UNIVERSITY OF ALBERTA

Topic: Local and Hybrid Cloud

A MINT 709 Capstone Project submitted to the Departments of Computing Science and Electrical and Computer Engineering of the University of Alberta in partial fulfillment of the requirements for the degree of

MSc Internetworking

By Syed Ahmed Hussain
2-28-2021

Acknowledgment

I am highly indebted to my supervisor, Mr. Juned Noonari in the preparation of this report, whose patience, time, kindness, and academic experience, have been invaluable to the completion of this report.

The encouragement and support from many of my family, relatives, and friends have been indispensable to the writing of this report. I would not have contemplated this road if not for my parents, who instilled within me a love of creative pursuits for an MSc in Internetworking, all of which finds a place in this report.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

1.0 ABSTRACT

The objective of this project is to study Local and Hybrid Cloud computing models and benchmark which of the cloud providers (or vendors) provide top-of-the-class services when comparing their components and services. The project will focus mainly on Scalability, Latency, Security, Distributed Computing, Virtualization, and Containerization. The project will also investigate the business model of globalization and localization of cloud services. Finally, the project is going to discuss how Multi-Cloud solutions can be used to bring the best features from multiple cloud providers and local cloud.

2.0 INTRODUCTION

Many businesses are going online and there is a growing trend towards digital transformation. Cloud computing is the buzzword that comes to the minds of individuals who are part of this digital transformation. Cloud is the delivery of IT resources such as servers, databases, storage, networking, software, and analytics over the internet as a service. There are several key reasons why a company may move its IT resources to the cloud:

- Cost: Cloud computing significantly reduces the capital costs for buying IT hardware and software. IT also reduces the operational costs of data centers such as electricity and staff for managing the infrastructure.

- Scalability: Cloud computing can elastically scale resources based on the requirement of a business.

- Reliability: Cloud Computing helps in making business services available round the clock as it can mirror data and services on multiple (datacenter) sites. [1]

Locally hosting a cloud on-premises provides businesses the similar benefits that of the publicly available cloud. The benefits are similar which include self-service, scalability, and elasticity. Besides, a private cloud provides better security and privacy, dedicated resources, cost control, and more that will be discussed in this report.

3.0 CHALLENGES FACED BY LOCAL AND HYBRID CLOUD

This project report focuses on the following challenges faced by cloud computing:

LATENCY

Network Latency is a major factor that affects cloud computing since all the services are hosted over the internet. This usually happens due to saturation or large data sets being transported. The bursty nature of network traffic makes latency a key performance issue. [2]

CLOUD SECURITY

CIA triad is significantly affected which can hinder the adoption of Cloud services. Several traditional attacks are possible such as DoS, side-channel, and keystroke timing attacks. But in addition to traditional attacks, there are also Cloud-specific attacks that exploit vulnerabilities in Hypervisor. EDoS is another type of attack that exploits utility pricing by fraudulent resource consumption. [3]

IMPLICATIONS OF GOING GLOBAL

There are critical issues related to the Legal, Privacy, Compliance, and Personal Information Protection and Electronic Documents Act (PIPEDA) that needs to be addressed when going global. [4]

RELIANCE ON SINGLE CLOUD PROVIDER

Differences in Features and related Costs between different Vendors is another issue that needs to be addressed. To optimize the cost of operations and make use of the best of both worlds. [5]

BUSINESS MODEL FOR CLOUD COMPUTING

It is still difficult to categorize SMB, SME and Large Enterprise need to migrate to the cloud and what type of cloud i.e. Public, Local, Community, or Hybrid cloud deployment models best fits their needs.

FEATURE AVAILABILITY

There is no clear benchmark about the features available from different local and hybrid cloud vendors such as Scalability, Virtualization, Containerization, etc.

4.0 METHODOLOGY

For the challenges posted before this section, the following methodology is used to analyze the project scope:

1. Literature Review: In this section, major features such as Scalability, Virtualization, Containerization, Compute, Storage, Networking, and more are studied, and this section is used as a reference to benchmark major cloud providers. This section deals with all the theories related to cloud and service models, private cloud, hybrid cloud, cloud security, distributed cloud, and multi-cloud technologies. The project will take a deep dive into the legal compliance and best practices published by NIST, (Canadian center of Cyber Security) ITSG-33, and PIPEDA. And will also analyze if Hybrid or Local Cloud Computing will be a better fit for some security-sensitive organizations.

2. Benchmarking local and hybrid cloud technologies: This section deals with benchmarking of the issues mentioned in the abstract which are as follow:

    a. Latency: To tackle this, Distributed Cloud computing technologies or services will be analyzed and benchmarked between top Private and hybrid cloud technologies.

    b. Cloud Security: Security services provided by top Cloud providers will be analyzed and benchmarked.

    c. Scalability: Types of scalabilities will be discussed and the support for this feature will be compared in this section

    d. Multi-cloud: Features for multi-cloud will be analyzed and compared using the literature review section as a reference.

    e. Distributed Cloud: Edge computing support will be compared and analyzed in this section.

    f. Containerization and virtualization: In this section Virtualization will be compared with containerization.

3. Business Model for Cloud Computing: This project will investigate which type of cloud computing model suits best for SMB, SME, and Large Enterprise based on an economical perspective.

5.0 LITERATURE REVIEW

5.1 WHY CLOUD COMPUTING

The primary reason to have a cloud-based infrastructure based on NIST definition:

> Cloud Computing is defined as a model that enables ubiquitous, convenient, on-demand access to a shared pool of configurable computer network resources that can also be rapidly provisioned or released with very little management. [6]

5.1.1 ESSENTIAL CHARACTERISTICS OF CLOUD

On-Demand Self-service: Consumers can self-provision the computing resources when needed automatically without human interaction.

Broad Network Access: The cloud is capable to run over the network and is accessible to different types of thin or thick clients such as mobile phones, laptops, workstations, and tablets.

Resource Pooling: Cloud makes use of the multi-tenant model and provides resources to multiple consumers. The virtual and physical resources are allocated dynamically based on demand. The consumer has no knowledge or control of the location of resources allocated but may be able to get a higher level of abstracted location such as country, state, zones, or datacenter.

Rapid Elasticity: Cloud can provide rapid elasticity of provisioned resources. The resources can (automatically) scale rapidly inwards or outwards based on the required demand. This capability may appear to have unlimited capacity.

Measured Service: Cloud systems can also optimize or control resource utilization automatically by using metering services. The usage utilization can be reported to the user or provider providing transparency. [6]

5.1.2 SERVICE MODELS

Software as a Service (SaaS): It is a cloud-based application service provided to the consumer. The consumer is not responsible for underlying infrastructure which includes network, server, OS, and storage. The only responsibility the consumer may hold is of the application configuration settings.

Platform as a Service (PaaS): It is a service in which the consumer can deploy their own application by programming or acquiring the license. The consumer is not responsible for underlying hardware such as network, server, OS, and storage. The consumer is completely responsible for the ins and out of the application which may include bug fixes, updates, performance, etc.

Infrastructure as a service (IaaS): Consumer is provided with an ability to provision and customize compute, storage, networks and install their preferred OS. The consumer is not in complete control of the underlying infrastructure. [6]

The figure from Microsoft displays with examples the service models as follows:

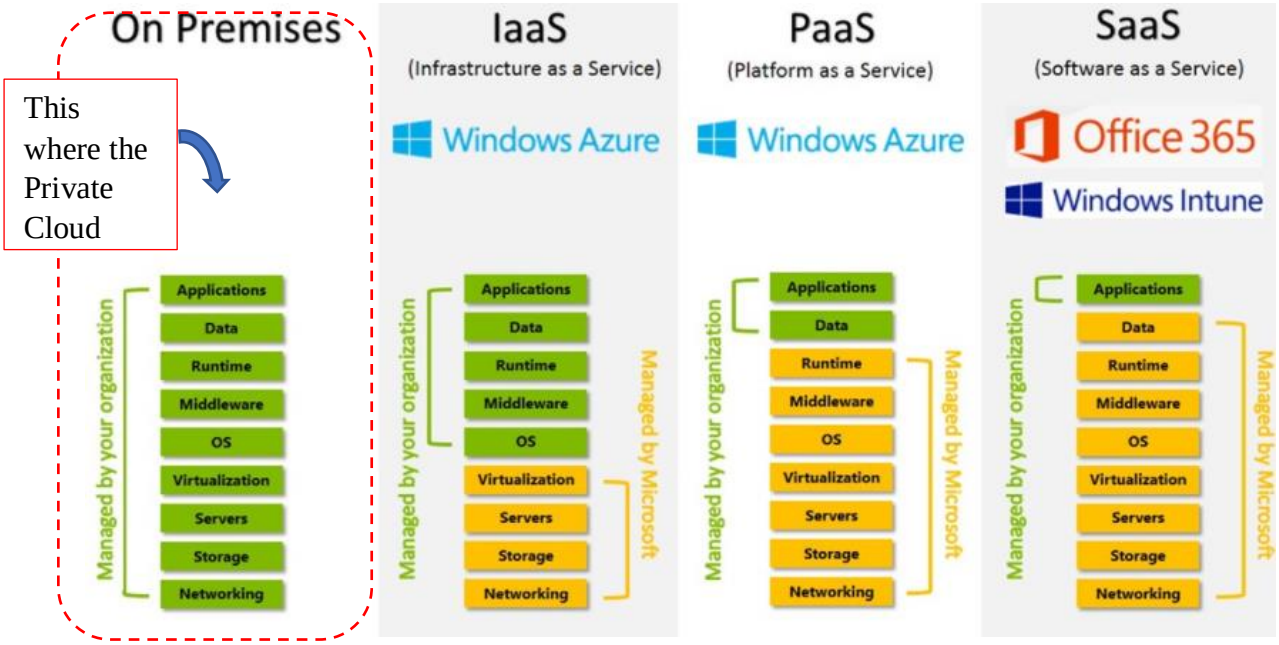

Figure 1: Service Models [7]

Since the focus of this report is on Private and Hybrid Cloud, which is on-premises infrastructure with an overlap of IaaS due to the presence of a hybrid component.

5.2 CLOUD COMPUTING DEPLOYMENT MODELS

The most common cloud models used in the industry are Private, Public, Hybrid Cloud, and Community Cloud models. National Institute of Standards and Technology (NIST) has defined these models and this report will discuss them as follows:

5.2.1 PRIVATE (LOCAL) CLOUD

According to NIST, it is a cloud-based infrastructure is designed specifically for an individual organization that may involve multiple business units. The IT infrastructure cloud is owned, operated, or managed by the organization itself, or another third-party service provider, or a combination of both. [6]

The following figure illustrates the simplified private cloud architecture:

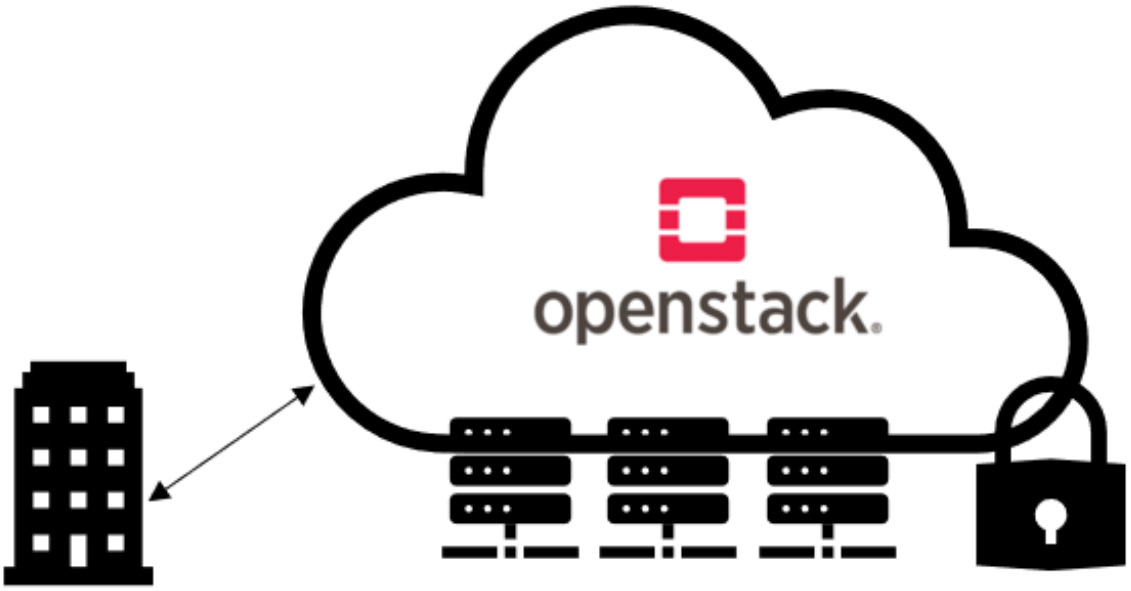

Figure 2: Private Cloud

5.2.2 PUBLIC CLOUD

This cloud infrastructure is the complete opposite of a Private cloud. NIST defines it as for open use to the public. This infrastructure may be owned, managed, or operated by an organization and it is present on the premise of the cloud service provider (CSP). [6]

The following figure illustrates the simplified public cloud architecture:

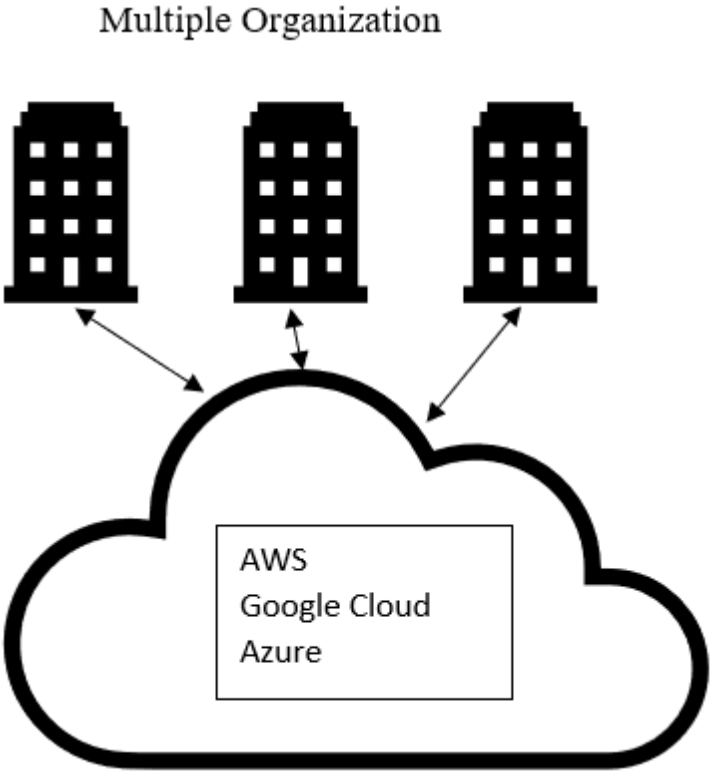

Figure 3: Public Cloud

5.2.3 HYBRID CLOUD

It is defined by NIST as a composition of multiple cloud infrastructure which may include private, public, or community cloud. They remain a separate entity and are connected by some proprietary or standardized technologies that allow data or application mobility. [6]
The following figure illustrates the simplified hybrid cloud architecture:

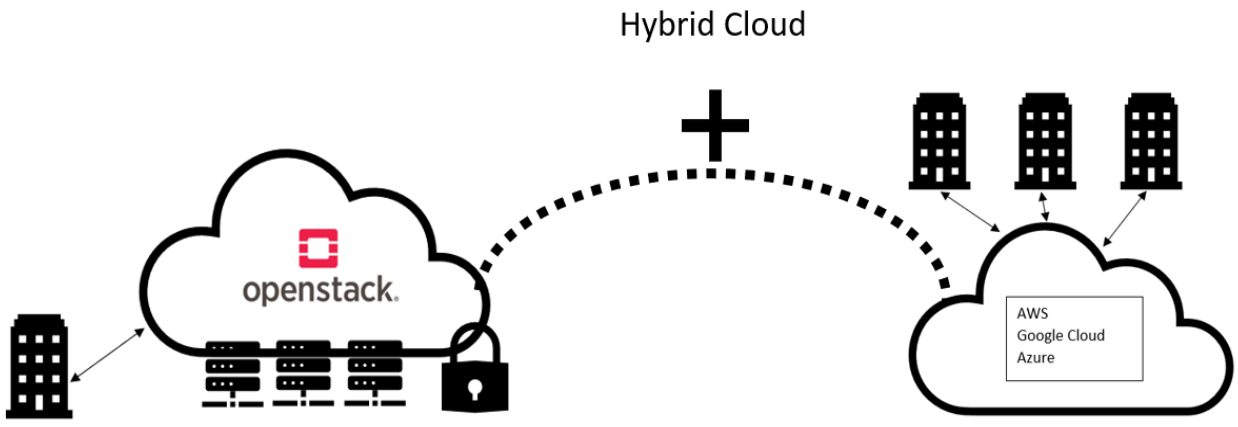

Figure 4: Hybrid Cloud

5.2.4 COMMUNITY CLOUD

A community cloud is an infrastructure that is reserved for organizations that have security concerns and have compliance issues with the governing bodies. It can be managed, operated, or owned by one or more organizations in the community such as the Health care industry. [6]

The following figure illustrates the architecture of the community cloud:

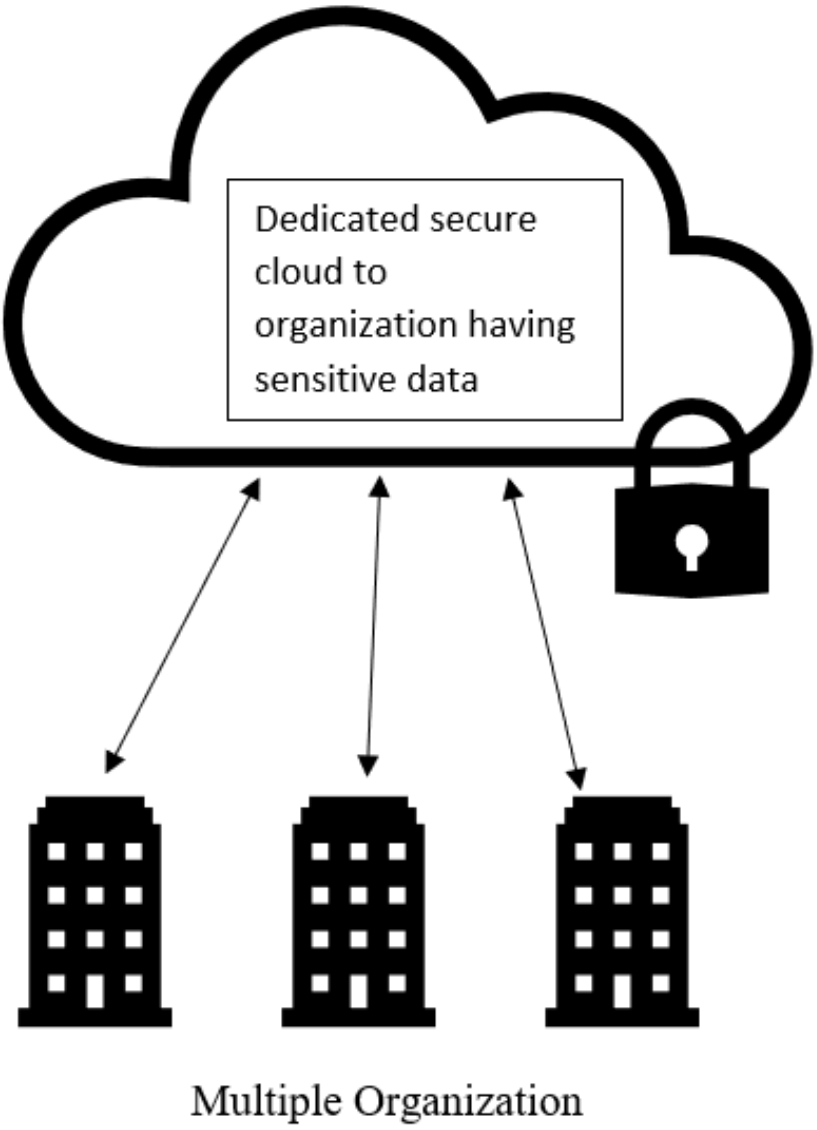

Figure 5: Community Cloud

5.3 CASE OF LOCAL AND HYBRID CLOUD

5.3.1 SERVICE-ORIENTED FRAMEWORK

This report will discuss the private cloud infrastructure based on Service-oriented infrastructure (SOI) which divides various layers in a very systematic way. As see in the figure below, there are a total of five layers.

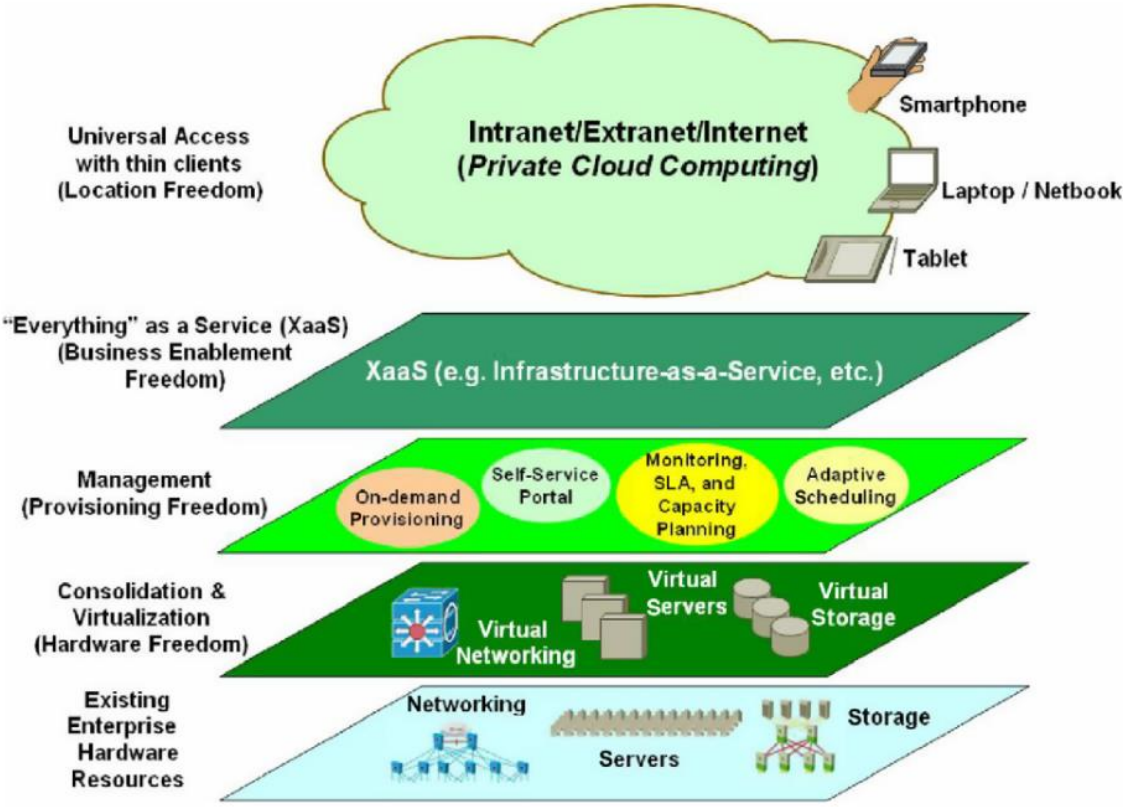

Figure 6: SOI framework [8]

The lowest layer is based on the physical hardware that consists of servers, storage, networking (e.g. routers and switches), firewalls, and other telecommunication hardware. The Consolidation and Virtualization layer consists of virtual networking components, servers, storage, and a firewall. This removes the limitation associated with physical hardware and allows optimization, flexibility, and efficiency of deployment. Then comes the management layer, this layer simplifies the provisioning of the bulk of resources in a simpler and GUI-based portal that contains On-demand provisioning, a self-service portal, metering (e.g. SLA and usage), and adaptive scheduling. [8]

5.3.2 CLOUD BURSTING DURING PEAK LOAD

The significant advantage of using hybrid cloud infrastructure is to compensate for the peak demand faced by private using the public cloud infrastructure. So, when the demand for an application increases at an instant, the extra load is burst to the public cloud. This in turn saves the cost of deploying additional hardware for seasonal traffic and any possible latency during peak time. [9]

5.3.3 CLOUD SECURITY

5.3.3.1 SECURITY DOMAINS

Security domains: They contain users, apps, and IT infrastructure that share trust requirements and the same Authentication and Authorization requirements. These security domains are as follows:

- Public: This refers to the internet which is an untrusted area
- Guest: This refers to the traffic between compute instances.
- Management (Control Plane): Refers to the area where services interact.
- Data Plane: This refers to an area containing storage services. [10]

The following figure displays the trust level of each security domain:

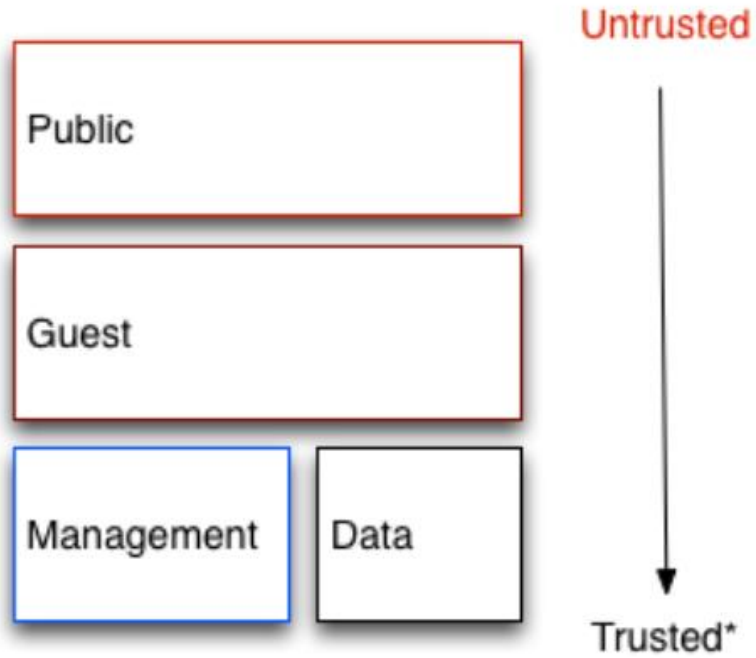

Figure 7: Security Domains in a cloud [10]

5.3.3.2 THREAT ACTORS

Cloud technologies are exposed to attacks by threat actors as follows:

- **Intelligence services**: These are the most capable adversary that can deploy a tremendous amount of resources to attack a target. Requires strong controls in place.

- **Serious organized crime:** Attackers that are highly capable and financially driven. They can fund the development of in-house exploits.

- **Highly capable groups**: Also known as Hacktivists that have a political motive to attack a cloud operator or service and are not commercially funded.

- **Motivated individuals**: Can be rouge or malicious employees and customers acting alone.

- **Script kiddies**: Make use of the automated tool for scanning and exploitation and can damage an organization's reputation. [10]
  The following figure summarizes the likelihood of an attack and the complexity of exploitation:

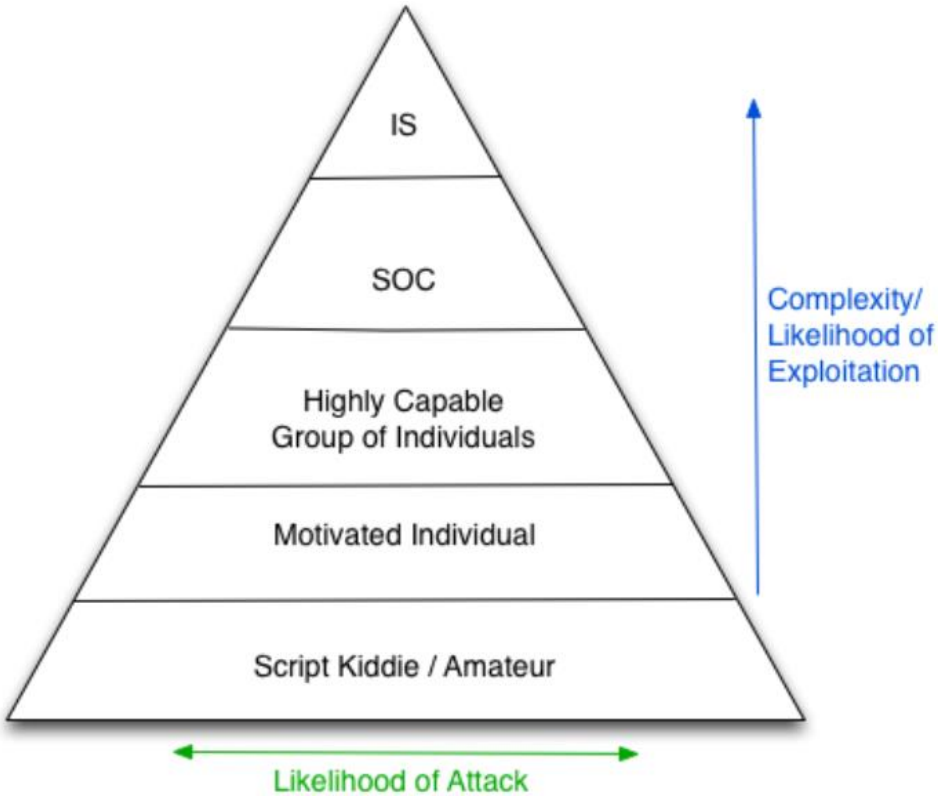

Figure 8: Likelihood of an attack vs complexity of exploitation [10]

5.3.3.3 ATTACKS ON CLOUD PLATFORMS AND REPUTATIONAL RISK

Cloud platforms are ideal for distributed denial of service attacks (DDoS) and brute force attacks. The reason behind this is due to a large number of resources present in it which results in a greater attack surface. And the risk is even greater when the private cloud is exposed to the public network to obtain hybrid or multi-cloud functionalities. It can damage the reputation of an enterprise especially if the cloud is used by attackers to host malicious software for launching attacks on other networks. Internal threats such as a rogue employee can leak sensitive data. [10]

The following figure displays the attacks that are possible by the threat actors:

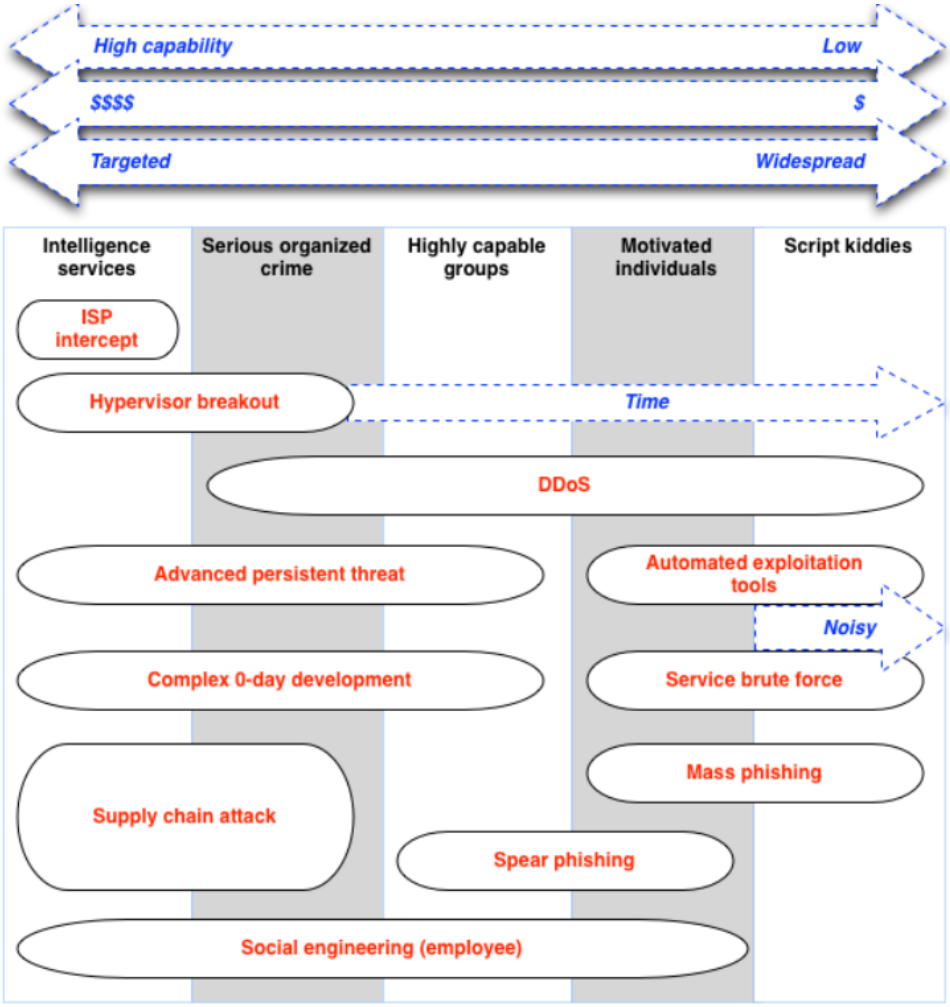

Figure 9: Attacks faced by types of groups [10]

Methods such as data loss prevention (DLP) systems, egress security groups, outbound traffic inspection. Employee training, customer awareness can help in mitigating most of the risks associated with these attacks. [10] Another thing to notice is EDoS attack is irrelevant to private or hybrid cloud as they do not rely on Public Cloud Provider.

5.3.4 GOVERNMENT OF CANADA ON DATA SOVEREIGNTY AND CLOUD

The Treasury Board of Canada Secretariat has carried out an extensive study to address aging IT infrastructure, emerging technology, and security issues surrounding cloud computing on data sovereignty.

The following were the conclusions from the paper:

- The government of Canada has adopted a cloud-first strategy and departments are required to consider using cloud computing services

- Up to Protected B level data is allowed over the Public cloud and the rest is to be stored locally. The security levels of sensitive government data are described below:

  - Protected A Level Data: Data that can cause injury to an individual, organization, or the government

  - Protected B Level Data: Data that can cause serious injury to an individual, organization, or the government

  - Protected C Level Data: Data that can cause extremely grave injury to an individual, organization, or the government [91]

- The government is currently working on developing a protected-cloud contract to allow access to public cloud services. [92]

All this stated information points to the importance of having a private or a hybrid cloud solution as it can meet both needs to secure Protected C Level Data and allows government departments to utilize modern IT infrastructure.

5.3.5 DEFENSE IN DEPTH FOR CLOUD-BASED SERVICES

Canadian Center For Cybersecurity has introduced a guide for cloud computing technology called ITSP.50.104 Guidance on Defence in Depth for Cloud-Based Services. In this document, there are several recommendations to implement a defense in depth on a cloud-based infrastructure. The following figure summarizes the defense-in-depth strategy proposed by the Canadian center for cybersecurity

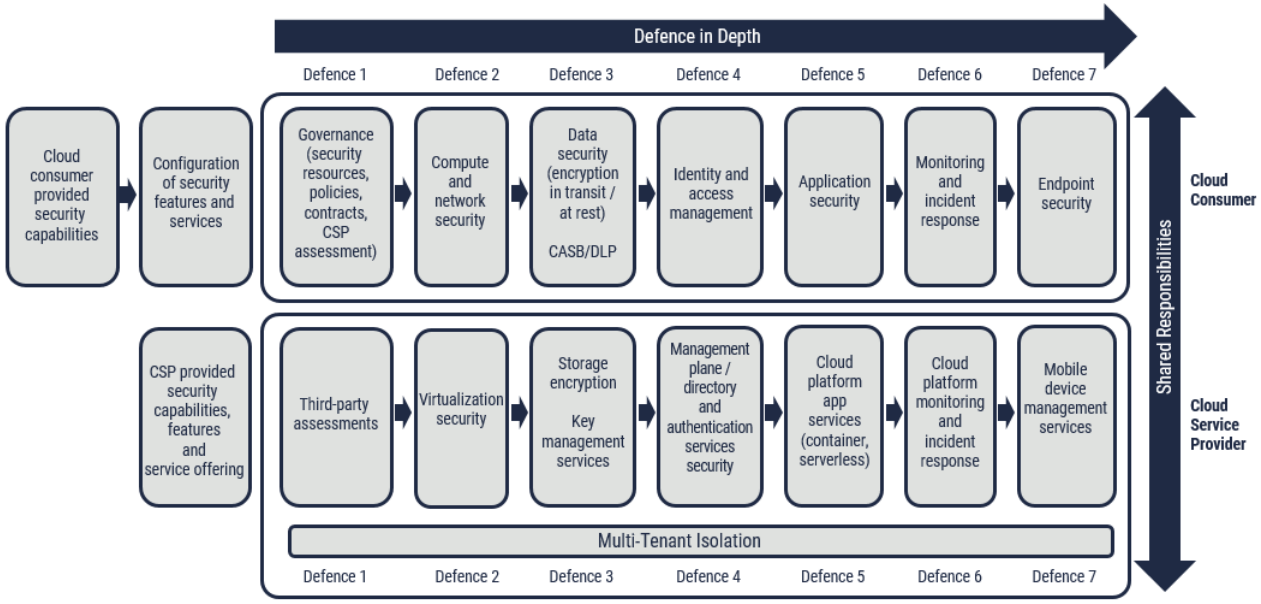

Figure 10: Cloud Defence In Depth Considerations ITSP.50.104 [93]

The concept identifies the importance of shared responsibility by both cloud service providers and consumers.

