planes and nodes. The control plane manages the overall state of the cluster, while the nodes are the worker machines that run containers. Kubernetes employs a declarative configuration model, meaning that administrators specify the desired state of the system in configuration files, and the platform takes care of making it happen. Resources like CPU, memory, and network bandwidth are allocated based on these configurations. Cross-cloud capabilities come into play by allowing these containerized applications to be managed consistently, irrespective of whether they are deployed in a private cloud, public cloud, or a combination thereof. This is achieved through a consistent API layer and a set of standard primitives like pods, services, and volumes that abstract the underlying infrastructure complexities.

Monitoring tools typically run continuously and scan incoming log data for patterns or events that match predefined rules or heuristics [27]. When a match is found, the system generates alerts that could lead to automated actions or manual intervention by system administrators. In hybrid cloud environments, real-time monitoring becomes crucial due to the highly dynamic nature of the resources. Events in one part of the hybrid cloud can influence operations in another, thus making real-time assessment valuable for maintaining a coherent security posture and ensuring uninterrupted service.

Audit and compliance requirements have increasingly become a compelling reason for robust log management. Regulatory frameworks such as the General Data Protection Regulation (GDPR) in Europe or the Health Insurance Portability and Accountability Act (HIPAA) in the United States mandate the secure storage and management of logs for specified periods. Centralizing log information assists in simplifying the audit process, as auditors can access a single repository of data to verify compliance against regulations. Audit trails can be generated from these centralized logs, showing a sequence of activities that led to particular events or changes. This facilitates accountability and traceability, which are fundamental elements of compliance frameworks.

Insights drawn from log analytics can offer perspectives on various aspects of system operations, including user behavior, system performance, and bottlenecks. Data visualization tools and dashboards often accompany modern log management solutions, providing real-time and historical data analyses. These tools allow organizations to drill down into the data and isolate variables, thus making it easier to identify trends or anomalies. This level of insight is crucial for making informed decisions, whether they relate to optimizing system performance, enhancing user experience, or mitigating security risks.

Storage is a significant concern in log management, particularly because modern hybrid cloud infrastructures can generate enormous volumes of log data. Storage not only involves the physical or virtual space where logs are kept but also entails management policies such as retention periods and data lifecycle. Some logs may need to be kept for extended periods for compliance, while others could be aggregated and summarized before being archived or purged. Organizations often use storage solutions like high-capacity disk arrays, or cloud storage services with adequate encryption and redundancy

features, to meet these diverse storage requirements.

Complexity and noise are significant challenges in log management and analytics. As systems grow more complex and logs proliferate, organizations may need specialized analytics tools capable of parsing large datasets quickly and efficiently. Similarly, "noise" refers to the irrelevant or less significant data that often crowds log files, making it difficult to pinpoint meaningful patterns. Filtering out noise is an ongoing challenge that often requires a blend of rule-based analytics and machine learning algorithms to hone in on the critical events or trends amidst a sea of mundane data.

Conclusion

Hybrid cloud combines the resources of both public and private cloud infrastructures, offering organizations a versatile platform for data storage, application deployment, and various computational needs. The applying of hybrid cloud systems brings several key advantages to an organization, among them being cost-efficiency, heightened security measures, and a flexible, scalable environment. By leveraging both public and private resources, organizations can allocate tasks and data storage in a manner that maximizes efficiency while minimizing costs. For example, sensitive data can be kept in a private cloud to ensure security, while less-sensitive tasks can be offloaded to the more cost-effective public cloud. Additionally, the scalable nature of hybrid cloud allows for rapid adjustments to infrastructure to meet the fluctuating demands of business operations.

Effective hybrid cloud management is crucial for reaping the maximum benefits of this infrastructure model. Managing a hybrid cloud environment involves resource allocation, performance monitoring, and ensuring regulatory compliance. Proper management practices enable IT departments to align the capabilities of the hybrid cloud with the specific requirements and goals of the business. This alignment is essential for optimizing the use of resources and for reducing risks associated with security breaches and data loss. Furthermore, proper governance ensures that the organization adheres to compliance standards, reducing the likelihood of legal complications that could arise from data mismanagement or non-compliance with industry regulations.