Multi-Tenant Isolation: This is the CSP's responsibility to ensure that a security breach to one tenant does not affect the CIA triad of others.

It also shows steps to harden the security posture of cloud services:

Defense 1: Focuses on the importance of Cloud Security Governance and Risk Management. This section of the guide discusses the importance of resource allocation, establishing policies and guidelines, third-party assessments and evaluations, and cloud contracts.

Defense 2: Focuses on the importance of network and compute security by CSPs. It discusses that CSPs should normally implement DDoS protection, micro-segmentation, IPS, network intrusion prevention systems (NIPS), and firewalls. CSPs are also responsible for making network security groups available to the consumers as it can help them to enforce ACL, network segmentation, and monitoring. Routing is mentioned as a primary responsibility of consumers. For connecting a private cloud to

a public cloud and turning it into a hybrid cloud, the guide recommends the use of a Connection Exchange Provider (CXP) as shown in the figure below:

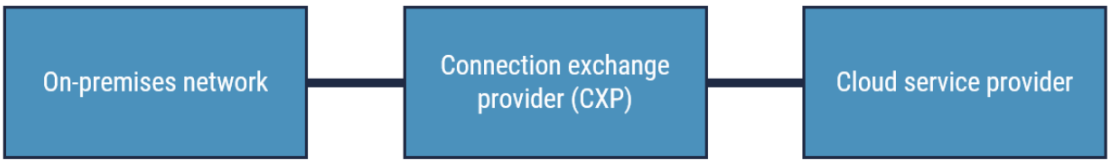

Figure 11: Dedicated private network connection to CSP via Connection Exchange Provider [93]

The reason for using CXP is to provide better security, dedicated bandwidth, and low latency. And the guide also recommends separating the network traffic from on-premises from off-premises cloud infrastructure. Bastion/Transit network can be used to implement the separation between the on-premises network from the off-premises network traffic. It provides locations where access control and other security tools can be implemented.

The following figure displays how Bastion or Transit Network can be implemented:

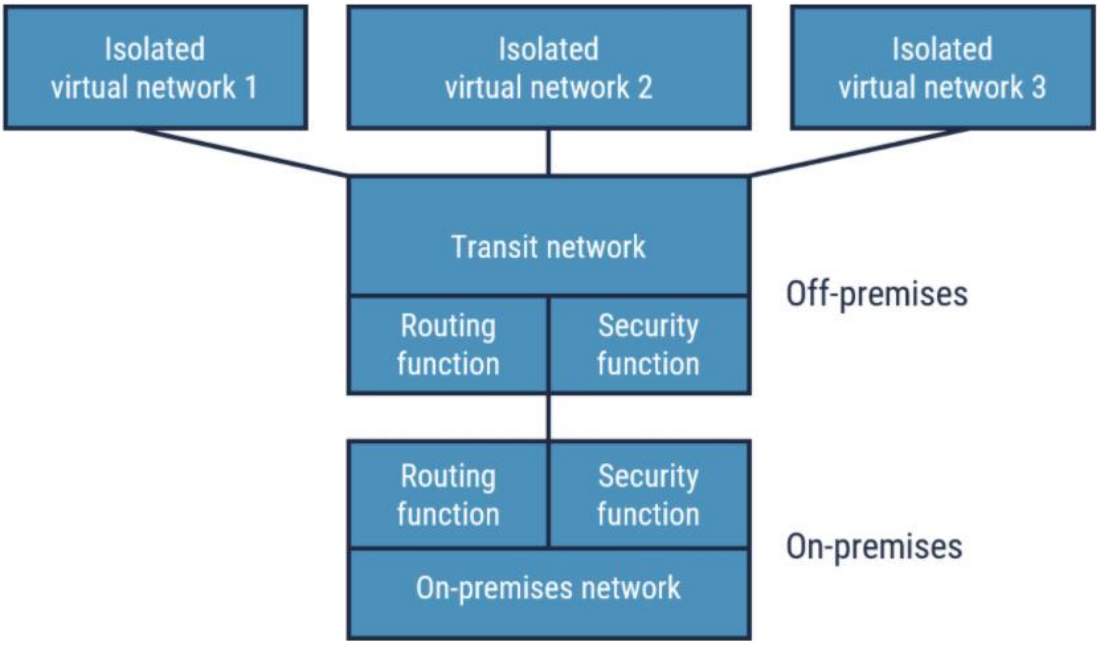

Figure 12: Bastion or Transit Network implementation [93]

Compute security involves (image and patch) management and security configuration of servers, processors, hypervisors, VMs, containers, and other serverless compute.

Defense 3: The focus of this defense revolves around data security. The guide addresses the responsibilities such as access control, encryption, key management, monitoring, lifecycle management, data migration, data remanence, data replication, data at rest, data in transit, and management plane security to ensure data is secured in the cloud.

Defense 4: This defense involves identity and access management (IAM), and it focuses on federation, strong authentication, use of standards and protocols such as LDAP, and access management such as role-based access control (RBAC).

Defense 5: Focuses on application security involving the development of secure apps, source code analysis, vulnerability testing, secure deployment, runtime vulnerability management, and threat protection.

Defense 6: This defense focuses on monitoring and incident response. The guide provides a set of recommendations to organizations so to ensure that they implement policies, procedures, and technology in order to prepare, detect, respond, mitigate, recover and learn from an incident.

Defense 7: The focus of this defense is on endpoint security, CPSs are supposed to provide capabilities to consumers to manage their endpoints using mobile device management (MDM).

5.4 DISTRIBUTED CLOUD

Distributed cloud brings cloud services geographically close to the end-users so to reduce latency and improve QoS. Distributed cloud is a public cloud computing service that helps to extend the cloud provider stack to where a customer needs it and can be deployed on-premise in multiple different locations. It provides a centralized management platform for these distributed micro-cloud satellites. [11]

The following figure shows the architecture of a distributed cloud:

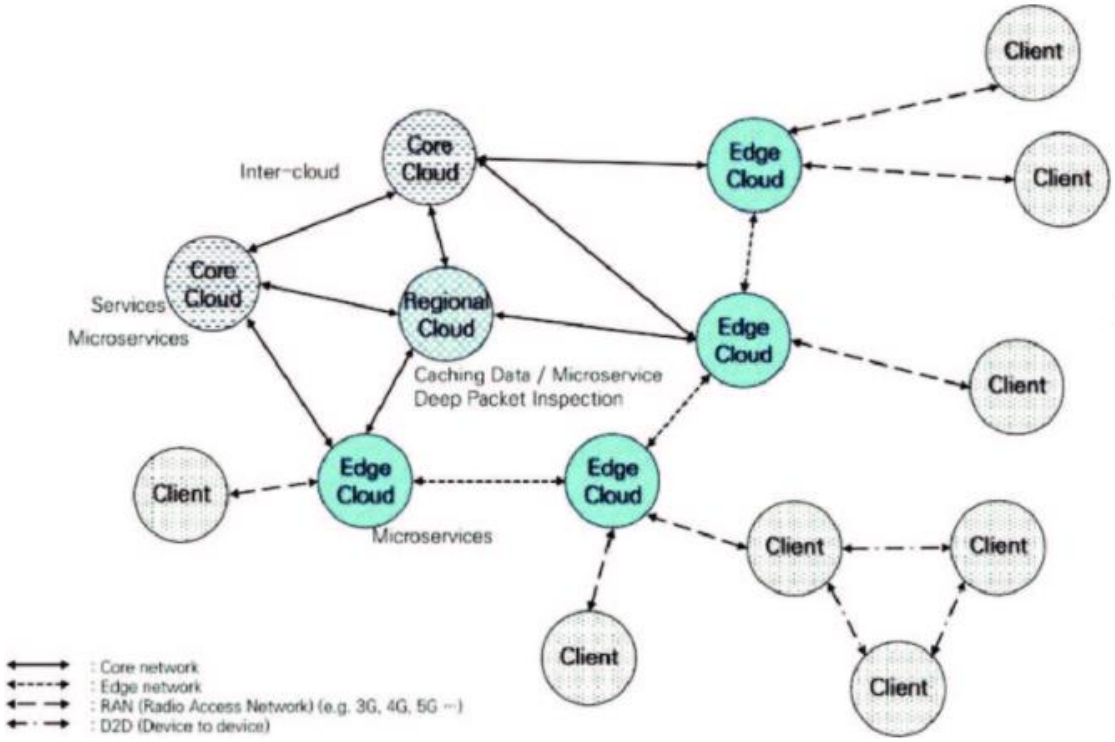

Figure 13: Distributed Cloud Concept [12]

In a study carried out by are 3 layers to the distributed cloud:

- **Core Layer**: It is the highest-level layer in the architecture of the distributed cloud, it operates like the traditional core cloud technology which involves management, provisioning of cloud services, and storage of data.

- **Regional Cloud**: It is a middle layer inside a distributed cloud architecture. It has capabilities of the proxy, self-deployment of cloud services, and caching data. Examples are Proxy Mobile IPv6 and content delivery network.

- **Edge cloud**: Final layer of distributed cloud architecture that is present close to the end-user. It supports microservices and provides real-time services. [12]

5.4.1 EDGE COMPUTING

The most important use of distributed computation is for edge computing as the demand for the Internet of Things (IoT), Artificial Intelligence (AI), and telecommunication (telco) is on the rise. As these applications require real-time processing of a huge amount of data, distributed computing is an ideal solution to quickly process the data and reduce latency. In edge computing, the application workloads are placed physically closed to where the data is generated.

Examples of edge computing include:

- Regions where the integration of smartphones or barcode scanners is high.
- Regions where IoT devices such as cameras and sensors are generating massive amounts of data.

In layman's term, edge computing can be defined as follows:

"Bring the math to the data"

The purpose of edge computing is to put the computation near where data is created and move the prevent the data from going to the centralized cloud. [11]

Use Cases for Edge Computing:

5.4.1.1 DATA COLLECTION AND ANALYTICS

It can help prevent sending data collected from IoT devices to the centralized data center for analytics which can be counterproductive due to added latency and excessive bandwidth utilization. A huge amount of data can be generated at the edge and taking the analytics closer to the source can help minimize costs of transportation by only sending the data that is necessary for a condensed form. [13]

5.4.1.2 SECURITY FOR EDGE DEVICES

Edge devices such as IoT sensors are expanding in number so is the threat of a cyber-attack against them. Edge computing can mitigate these kinds of attacks by bringing the security elements closer and enabling better performance of security applications to respond quickly to these breaches. [13]

5.4.1.3 COMPLIANCE REQUIREMENTS

Compliance requirements that edge computing infrastructure covers include geofencing, data sovereignty, and copyright enforcement. It helps with restricting access to data based on geographical and pollical boundaries. [13]

5.4.1.4 NETWORK FUNCTION VIRTUALIZATION (NFV)

Telco operators can maximize efficiency and lower costs by running virtual network functions for service delivery on top of edge computing infrastructure. [13]

5.4.1.5 REAL-TIME APPLICATIONS

Applications that require real-time computation such as AR/VR, smart cars, tactile internet Industry 4.0, telemedicine, and smart cities are time-sensitive applications. These applications cannot tolerate any latency or jitter. These automated applications require high availability and edge computing infrastructure has become a necessity. The figure below displays applications that may require edge computing in the future:

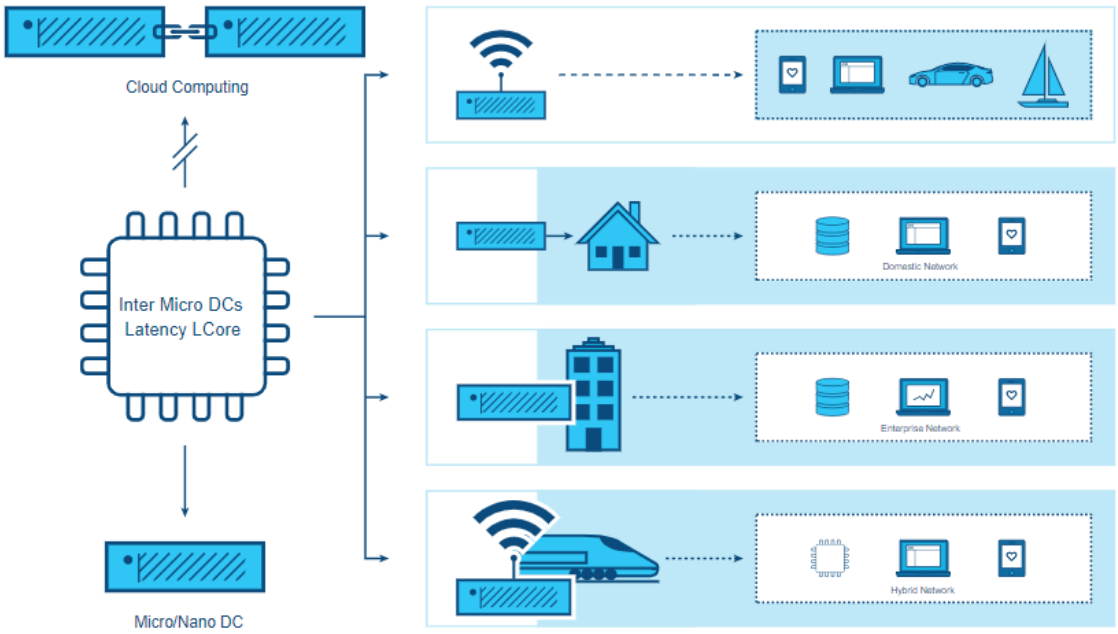

Figure 14: Applications of Edge Computing [13]

5.4.1.6 SELF-CONTAINED AND AUTONOMOUS SITE OPERATIONS

Provides semi-autonomous functionality to sites that have poor connectivity in remote areas such as oil rigs, mines, and transportation routes (for ships, buses, or even planes). It also allows the point of sales (POS) and offices to run even when the connectivity is down. [13]

5.4.1.7 PRIVACY

Edge computing can help medical applications anonymize personally identifiable data (PII) and personally health identifiable (PHI) data. [13]

5.5 CONTAINERIZATION AND VIRTUALIZATION

This section compares to key technologies that are virtual machines and containers used for running applications.

5.5.1 ARCHITECTURE FOR CONTAINERS AND VMS

5.5.1.1 CONTAINERS

According to Microsoft, container architecture comprises of application and some lightweight APIs and services that run on top of the OS kernel which runs on top of the physical hardware shown in the following figure: [14]

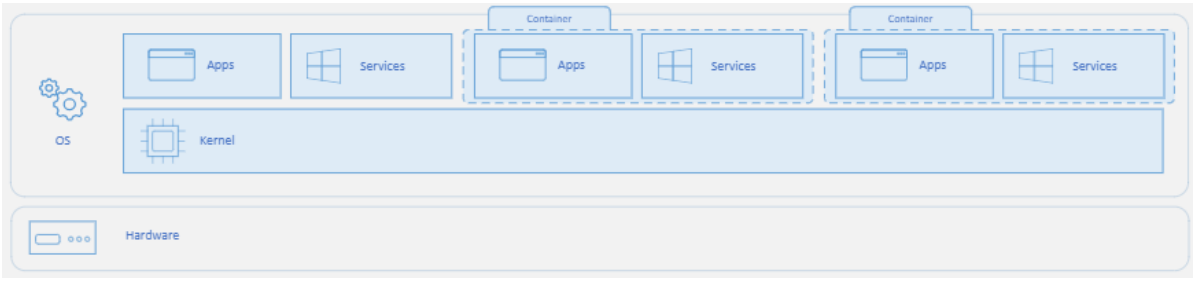

Figure 15: Container Architecture [14]

5.5.1.2 VIRTUAL MACHINES

The Virtual Machines (VMs) run an entire Operating System (OS) and separate kernel for each application and service on top of the physical hardware. [14]

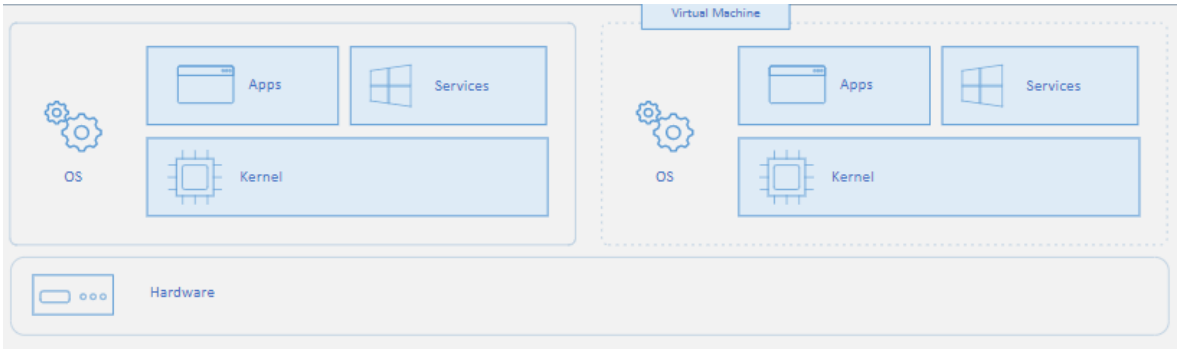

Figure 16: VM Architecture [14]

5.5.2 DEEP DIVE INTO CONTAINERIZATION

The name container is derived from shipping companies as it helps to physically isolate shipments from multiple different cargos. So a purpose of a container is to bundle up a software package that contains all the related configurations, dependencies, and libraries that are required to operate an application. This helps software developers and DevOps engineers to quickly and easily deploy apps seamlessly in multiple environments.

Containers virtualize the operating system to perceive to the application that it is present in an ecosystem inside the OS which is having a CPU, RAM, storage, and a network to connect to. Containers make use of a shared operating system which makes them fast, efficient, and lightweight when compared to VMs because they do not require to boot an entire OS, load libraries separately, and the overall overhead is significantly reduced when we account for maintenance, patching, and updates. There are some problems that containers face since that the containers are tied to the operating system they are portable to that same type of OS. That means if a container is deployed in a Windows OS then it can only be ported to another Windows OS. [15] [88]

One of the other issues that containers have when compared to VMs is poor isolation. Weaker isolation leads to less secure infrastructure.
The following figure from Palo Alto displays the weak isolation issue:

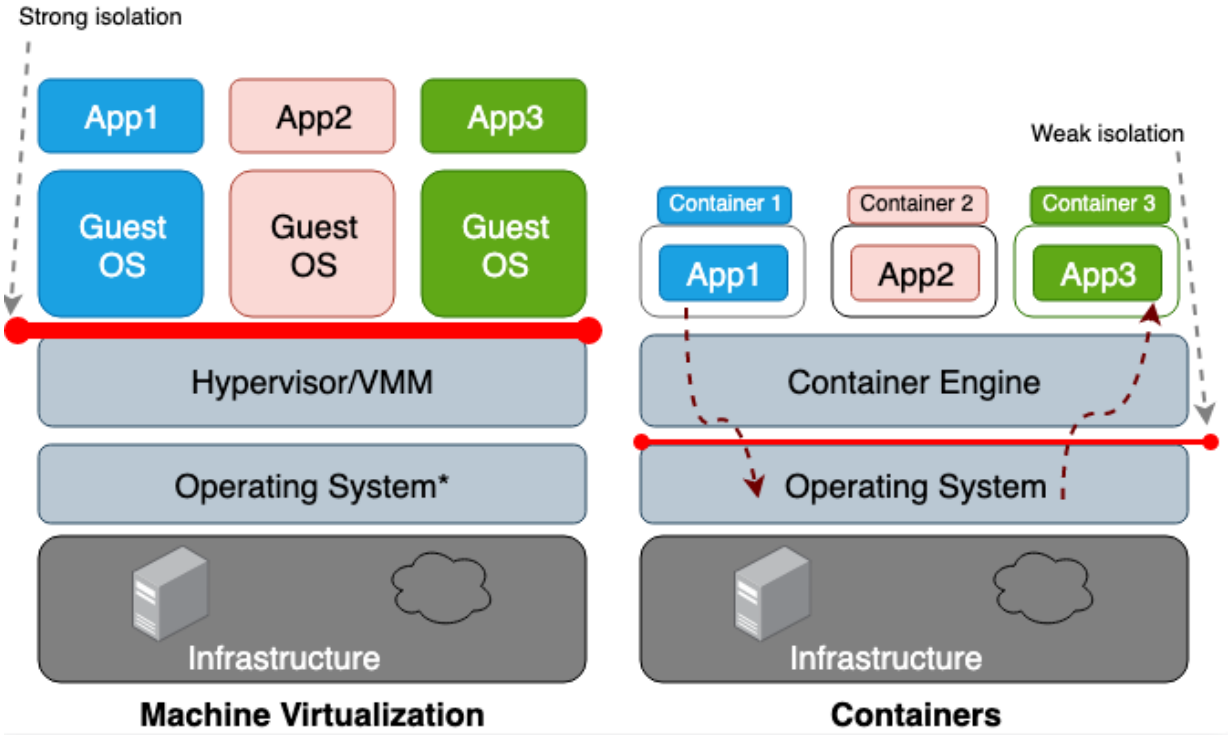

Figure 17: Hypervisor provides the stronger isolation between each guest OS. In containers, the host operating system provides the weak isolation between each container so it is less secure [16]

The key difference between the VM and container design is that VM has hardware-level virtualization whereas the containers make use of OS-level virtualization. OS-level virtualization results in sharing the OS resources which makes containers less secure. There are multiple CVE that lead to container breakouts such as CVE-2014-3519, CVE-2016-5195, CVE-2016-9962, CVE-2017-5123, and CVE-2019-5736. [16]

5.5.3 CONTAINER ORCHESTRATION: KUBERNETES

Developers are increasing the use of container micro-services day by day for modern app development. As the number of applications grows, more containers are introduced to the server and their management and operations become difficult. To reduce the complexity of management of these containers an opensource API called Kubernetes is used by developers and IT staff.

Key benefits of Kubernetes:

- It provides service discovery by exposing containers to DNS and provides load balancing between containers
- It allows storage orchestration in which storage of our choice can be mounted.
- It supports automated rollouts and rollbacks, this can involve changing the state of a container, creating one, or even removing the existing one using automation. (High Scalability)
- It has an automatic bin packing feature that allows us to allocate CPU and Ram to a cluster of nodes.
- It is designed for cloud-native applications
- It can self-heal a failed container by replacing it which provides high availability and no downtime (Disaster Recovery) [17]

Components of Kubernetes:

- Node: There are 2 types of Nodes present in Kubernetes:
  - Master Node: Manages schedules and resources of the containers or Pods that are running inside Worker Nodes. It also contains etcd which is a database of all the Worker Nodes resources and state changes.
  - Worker Node: Contains multiple Pods or containers.
- Pods:
  - The smallest unit of Kubernetes service is a Pod.
  - It runs as an abstraction layer over containers. It usually runs a single application in a container inside a pod.
  - Each Pod has its own IP address.
  - They can contain an application or a database.
- Service: If a Pod fails it is replaced by a new Pod with a new IP address and so Kubernetes makes use of the service as it needs to maintain the connection between applications or databases. Each service is attached to each pod and even if the Pod fails, the service stays and it has its own permanent IP address. The lifecycle of a Pod is independent of the lifecycle of a service attached to it. Service also acts as a balancer for replicas of the Pods. There are 2 types of services for Pods:
    - External Service: For applications that are accessible for example through a browser. This is usually managed through an API object

called Ingress. Ingress provides load balancing, SSL termination, and name-based virtual hosting for an application inside a Pod.
    - Internal Service: For a database that is not exposed externally requires the use of internal service.
- ConfigMap: An API that manages external configuration related to a pod such as URL, environmental variables, and command-line argument in a volume. It with changing external configurations without affecting the Pod and allows portability of the applications. It is not meant to store sensitive data such as passwords.
- Secret: Stores and manages sensitive data such as passwords, keys, certificates, and tokens.
- Volumes: Allows Pods to create persistent storage inside or outside a Kubernetes cluster. It can be on an on-premise hard drive or in a public cloud. Kubernetes does manage the persistence of data and there can a loss of data when a container fails.
- StatefulSet: It is a workload API that deals with the management deployment and scaling of stateful applications such as Databases. That means it can replicate reads and writes on Database replica to maintain consistency across other Database replicas.
- Container runtime: Process that allows containers to run inside a Node.
- Kubelets: It acts as an interface between the container inside a pod and the Node. It is responsible for running a Pod inside a Node.

5.6 MULTI-CLOUD

Multi-cloud is a broad term and can involve a mix of multiple public cloud service providers and private cloud. The following shows the benefits and underlying issues for multi-cloud:

- **Vendor lock-in**: It is the first fundamental reason to consider a multi-cloud strategy. It makes use of multiple cloud platforms to avoid vendor lock-in and makes use of vendor diversification to improve security posture. In simple words, it can be defined as putting all your eggs in one basket.

- **Maintaining workload portability**: When using the functionality of each cloud platform it is important to consider workload portability. For example, serverless compute from AWS Lambda is not portable to other Cloud computing platforms which makes it difficult to migrate an application workload to another cloud provider such as Azure Functions and Google Cloud Functions.

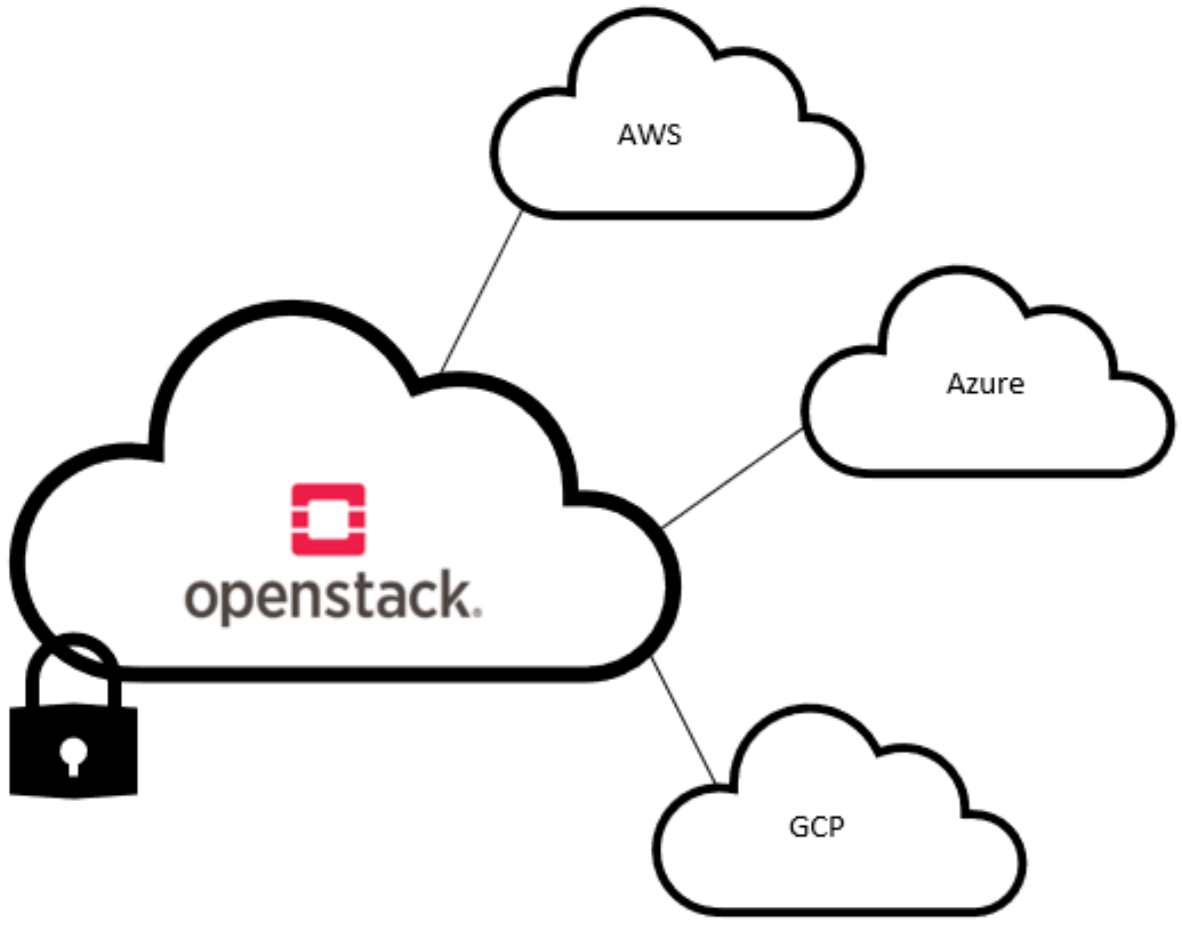

Figure 18: Multi-Cloud makes use of a combination of CSPs and private cloud

- **Reliability**: This also provides high availability. In case the primary cloud provider is experiencing downtime, the secondary cloud provider can take over.

- **Price-competitive**: This benefit is clear as the power of bargain is transferred to the customer as it removes the ability of a vendor to overcharge and allows the enterprise to have a better return on investment (ROI).

- **Optimal application environment**: It gives the customer to pick and utilize the best tools from multiple cloud providers to achieve optimal performance for an application. For example, it helps in providing multiple availability zones from different providers. The one nearest to the end-user will provide the lowest latency and better performance.

- **Regulatory compliance**: Having a diversified portfolio of public cloud vendors and the private cloud helps with the changing regulatory requirements set by Governments around the world. The data protection and sovereignty laws require businesses to keep the local data in certain geographic locations. [84]

5.7 PRIVATE AND HYBRID CLOUD PROVIDERS

5.8 OPENSTACK

5.8.1 OVERVIEW

OpenStack is an open-source cloud technology platform that allows businesses to control and manage their resources such as compute, storage, and networking hardware in the data centers. The figure below shows the modular components that are used for providing core services which are Horizon, Nova, Cinder, Neutron, Glance, Swift, and Keystone which will be discussed in detail in this report:

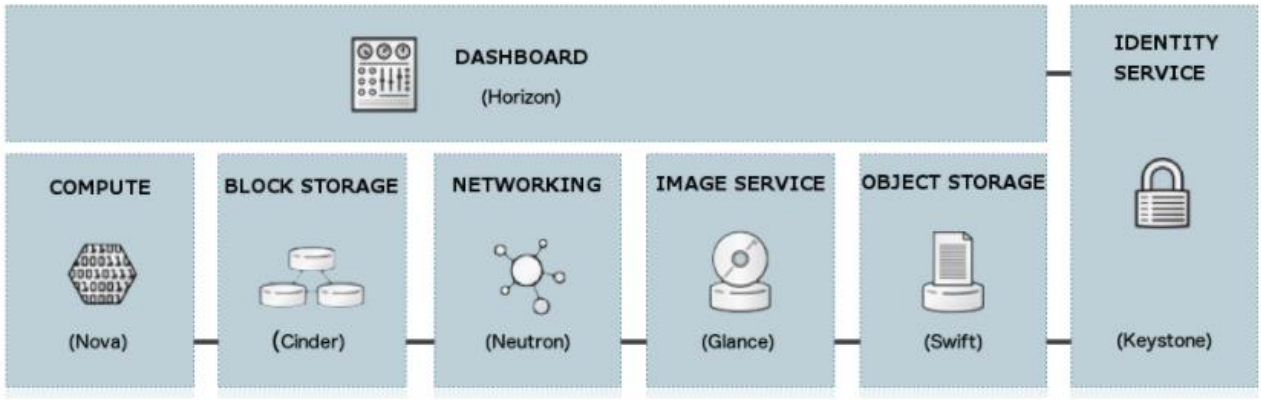

Figure 19: Services in OpenStack [18]

Other than IaaS functionality, OpenStack also provides services such as orchestration, fault management, and service management as shown in the landscape constructed by OpenStack in the following figure:

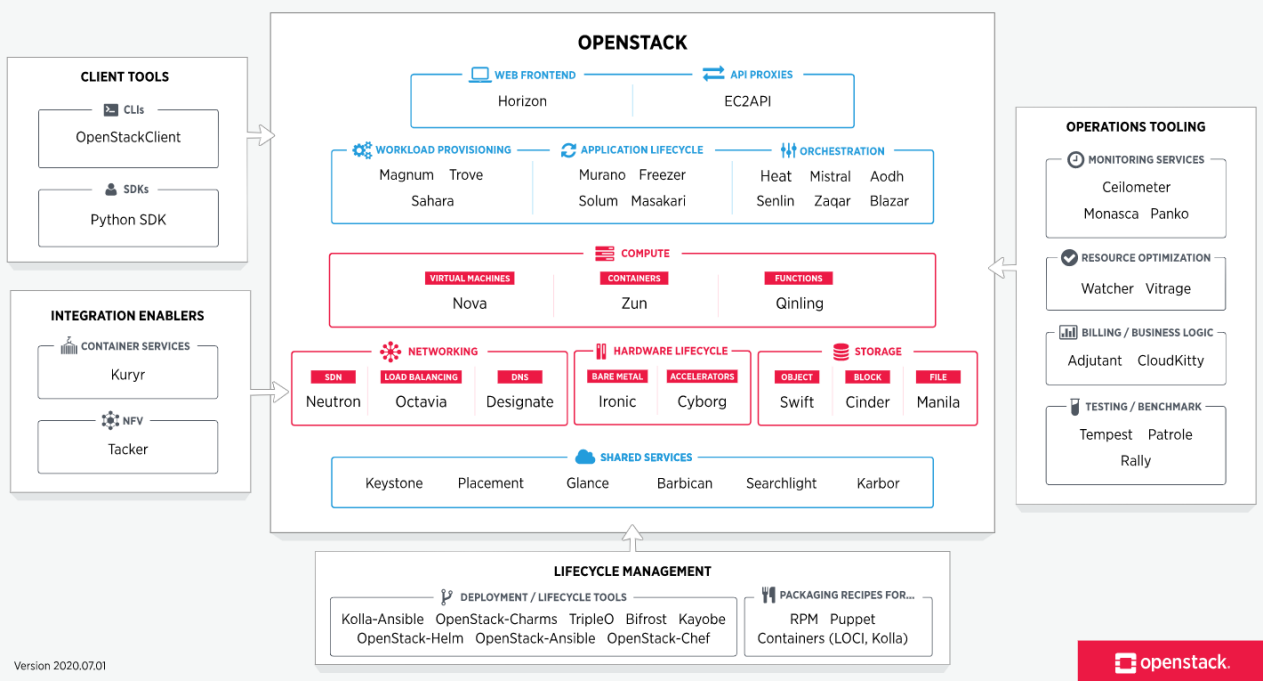

Figure 20: OpenStack Map [19]

A brief overview of the services present inside the landscape are listed in the tables below:

Table 1: Services in OpenStack [20]

Projects	**Compute**
Nova	Provides scalable, on-demand, and self-service access to computing resources.
Zun	This project provides API for running and managing containers which are supported by multiple different technologies
Qinling	Provides capabilities that support serverless compute.