The study explored into various aspects of hybrid cloud configuration strategies management, its advantages and challenges. and. Policy-based resource management is a systematic approach to overseeing resources in a hybrid cloud environment through the application of predefined rules or policies. These policies cover a broad range of operations such as access control, computational resource allocation, and compliance protocols. One significant advantage of this method is improved resource utilization. By standardizing the allocation based on policies, resources can be used more efficiently, thereby reducing waste. The automated governance aspect ensures that security and compliance measures are automatically enforced, thus reducing manual oversight and potential human errors. Cost optimization is another compelling aspect, as the system can allocate resources based on usage patterns, which can substantially reduce operational costs. However, this approach is not without its drawbacks. The formulation and maintenance of these policies require specialized expertise, adding a layer of

complexity to the system. Furthermore, the rigid nature of policies may not accommodate exceptional or ad-hoc scenarios easily, thus potentially hampering flexibility. Administrative overhead can also be a challenge, as policies may require regular updates.

Cross-cloud load balancing serves to distribute network traffic across various cloud environments and possibly on-premises infrastructure. By doing so, it enhances application performance, reliability, and fault tolerance. Scalability setup can dynamically adapt to fluctuations in demand. Moreover, performance can be optimized by routing traffic to the server locations that can best accommodate the load. However, implementing cross-cloud load balancing introduces additional complexity because managing multiple cloud providers necessitates a more intricate operational setup. There can also be additional cost implications, especially when data must be transferred between different cloud providers. Consistency in configurations across different environments can be challenging to maintain, given the variations in cloud service offerings.

A hybrid cloud service mesh is designed to facilitate secure and seamless communication between services and applications across multiple cloud environments and potentially on-premises infrastructures. Unified monitoring offers centralized control and visibility across diverse services and platforms. Enhanced security measures, including robust encryption and identity-based access controls, can be readily implemented. Additionally, the architecture enables interoperability between various services, applications, and infrastructure components. On the downside, the configuration and ongoing maintenance of a service mesh demand specialized expertise. Latency is another concern, as the additional network hops required for communication can slow down service response times. Vendor lock-in is also a potential limitation if the service mesh technologies used are specific to a single provider.

Cross-cloud container orchestration involves overseeing containerized applications across multiple cloud providers and on-premises infrastructures. It offers a high degree of flexibility, allowing for quick deployment and scaling of applications. Portability is another key advantage, as applications can be readily moved between different cloud environments. Consistency in the management and monitoring of containers is also ensured, thereby simplifying administrative tasks. Nevertheless, this approach also brings its challenges. The management of multiple orchestrators or configurations can be complex and requires a deep understanding of the underlying systems. Security is another area of concern; if not configured properly, the system could be vulnerable to threats. Additionally, orchestrating containers across different environments may consume additional computational resources, which could have an impact on performance and costs.

Log management and analytics involve the collection, monitoring, and analysis of log data from various sources in a hybrid cloud environment. Real-time monitoring enables immediate identification of security threats or performance bottlenecks. It also simplifies the audit and compliance processes by centralizing all log information, making it easier to produce

reports for regulatory bodies. Furthermore, the analytics component can provide valuable insights into user behavior and system performance. However, storing large volumes of log data can be a challenge, often necessitating significant storage capabilities. Another drawback is the specialized skill set required for effective log analysis. Lastly, the potential for 'noise' in the log data—irrelevant or less important information—can make it difficult to identify meaningful patterns and insights.

References

[1] G. Boss, P. Malladi, D. Quan, L. Legregni, and H. Hall, "Cloud computing," *IBM white paper*, vol. 321, pp. 224–231, 2007.

[2] W. Voorsluys, J. Broberg, and R. Buyya, "Introduction to cloud computing," *Cloud computing: Principles*, 2011.

[3] W. Kim, "Cloud computing: Today and tomorrow," *J. Object Technol.*, vol. 8, no. 1, p. 65, 2009.

[4] B. Hayes, "Cloud computing," *Commun. ACM*, vol. 51, no. 7, pp. 9–11, Jul. 2008.