Projects	**Hardware Lifecycle**
Ironic	This a Bare Metal Provisioning service, it provides high scalability, on-demand, self-service access to: • Computer Resources • Virtual Machines • Containers
Cyborg	Provides lifecycle management acceleration by making use of general-purpose management framework for accelerators such as Field Programmable Gate Arrays (FPGAs) and GPUs.

Projects	Storage
Swift	Swift is an object storage service, it provides high availability, distribution, and consistency to the object or blob storage. It helps organizations save data in a very efficient and cost-effective manner. It is also durable and scalable. Swift is suitable for storing unstructured data.
Cinder	This a Block storage service in OpenStack. It makes use of virtualization of block storage devices and makes use of API to request or use the resources.
Manila	Provides access to file-sharing or distrusted file system.

Projects	Networking
Neutron	It is a software-defined networking (SDN) service that provides networking-as-a-service (NaaS) using virtualization.
Octavia	Open source load balancing service for Virtual machines, containers, and bare metal servers deployed in OpenStack. It is also horizontally scalable.
Designate	It is a DNS-as-a-service designed for OpenStack

Projects	Shared Services
Keystone	Keystone is an identity service in OpenStack used for API client authentication and service discovery. It provides built-in support for LDAP, OAuth, OpenID Connect, SAML, and SQL
Placement	It is an HTTP API that helps in tracking resource usages for the cloud.
Glance	Image service used for discovery, registration, and retrieval of VM images.
Barbican	It is a Key Manager service that securely stores, provisions, and manages secret data such as passwords, encryption keys, and certificates.
Karbor	Protects application data and metadata against loss or damage. It is not an application security or (data loss prevention) DLP service.

Projects	Orchestration
Heat	Orchestration tool that helps manage infrastructure resources for cloud applications using templates.

Senlin	A service for the creation and operation of clusters of similar objects in the OpenStack cloud.
Mistral	Service that manages the workflow of tasks or processes in a distributed business environment. It carries out state management, orderly execution, parallelism, synchronization, and high availability.
Zaqar	It is a service used for cloud messaging for mobile and web development. It makes use of RESTful API and it can be used to send messages between different components of SaaS and Apps.
Blazar	A service that allocates reservation of resources for a specified amount of time.
Aodh	Service that triggers alerts and alarms based on the conditions and rules set for events.

Projects	**Workload Provisioning**
Magnum	Magnum provisions container orchestration engines for example Docker Swarm, Kubernetes. And Apache Mesos. It makes use of the Heat service for orchestration of Operating System images that have orchestration engines inside them and runs the image in VM or on bare metal inside a cluster.

Sahara	A project which allows users to run or provision Big Data Framework Processing such as Hadoop, Spark, and Storm,
Trove	It acts as a database-as-a-service and helps with deploying relational and non-relational database engines on OpenStack.

Projects	**Application Lifecycle**
Masakari	It is called the instance high availability service that can recover KVM-based virtual machines during a failure, process down, provisioning process down, and nova-compute host failure.
Murano	It is an application catalog where developers and administrators can publish cloud-ready applications. It allows cloud application users to easily use these applications. Murano makes use of Heat service for the orchestration of infrastructure resources for these applications.
Solum	Helps with application development process using software development lifecycle automation

Freezer	A service that provides distributed data backup, restore, and as well as disaster recovery. It supports block-based backups, incremental backup, point-in-time actions, and backup sync over several nodes.
API Proxies	
EC2API	A project that helps in providing EC2 API proxy to Nova service in OpenStack.
Web Frontend	
Horizon	It is a service that provides an interface to OpenStack services with the help of a web-based dashboard.

5.8.2 NOVA

Nova is the compute service that helps with the management and automation of virtual machine (VM) instances that run applications, provide development and testing environment, and other high-performance computing such as Big Data. It works will diverse virtualization technologies and bare metal technologies.

It is responsible for the creation, deletion, scheduling, and sizing of an instance of the virtual machine. It has the ability to scale the VM horizontally. It is a distributed application with multiple components or services which has an architecture as shown in the following figure:

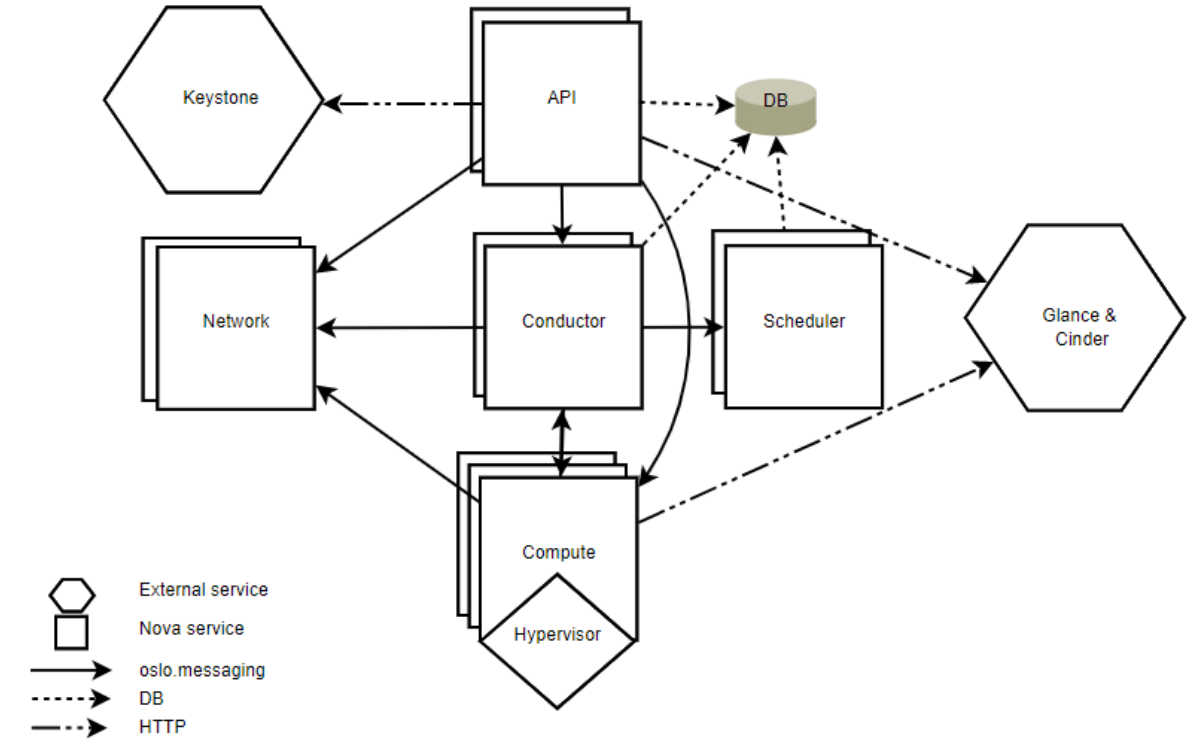

Figure 21: Nova Architecture [21]

- **Nova-API**: The purpose of this component is to obtain HTTP requests, conversion of commands, and communicate with other components in the architecture using Oslo.messaging queue or via HTTP. It is based on REST (Representational State Transfer) API.

- **DB**: It is a database service based on SQL for storing data.

- **Nova-Scheduler**: It is responsible for deciding and allocating instances for each host.

- **Nova-Network**: Deals with IP forwarding, VLANs, bridges, layer 2 traffic, and this service can also be handed over Neutron which is another core component of OpenStack.

- **Nova-Compute**: Is responsible for managing communication that takes place between the hypervisor and the virtual machines.

- **Nova-Conductor**: Manages requests that require coordination such as build and resizing and performs as a database proxy. [21]

The following shows how the Nova-network is replaced by Neutron core service:

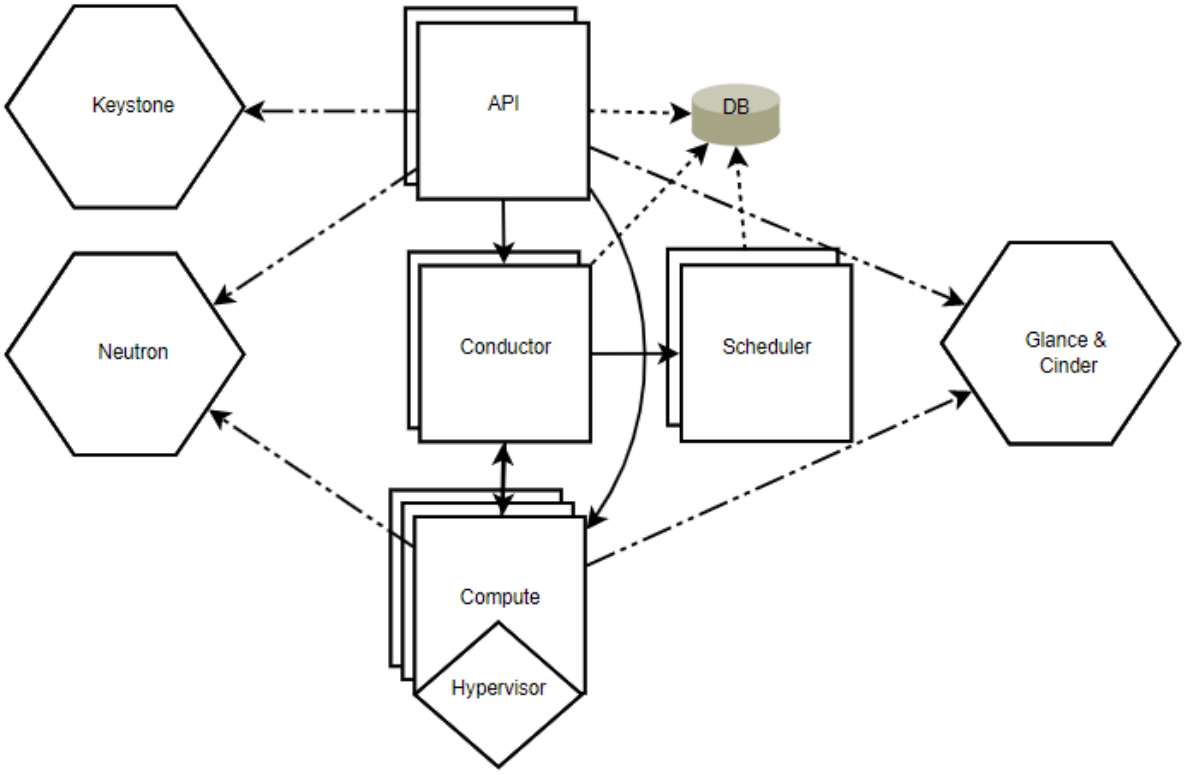

Figure 22: Nova Architecture Using Neutron [21]

Since Neutron is the core network service in OpenStack, it is preferable to use it instead of nova-network.

Certain requirements need to be fulfilled to launch an instance:

- **Image**: Pre-build Images from Glance so to quickly boot
- **Network**: To determine where to attach an instance in a network
- **Flavor**: To determine the size of the instance (e.g. vRAM, vCPU, and vStorage).

  The following image shows the terminal with a list of default flavors available in OpenStack

```
[root@localhost ~(keystone_admin)]# openstack flavor list
+----+-----------+-------+------+-----------+-------+-----------+
| ID | Name      |   RAM | Disk | Ephemeral | VCPUs | Is Public |
+----+-----------+-------+------+-----------+-------+-----------+
| 1  | m1.tiny   |   512 |    1 |         0 |     1 | True      |
| 2  | m1.small  |  2048 |   20 |         0 |     1 | True      |
| 3  | m1.medium |  4096 |   40 |         0 |     2 | True      |
| 4  | m1.large  |  8192 |   80 |         0 |     4 | True      |
| 5  | m1.xlarge | 16384 |  160 |         0 |     8 | True      |
+----+-----------+-------+------+-----------+-------+-----------+
```

Data at rest feature called Ephemeral disk encryption allows VMs to encrypt data used as a temporary workspace. The vestigial remanent data may stay on the disk when a VM is unmounted resulting in a data privacy issue. [22]

5.8.3 GLANCE

This core service is used to store virtual machine disk images. This service allows the running of prebuilt images to spin up the instance to Nova. A feature is also incorporated in OpenStack which prevents unverified VM images from the Glance service to run. Glance service carries out the verification by retrieving the certificates from the key manager. It is also based on the REST API and makes use of Swift to store these images as a backend. The following figure shows the architecture that Glace utilizes:

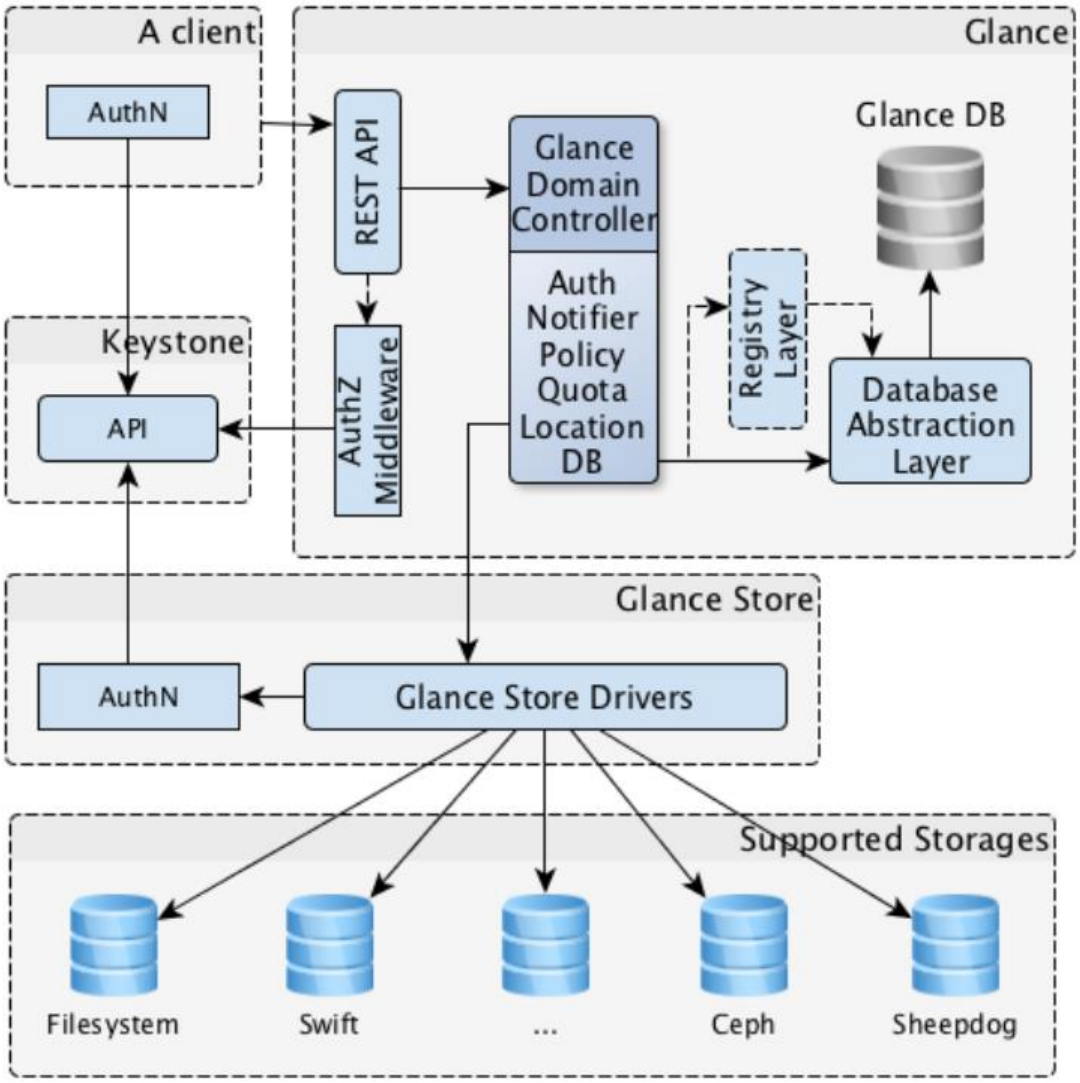

Figure 23: OpenStack Glance Architecture [23]

The following components are part of the architecture:

- **Client**: It is an application that utilizes the Glance server.
- **REST API**: Glance uses REST API to expose the functionalities
- **Database Abstraction Layer (DAL)**: It is an API the unifies the communication that takes place between DB and Glance.
- **Glance Domain Controller**: Manages the operation of the internal server such as authorization, notifications, policies, and database connections.
- **Glance Store**: Manages the integration that takes place between different databases and Glance.
- **Registry Layer**: An extra optional layer that provides added secure communication between Database Abstraction Layer and the Glance Domain Controller. [23]

5.8.4 NEUTRON

A core service that provides virtual network connectivity using routers, subnets, ports, and abstractions. The virtual network topologies also include network security groups, firewalls, load balancers, and virtual private networks. OpenStack also calls it 'networking as a service'. It supports network models such as Flat, VLAN, VXLAN, GRE, static IP, and DHCP. Plus, it can also make use of SDN platforms. [24] [25]

The networking architecture of OpenStack is said to be a standalone service that operates multiple processes over several nodes. The main process is the Neutron-Server. The following figure shows the networking architecture of OpenStack:

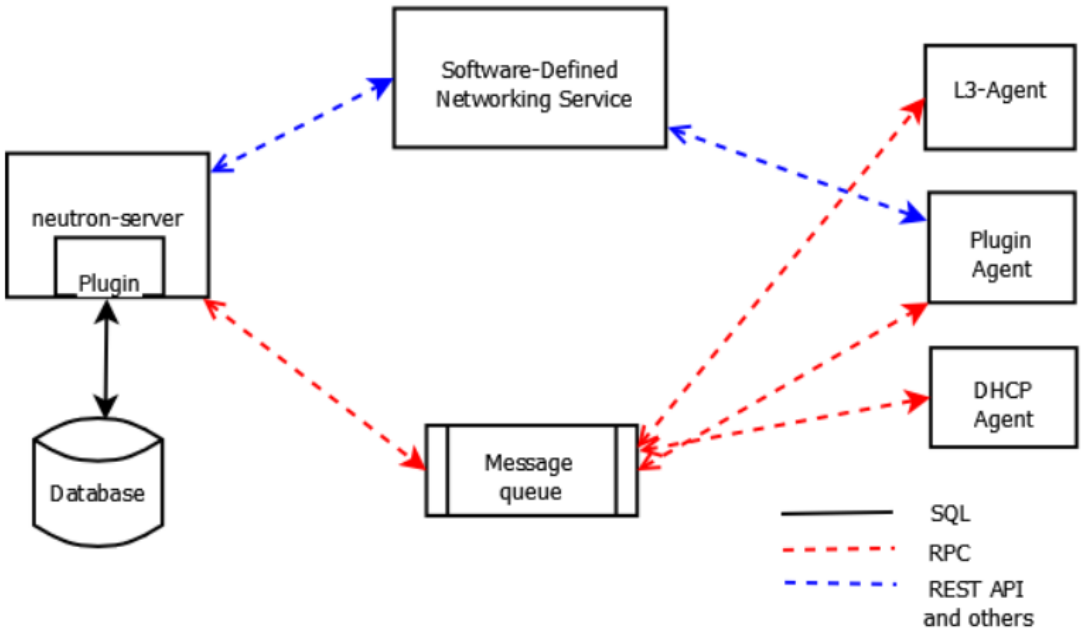

Figure 24: Neutron Architecture [25]

The components of this architecture are explained as follows:

- **Neutron-server + plugin**: Runs on the network node and provides service for Networking API and extensions. It is responsible for IP addressing and network model. A neutron-plugin allows access to the persistent database using AMQP (Advanced Message Queuing Protocol).

- **Plugin agent (Neutron-agent):** Manages configuration of switches (virtual switches) on each compute node. The plugin in use decides which agents to run. It requires message queue access.

- **Neutron-DHCP-Agent**: Responsible for DHCP configuration and services. It requires access to the message queue.

- **Neutron-L3-Agent**: Provides virtual machines with access to the external network via NAT (network address translation) and this service requires access to the messaging queue.

- **SDN (Software Defined Network) server/service**: Provides tenant network with additional services and interacts with neutron-server, neutron-plugin, and plugin-agents using REST API.

The OpenStack networking service has a placement on physical servers as shown in the following figure:

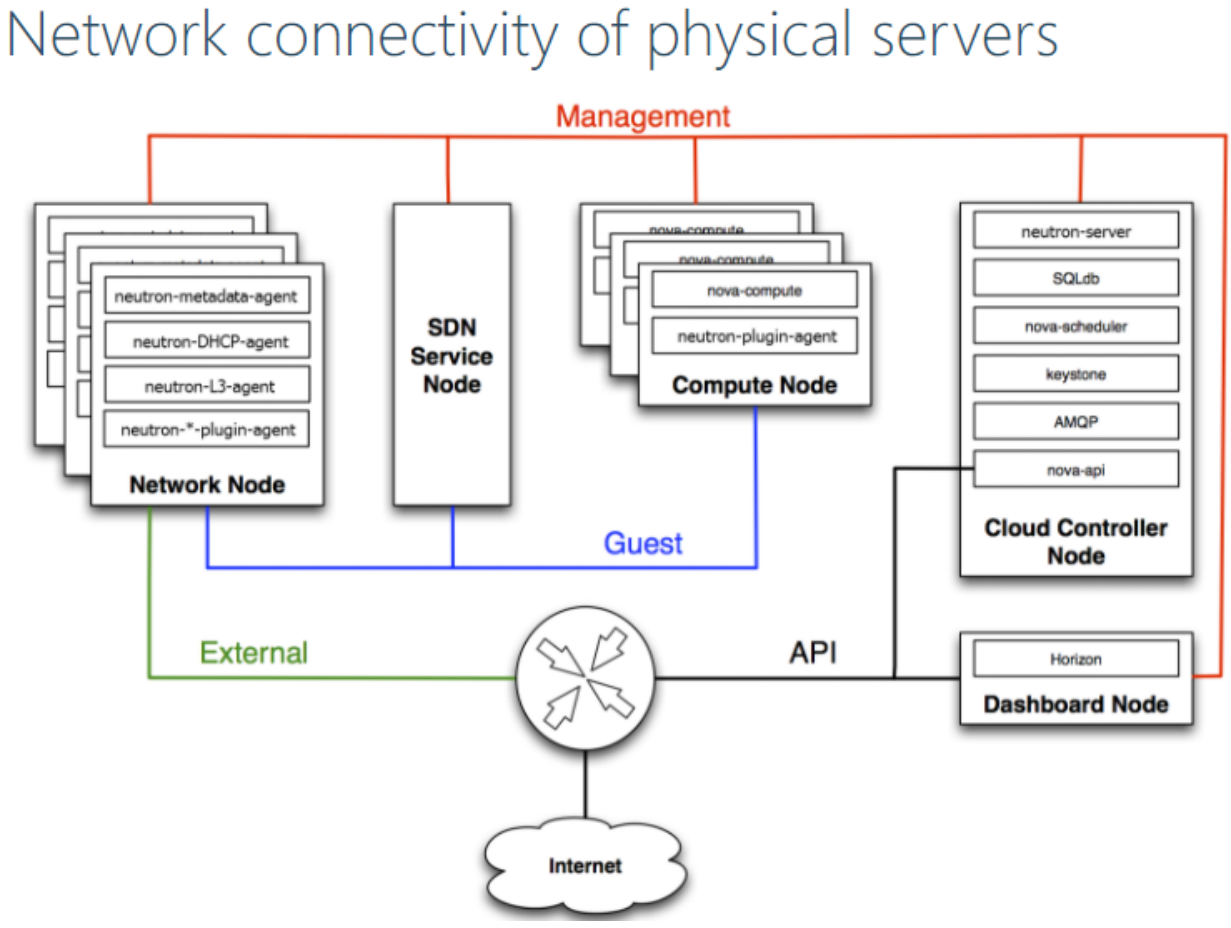

Figure 25: Neutron Connectivity to Physical Servers [25]

As seen in the figure, there are four different networks:

- **Management Network:** Provides internal communication between the components of OpenStack.

- **Guest Network**: Provides communication between VMs inside the cloud deployment.

- **External Network**: Provides access to a public network or the internet for some of the VMs.

- **API Network:** Helps in exposing the OpenStack API and Networking API to tenants. And the IP address is made reachable to the public network. [25]

5.8.5 STORAGE TYPES

Cloud environments offer multiple different types of storages:

Ephemeral Storage: It is temporary storage that is deployed when a Nova compute service VM is started and disappears when it is terminated.

Persistent Storage: This is a permanent solution and it remains available even if the VM is terminated or is in any other state. There are 3 types of persistent storages which are as follows:

- **Object Storage**: Used for storing a large amount of unstructured data in the cloud. It can store VM images and has better support for distributed deployments across different datacenters with the help of asynchronous replication which makes the data resilient to disaster. It is not suitable for applications that demand high performance. It is mostly used for active archiving, BI (business intelligence), or data warehousing and backup or disaster recovery solutions. [26]

- **Block Storage**: Breaks the data into blocks and stores them separately wherever it is efficiently stored on SAN (Storage area network). Preferred by the developer due to low latency, high redundancy, and high efficiency in data transfers. Block storage can be used for boot volume for VMs, and databases that require high performance. These storages are also ideal for containers due to the high IOPS. [27]

- **File-Based Storage**: Storage that is usually associated with NAS (Network-attached storage). Organized as a directory tree with folders and files. Data access is dependent on a path. It is highly scalable, has simultaneous read and writes, and is accessible to multiple runtimes. This type of storage can be used for a mix of structured and unstructured data, and when multiple users are accessing a shared file at once. [27]

5.8.6 CINDER

Cinder is an example of block-level storage which is persistent storage used for Nova compute instance. It can be attached as secondary storage or as root storage to boot up instances. It is capable of managing snapshots of instances. Supports SSDs and old magnetic drives. It is usually attached to Nova compute nodes using iSCSI. It supports Volume lifecycle management such as create, delete, extend, attach, and detach volumes. It is suitable for database storage. It also makes use of volume encryption functionality for data at rest via Castellan (a generic Key Manager) and key creation can be requested by Cinder when a user decides to encrypt a volume.

The following figure shows how Cinder architecture operates:

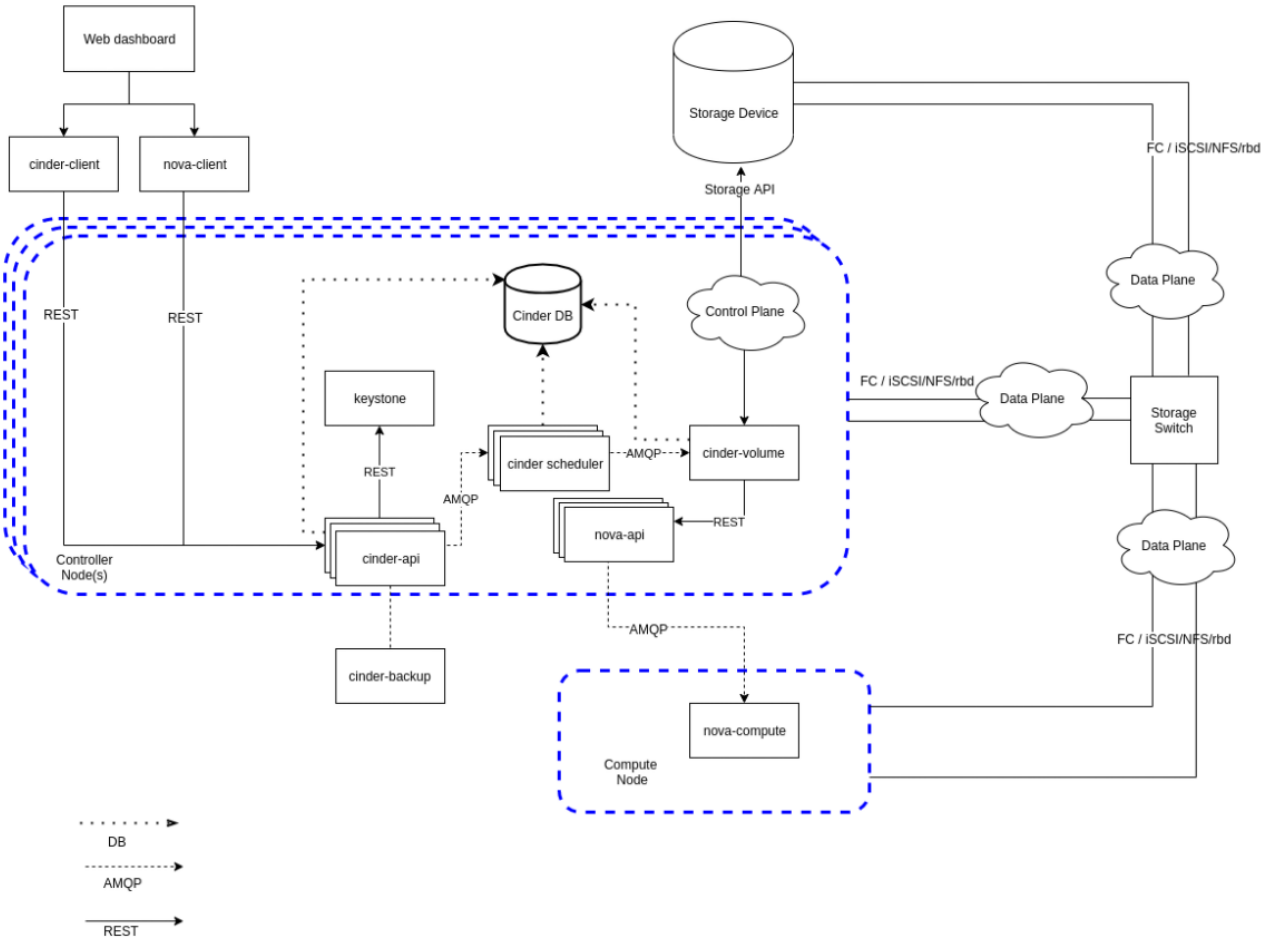

Figure 26: Cinder System Architecture

- **Cinder database**: It is a SQL database storage used by components present in Cinder architecture.

- **Web Dashboard**: It is an external component that communicates with Cinder using an API

- **Cinder API**: Responsible for receiving HTTP requests, converting commands, and communicating with other components in the Cinder architecture using queue or HTTP.

- **Auth Manager**: It manages users, projects, and roles and can utilize DB or LDAP.

- **Cinder Scheduler**: Allocates the volume based on each host.

- **Cinder Volume**: Responsible for management of attachable block devices

- **Cinder Backup**: Responsible for management of backup of block storage devices. [28]

5.8.7 SWIFT

It is a simple object storage service that works on REST API. It is highly scalable, easily distributed, and replicated. It supports high concurrency when the storage is accessed by multiple users. It supports eventual consistency which is the reason it has great scalability. Swift stores 3 replicas by default which can be changed to create a balance between cost and durability. The storage also stores the MD5 checksum for each object to prevent files from getting corrupted. [29]

Data Placement in Swift is done as follows:

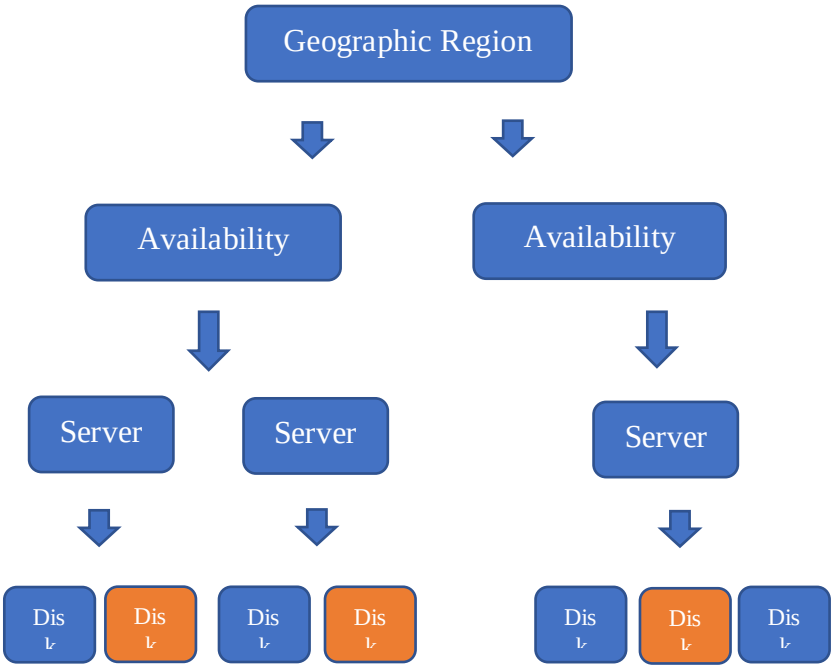

Figure 27: Data Placement in Swift

Swift makes 3 replicas by default and places them in separate zones whenever possible. Wider the deployment the more spread-out the data placement will occur as shown in the above figure. The orange shade is given to the data placement.

Swift object storage architecture is quite modular as it is designed for high availability, durability, and concurrency. The architecture contains the following components:

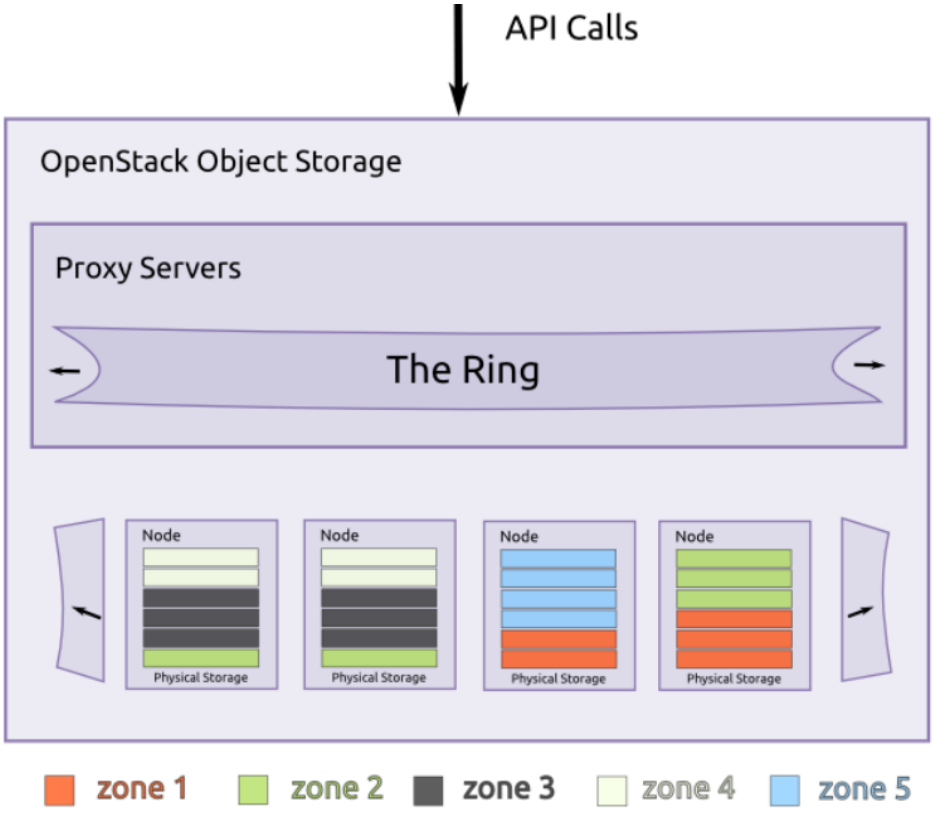

Figure 28: Object Storage building blocks [29]

- **Proxy servers**: This component is responsible for incoming API requests. Examples include file upload, modifications to metadata, or the creation of containers. Can also improve performance by utilization of cache.

- **Zones**: are responsible for the isolation of data. If there is a failure in one zone, the data in the other zone is unaffected as the data is replicated across different zones. [29]

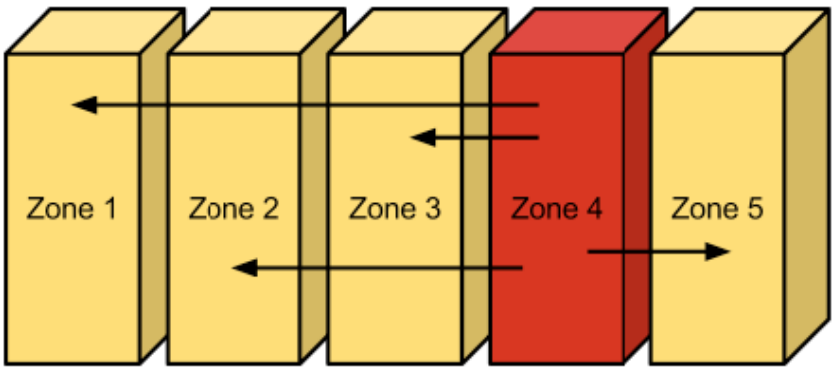

Figure 29: Zones in Object Storage [29]

- **Rings**: Carry out logical mapping of the names of data to disks on specified locations. There are used by a proxy server to map data. The following figure shows the workflow of Ring inside Object storage:

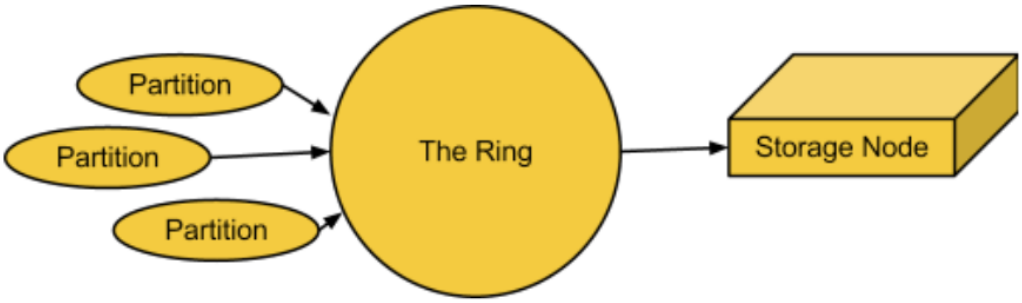

Figure 30: Ring in an Object Storage [29]

- **Accounts and Containers**: In order to track the object data locations, the account database references a list of containers in that account and the container database references a list of objects (data itself).

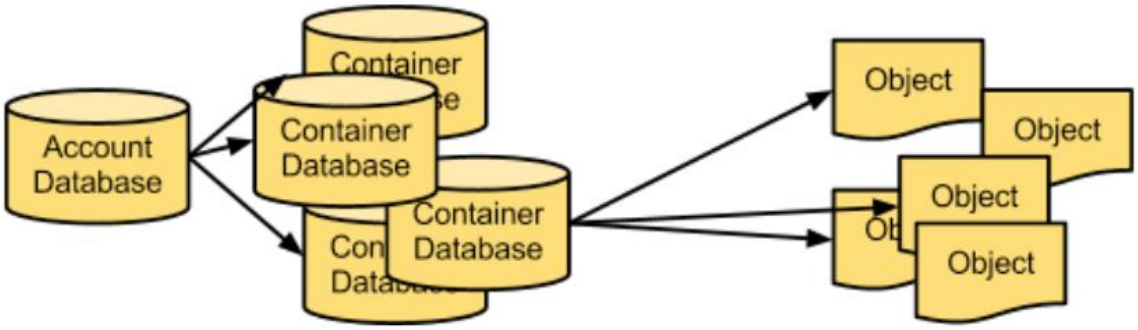

Figure 31: Accounts and Containers [29]

- **Partition**: It stores account database, container database, and objects. And help in the management of the location where the data reside. [29]

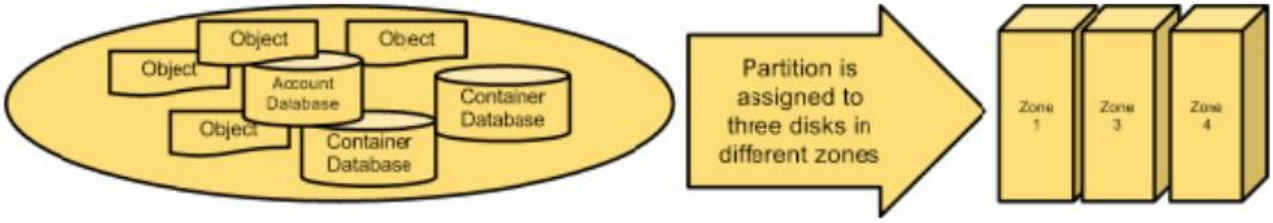

Figure 32: Partitions [29]

5.8.7.1 REPLICATORS

Replicators are responsible for examining each partition continuously. It compares each replica to the replica in other zones to monitor if there is any difference.

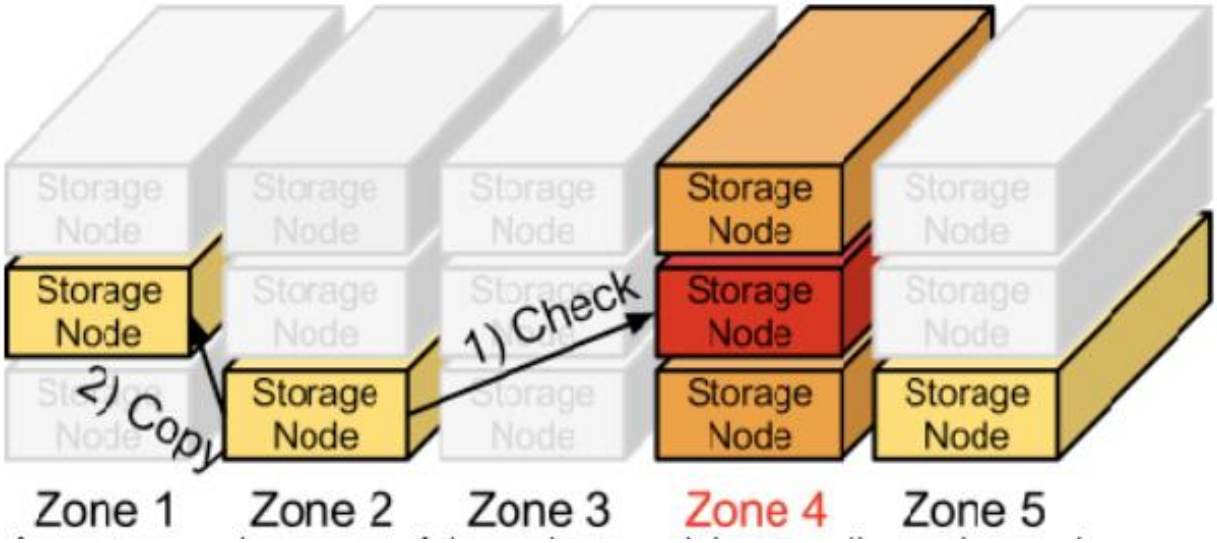

Figure 33: Replication [29]

Replicator makes use of Hashes for comparison and it stores a hash file inside each partition for each directory. If it detects any differences in hashes between partitions, it replicates those directories over. If any of the zones go down due to failure, nodes containing replica nodes and copies of data will take over. [29]

5.8.8 KEYSTONE

Keystone is a core service in OpenStack that provides API client authentication, service discovery, and authorization for distributed multi-tenant by utilizing Identity API. It also supports OAuth, SAML, LDAP, and OpenID Connect. This service makes use of tokens for authentication of requests.

The following figure shows the architecture of Keystone:

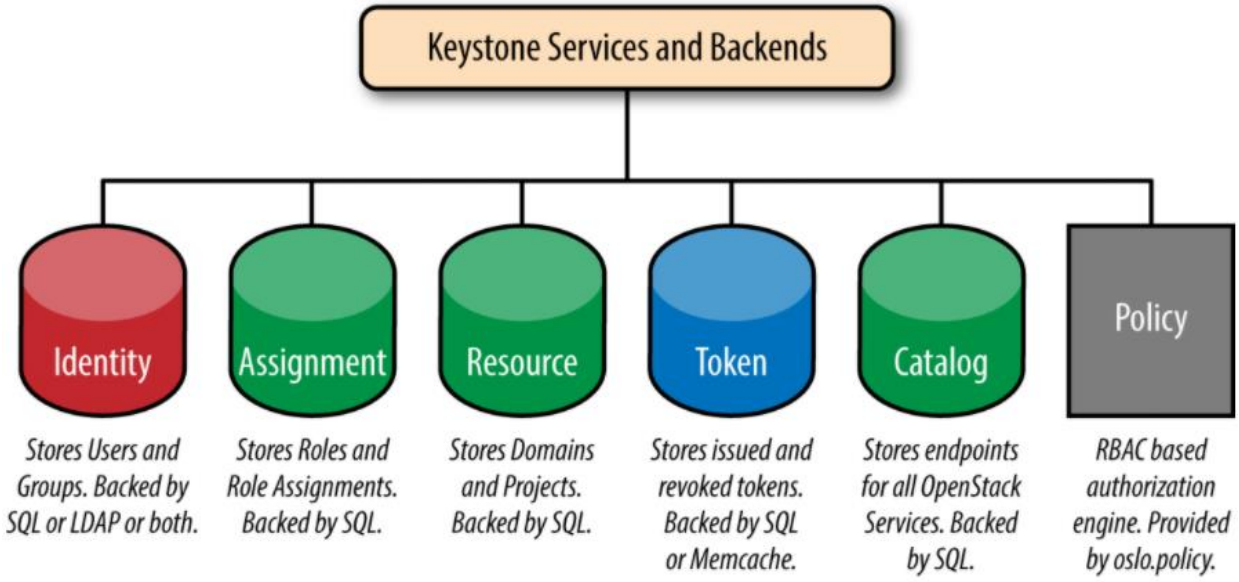

Figure 34: Services and backends inside Keystone [30]

Identity: This is a service responsible for auth credential validation and contains data about *users* and *groups*. This service runs as a frontend and may use backends such as SQL and LDAP.

Resource: A service that provides information about *domains* and *projects*. The domains and projects are explained as follows:

- **Projects**: they are represented as the base unit for ownership in OpenStack. And each project is inside a specific domain which makes the name of each project unique inside a domain but not globally.

- **Domains**: It is a high-level container with a namespace that contains *projects*, *users*, and *groups*. They can be used to entrust the management of OpenStack resources.

Assignment: It is a service responsible for providing information about *roles* and *role assignments*. A *role* is a level of authorization and *role assignment* refers to 3-tuple containing role, resource, and identity.

Token: A service that provides validation and management of tokens that are used for authentication of requests after the credentials of a user is verified.

Catalog: It is a service that contains URLs and endpoint information for different cloud services. It helps in routing the requests to create VMs and store data. [31]

5.8.9 HIGH AVAILABILITY SUPPORT

High Availability is also supported by grouping computing nodes into different availability zones, regions, and host aggregates. The following figure shows how the nodes can be separated into different regions and availability zones. [32]

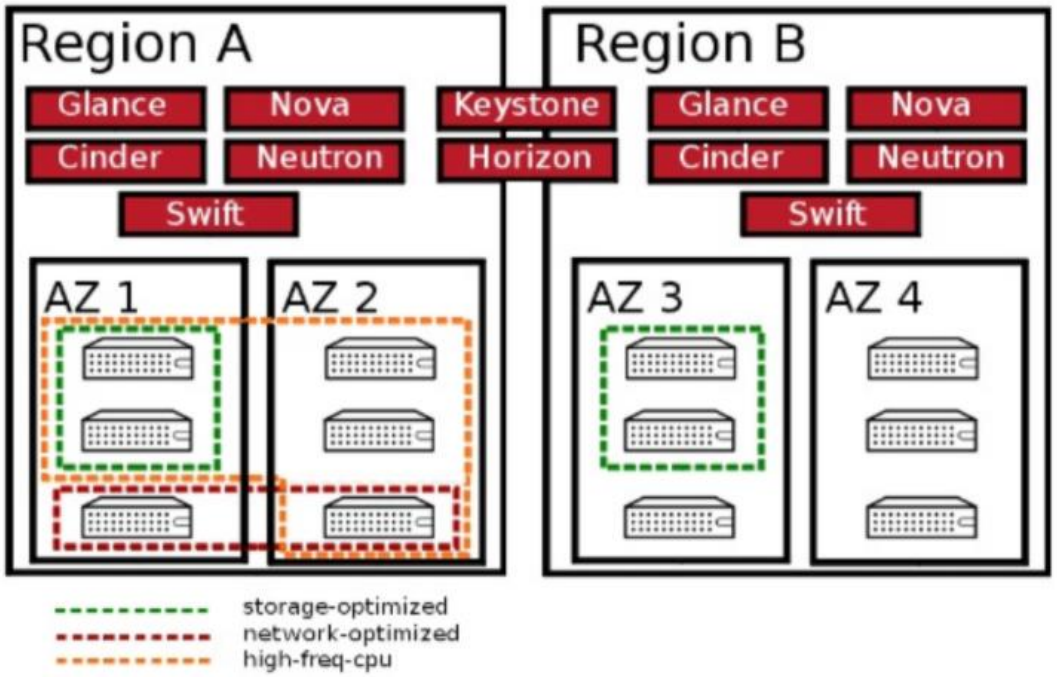

Figure 35: OpenStack Availability Zones [32]

5.8.10 AUTO-SCALING SUPPORT

OpenStack supports Auto-scaling and defines it as a capability of Cloud infrastructure to detect conditions related to the load on the compute workloads automatically and respond to it without the intervention of an IT staff. Auto-scaling can horizontally scale up and scale down the resources based on load.

There are several methods by which OpenStack supports auto-scaling of clusters such as a Heat AutoScalingGroup and Senlin Clusters. [33]

The following figure shows the conceptual diagram that is used as a reference when deploying Auto-Scaling in OpenStack:

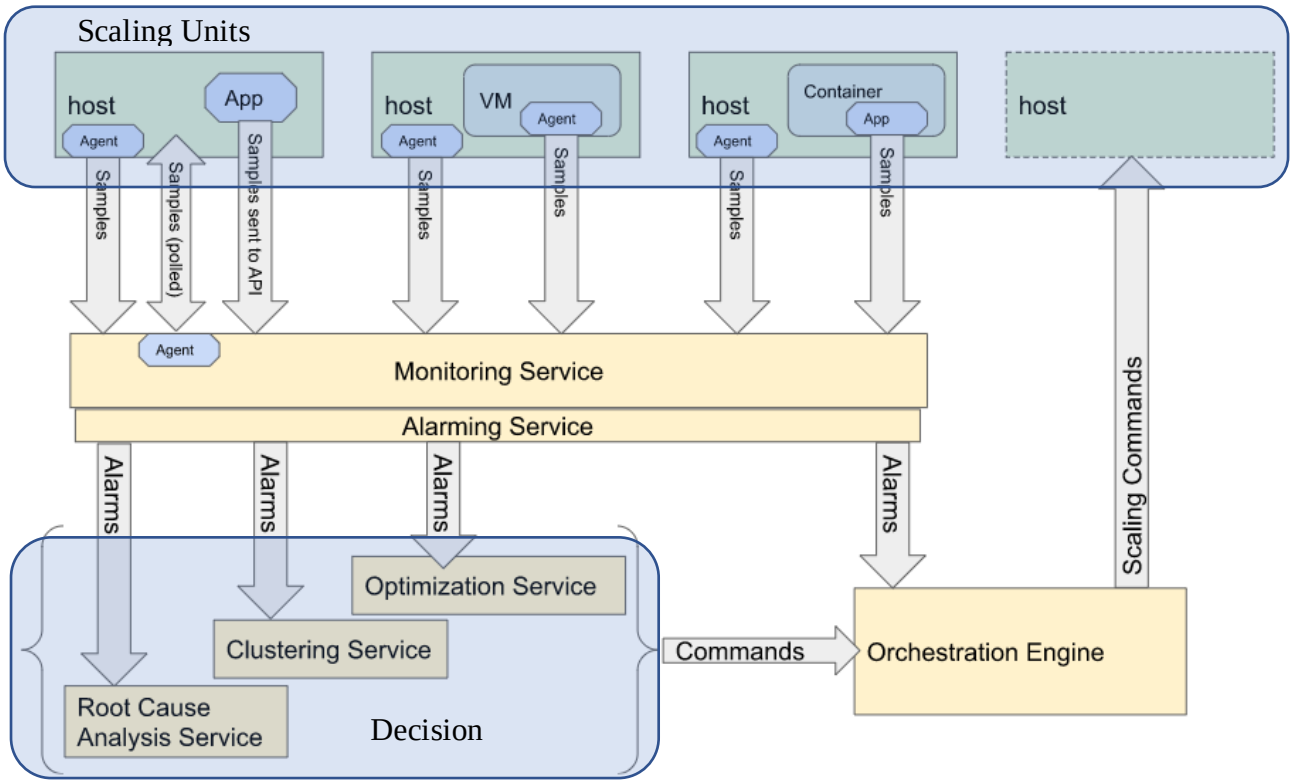

Figure 36: Conceptual diagram on Auto-Scaling [33]

Components of Auto-Scaling in OpenStack:

- **Scaling Units**: OpenStack allows to control the following units using Auto-Scaling:
  - VMs running on Compute Instance
  - Compute instances

- Containers running on a Compute instance
- Network Attached Storage (NAS)
- Virtual Network Functions

- **Monitoring Service**: It allows to automatically sense conditions related to workload and can make use of polling or an agent installed on the Scaling Units which can be as follows:

  - Monasca: It is a solution that is an open-source multi-tenant, scalable, high-performance, fault-tolerant monitoring-as-a-service.

  - Ceilometer from the Telemetry project: This service helps with the collection of data on physical and virtual resource utilization based on set triggers.

  - Prometheus: Display a chart in OpenStack and provides a time-series database and monitors components of OpenStack-Helm.

- **Alarming Service**: Create alerts based on set thresholds. The following can be used for setting these alarms:

  - Monasca: Makes use of services such as built-in alarm threshold and notification.
  - Aodh service from Telemetry project.

- **Decision Services**: Services that are responsible for interpreting the metrics and alarms based on logic and send commands to the Orchestration Engine.

- **Orchestration Engines**: Services such as Heat, Senlin, and Tacker can be used to orchestrate auto-scaling in OpenStack [33]

5.8.11 SELF HEALING SUPPORT

Self-healing support makes use of the same services and technologies as Auto-scale service, but it is used for healing failures in the OpenStack Cloud.

5.8.12 HORIZON

Horizon is a web-based OpenStack dashboard that helps manage all core services such as Nova, Glance, Cinder, Neutron, and Swift. It is designed and architected based on the following key values:

- **Core Support**: It provides simplified out-of-the-box support for OpenStack projects. Contains 3 different dashboards: User Dashboard, System Dashboard, and Settings Dashboard.

- **Extensible**: Provides ease for developers as it makes use of Dashboard class which makes the API consistent.

- **Manageable**: It makes the core codebase easier to navigate and makes code easy to find using granular breakdown.

- **Consistent**: Applications have consistency in visual (base forms, widgets, tags, and views) and interaction paradigms.

- **Stable**: The changes made are backward compatible by making use of reliable API.

- **Usable**: Provides the best possible user experience when using the interface. [34]

5.9 VMware CLOUD FOUNDATION

5.9.1 OVERVIEW

VMware Cloud Foundation is a private and hybrid cloud platform for managing VMs and containers that makes use of HCI (Hyperconverged Infrastructure) which is a software-defined data center stack and it combines compute, storage, cloud management, and data networking. HCI is a replacement of traditional three-tier (separate compute, storage, networking, and management) architecture, which was expensive to build, and difficult to operate and scale. [35]

The following figure shows the core components of VMware Cloud Foundation:

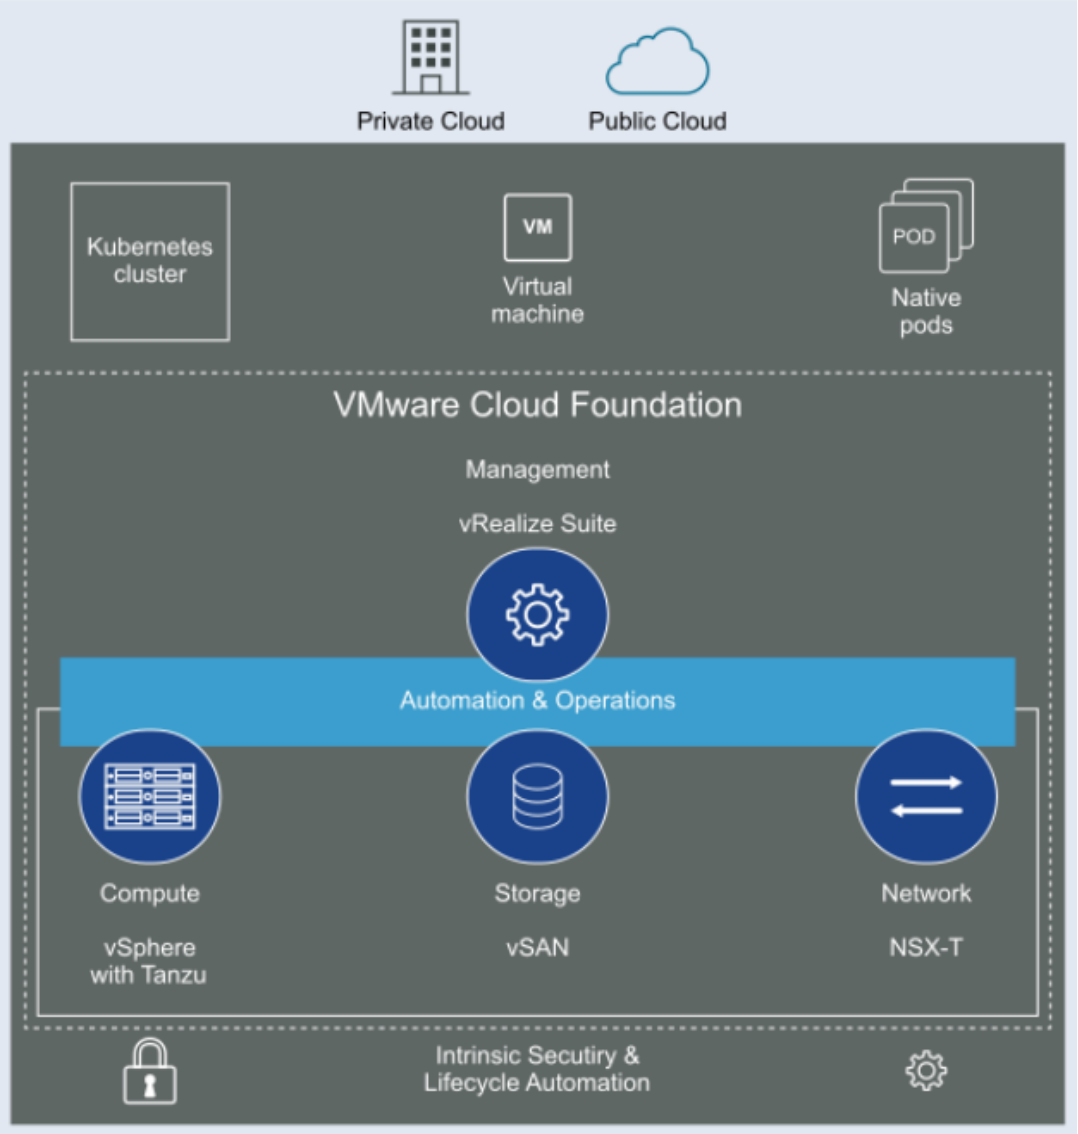

Figure 37: VMware Cloud Foundation core components [36]

VMware Cloud Foundation (VCF) makes use of core services such as VMware vSphere with Tanzu for Compute, vSAN for storage, NSX-T for data networking, and vRealize Suite for automation as shown in the figure above. It provides additional features such as the Kubernetes cluster for container orchestration, network security, and other cloud management tools.

VCF makes use of automation using the Cloud Builder appliance to bring-up the complete software-defined stack. It then uses SDDC (Software-defined Data Center) Manager to automate configuration and provisioning. Another purpose of SDDC is to automate the lifecycle management of this stack. [35]

5.9.2 SOFWARE-DEFINED DATA CENTER MANAGER

The purpose of the SDDC manager is to automate the complete lifecycle of the VCF.

It is responsible for:

- Initial bring-up
- Configuring
- Provisioning
- Managing
- Operating
- Monitoring
- Upgrading and patching the logical and physical resources of VCF

The configuration, operation, and management functions are carried out by:

- The abstraction and aggregation of available physical resources into a logical entity.

- Resource management of physical resources by addition or removal of hosts or switches, failure management, and their maintenance.

- Helping with the orchestration of booting-up and shutdown of logical software and other VCF components such as ESXi, vCenter Server, vRealize operations, NSX, and vSAN.

- Generation of logical resource structural mapping which is based upon the physical profiles, events, and operations.

- Interaction with VCF components for example NSX for network management, vRealize for monitoring health, etc.

SDDC manager uses a multi-threaded execution engine as follows:

Service Engine: Allows SSDC manager to perform management plane functions. It makes use of the Java executor service framework which is initialized with multiple runnable threads and scheduler threads.

Physical Resource Manager (PRM): It is responsible for managing the physical components of a rack and maintaining the software associated with it.

Event Engine: Displays events information on the UI of SDDC manager and pushes them to vRealize Log Insight. [37]

The following figure illustrates how the SDDC manager is linked to the management and infrastructure plane:

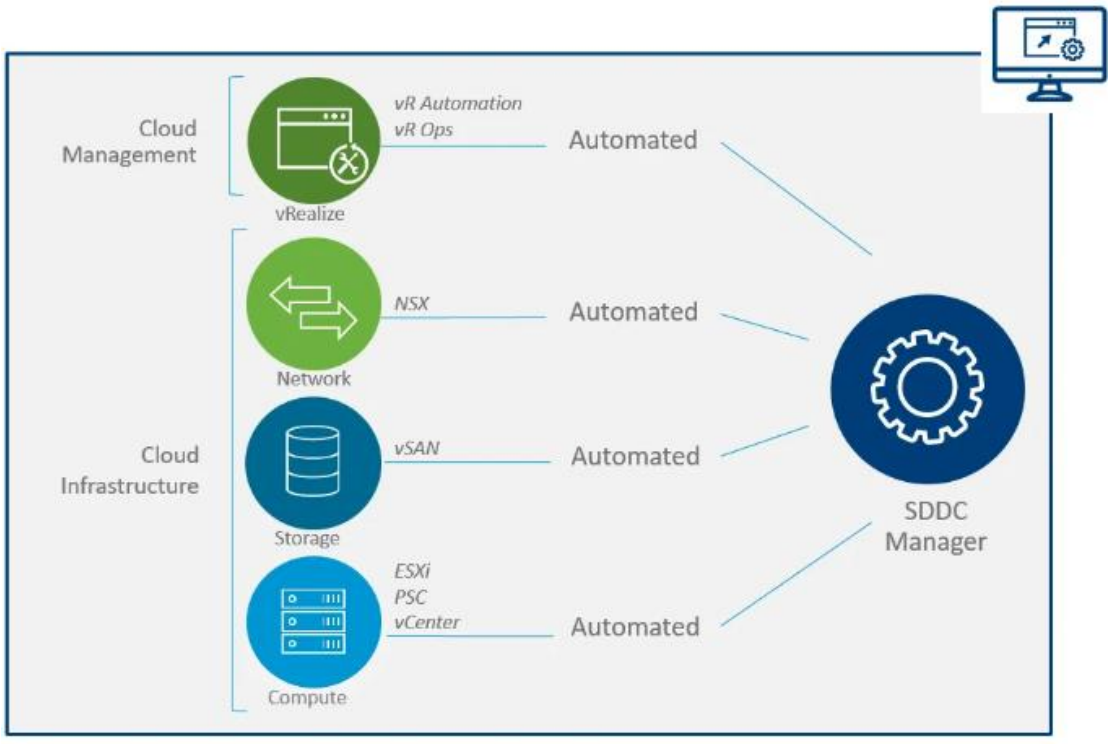

Figure 38: SDDC Manager [38]

5.9.2.1 CLOUD FOUNDATION BUILDER

It is a service that automates the deployment of complete VMware validated designs which are distributed as an Open Virtualization Appliance (OVA). The VMware validated design bundle containing VMware binaries and signed certificates are required to be uploaded.
The following shows the workflow based on which the Cloud Builders needs to be performed :

1. Get the environment perquisites based on the Deployment Guides document ready as follows:
    a. Get the ESXi hosts prepared
    b. Get the UMDS VM prepared
    c. Get the vRealize Automation IaaS Master VM prepared
    d. Get the SQL server prepared for v Realize Automation
    e. Get the Site Recovery Manager ready
    f. Then generate the signed certificates

2. List or fill the deployment parameters in an XLS file format

3. Now download and deploy the Cloud Builder appliance and:
    a. The software bundle should be uploaded
    b. And the signed certificates

4. Deployment of SDDC:
    a. The JSON file will have to be generated
    b. Auditing needs to be carried out of the parameters and target environment before the deployment of the SDDC
    c. Deploy the SDDC

5. In the end, manual post-deployment procedures are carried out. [39]

5.9.3 VSPHERE

VMware vSphere makes use of virtualization to change separate data centers into aggregated IT infrastructure that makes use of computing, storage, and data networking resources. It manages these resources as a unified environment and provides tools to control the data centers present inside. [40]

There are 2 major core components of VMware vSphere illustrated in the architecture which is displayed in the following figure as ESXi and vCenter:

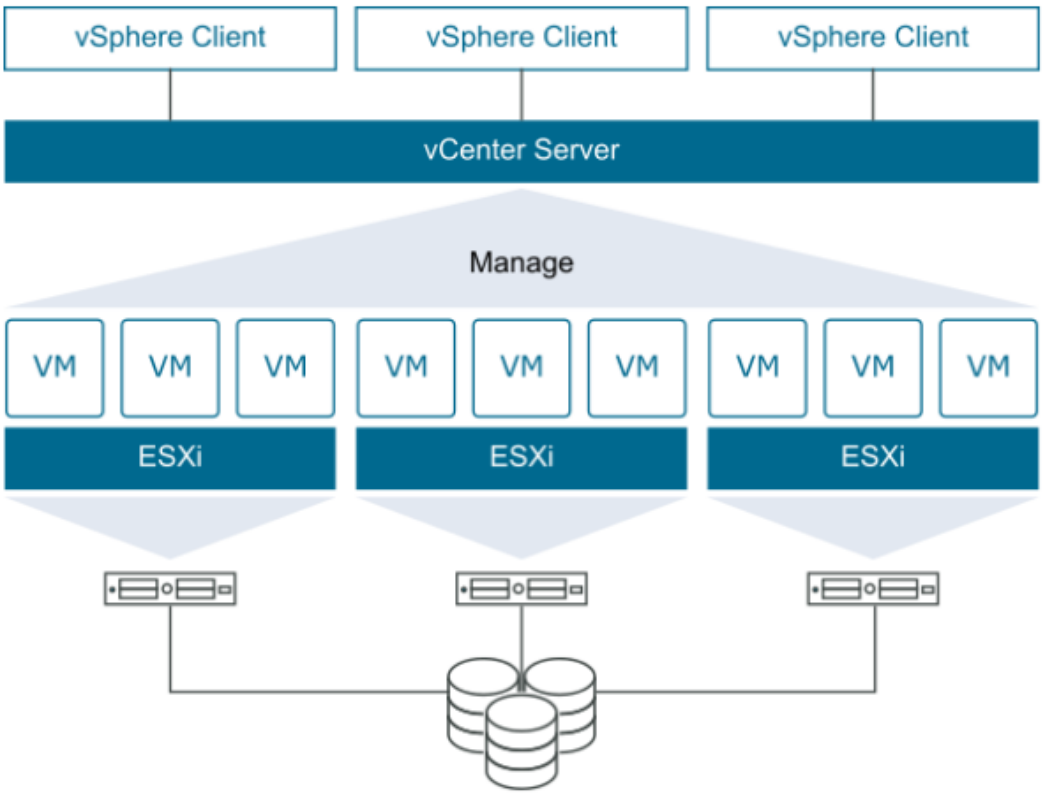

Figure 39: Major Components of vSphere [40]

5.9.3.1 ESXi

ESXi (Electric Sky X integrated) is a type-1 hypervisor that is directly installed on bare metal servers. It provides a virtualization platform where VMs and other virtual appliances can be created and operated. Type 1 hypervisor has better security, reliability, and management as compared to type 2 as it is independent of the OS.

The following figure shows the architecture of the ESXi host:

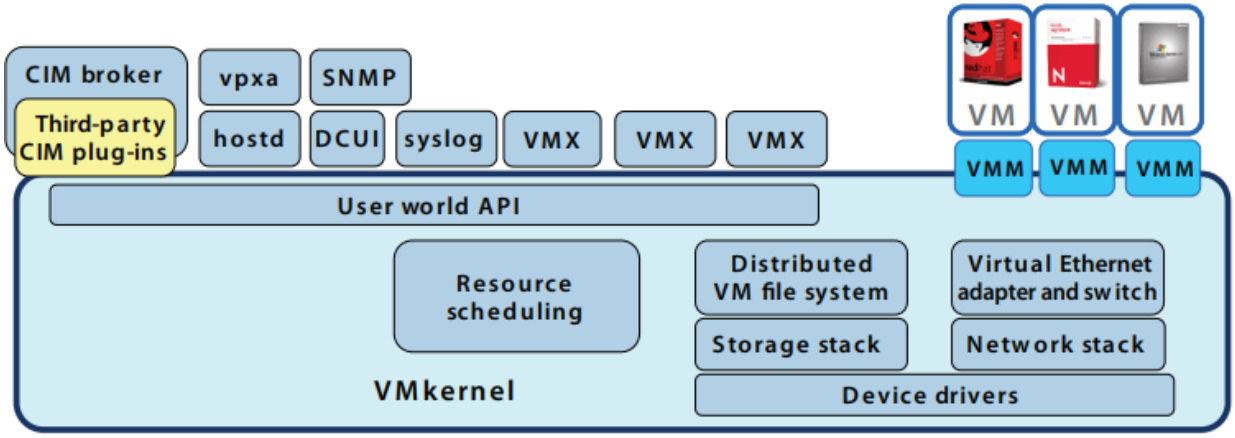

Figure 40: Architecture of ESXi [41]

ESXi runs on top of its kernel called VMkernel. The main components that run on top of VMkernel are:

- **Direct Console User Interface (DCUI)**: It is used for initial configuration and is defined as low-level management and a configurational interface that can be accessed through the console.

- **Virtual Machine Monitor (VMM)**: Helps with the execution of VMs and an as a helper process called VMX

- **Common Information Model (CIM)**: CIM is a system that helps with hardware management through remote applications. [41]

5.9.3.2 VCENTER SERVER

It pools and manages several ESXi hosts and resources connected over a network. It also acts as a central administrator for ESXi hosts.

It serves the following three functions as shown in the following figure:

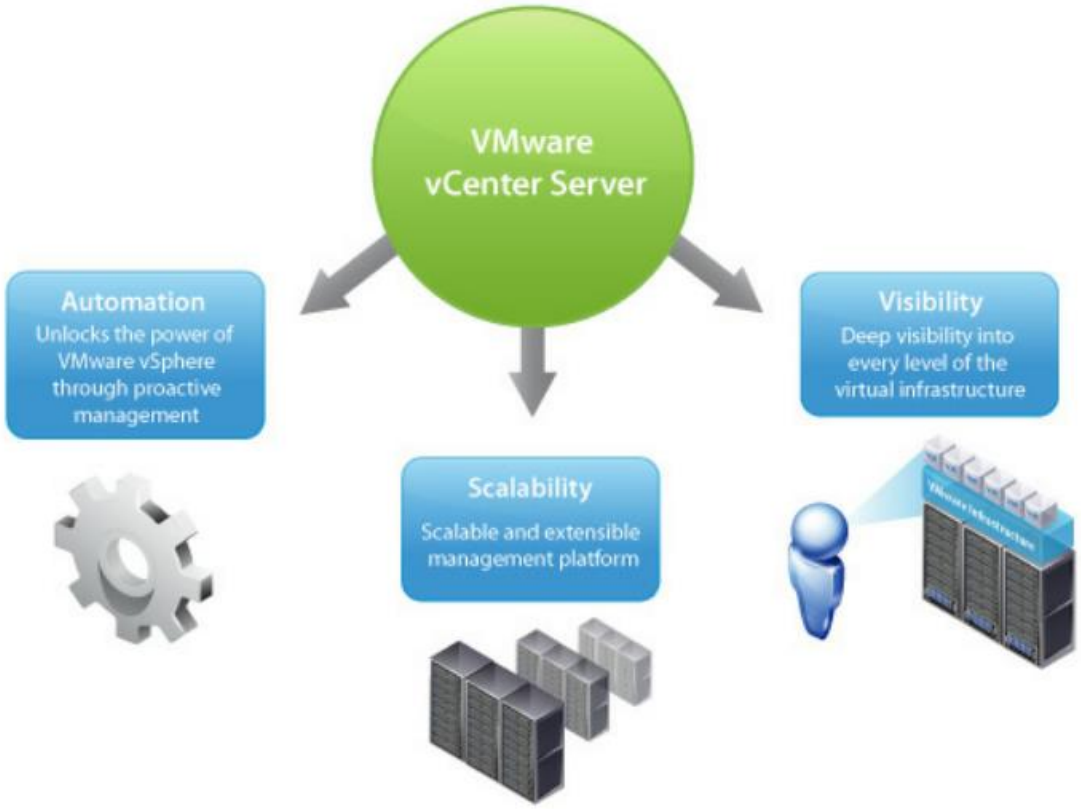

Figure 41: vCenter Server [42]

- **Visibility**: It configures ESXi hosts and VMs and monitors their performance using events and alerts.