[5] L. Wang *et al.*, "Cloud Computing: a Perspective Study," *New Generation Computing*, vol. 28, no. 2, pp. 137–146, Apr. 2010.

[6] M. Vaishnnave, K. S. Devi, and P. Srinivasan, "A survey on cloud computing and hybrid cloud," *Int. J. Eng. Res. Appl.*, 2019.

[7] G. Lackermair, "Hybrid cloud architectures for the online commerce," *Procedia Comput. Sci.*, vol. 3, pp. 550–555, Jan. 2011.

[8] J. Weinman, "Hybrid Cloud Economics," *IEEE Cloud Comput.*, vol. 3, no. 1, pp. 18–22, Jan. 2016.

[9] F. Luo, Z. Y. Dong, Y. Chen, Y. Xu, and K. Meng, "Hybrid cloud computing platform: The next generation IT backbone for smart grid," *2012 IEEE Power and*, 2012.

[10] R. Raju, J. Amudhavel, and N. Kannan, "A bio inspired Energy-Aware Multi objective Chiropteran Algorithm (EAMOCA) for hybrid cloud computing environment," *on green computing …*, 2014.

[11] A. Dubey, G. Shrivastava, and S. Sahu, "Security in hybrid cloud," *Global Journal of Computer Science*, 2013.

[12] J. K. Wang and X. Jia, "Data security and authentication in hybrid cloud computing model," *2012 IEEE Global High Tech Congress on*, 2012.

[13] M. I. Tariq, "Agent based information security framework for hybrid cloud computing," *KSII Transactions on Internet & Information Systems*, 2019.

[14] P. Lu, Q. Sun, K. Wu, and Z. Zhu, "Distributed online hybrid cloud management for profit-driven multimedia cloud computing," *IEEE Trans. Multimedia*, vol. 17, no. 8, pp. 1297–1308, Aug. 2015.

[15] Y. Khmelevsky and V. Voytenko, "Hybrid cloud computing infrastructure in academia," in *WCCCE 2015-the 20th Western Canadian Conference on Computing Education*, 2015, pp. 8–9.

[16] M. I. Tariq, "Analysis of the effectiveness of cloud control matrix for hybrid cloud computing," *Int. J. Future Gener. Commun. Netw.*, vol. 11, no. 4, pp. 1–10, Jul. 2018.

[17] T. V. N. Rao, K. Naveena, and R. David, "A new computing environment using hybrid cloud," *and Computing …*, 2015.

[18] S. Goyal, "Public vs private vs hybrid vs community - cloud computing: A critical review," *Int. J. Comput. Netw. Inf. Secur.*, vol. 6, no. 3, pp. 20–29, Feb. 2014.

[19] R. Balasubramanian and M. Aramudhan, "Security issues: public vs private vs hybrid cloud computing,"

International Journal of Computer, 2012.

[20] G. Mateescu, W. Gentzsch, and C. J. Ribbens, "Hybrid Computing—Where HPC meets grid and Cloud Computing," *Future Gener. Comput. Syst.*, vol. 27, no. 5, pp. 440–453, May 2011.

[21] C. Zou, H. Deng, and Q. Qiu, "Design and implementation of hybrid cloud computing architecture based on cloud bus," *conference on Mobile ad-hoc and …*, 2013.

[22] A. Gordon, "The hybrid cloud security professional," *IEEE Cloud Comput.*, vol. 3, no. 1, pp. 82–86, Jan. 2016.

[23] D. S. Linthicum, "Emerging hybrid cloud patterns," *IEEE Cloud Comput.*, vol. 3, no. 1, pp. 88–91, Jan. 2016.

[24] K. A. Beaty *et al.*, "Managing sensitive applications in the public cloud," *IBM J. Res. Dev.*, vol. 60, no. 2–3, p. 4:1-4:13, Mar. 2016.

[25] J. Thaler, W. Shin, S. Roberts, and J. H. Rogers, "Hybrid approach to hpc cluster telemetry and hardware log analytics," *2020 IEEE High*, 2020.

[26] M. N. Birje and C. Bulla, "Cloud monitoring system: basics, phases and challenges," *International Journal of Recent Technology Engineering (IJRTE)*, vol. 8, no. 3, pp. 4732–4746, 2019.