- **Scalability**: The vCenter visibility is scalable across multiple ESXi servers and VMs.

- **Automation**: The alerts generated in vCenter can be used for triggering automated actions using orchestration. [42]

A single vCenter instance contains the following components:

- **vSphere Client and Web Client**: These are tools that are used for the management of ESXi hosts of a vCenter Server.

- **vCenter Server DB**: Used for storage of inventory items, security roles. Resource pools etc. Supports MS SQL Server and Oracle databases.

- **vCenter Single Sign-On (SSO)**: Multiple repositories can be authenticated using a single credential such as open LDAP or MS AD.

- **Managed hosts**: ESXi hosts and the VMs inside them.

Tanzu: VMware has introduced another service called Tanzu as a part of vSphere which has transformed it for running Kubernetes container workloads directly on the hypervisor or the ESXi hosts. [43]

vCenter provides great features to vSphere such as:

- **vSphere High Availability (HA)**: Provides failover protection from OS or hardware outages. High availability is maintained as follows:

    - Monitoring and detection of any hardware and OS failure inside hosts and VMs.

    - Helps with restarting the VMs during outages without manual intervention.

    - Helps with reduction of downtime during outages.

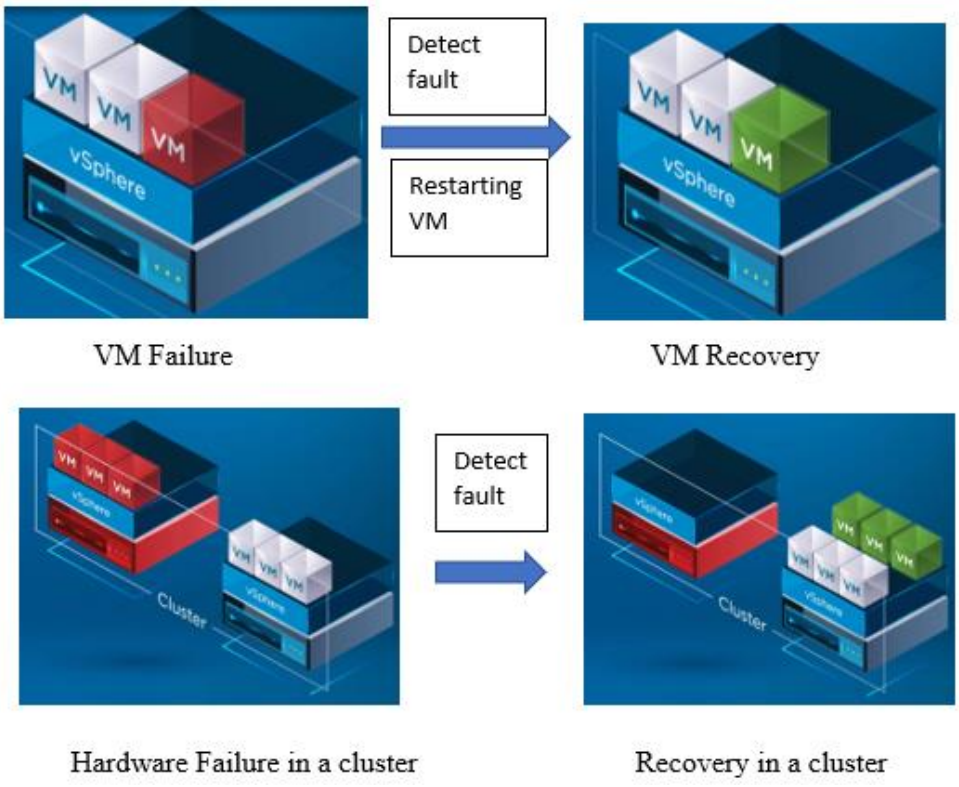

Figure 42: vSphere HA [44]

It provides automated protection when a hardware failure is detected in a cluster with no modification or changes to VM.

According to VMware, vSphere HA provides scalability, reliability, and usability:

- Scalability: vCenter works on the Master-Slave relationship instead of the Primary-Secondary relationship which eliminates the need for administrators to know where the primary and secondary nodes are located. vCenter also supports IPV6 which addressing the need for large address space.

- Reliability: vCenter ensures high reliability by eliminating external dependencies such as DNS resolution. It also relies on multiple communication paths for redundancy. And it prevents disruption of VM-VM anti-affinity rules defined by Distributed Resource Scheduler which will be discussed later in this section.

- **vSphere Fault Tolerance**: Creates shadow instance which is a replica of the primary VM in another server. If a failure is detected in the primary the vSphere Fault Tolerance feature triggers failover to the shadow VM. The shadow VM is then changed into the primary and a new shadow instance is created in some other server. [44]

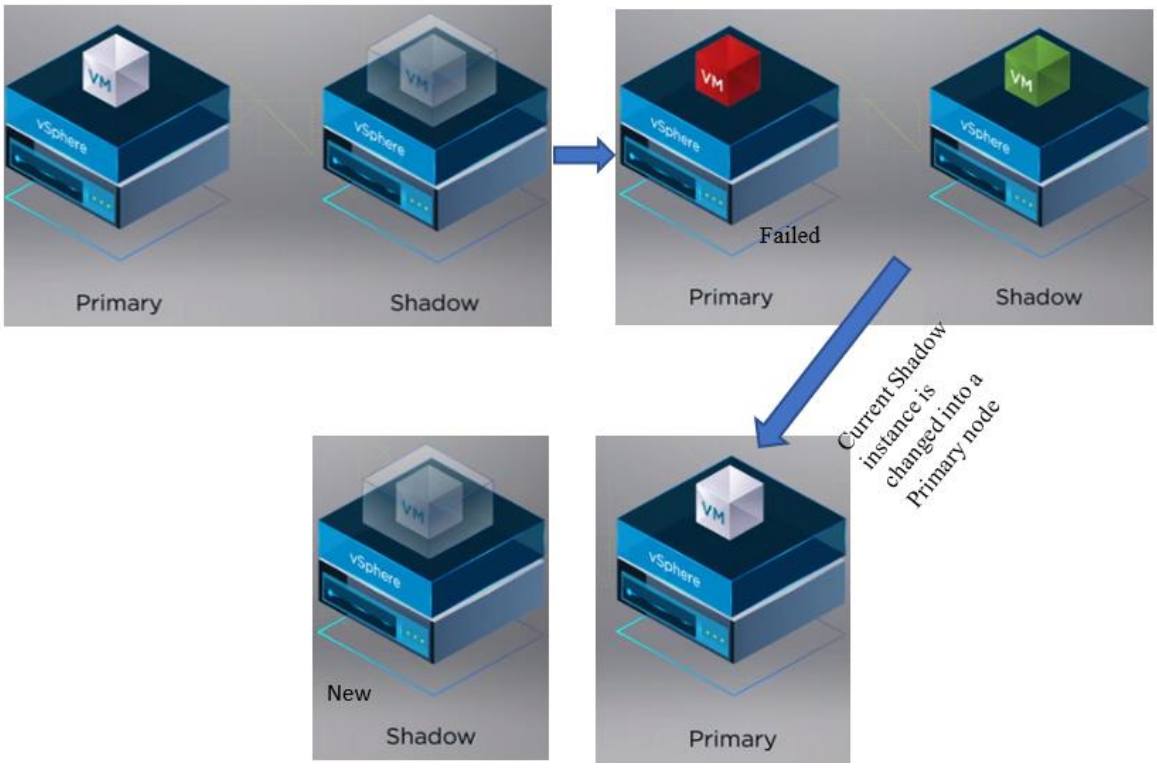

Figure 43: fault tolerance

- **vSphere vMotion**: It is a service that allows live migration of virtual machines from one physical server to another while it is running avoiding any downtime during maintenance, upcoming natural disaster, or during moving an office. The vMotion technology works by using the 3 technologies, the first involves encapsulation of the entire state of VMs by a set of files that are stored on shared storage which can be a Fiber channel or SAN (storage area network) or iSCSI or a NAS (network-attached storage). A clustered file system in VMware called VMFS (Virtual Machine File System) makes multiple installations of the ESXi server to allow access to the VMs concurrently. The second technology involves the fast transfer of active memory and precise execution state of VM over a high-speed network. This allows the VM instantaneously to move from one ESXi host to another. And the final third technology involves the usage of virtual networks in ESXi hosts which allow the preservation of network identity and network connections. vMotion works on managing the virtual MAC address of the VM and sends ping messages to the router so to make it aware of the new location. All these three technologies work together to prevent any downtime during the migration of the VMs. [45]

  The following figure illustrates the vMotion technology in vSphere:

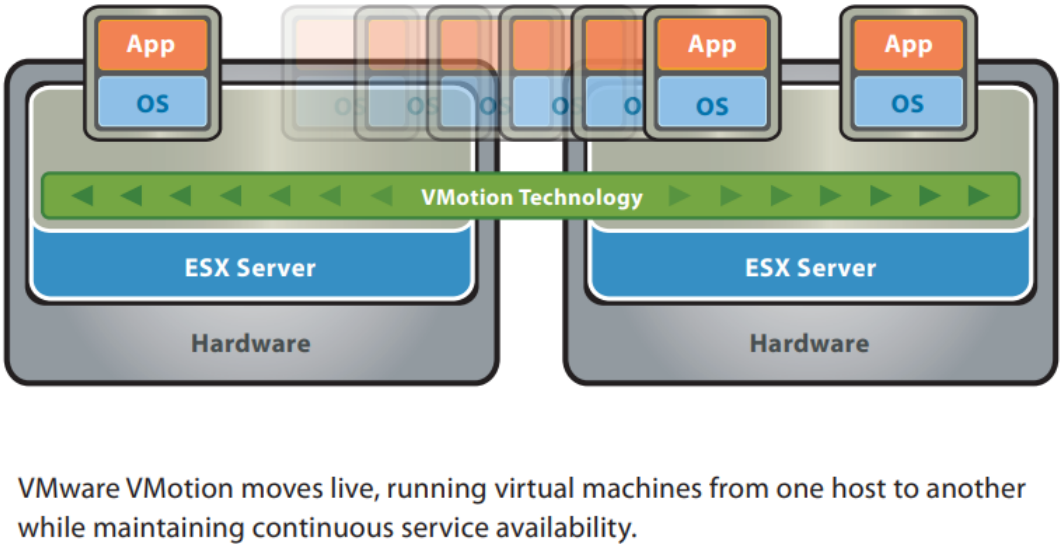

Figure 44: VMware vMotion [45]

- **vSphere Storage vMotion**: It is a component that responsible for the live migration of storage connected to VM from one storage system to another. It operates by first copying metadata to an alternate location, then the disk file also called Virtual Machine Disk File (VMDK) is replicated to this new location by making use of vSphere's Changed Block Tracking (CBT) feature to maintain data integrity. In the third step, CBT is again queried for making 2nd copy involving only blocks of storage

that are changed with respect to the first replication to another new location. This step continues until the copies are identical to each other. The old VM is stopped and the pointer is moved to the new virtual disk image. Finally, before starting the new virtual machine, the last changed blocks of the source disk are replicated, and the old disk is removed. [46]

The following figure illustrates the process that occurs during Storage vMotion inside vSphere:

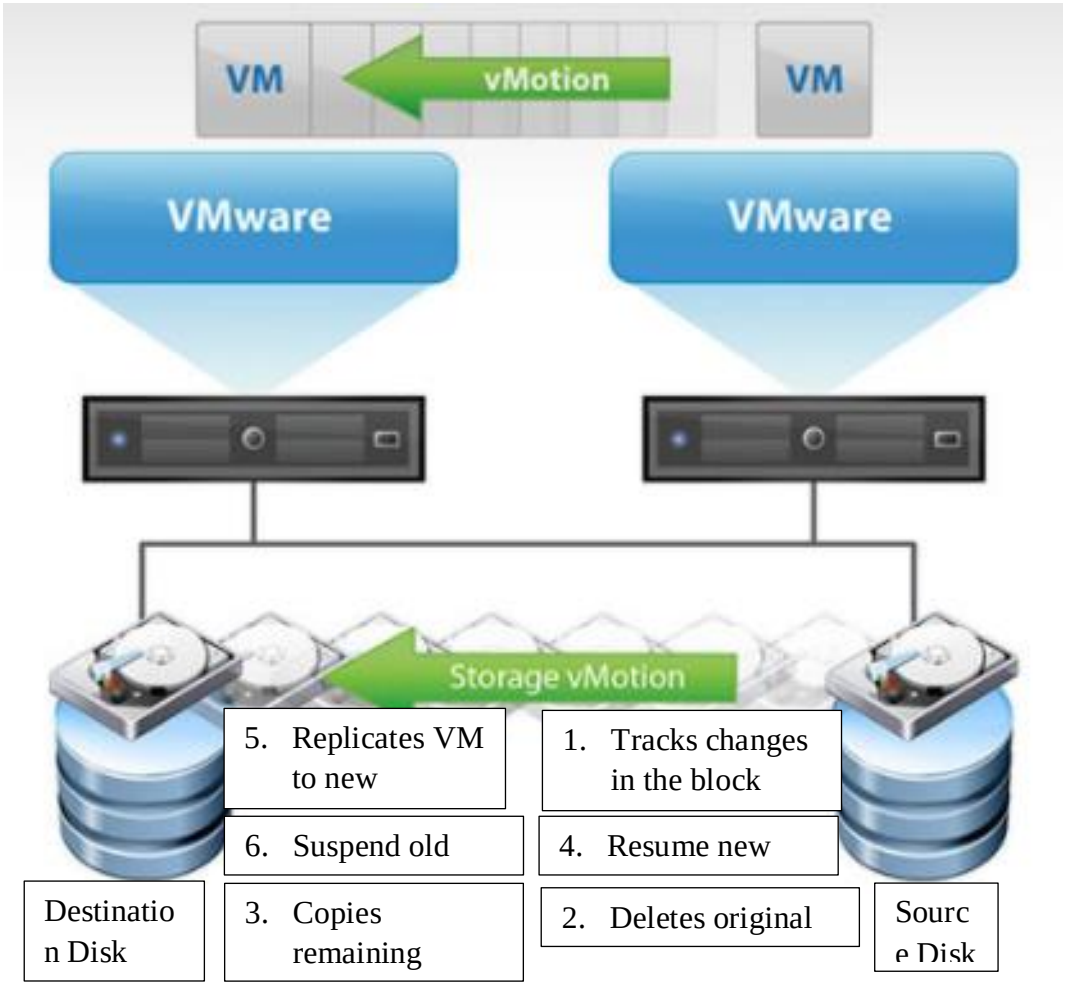

Figure 45: vMotion Storage [47]

- **vSphere Distributed Resource Scheduler (DRS)**: DRS is responsible for load balancing VMs based on available resources via vMotion. The rules are user-defined for automation and can also be carried out manually. Anti-affinity rules and DRS groups are used for separation for the identification of redundant VMs. [48]

The process of workload balancing is shown in the following figure:

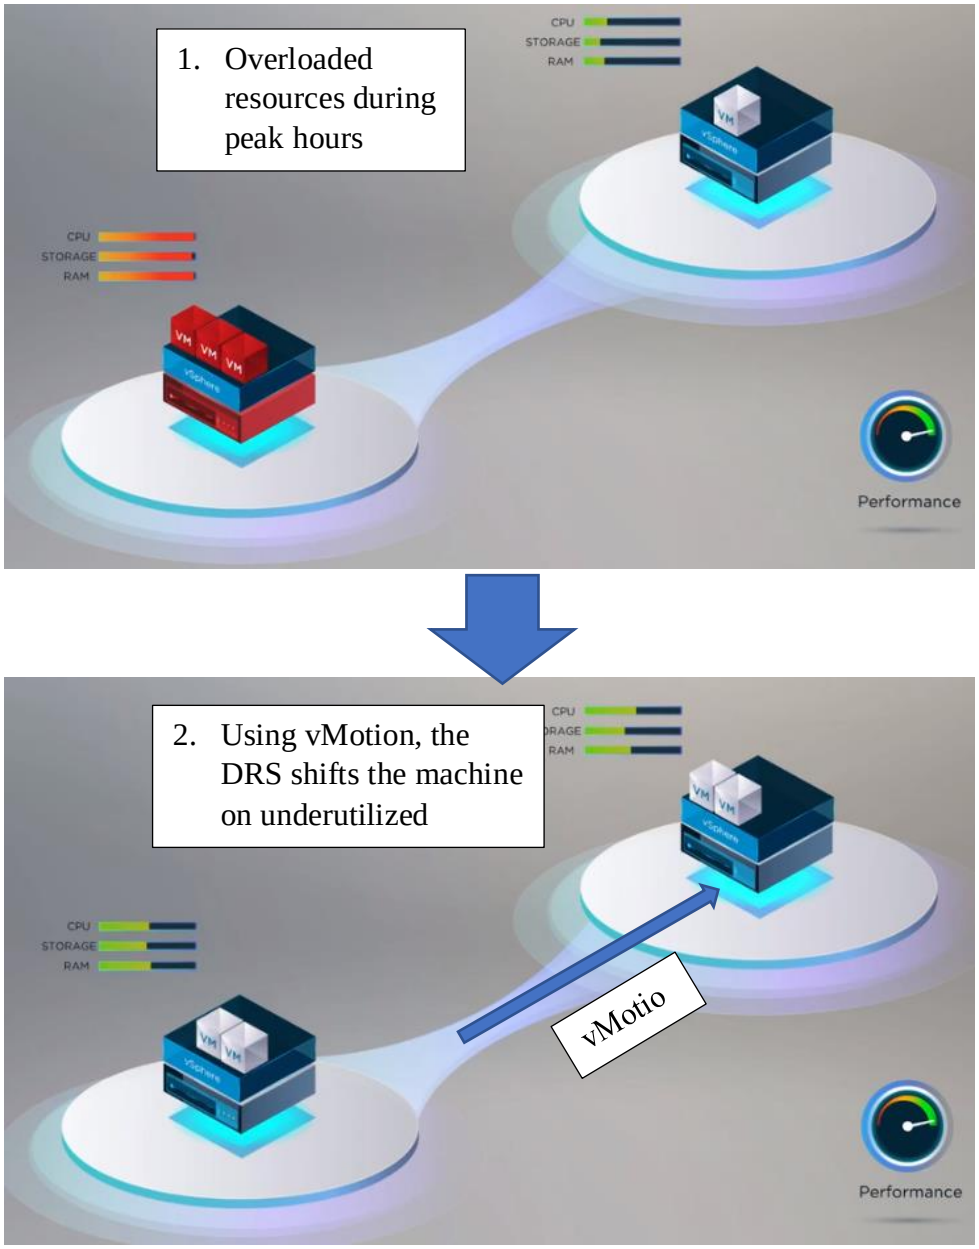

Figure 46:vMotion for DRS [49]

This methodology allows high availability, optimal performance, and scalability. DRS also optimizes power usage also called Distributed Power Management (DPM) by removing or shutting down underutilized resources and shifting the VMs using vMotion during low demand.

It has a new functionality such as machine learning to predict the usage and optimize power usage based on expected demand. It has the capability of prioritization resources based on the rules and policies set by the admin.

5.9.3.3 GPU-AS-A-SERVICE

VMware Cloud Foundation enables organizations to run share high-performance GPU over the network infrastructure for Artificial Intelligence (AI) and Machine Learning (ML). It prevents wastage of underutilized GPU resources by sharing them with data scientists and developers when needed and avoids dedicated allocation to a single entity. NVIDIA GPUs are utilized in the industry for computing horsepower to carry out deep learning and train neural network-based models.

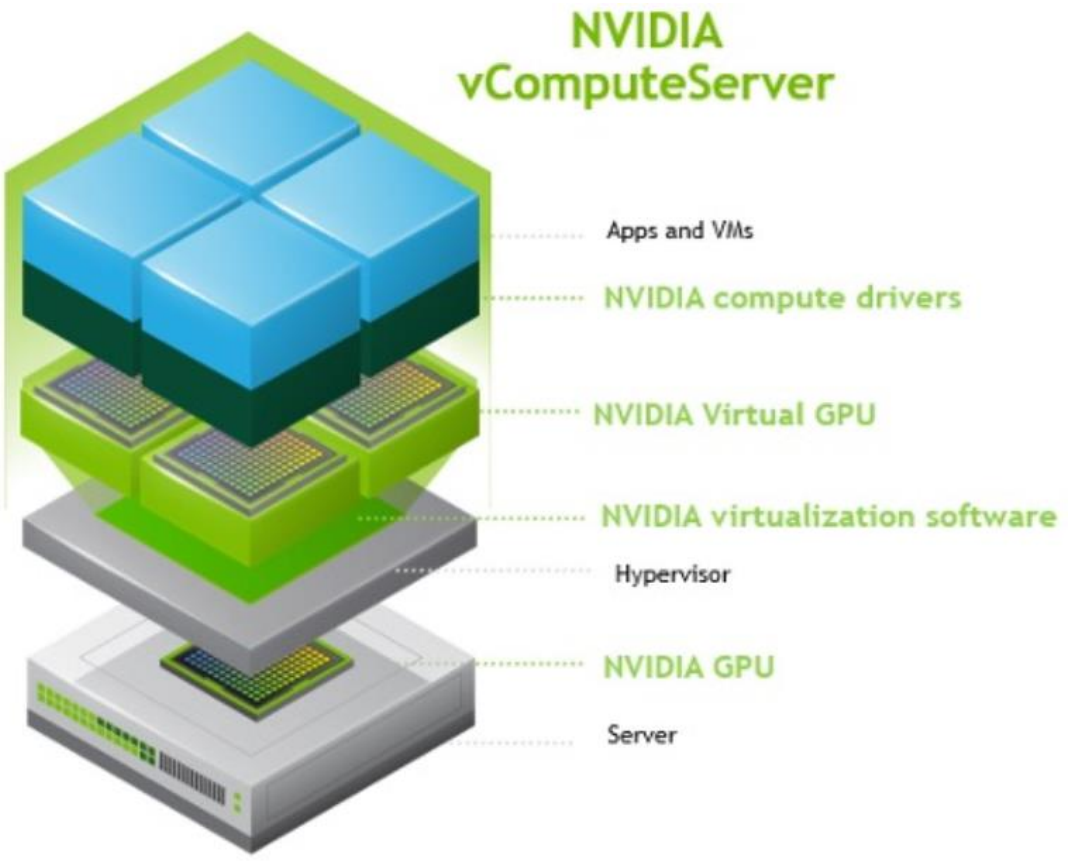

Figure 47: Layered model showing NVIDIA vGPU components [50]

NVIDIA vComputeServer allows sharing of vGPU among multiple VMs. This allows running more than 60 AI applications. This drastically reduces operational costs. The vGPU can be managed from vCenter inside vSphere. [50]

High-Speed Networking with PVRDMA and RDMA

VMware vSphere supports Remote Direct Memory Access (RDMA). RDMA reduces latency in the network by providing direct access to memories of the host and bypassing

the OS and CPU. Making use of this technology reduces the load on the CPU and optimizes bandwidth utilization. vSphere can utilize VMware's para-virtualized RDMA (PVRDMA) technology which enables virtualization compatible with RDMA.

The following figure displays how RDMA is implemented using a PVRDMA stack:

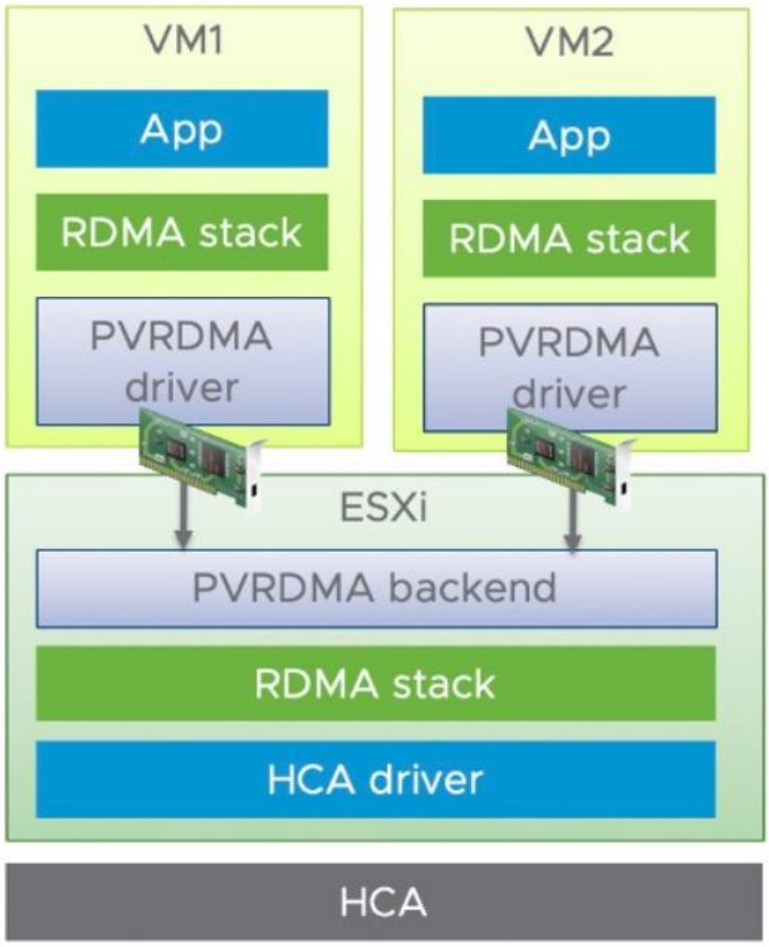

Figure 48: The end to end PVRDMA stack [50]

vSphere also supports RDMA over Converged Ethernet (RoCE) which allows RDMA to operate over an Ethernet network. [50]

5.9.4 VSAN

A core VMware Hyperconverged Infrastructure (HCI) technology that is responsible for the aggregation of data storage devices that can be directly attached or locally present into a shared single storage pool of hosts in a single cluster. [51]

The use of vSAN simplifies the management of private and hybrid cloud infrastructure by the elimination of the specialized skill set that was required for traditional infrastructure deployment. It can run on x86 servers from multiple OEMs. It supports both traditional and modern container-based applications. [52]

The following figure illustrates how vSAN changes different storage servers into abstraction as a single cluster:

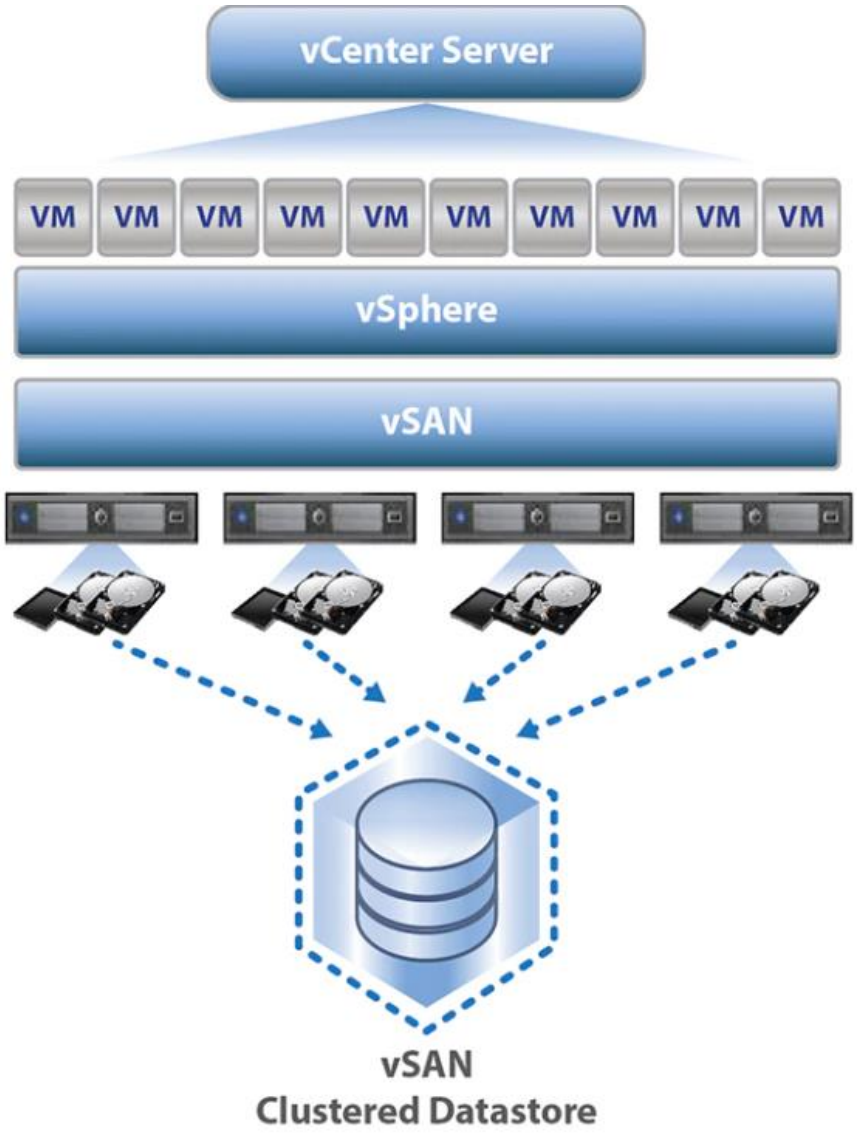

Figure 49: vSAN Architecture [53]

5.9.4.1 VSAN ARCHITECTURE

vSAN is software-defined storage designed for vSphere. It can be provisioned and managed through vCenter or vSphere Client. vSphere and vSAN are embedded on the hypervisor delving storage and compute from the same x86 server. Deployment in a cluster can range from 2 nodes up to 64 hosts.

Servers with Local Storage: There are two possible ways storage can be designed in the infrastructure; one is the all-flash configuration and the other can be a hybrid approach containing a mix of magnetic disks and flash drives. To improve performance, faster flash drives are used for Cache and slower ones such as magnetic drives are used for capacity. The writes are carried out at the cache layer and later transferred to capacity when the blocks are cold, or when the cache space is running out making the storage cost-effective and fast. [52]

The following figure shows how data can be cached in high endurance drives:

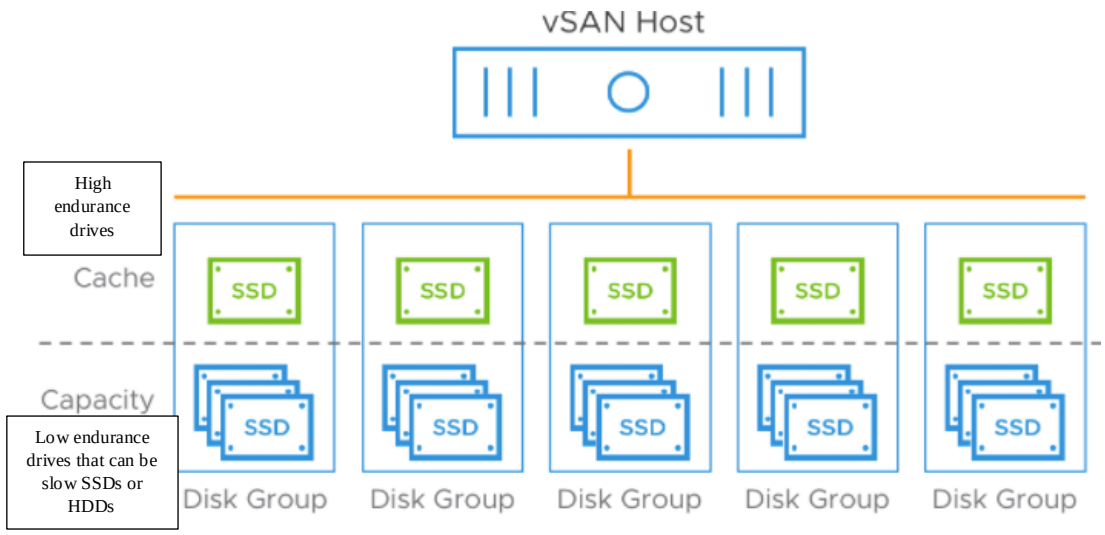

Figure 50: vSAN using high endurance drives for caching

Traditional Storage Controller Virtual Appliance vs vSAN

A Storage Controller Virtual Appliance requires high computational resources such as CPU and RAM with dedicated storage arrays. Most Hyperconverged Infrastructure (HCI) makes use of them. They are used for reducing VM workload, but it increases the cost of operation due to additional resource requirements. Usage of storage virtual appliance also result in increased latency due to additional steps required to handle operations when writing data as shown in the figure below:

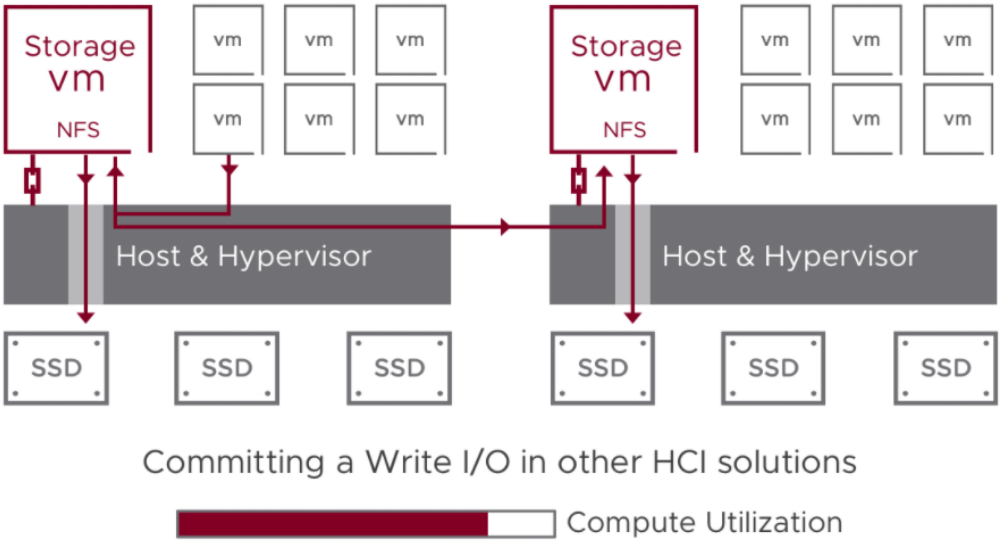

Figure 51: High latency as traditional storage requires multiple steps [52]

Whereas vSAN has lower latency as it does not require multiple steps to do the same operation when writing data. Since it is a native technology in vSphere, it uses 10% lower computational resources when compared to a storage virtual appliance. [52]

The architecture of vSAN operations is illustrated in the figure below:

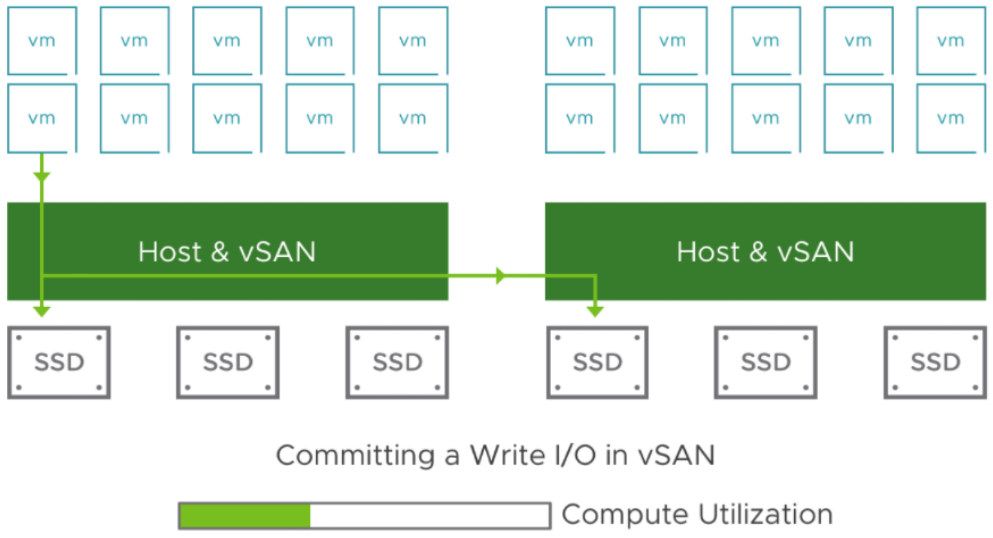

Figure 52: Lower latency as it does not require multiple steps [52]

The above figure shows the much simpler design with shorter I/O paths and the removal of storage virtual appliance which lowers the use of resources which results in lower costs of ownership. [52]

5.9.4.2 vSAN CLUSTER TYPES

Standard Cluster Topology

A standard cluster may contain at least three physical nodes, and which can also be scaled up to 64 nodes. The nodes are usually located in a single location. The link of 10Gb or higher is used for connectivity for all-flash configuration and is recommended speed for hybrid configuration. [52]

Two-Node Cluster Topology

2-node cluster topology involves 2 physical nodes in a single cluster and location. These clusters are connected directly or may be connected via switch in the same network.

The following figure shows the architecture of a 2-node cluster:

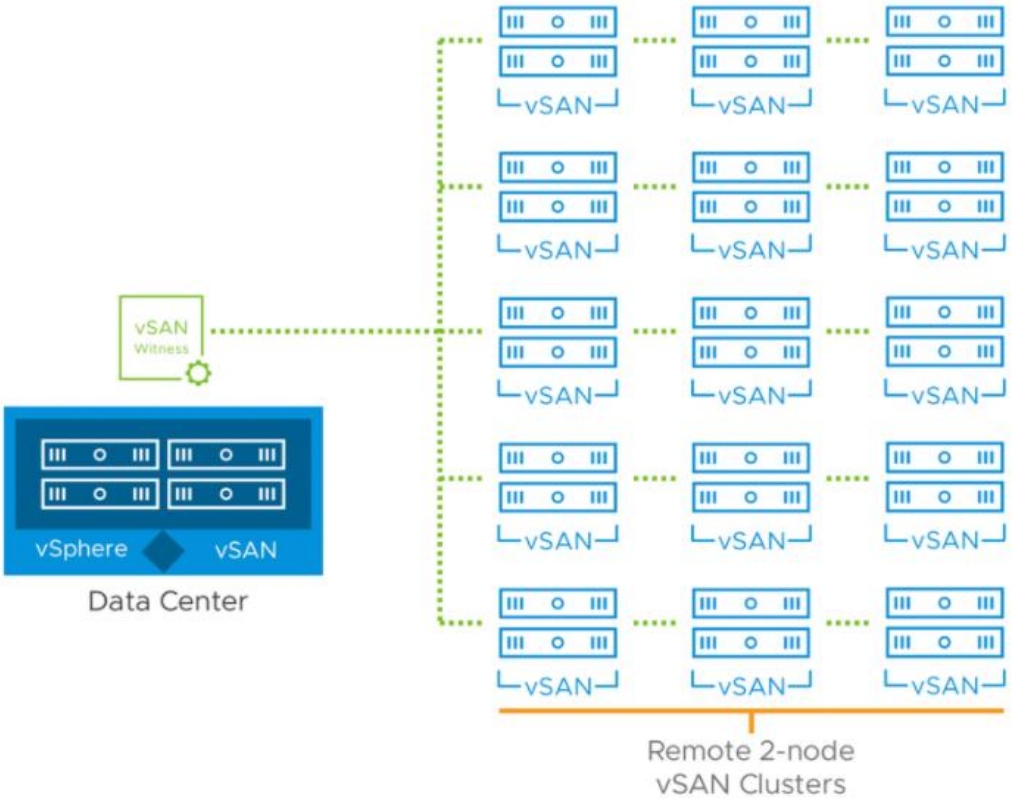

Figure 53: Two-Node Cluster Topology [52]

Direct connection helps in the reduction of the cost of deployment by the removal of a network switch from a cluster. vSAN Witness Host is used for establishing a quorum on failure for the 2-node topology. [52]

Stretched Cluster Topology: vSAN has the capability to run stretched cluster topology which an active-active disaster recovery solution as shown in the following figure:

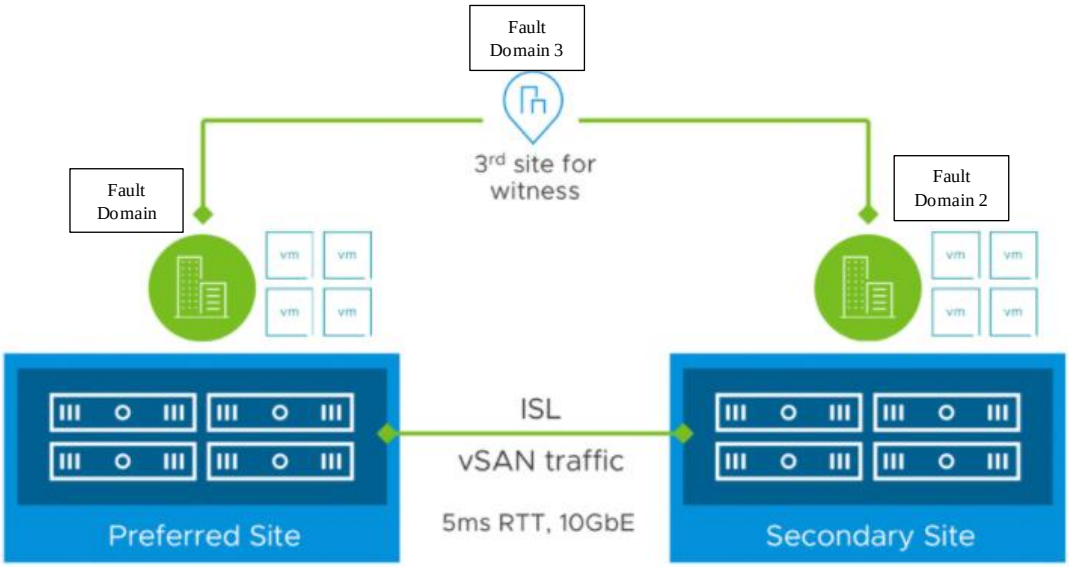

Figure 54: Stretched Cluster Topology [52]

Replication happens between the 2 sites A and B through the L2 network. The distance between active sites should not allow the round-trip time (RTT) latency to exceed 5ms and a 10 Gbps link is required between the sites. The RTT between Witness and other sites can go up to 200ms and a minimum of 100 Mbps link can be used between them. The number of hosts in a cluster can be a maximum of 30 which is stretched up to 15 at each site. This topology is well integrated with vSphere High Availability. That is, if a site goes down, vSphere HA will automatically restart VMs at the Secondary site in a matter of seconds. The Stretched Cluster consists of three fault domains. [52]

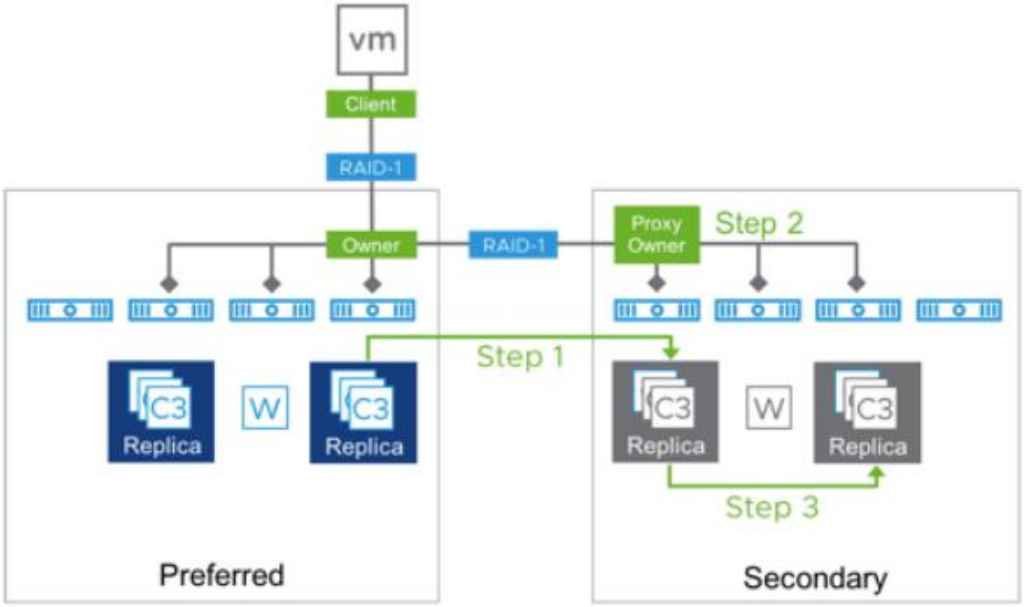

Figure 55: vSAN provides redundancy by storing replicas on secondary sites [52]

In case of failure which causes an incomplete dataset, it is efficient to only copy the part of that dataset instead of a full copy. vSAN can carry out partial synchronization. In large failures, as shown in the figure below, vSAN will failover to the Secondary, and the Preferred site will not become the authoritative site until connectivity between the Witness and Secondary sites is restored which prevents the Preferred site to report as available with stale data. [52]

The following image shows the scenario when a site goes down in a stretched cluster topology:

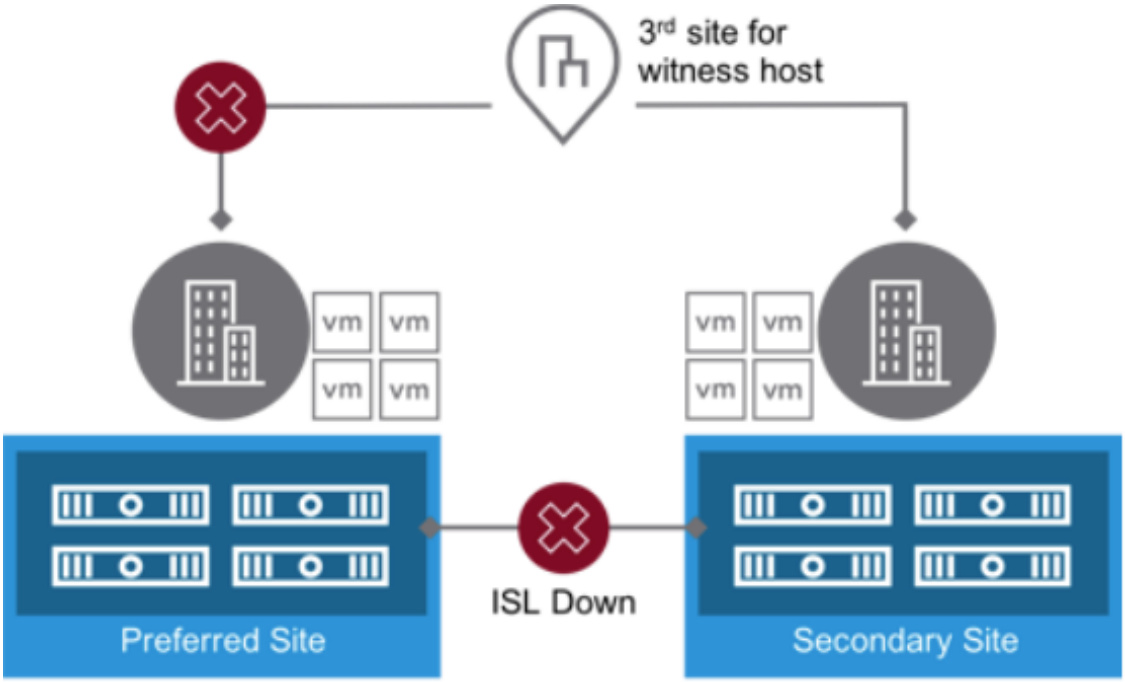

Figure 56: Failure in the Stretched Cluster Topology [52]

HCI (Hyperconverged Infrastructure) Mesh Cluster: In this vSAN topology, the local storage is polled from hosts in a cluster and merged into a single vSAN datastore. This topology provides more flexibility in borrowing or sharing the capacity over the vSAN cluster. This capability also allows cross-cluster storage consumption. It allows the administrator to connect vSAN datastore from a server cluster to client clusters, which helps with extending server capacity to the clients. HCI Mesh makes it easier to migrate a VM using vMotion without migrating the storage. It also improves scalability for both compute (VM) and storage. It makes use of hub and spoke topology to associate the server cluster to client clusters. [52]

The following figure shows the architecture of the HCI Mesh Cluster:

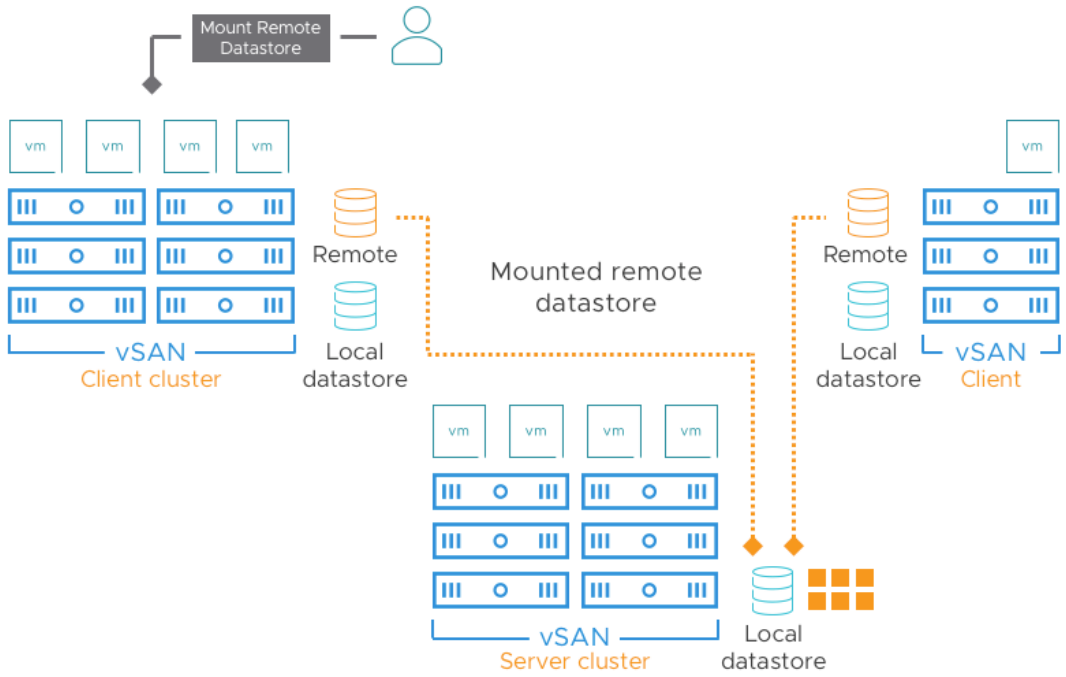

Figure 57: Architecture of HCI Mesh Cluster [52]

5.9.4.3 vSAN DATA SERVICES

Space Efficiency: This feature helps in the reduction of the total cost of ownership of (flash) storage by the following technologies:

- *Deduplication and Compression*: Reduces the usage of physical storage by 7 times. VMs with similar data and OS benefit the most by Deduplication. Compression works great for files such as text, bitmap, and other program files. This feature only works on the all-flash configuration. [52]

- *Compression-only*: A subset of Duplication and Compression and is only used for the purpose of compression. It is a cluster-wide setting that is implemented at the capacity device level. This feature is carried out at the device level which helps in the reduction of failure domains when compared with Deduplication and Compression.

  The following figure illustrates the device level compression:

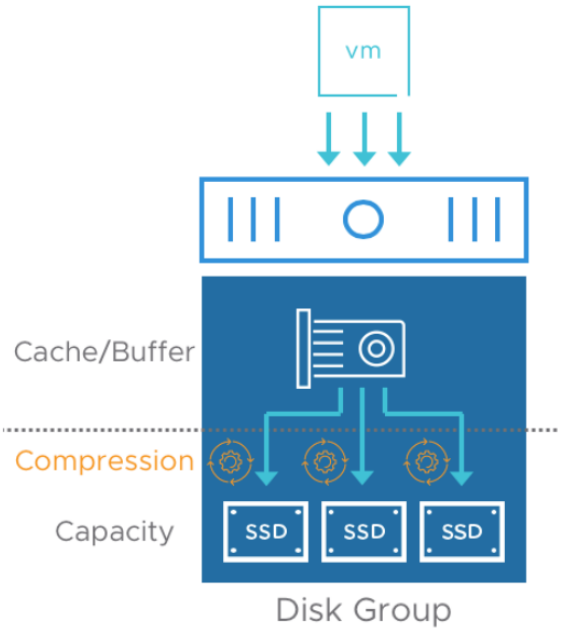

Figure 58: Device-level compression [52]

- *Erasure Coding*: Another space efficiency feature for all-flash configuration. This feature provides redundancy that is equivalent to mirroring with reduced capacity usage. Erasure coding is a methodology that involves breaking data into multiple parts and dispersing it across multiple devices. Parity is then added to the data so that recreation will be possible in case of pieces get corrupted or lost. [52]

  The following figure shows how redundancy is created using erasure coding:

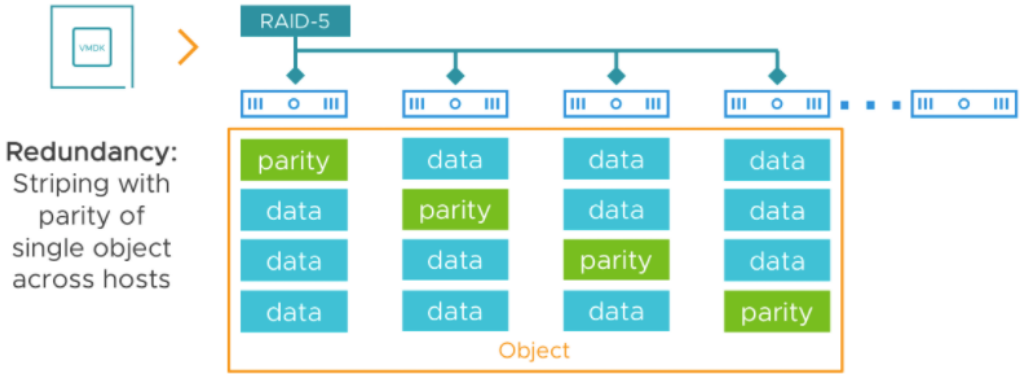

Figure 59: Erasure Coding [52]

- *TRIM/UNMAP*: The TRIM/UNMAP are commands for ATA and SCSI protocols that help make OS more efficient with storage usage.

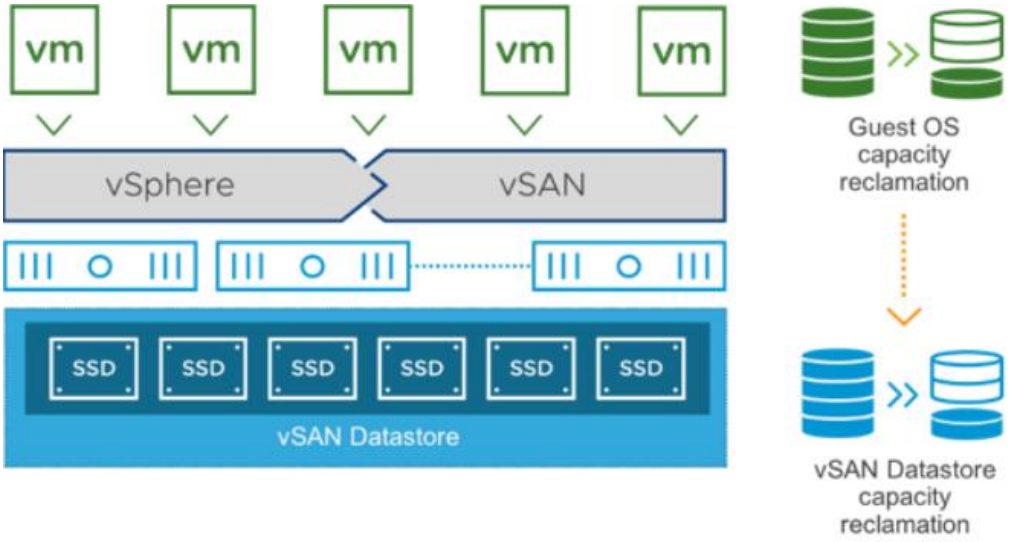

Enabling this feature makes vSAN aware of TRIM/UNMAP commands coming from operating systems and helps vSAN reclaim the unused storage. The figure above shows the mechanism of the TRIM/UNMAP methodology. [52]

- *iSCSI Target Service*: This service makes use of the iSCSI protocol so to provide Block storage to physical workloads. This simplifies the process and reduces capital costs incurred due to external storage arrays.

- *IOPS Limits*: vSAN has the capability to limit the IOPS that a VM may generate which may negatively affect the performance of other VMs. Limiting IOPS for specific VMs becomes advantageous to other VMs as it reduces the workload monopoly. The following figure displays the graph on the impact of IOPS limit to a VM:

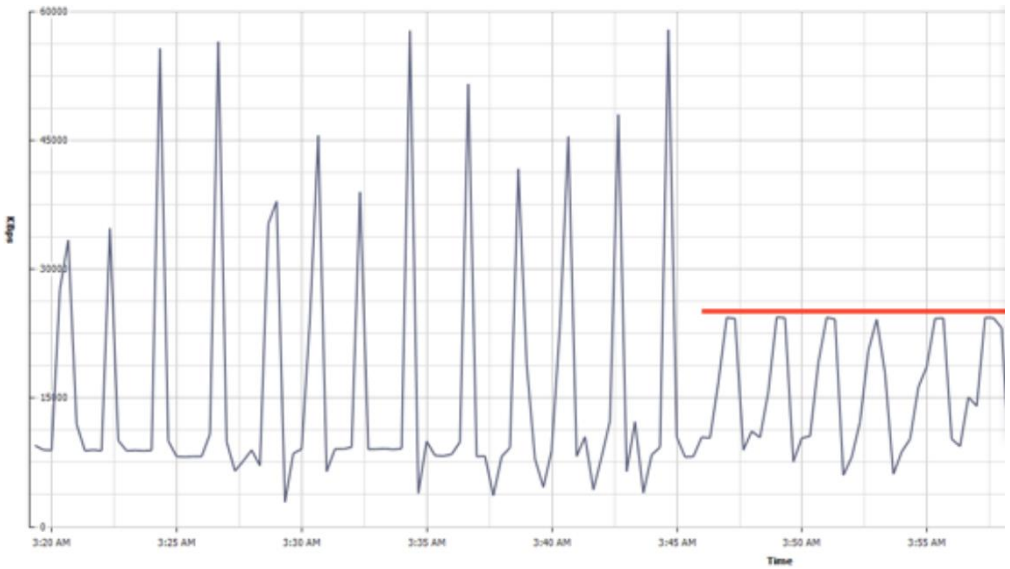

- *Native File Services (NFS)*: vSAN native file sharing service which supports NFSv3, NFSv4.1, SMB versions 2.1 and 3. When NFS is enabled, a container set is deployed for each host. The containers are responsible for provisioning the file system. A layer called Virtual Distributed File System (VDFS) is responsible for providing scalable filesystem by aggregation of objects in vSAN in the form of a file server and a control plane. This file share service is connected to vSAN Storage Policy management. It also has support for MS AD for authentication. [52]

  The following figure shows NFS architecture:

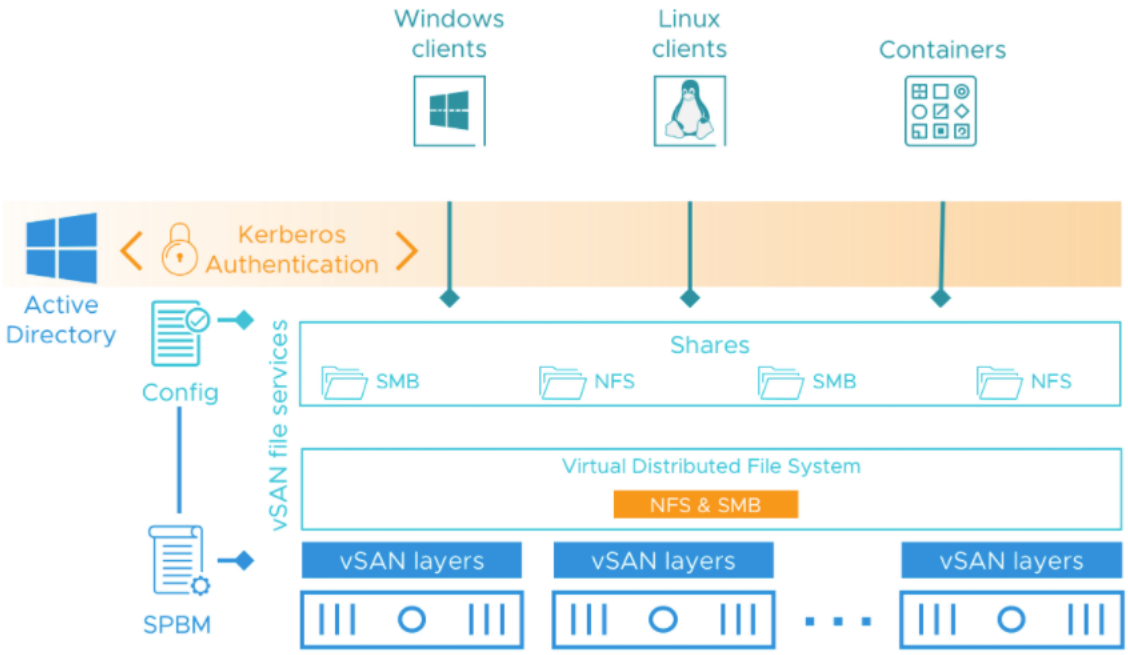

Figure 60: NFS architecture

5.9.4.4 VMWARE SECURITY

Native VMkernel Cryptographic Module: VMware makes use of a FIPS 140-2 validation which is a part of the Cryptographic Module Validation Program (CMVP) designed by NIST and Communications Security Establishment (CSE).

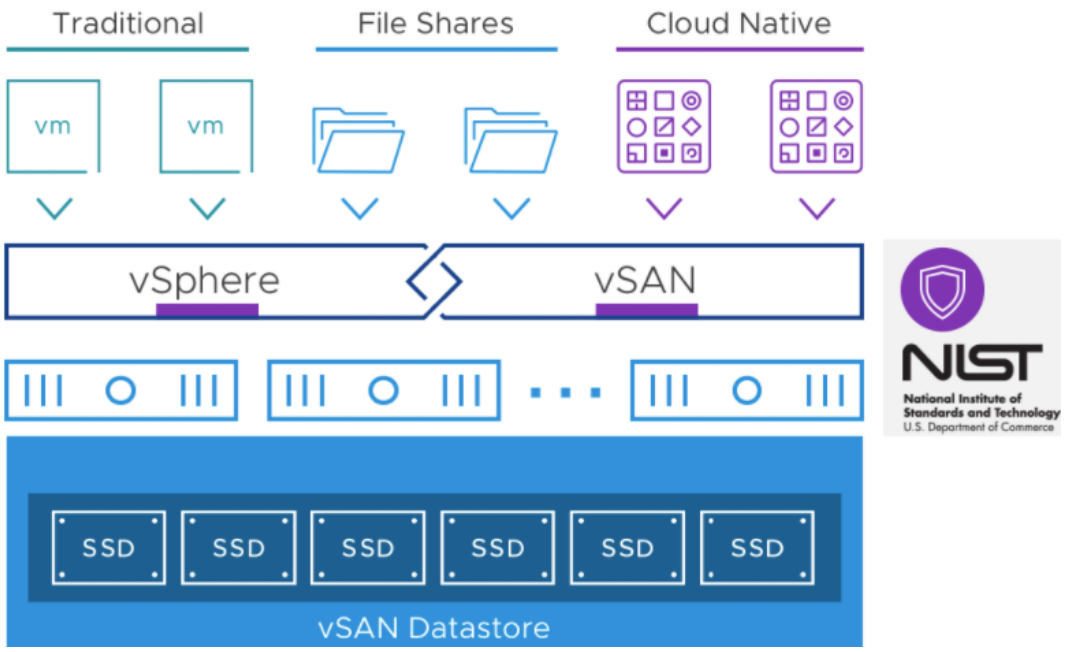

Figure 61: FIPS 140-2 certification has satisfied all the 11 requirement areas of Cryptographic Module Standards [52]

FIPS 140-2 certification has satisfied all the 11 requirement areas relating to design and implementation set by Cryptographic Module Standards. This module has been tested with algorithms and operational testing. An award certificate number 3073 is given to VMware by CMVP.

Key Management: For encrypting data-at-rest and VM encryption, key management is a core requirement. Key Management Solution makes use of Key Management Interoperability Protocol (KMIP) 1.1.

vSAN Encryption: vSAN carries out encryption for Data-at-rest and Data-in-transit using FIPS 140-2 cryptographic modules. [52]

.

5.9.5 NSX DATA CENTER

NSX is a core service that provides data networking and security functionality to vSphere inside the VCF platform. The networking is abstracted from the physical infrastructure where the VCF is deployed. It provides ease of control on managing data networks and applying security policies from a software-defined approach. Network virtualization helps to provide the same flexibility that is achieved using a VM, i.e. IT staff can create snapshots, save, move, or restore the Datacenter network configuration. NSX data center expands the network security functionality to applications and heterogeneous environments which is implemented across any cloud deployment models. [54]

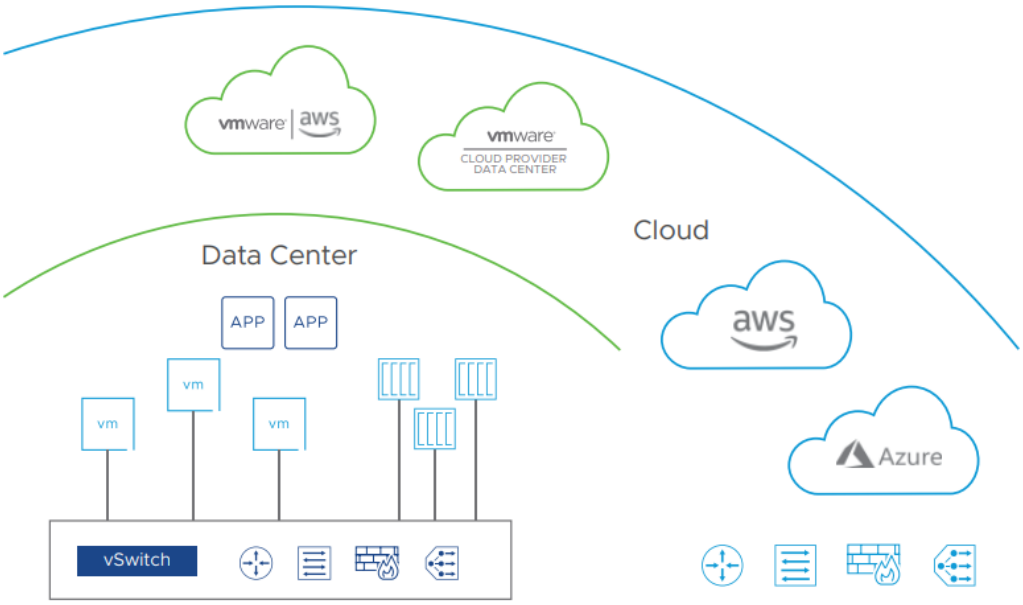

Figure 62: NSX can be implemented across multiple Cloud vendors so a good option for multi-cloud [54]

NSX provides core networking capabilities such as:

- Logical switching
- Logical routing
- Logical firewalling
- Logical load balancing
- Virtual Private Networks (VPN)
- Quality of service (QoS)
- Monitoring

All the stated services are non-disruptive to the underlying infrastructure which includes public cloud, private cloud, containers, and bare metal servers as these all are deployed in a virtual network.

Key Features of NSX

Table 2: Features of NSX [55]

Switching	NSX can enable logical Layer 2 extension across a Layer 3 network. And supports the use of VXLAN and GENEVE-based networks.
Routing	Enables use of dynamic routing in the hypervisor kernel between virtual networks and is also capable of providing scale-out routing with active-active failover. Both static and dynamic routing protocols are supported. IPv6 can also be configured.
Gateway Firewall	Provides the capability of stateful firewalling including layer 7 which includes identification of application and distributed FQDN allow listing. It is distributed across the entire network with centralized management and policy.
Distributed Firewall	Capable of integrating with cloud-native platforms which include Kubernetes and other cloud service providers such as Azure and AWS. Also provides the capability of stateful firewalling including layer 7 which includes identification of application and distributed FQDN allow listing.
Load Balancing	NSX provides Layer 4 – 7 load balancing capabilities with SSL pass-through offloading. Provides server health checks and application rules.
VPN	Provides site-to-site VPN, unmanaged VPN, and remote-access capabilities.
NSX Gateway	It is a gateway for bridging VLANs in the network.

NSX Intelligence	NSX Intelligence can provide automated recommendations for security policies. It continuously monitors the network traffic flow which enables auditing the security posture.
NSX Distributed IDS/IPS	NSX has a purpose-built threat detection engine to prevent lateral threat movement in traffic. Replaces the use of separate applications and helps achieve regulatory compliance.
Federation	Provides simplicity in enforcing centralized policies across the entire network.
Virtual Routing and Forwarding (VRF)	NSX supports the use of VRF which helps in the separation of data plane containing separate routing table and firewall.
NSX Data Center API	Supports use of RESTful API for integrating with cloud and DevOps tools.
Operations	NSX makes use of native operations such as CLI, Traceflow, SPAN, and IPFIX for troubleshooting and monitoring the virtual infrastructure.
Context-Aware Micro-Segmentation	NSX provides the ability to dynamically and automatically create security groups and policies that can also include elements such as tags, OS, and machine names for better micro-segmentation on top of attributes such as IP addresses, ports, and protocols.
Automation and Cloud Management	NSX is well integrated with vRealize Automation Cloud, OpenStack, and more.

Third-Party Partner Integration	NSX supports 3rd part vendor firewalls, switching, IDS, IPS, and more with their management, control, and data plane
Multi-Cloud Networking and Security	Allows uniform networking and security implementation across different cloud platforms and is not dependent on the underlying topology.
Container Networking and Security	The NSX has support for Kubernetes/Cloud Foundry application instances and Kubernetes Networking policy. It is well integrated with the VMware Tanzu Kubernetes grid.
VMware Carbon Black Technology	Provides protection for workloads using Real-time Auditing and Next-Generation Antivirus (NGAV). [56]
NSX Advanced Load Balancer with Web Application Firewall	Makes use of automation and deep understanding of applications to prevent and lower the chances of false positives from occurring. [56]

5.9.6 VREALIZE SUITE

vRealize Suite consists of three major components which are as follows:

vRealize Log Insights: It is a log management system that is highly scalable. It contains dashboards, analytics, and other third part extensibilities. It can provide in-depth operational visibility and quicker troubleshooting of physical, virtual, and cloud infrastructure.

vRealize Log Insight

Improve troubleshooting and security with centralized log management, visibility and analytics.

vRealize Automation: It is a cloud automation tool that helps speed up the deployment of IT services through pre-defined policies. This provides developers with great agility and flexibility. This also allows the IT team to govern and control seamlessly.

vRealize Automation

Increase agility, productivity and efficiency through self-service automation.

vRealize Operations Manager: This component provides intelligent operations management with visibility involving application to storage in the virtual, cloud, or physical environment.

vRealize Operations

Optimize, plan and scale SDDC and cloud deployments with self-driving operations.

vRealize Suite Lifecycle Manager: Helps in delivering a complete application lifecycle management. [57] [58]

5.10 MICROSOFT AZURE STACK PORTFOLIO

5.10.1 OVERVIEW AZURE STACK PORTFOLIO

Azure Stack portfolio is an extension to Azure public cloud so as to allow consistency between them. It allows organizations to run and build hybrid applications on their data centers on-premise, edge locations, and cloud. The following figure shows the Azure Stack portfolio:

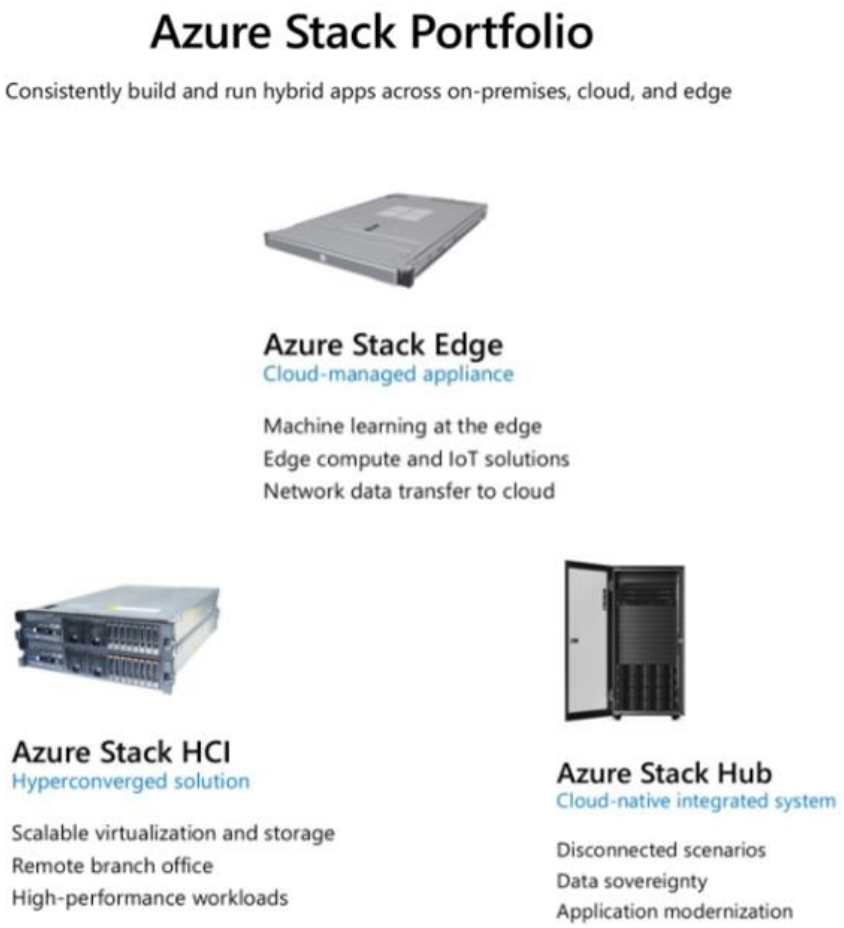

Figure 63: Azure Stack Portfolio [59]

The portfolio consists of three major services, Azure Stack Edge, Azure Stack HCI, and Azure Stack Hub. They are explained in the upcoming sections of this report. [59]

Both Azure Stack Hub and Stack HCI make use of industry-standard hardware on-premise and both have Hyperconverged computing, storage, and networking capability. Azure Stack Hub has additional features which allow it to run cloud apps on-premises and as well as the public cloud. It can also make use of Azure public cloud features such as Azure Portal, API, IaaS, PaaS, and other cloud platform admin tools.

The following figure shows the overlap in the features of the components:

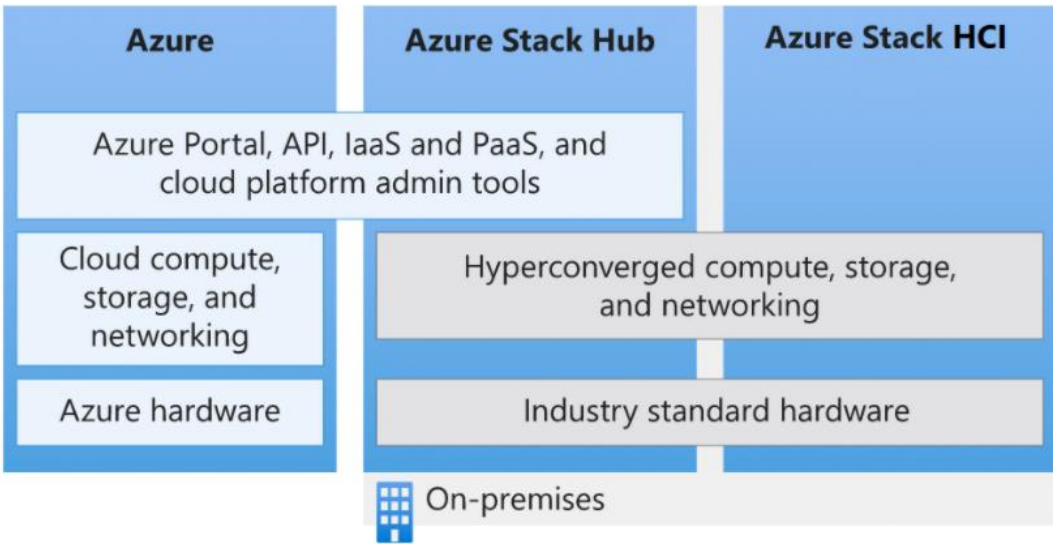

Figure 64: Comparing components of the portfolio [60]

5.10.2 AZURE STACK EDGE

Azure Stack Edge is a Microsoft-managed Hardware-as-a-service appliance that delivers services such as compute, storage, and intelligence to the edge. It is a subscription-based service with no upfront cost which delivers hybrid cloud services on-site. It supports services such as artificial intelligence (AI), Hyper-V VM Compute, FPGA compute, storage, Kubernetes, high availability, and GPU.

Figure 65: Rugged Edge Device [59]

Microsoft provides rugged edge devices that are read for harsh environments and are battery power.

Use cases:

- **Inference with Azure Machine Learning (ML)**: The edge devices can run ML models so to obtain fast results without sending data to the public cloud.
- **Preprocess data**: The data can be preprocessed before sending so to create a more actionable dataset.
- **Transfer on-premise data over the network to Azure:** Edge devices can be used to transfer data quickly for further computation or for archival purposes. [59] [61]

The following figure shows the architecture of Azure Stack Edge:

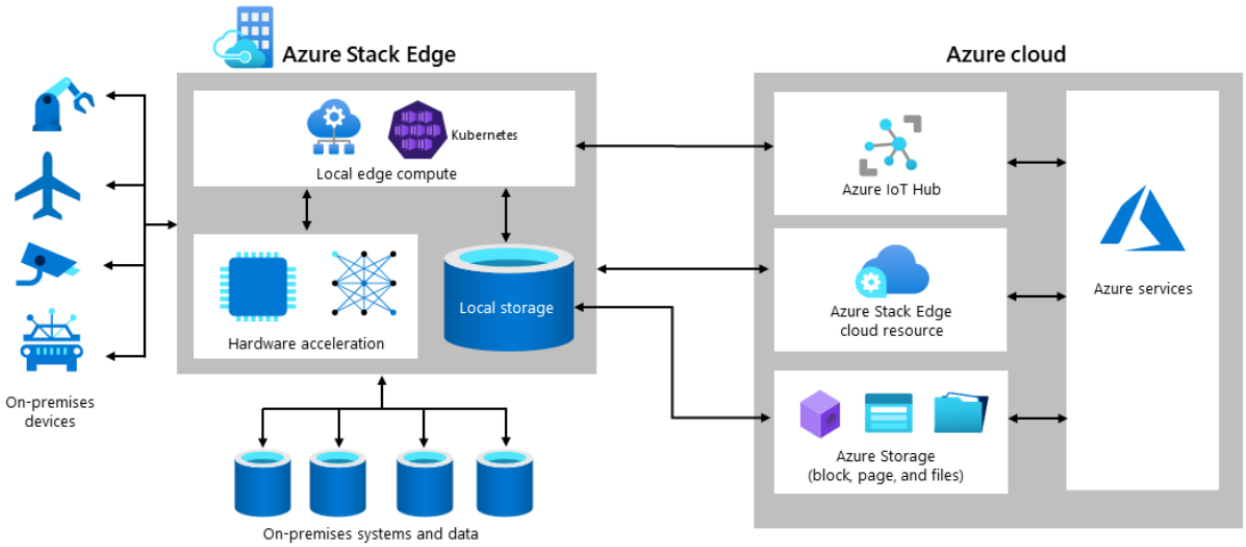

Figure 66: Architecture of Azure Stack Edge [61]

Table 3:Azure Stack Edge components [61]

Components	Description
Azure Stack Edge physical appliance	It acts as an edge location on-premise and functions as a network storage gateway. It offers functionality such as GPU and FPGA compute which allows an organization to accelerate AI inferencing. The local edge Compute makes use of containers that are orchestrated via Kubernetes. Container management can be carried out using Azure IoT Hub.
Azure Stack Edge resource accessible via the Azure portal	Helps with administration and monitoring of Azure Stack Edge appliances, which involves the management of local share hosting data that is being processed and transferred to public cloud storage called Azure Store.
Azure Stack Edge local web user interface (UI)	Helps with connecting the Azure Stack Edge appliance and facilitates initial installation. It is also used for the management of appliances such as restarting or copying logs.

5.10.3 AZURE STACK HCI

A Microsoft's Hyperconverged infrastructure cluster subscription-based technology that is capable of hosting Windows and Linux Hyper-V VM workloads, storage, Kubernetes clusters, and managing networking resources using hybrid cloud deployment model (a combination of hosting in local on-premise with validated partner hardware and public cloud). Built-in hybrid services allow cloud-based Monitoring, Site Recovery, Virtual machine backups, and a centralized view of all Azure Stack HCI (public and private) deployments through Azure portal connected through Azure Arc. It also has built-in support for standard tools such as Windows Admin Center and PowerShell.

The following figure displays how the Azure Stack HCI is structured:

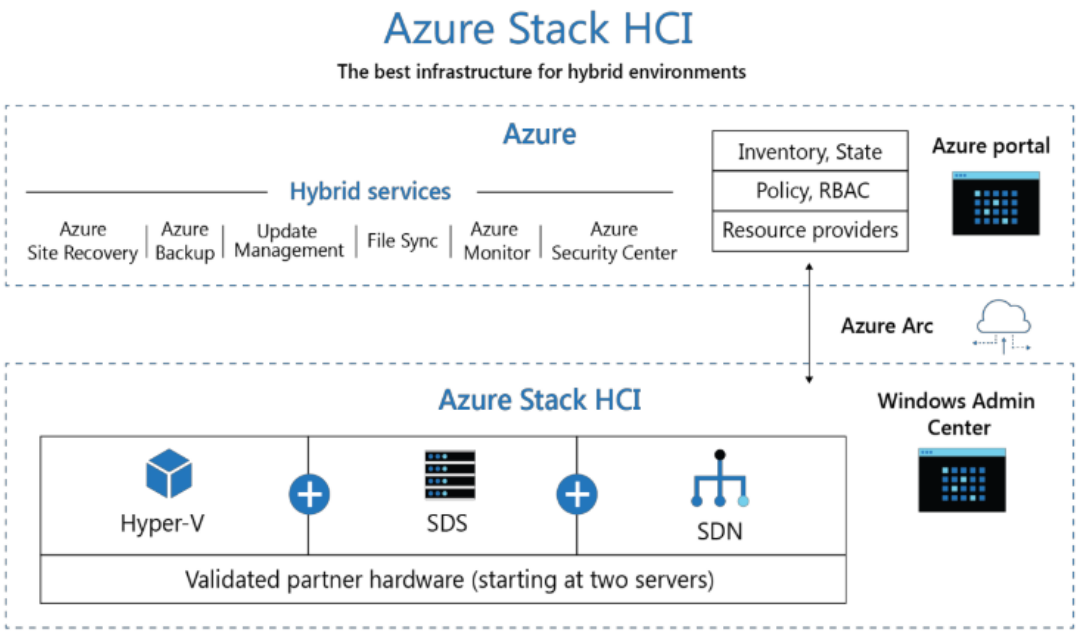

Figure 67: Azure Stack HCI Architecture [60]

Azure Stack HCI is significantly better than a simple Virtualization-based Windows Server and makes use of Storage Spaces Direct (SDS) which provides great storage price-performance. Storage Spaces Direct supports technologies like stretch clustering across sites for automatic failovers during disaster recovery. [60]
There are certain requirements that need to be met so to start an Azure Stack HCI:

- A cluster containing two or more server are required which are from Microsoft hardware partners.

- Azure subscription.

- A connection to the internet for each of the servers in the cluster via outbound HTTPS traffic. This connection should be made available to Azure endpoints at least every 30 days.

- Stretched clusters across sites should have a minimum link of 1 GB and average round trip latency of 5 ms for synchronous replication.

- For Software-Defined Networking, a virtual hard disk (VHD) will be needed in Azure Stack HCI to create Network Controller VM. [60]

5.10.3.1 AZURE STACK HCI HYBRID SERVICES

Azure Site Recovery: Help to provide high availability and disaster recovery as a service (DRaaS) by replicating the VMs to the public cloud. Reduces the cost of operating or owning a Hot site or Warm site. Azure Site Recovery manages the replication of VMs between Azure regions and On-premise VMs. It also helps to keep the desired recovery time objectives (RTO) and recovery point objectives (RPO). The value of RTO can go as low as 30 sec. The Azure Site Recovery service has a built-in mechanism to test the disaster recovery without disruption and the service complies with industry regulations such as ISO 27001. [62] [63]

Azure Monitor: AI-based centralized hub which helps to track apps and infrastructure using advanced analytics. This service aggregates and saves the telemetry into a log. Data collected is divided into metrics and logs. Metrics are lightweight numerical values that describe real-time scenarios. Whereas logs contain data that is organized into records with properties.
The following figure shows the architecture of Azure Monitor:

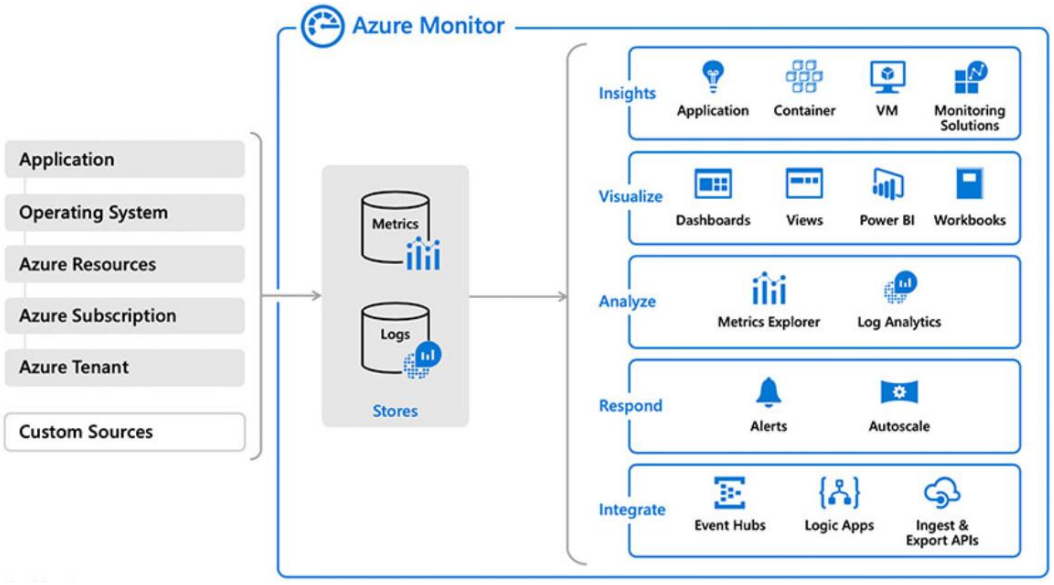

Figure 68: Azure Monitor [60]

Analyzed data from the log is used to generate insights, create alerts, display visuals in a dashboard or Power BI, and the logs are used by machine learning algorithms to identify and resolve problems. It can also be used for integrating applications using logics apps and more. [60]

The following diagram illustrates the failover mechanism:

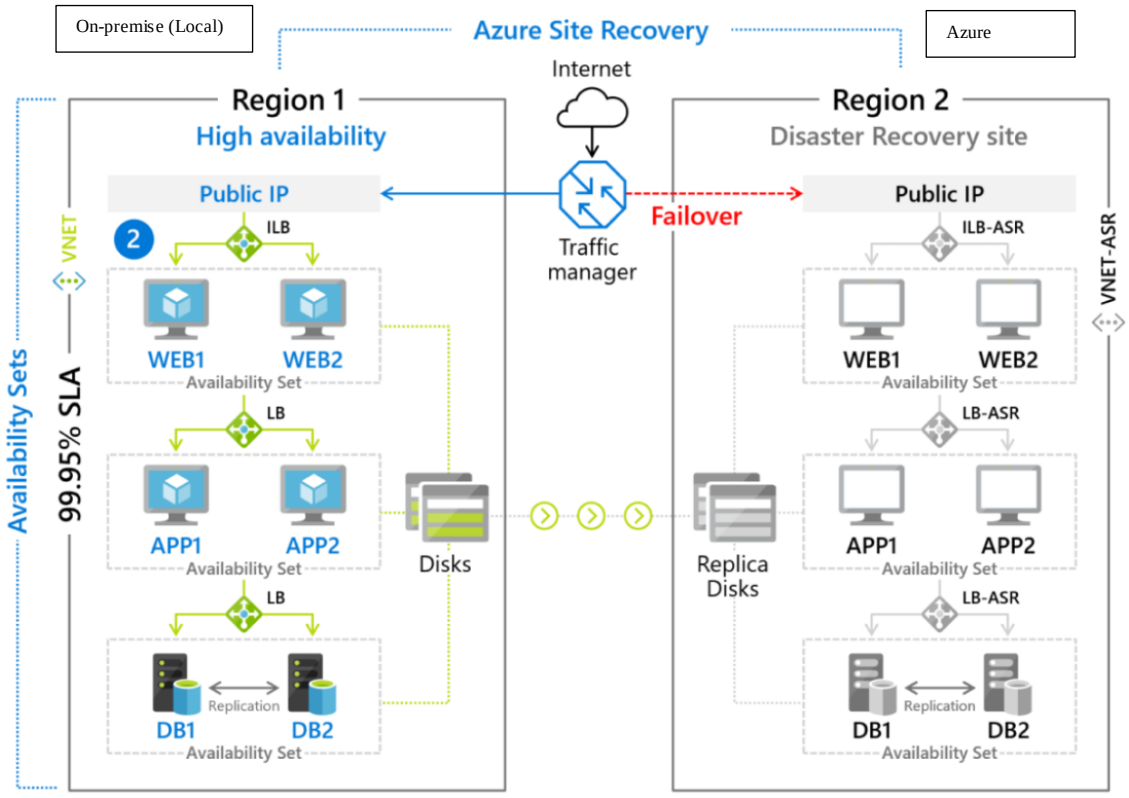

Figure 69: Azure Site Recovery in Public Cloud [62]

Cloud Witness: It is a service that is a lightweight tiebreaker for the cluster quorum. It is a replacement of 3rd datacenter used as a witness in multi-site stretched failover cluster quorum and hence reduces additional operational cost.

There is a total of 5 votes and only 3 are needed for the cluster to keep on operating, so a witness acts as 1 extra when either of the clusters fails. Cloud Witness makes use of Azure (Public Cloud) as the arbitration point and makes use of Azure Blob Storage in case of split-brain resolution. It is a good solution as most organizations do not have a third Datacenter that can host a File Server backing the File Share Witness. [64]

The following figure displays the Cloud Witness topology that can be deployed in Azure Stack HCI:

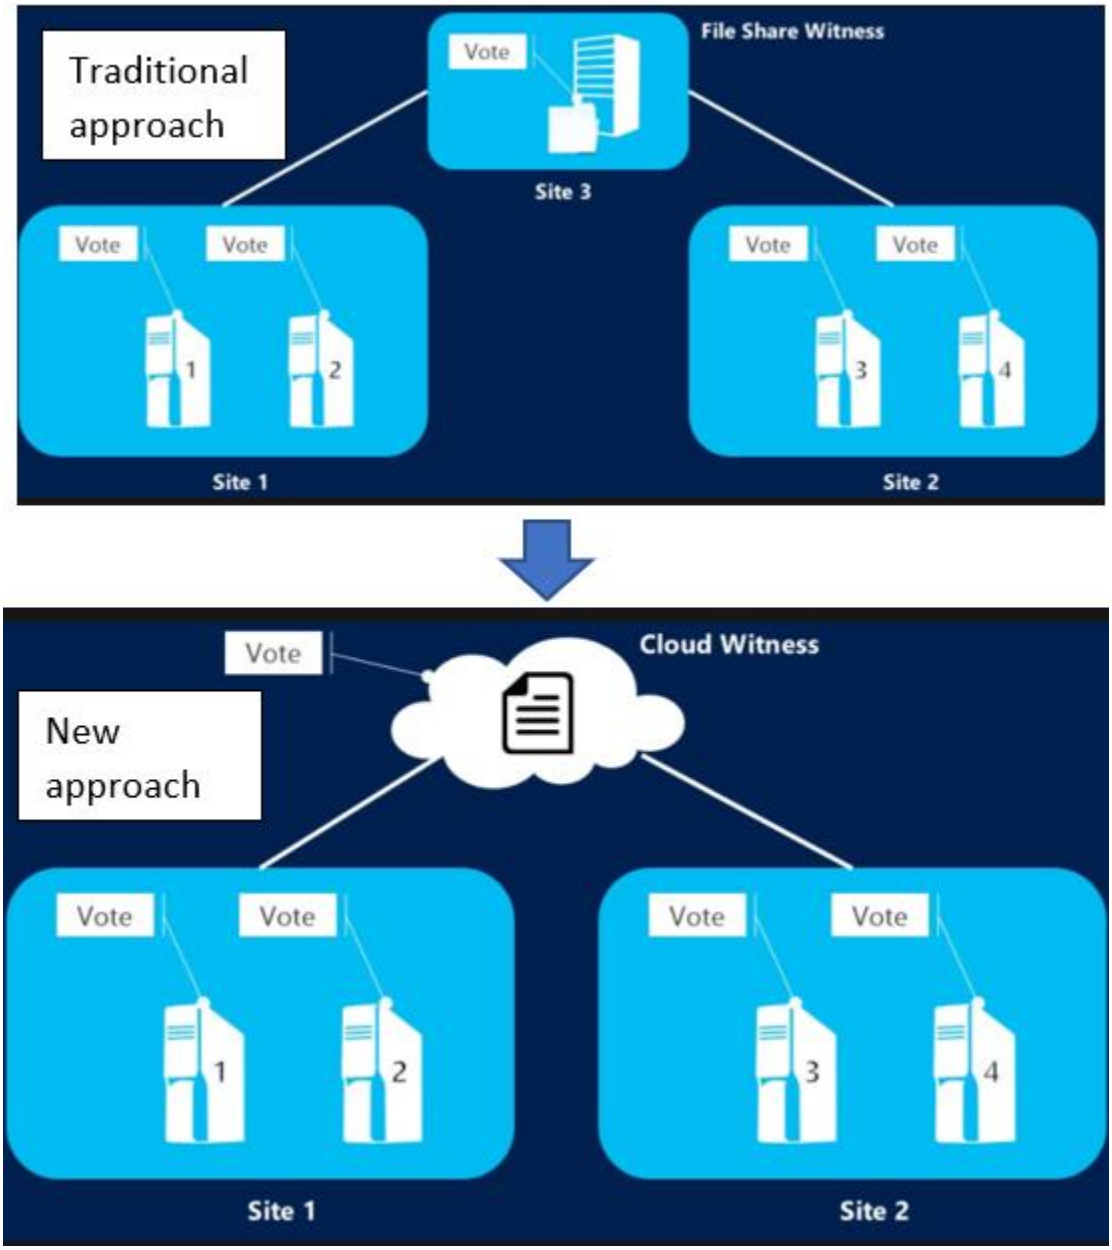

Figure 70: Multi-site stretched clusters with Cloud Witness as a quorum witness [64]

- **Azure Backup**: Provides backup of files folder and system state on the public cloud. Also offers offsite data protection and protection against ransomware attacks. Azure Backup makes use of Microsoft Azure Recovery Services (MARS) agent for protecting the on-premise Hyper-V VMs, VMware, and other workloads. It has three types of replication for data or storage for high availability:

- **Locally Redundant Storage (LRS)**: Replicates the data 3 times and stores it in a Datacenter in the same region.

- **Geo-redundant Storage:** Geo-redundant Storage is a default configuration and is also recommended by Microsoft. This type stores data replicates into secondary regions.

- **Zone-redundant Storage**: This replication type stores data in the availability zones which provides zero downtime.

The following figure illustrates the way the Azure Backup Services operates:

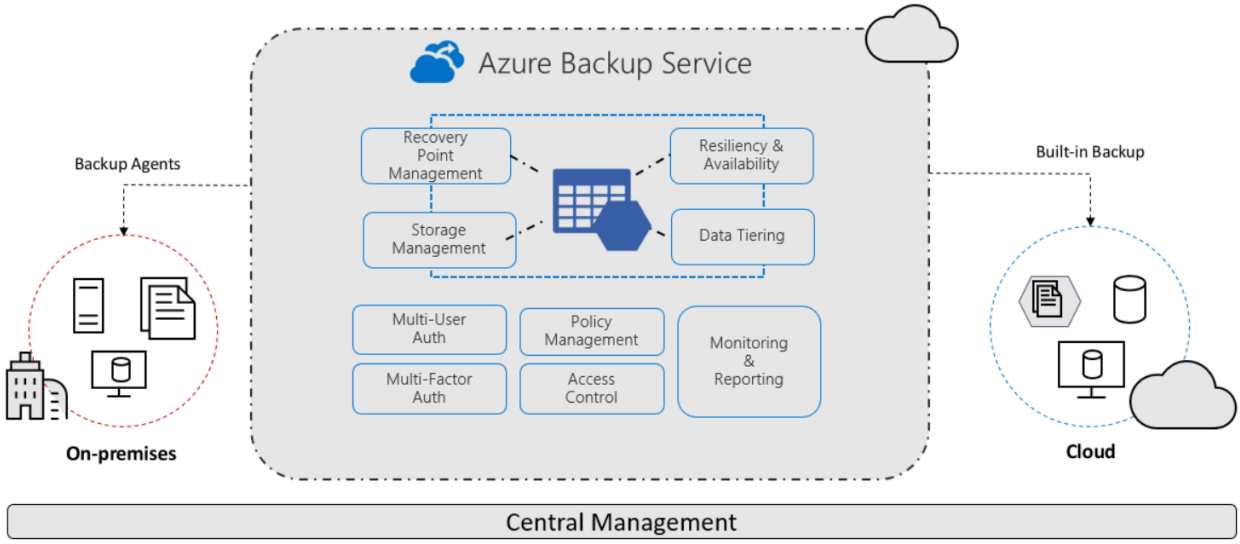

Figure 71: Azure Backup service [65]

- **Azure Update Management** carries out update deployment and assessment of VMs running Windows and Linux. It is a part of the Azure Automation service. It provides information regarding updates for the OS running on the VMs.

The following figure shows how Update Management operates to apply updates to OS in VMs:

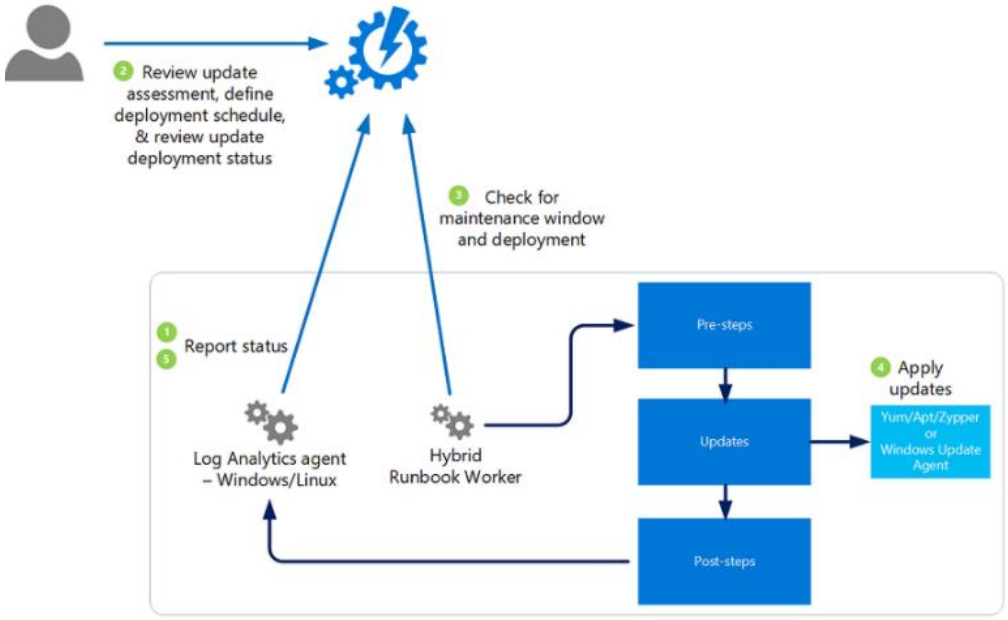

Figure 72: Update Management [66]

- **Azure Network Adapter**: Allows the use of point-to-site VPN. point-to-site VPN requires complex procedures to setup, certificate management, and even an infrastructure setup may be required such as an express route. Azure Network Adapter makes this procedure easy with a single click and automates Azure Virtual Network Gateway configuration and the VPN client. [67]

- **Azure File Sync** helps to sync the on-premise file server to the Azure public cloud.

- **Azure Security Center**: Azure Stack HCI can utilize an online Azure Security Center Azure for monitoring the VMs on-premise. It has built-in support for Azure Defender which provides threat protection for VMs running on-premise.
    - It detects any unusual attempts that are made to access the storage.
    - Scans containers for vulnerabilities and provides protection for Azure Kubernetes Service instances.
    - Protects Windows and Linux servers and clients. Scans for vulnerabilities in the VMs.
    - Helps monitor the state of security of hybrid cloud workloads from a single console.
    - Makes use of AI and automation using advanced analytics and machine learning to identify attacks and zero-day exploits.
    - Connects with existing tools such as SIEM (security information and event management) systems. [68]

5.10.3.2 WINDOWS ADMIN CENTER

Windows Admin Center provides administration in Azure Stack HCI. It is a browser-based tool that can be used to monitor and manage Windows server infrastructure and Azure HCI clusters. Windows Admin Center provides central management of Compute, network, storage, and other virtual resources. It can also administer features installed on machines present in the infrastructure. It can provide cluster-wide monitoring and alters, for example, CPU utilization, storage capacity, IOPS, throughputs, and latency. It also provides SDN support which involves monitoring and management of virtual networks, subnets, and connection of VMs to virtual networks. A separate administration tool called Microsoft System Center 2019 can be used for functions that cannot be performed using Windows Admin Center such as managing heterogeneous networks such as VMware VMs and Linux servers. [69]

5.10.3.3 AZURE STACK HCI VIRTUALIZATION: HYPER-V

Hyper-V is the hypervisor and the hardware virtualization system that is used in Azure Stack HCI. It can run Windows or Linux OS. It is designed based on a thin, micro-kernelized Type 1 hypervisor that can run directly on the bare-metal server. The following figure provides the architecture of Hyper-V:

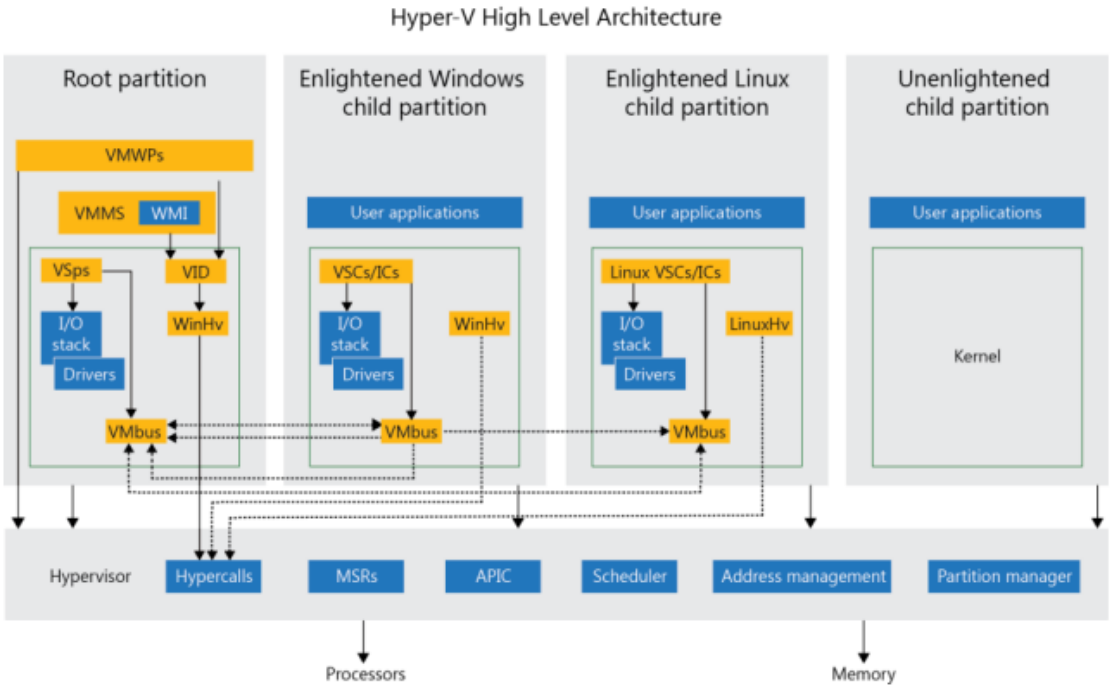

Figure 73: Hyper-V Software Architecture [69]

Figure 5. Hyper-V software architecture

Components that makeup Hyper-V are as follows:

Root partition: This partition runs the Hosts OS and is responsible for managing functions for the machine that include drivers, power management, and removable devices. This partition has access to physical memory and devices.

Child partitions: This partition is responsible for hosting the guest OS. It can access physical memory via Virtual Machine Bus (VMBus) or through the Hypervisor.

Virtual Machine Bus (VMBus): It is a communication channel for inter-partition commutation with multiple virtualized partitions.
Enlightened guest: An instance that is virtualization-aware of an OS and can communicate directly via VMBus.

Hyper-V Virtual Machine Management Service (VMMS): VMMS is responsible for the management of the state of VMs in child partitions.

Virtualization WMI provider: APIs exposed by VMMS that help with management and control of VMs.

Virtualization Service Provider (VSP): feature as a part of the root partition and helps provide support to child partition on a VMBus.

Virtualization Infrastructure Driver (VID): A driver that is responsible for providing partition management services, virtual CPU management services, and memory management services for partitions.

Virtualization Service Client (VSC): It is a synthetic device instance inside a VM. They are responsible for communicating with VSPs in parent partition through the VMBus for child partition's device I/O request.

Integration component (IC): Helps to communicate a child partition with other partitions and hypervisor.

Hypercall: Communication interface for the hypervisor. Provides access to optimizations provided by the hypervisor.

Advanced Programmable Interrupt Controller (APIC): Assigns priority levels to interrupt inputs.

Virtual Machine Worker Process (VMWP): This process is created for individual VMs to the management of services from OS in parent to guest partition.

Memory service routine (MSR): It is a process that the processor uses for reading or writing data to the memory.

Windows Hypervisor Interface Library (WinHv): It is a bridge designed for drivers for OS in guest partitions to communicate with the hypervisor by using standard Windows calling conventions.

Linux Hypervisor Interface Library (LinuxHv): It is a bridge designed for drivers for OS in guest partitions to communicate with the hypervisor by using standard Linux calling conventions. [69]

5.10.3.4 AZURE STACK HCI NETWORKING

Software-Defined Networking (SDN) is an optional service that can be used in Azure Stack HCI. It is a service that is capable of centrally segmenting, configuring, and managing the data network via software. This service can be used to design and provide network functionalities using virtualization of the following components:

- Routers
- Switches
- Firewalls
- Load Balancers

The SDN technology in Azure Stack HCI is the same technology that governs SDN in the Azure public cloud. There are some Windows Server roles that are utilized for SDN in Azure Stack HCI which is as follows:

- **Network Controller**: It is the server inside the SDN stack that acts as a "brain" that can be programmed and hence allows the IT team to managing, configuration, monitoring, and troubleshooting virtual and physical infrastructure for the Datacenter. By using Windows PowerShell, REST API, and a management application, components like Hyper-V virtual machines, virtual or physical switches, routers, firewalls, VPN gateways, and load balancers can be managed via the help of a Network controller. Network policies can be defined using JavaScript Object Notation (JSON) which is passed to the Network Controller via RESTful API. These policies are implemented on Hyper-V hosts virtual switch through SDN host agent.

- **Network Virtualization**: This is a server role that allows abstraction to separate the VM workloads from the physical network. This role provides a means for multitenant isolation on a shared IT infrastructure and hence improving security.

- **Remote Access (RAS)**: It is a server role that has a RAS gateway component, helps in routing the network traffic between the physical and virtual network resources.

- **Software Load Balancer (SLB)**: Help in evenly distributing the network traffic in Azure Stack HCI among different virtual networks. This allows to host multiple servers running the same workload, then provides high availability and scalability. SLB makes us of the multiplexer (MUX) that provides a virtual IP (VIP) to the external network via Border Gateway Protocol (BGP) and allocates dynamic IP (DIP) addresses among the VMs deployed in the network. [69]

The following figure illustrates the SDN architecture in Azure Stack HCI:

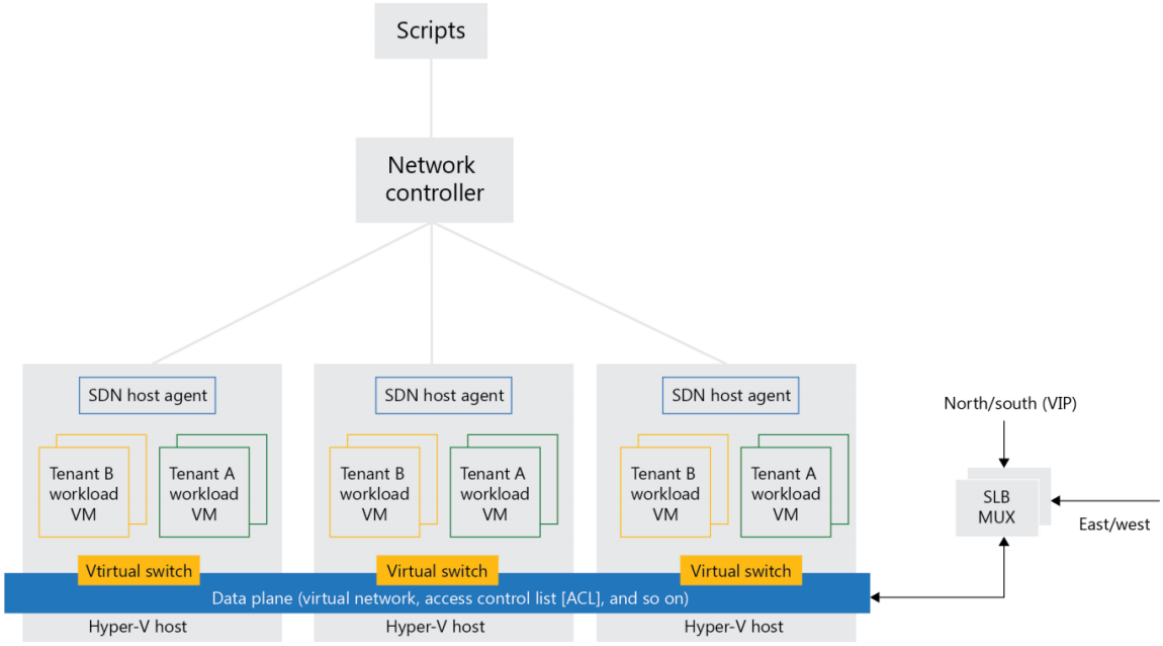

Figure 74: SDN Architecture in Azure Stack HCI [69]

- **Remote Direct Memory Access (RDMA) Support**: Azure Stack HCI supports the use of RDMA as it reduces network latency and increases throughput by significantly reducing CPU resources by bypassing the operating system processes. [70]

The following figure illustrates how RDMA can improve the performance of a network:

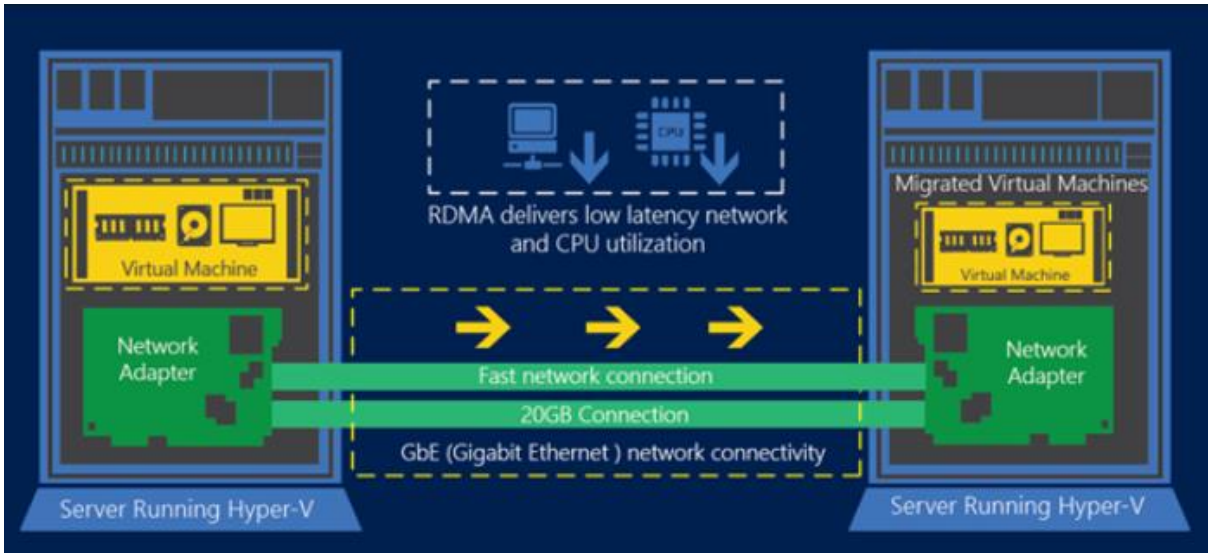

Figure 75: RDMA [71]

- **Dynamic Virtual Machine Multi-Queue (VMMQ) Support**: Makes use of multiple queues of modern network adapters to distribute the traffic on the NIC and prevents overloading of the CPU and, hence reducing latency.
    - Dynamic VMQQ optimizes the efficiency of the CPU of the host
    - Helps with automatically tuning of network traffic to the CPU cores so that VMs deliver expected throughputs
    - Bursty workloads are made to receive the right amount of traffic. [70]

5.10.3.5 AZURE STACK HCI STORAGE LAYER: STORAGE SPACES DIRECT

Storage Spaces Direct is defined by Microsoft as software-defined and shared-nothing storage. This technology makes use of standard servers that are used in the industry and it combines the locally attached drives to form a pool that has high availability and scalability and it lowers the cost of ownership as expensive SAN and NAS arrays can be avoided. It also provides easy management and provides a performance boost by making use of features such as remote direct memory access (RDMA) networking and persistent read/write drives. [69]

Storage Spaces Direct boost performance by allocating and distributing the fastest NVMe drives for reading and writing cache for each node in an Azure Stack HCI cluster as shown in the figure below:

Figure 76: Storage Spaces Direct NVMe drives for fast read/write caching [69]

In Storage Spaces Direct, multiple features are directly pulled from Windows Server which includes failover clustering, Cluster Shared Volume (CSV) file system, Server Message Block (SMB) 3, and Storage Space. A new addition to this technology includes Software Storage Bus which allows servers visibility of disks attached to each node in a cluster. [69]

Components of Storage Spaces Direct are as follows:

- Networking hardware: 10 Gbps or faster speed links are used between servers and Storage Spaces Direct makes use of SMB3 for communication between them.

- Storage Hardware: Two to sixteen servers containing local storage which can contain any SATA, SCSI, or NVMe but each server should have a minimum of 2 SSD drives out of the minimum total of six drives.

  The following figure displays the architecture behind Storage Spaces Direct:

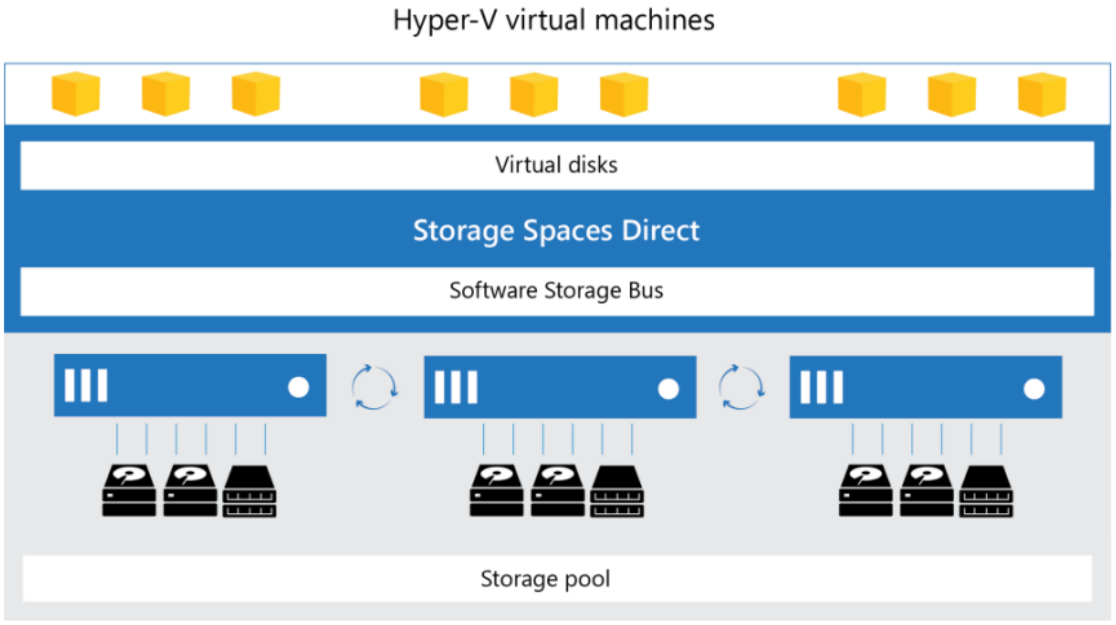

Figure 77: Architecture behind Storage Spaces Direct [69]

- Failover Clustering: It is a Windows Server built-in clustering feature that is which is used for connecting the server.

- Software Storage Bus: It is a software-defined storage fabric that spans the cluster and helps servers with the visibility of other attached local drives.

- Storage Bus Layer Cache: helps with dynamically binding of fast drivers such as SSDs to the drives with slower speed for example HDDs. This provides faster read/write operation as SSDs are used for caching.

- Storage Pool: When multiple drives joint together to form Storage Spaces Direct.

- Resilient File System (ReFS): It helps in acceleration in operations such as creating, expanding, and merging checkpoints in a .vhdx file. It also provides a built-in checksum for the detection or correction of an error in data. It also supports real-time rotation of data between Hot and Cold storage based upon the usage. [69]

5.10.4 AZURE STACK HUB

Azure Stack Hub is an Azure for private on-premise or local cloud environment and can do all that Public Azure cloud services can do.

It is capable of running:

- VMs
- App service web apps
- API apps
- Functions (Serverless Compute)
- Databases such as SQL and MySQL
- Containers
- Event Hubs
- IoT Hubs
- Service Fabric clusters
- Kubernetes clusters

The following figure shows the overview of Azure Stack Hub architecture:

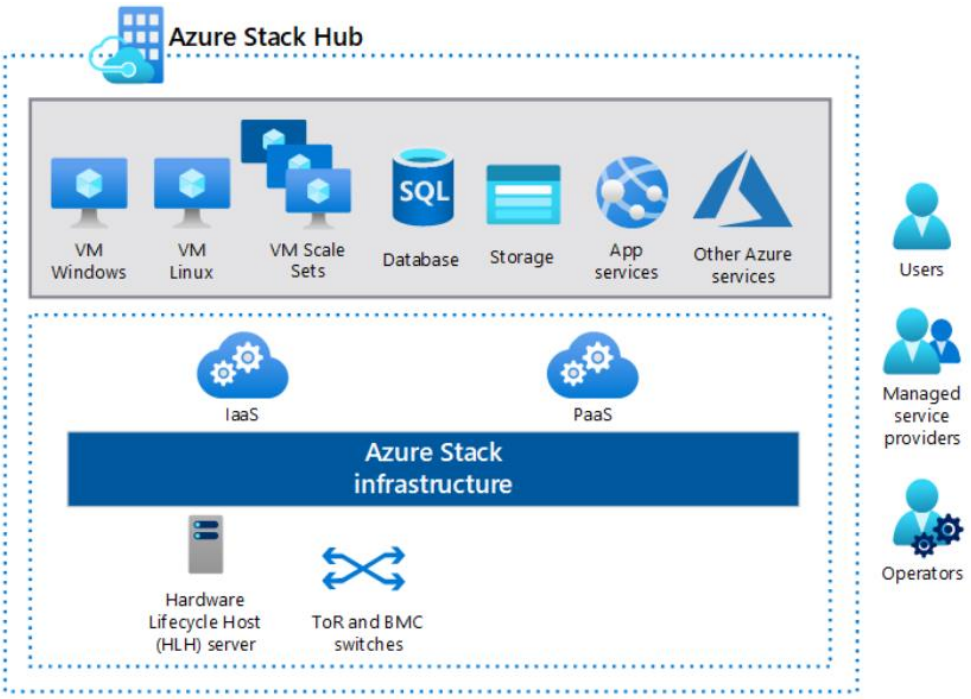

Figure 78: Azure Stack Hub architecture [72]

Azure Stack Hub does not work on custom infrastructure, it relies on preconfigured systems in racks of 4-16 servers (called scale unit) which are available from hardware vendors validated and certified by Microsoft only. The operating system that operates these servers is inaccessible and only the interfaces required to operate and implement workloads are accessible. The isolation of this architecture from the public network helps with satisfying regulatory requirements. [72]

The following figure shows the architecture of scale unit deployment in Azure Stack Hub:

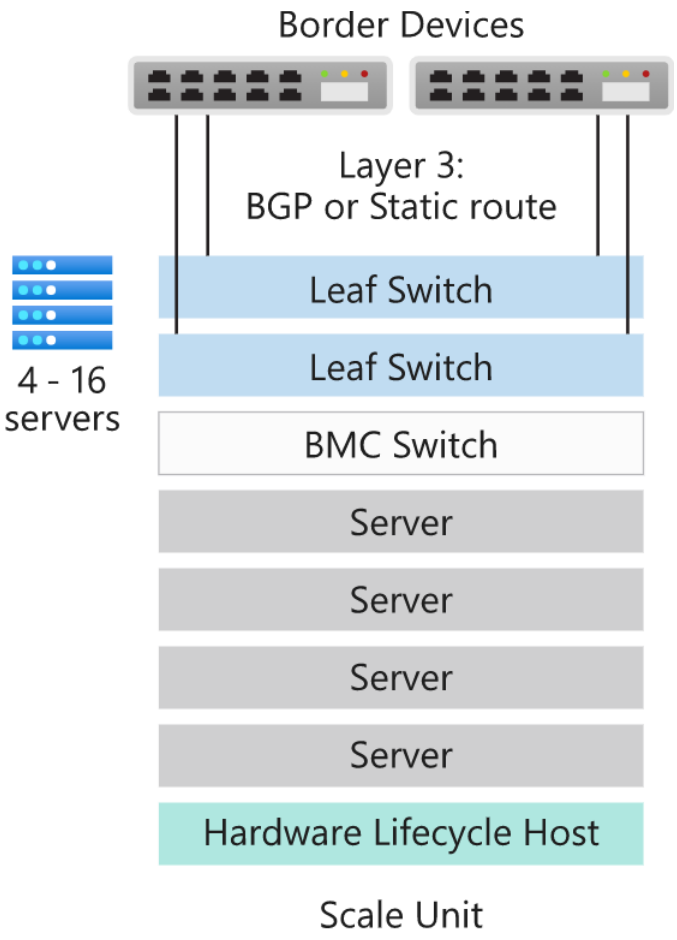

Figure 79: Azure Stack Hub integrated system [73]

Azure Stack Hub can make use of Azure Active Directory (Azure AD) as well as Active Directory Federation Services (AD FS). Since Azure Stack Hub is designed for disconnected private cloud deployment so in that scenario only AD FS. [73]

5.10.5 AZURE STACK RESOURCE PROVIDERS

IaaS and PaaS services that run on Azure Stack Hub are fundamentally dependent on the web services that are called resource providers.

Microsoft has divided the components that provide services as follow:

5.10.5.1 FOUNDATIONAL RESOURCE PROVIDERS

Compute Resource Provider: It can be used for providing an ability to create VMs and VM extensions in Azure Stack Hub.

Network Resource Provider: It can provide software-defined networking (SDN) and Network Function Virtualization for Azure Stack Hub. Resources like load balancers, public IP addresses, network security groups, and virtual private networks can be deployed using Resource Provider.

Storage Resource Provider: This Resource provider is responsible for delivering Azure-consistent storage services as follows:

- Blob storage: It is an object storage solution provided by Microsoft which helps in storing a large amount of unstructured data.

- Queue storage: It is used for storing messages so that they can be processed asynchronously from a store list. The figure shows an example of a Queue Storage: [74]

Figure 80: Queue Storage [74]

- Table Storage: This storage is a part of Azure Cosmos DB and is used for storing non-relational structured data (NoSQL). It is fast, efficient, and cost-effective than traditional SQL data storage. The following figure displays the Table Storage components: [75]

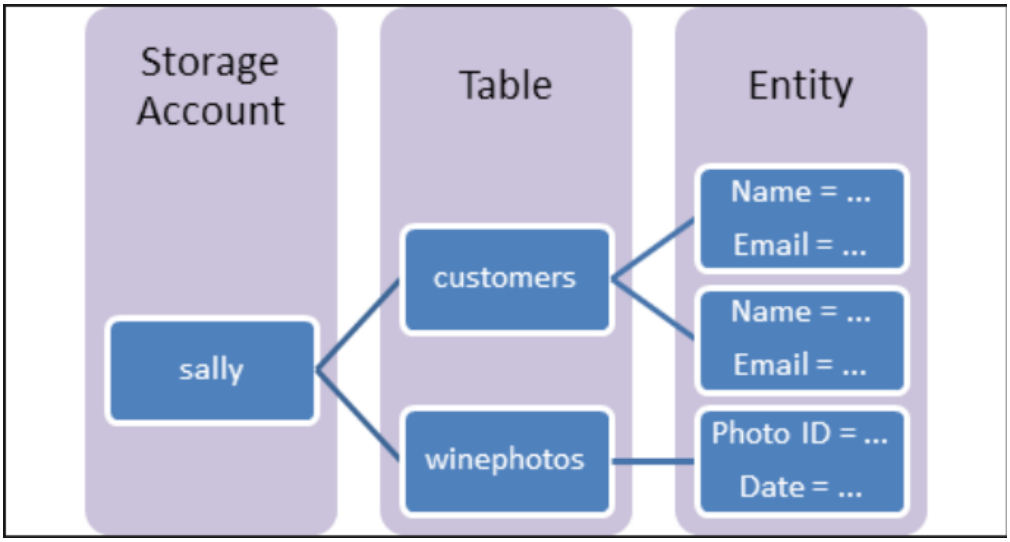

Figure 81: Azure Table storage [75]

- Key Vault: Key is responsible for secrets, key, and certificate management as follow:

    - Secrets Management: It allows Key Vault to store and control access to passwords, tokens, certificates, API keys, and more.

    - Key Management: Simplifies encryption key creation and control.

    - Certificate Management: Help and simplifies provisioning, management, and deployment of public and private Transport layer Security (TLS) based certification. [76]

5.10.5.2 OPTIONAL RESOURCE PROVIDERS

Other than the foundational resources, Azure Stack Hub offers PaaS resource providers ad follows:

- App Service: Allows the creation of the web, API, and Azure Functions for any platform for internal and external customers. It can help integrate the user applications with on-premise applications and help automate the business processes. The Azure App Service provides some key services as follows:

    - Multiple programming languages and frameworks: Support for ASP.NET, Node.js, Java, PHP, and Python. Also, support the use of PowerShell and other scripts.

    - DevOps optimization: Helps with setting up CI/CD (continuous integration and continuous development) using GitHub, local Git, and BitBucket. It also helps with the promotion of updates via testing, staging environment, and management of apps via CLI (command-line interface) tools such as PowerShell.

    - Visual Studio integration: Programming IDE such as Visual Studio are well integrated to program and deploy applications.

    App Services can help in the deployment of many types of applications that may require a specific type of workload:

    - Web Apps: Used for hosting web pages and applications
    - API Apps: Used for hosting RESTful APIs.
    - Azure Functions: Used for hosting serverless computer-based workloads. Azure functions allow Azure Stack Hub to run or execute a small part of

code without dealing with the configuration of the VMs or underlying infrastructure. [77]

- SQL Server: This resource provider allows Azure Stack Hub to run SQL databases as a service. Databases can be created for cloud-native applications, websites, and workloads.

- MySQL Server: This resource provider allows Azure Stack Hub to run MySQL databases as a service. [78]

5.10.6 SECURITY CONTROLS FOR AZURE STACK HUB

The security posture of Azure Stack Hub is designed based on two fundamentals:

- Assume Breach: Azure Stack Hub has assumed that the system is already compromised and so the objective is to detect and limit the impact of the breach.

- Hardened by Default: The security features and configuration on Azure Stack Hub are enabled by default. [79]

5.10.6.1 DATA AT REST

Microsoft has implemented data encryption for data at rest using BitLocker. BitLocker makes use of 128-bit AES encryption by default and can be configured to 256- bit AES. It satisfies the compliance standards such as PCI-DSS, FedRAMP, and HIPAA. [80]

5.10.6.2 DATA IN TRANSIT

Azure Stack Hub makes use of Transport Layer Security (TLS) 1.2 to carry out communication between Infrastructure components. The encryption certificates are managed by the infrastructure itself. All external endpoints also support TLS 1.2. [81]

5.10.6.3 SECRET MANAGEMENT

Azure Stack Hub internal service accounts rotate passwords and secrets every 24 hours using group Managed Service Accounts (gMSA), which is a domain account managed by Domain Controller. Azure Stack Hub uses 4096-bit RSA keys in all internal certificates.

5.10.6.4 WINDOWS DEFENDER APPLICATION CONTROL

Azure Stack Hub utilizes a Windows Security feature called Windows Defender Application Control (WDAC). It helps in applying executables filters and allows authorized code to run on the Azure Stack hub infrastructure.

5.10.6.5 CREDENTIAL GUARD

A feature that is a part of Windows Server and is utilized in Azure Stack Hub to protect the infrastructure credentials against Pass-the-Hash and Pass-the-Ticket attacks.

5.10.6.6 ANTIMALWARE

Hyper-V Virtual Machines and all other components in Azure Stack Hub are protected using Windows Defender Antivirus. If the Azure Stack Hub is connected the updates to antivirus is applied daily and if disconnected the updates can be applied monthly as a part of Azure Stack Hub updates

5.10.6.7 SECURE BOOT

Secure boot is enforced in Azure Stack Hub during boot-up of Hyper-V VMs which prevents malicious attacks. [82]

5.10.7 AZURE ARC

Azure Arc makes it easy to carry out governance and management by providing a consistent Azure-based Multi-cloud (public and private on-premises) management and edge platform. It helps with the management of the entire Multi-cloud environment via a single tool called Azure Resource Manager. It also provides a means to manage resources such as VMs, Kubernetes clusters, and databases just the way it is managed on the Azure public cloud. Azure Arc enables a business to utilize their traditional ITOps and at the same time allows integration of DevOps for supporting cloud-native patterns in the Multi-cloud environment. [83]

The following figure displays the architecture of Azure Arc:

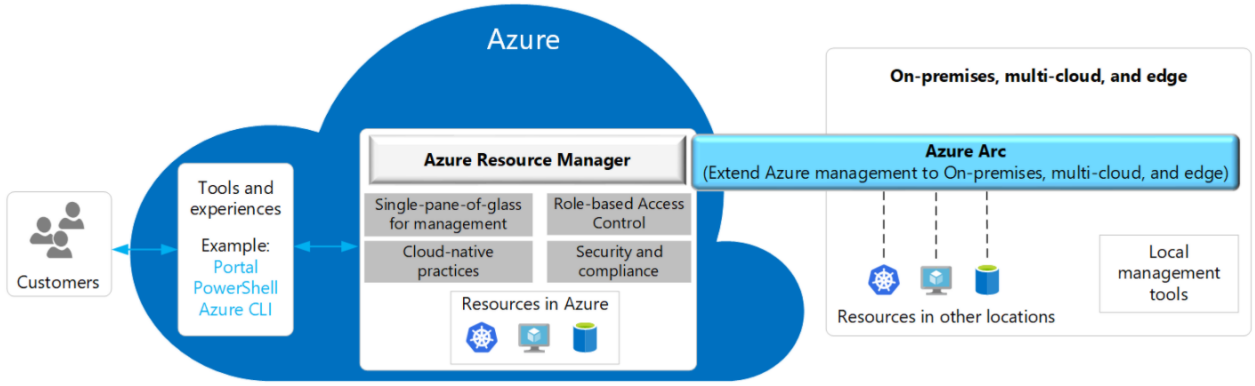

Figure 82: Azure Arc [83]

6.0 BENCHMARKING LOCAL AND HYBRID CLOUD TECHNOLOGIES

6.1 LATENCY

Since the focus of this report is on Local and Hybrid Cloud technologies, latency in the local cloud is dependent on the performance of the network and hardware installed whereas extension to Public Cloud via Hybrid Cloud technology is more dependent on the usage of WAN technologies such as SDWAN. By the analysis carried out, it shows Microsoft Azure Stack Portfolio outperforms as compared to the rest based on the inbuilt technologies.

The following table compares the how latency issue is tackled by each vendor:

Table 4: Benchmarking Latency

OpenStack	• It is an open-source cloud platform that can be used to design technologies for improved performance. • Storage Performance Development Kit driver can be used to design to run SSDs in polled mode, asynchronous and lockless- NVMe and RDMA. • There is no prebuilt technology like VMware Cloud Foundation and Microsoft Azure Stack Portfolio.
VMware Cloud Foundation	**vSAN Architecture:** • The hybrid approach of allocating drives of local storage of servers. • Makes use of shorter I/O paths and removal of storage virtual appliance • Makes use of IOPs Limits that reduces the monopoly by individual VMs **vSphere PVRDMA Support:** • Supports RDMA technology by making use of PVRDMA which reduces the network latency in the Cloud infrastructure.

Azure Stack Portfolio	**Azure Stack Edge:** • Makes use of offers functionality such as GPU and high-performance FPGA compute which allows an organization to accelerate AI inferencing and hence reduce latency at remote locations. The data can be preprocessed before sending so to create a more actionable dataset so to reduce latency. **Storage Spaces Direct (Azure Stack HCI):** • Provides performance boost by making use of features such as remote direct memory access (RDMA) networking and persistent read/write drives which reduces latency. • Boosts performance by allocating and distributing fastest NVMe drives for reading and writing cache for each node in Azure Stack HCI cluster and hence reduces latency **Azure Stack HCI supports the use of RDMA** • Reduces network latency and increases throughput by significantly reducing CPU resources by bypassing the operating system processes. The following figure illustrates how RDMA can improve the performance of a network **VMMQ Support** • Makes use of multiple queues of modern network adapters to distribute the traffic on the NIC and prevents overloading of the CPU and, hence reducing latency.

6.2 CLOUD SECURITY

Security is a key component of Cloud technologies and when Multi-cloud and Hybrid Cloud is integrated with Private Cloud, it can expose it to Public networks which exposes the IT infrastructure to multiple attacks. It was noticed that VMware Cloud Foundation had better built-in network security overall and OpenStack had a better built-in Identity service. Azure Stack HCI does provide advanced security monitoring but at the expense of exposing the network to public network as it requires the internet to operate advanced AI.

The following table displays what risk mitigation is in placed by the mentioned Cloud providers:

Table 5: Benchmarking Security

OpenStack	**Keystone** • Keystone is an identity service in OpenStack used for API client authentication and service discovery. It has provided built-in support for LDAP, OAuth, OpenID Connect, SAML, MFA, and SQL. • Supports all the top-of-the-line encryption technologies such as TLS 1.2 and TLS 1.3. • Horizon Dashboard is designed on the Django web framework for the deployment of best security practices. • Nova compute can make use of centralized logging for known security issues. • More security features can be programmed such as smart monitoring, alerting, and reporting using ELK stack, Elasticsearch, and more. **Data in-transit encryption**. • Supports all the top-of-the-line encryption technologies such as TLS 1.2 and TLS 1.3. **Image signature verification** • OpenStack prevents unverified VM images from the Glance service to run. Glance service carries out the verification by retrieving the certificates from the key manager.

	Data-at-rest Encryption • Provides volume encryption functionality for data at rest via Castellan (a generic Key Manager) and key creation can be requested by Cinder when a user decides to encrypt a volume. • Another data at rest feature called Ephemeral disk encryption allows VMs to encrypt data used as a temporary workspace. The vestigial remanent data may stay on the disk when a VM is unmounted resulting in a data privacy issue.
VMware Cloud Foundation	**Data in transit encryption** • Supports all the top-of-the-line encryption technologies such as TLS 1.2 and TLS 1.3. **Native VMkernel Cryptographic Module** • This module makes use of a FIPS 140-2 certification which is a part of the Cryptographic Module Validation Program designed by NIST and Communications Security Establishment. vSAN carries out encryption for Data-at-rest and Data-in-transit using FIPS 140-2 cryptographic modules. VMware has satisfied all the 11 requirement areas relating to design and implementation set by Cryptographic Module Standards. **NSX Security** provides the support for following: • Gateway Firewall: stateful firewalling with centralized management and policy. • Distributed Firewall: Capable of integrating with cloud-native platforms and applications. • NSX Intelligence: Monitors the network traffic flow and provides automated recommendations. • NSX Distributed IDS/IPS • Context-Aware Micro-Segmentation allows to automatically create security groups and policies.

	VMware Carbon Black Technology • Protects workloads using Real-time Auditing, Next-Generation Antivirus, Endpoint Detection & Response (EDR) solutions. **NSX Advanced Load Balancer with WAF** • Provides VMs and workloads with web application firewall and it makes use of automated learning so to reduce false positives.
Azure Stack Portfolio	**Data in transit encryption** • Supports all the top-of-the-line encryption technologies such as TLS 1.2 and TLS 1.3. **Windows Defender Application Control** • Provides executables filters and allows authorized code to run on the Azure Stack hub infrastructure. **Credential Guard** • Protects credentials against Pass-the-Hash and Pass-the-Ticket attacks on Azure Stack Hub **Antimalware** • All Azure Stack Hub components are protected using the Windows Defender Antivirus solution. **Network controls** • Simple network controls such as ACLs are supported by Azure Stack Hub on all rack switches, software-defined networks, and VMs. **Data at rest Encryption** • Uses specialized technology for encryption such as BitLocker for data at rest which complies with standards such as FIPS and HIPPA. **Azure Stack Hub (Private Cloud)** • Built-in FortiGate Next-Generation Firewall VM can be added to the Azure Stack Hub which provides solutions such as VPN and network virtual appliance.

	• The security posture of Azure Stack Hub is designed based on two fundamentals approaches that assume breach and are hardened by default. This helps them meet the major compliance standards set by the industry. • Secure boot is applied at bootup for all the VMs inside Azure Stack Hub. **Azure Stack HCI (Hybrid Cloud)** Azure Stack HCI can utilize an online Azure Security Center Azure for monitoring the VMs on-premise. It also has built-in support for Azure Defender which provides threat protection for VMs running on-premise. Azure Security Center can: • It supports BitLocker Drive Encryption for storage which complies with standards such as FIPS 140-2 and HIPAA. • Look for any anomalies and attempts that are made to access the storage. • Scan for vulnerabilities in containers and provides protection for Azure Kubernetes Service instances. • Scan for vulnerabilities in the VMs. • Help monitor the state of security of hybrid cloud workloads from a single console. • Use advanced technologies such as AI and automation using analytics and machine learning to identify attacks and zero-day exploits. • Supports tools such as SIEM systems.

6.3 SCALABILITY

In cloud computing, there are 2 types of scaling:

- **Vertical Scaling**: Involves the addition of or removal of vCPU, vRAM, or storage resources to the existing capacity of an instance based on the demand.

- **Horizontal Scaling**: Involves the addition or removal of instances to existing infrastructure based on demand.

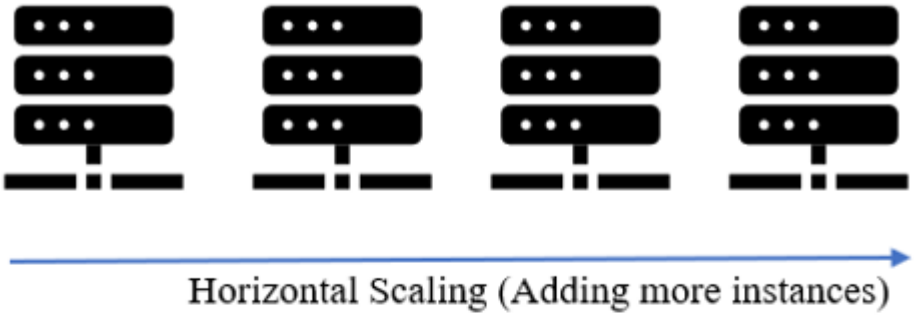

Horizontal Scaling (Adding more instances)

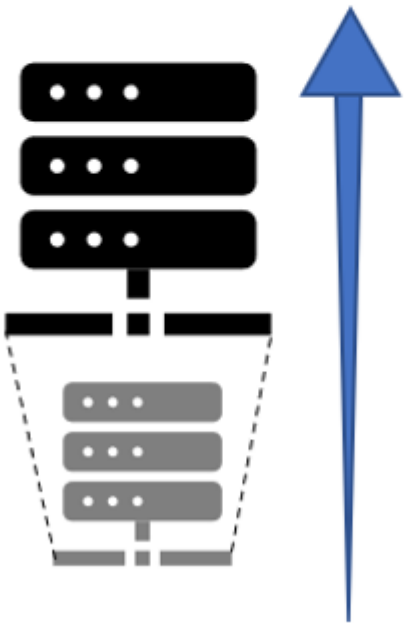

Vertical Scaling (Making the instances more powerful)

Figure 83: Vertical vs Horizontal Scaling

Capacity planning is the key to scaling the Private cloud infrastructure. But if the private cloud is integrated with public cloud such as Azure Arc, it makes use of the cloud bursting feature for auto-scaling. Azure Stack Portfolio used technology called Virtual Machine Scale Sets which was able to perform horizontal scaling. OpenStack named horizontal scaling as

Auto-scaling and Octavia. VMware Cloud Foundation made use of vSphere Distributed Resource Scheduler (DRS). DRS used vMotion and it was a vertical scaling technology. Horizontal scaling in VMware Cloud Foundation can only be utilized using the public cloud.

The following table compares scalability between the three vendors:

Table 6: Benchmarking Scalability

OpenStack	**Auto-Scaling** • OpenStack supports horizontal auto-scaling of clusters such as a Heat AutoScalingGroup and Senlin Clusters. **Octavia** • Open source load balancing and horizontally scalable service for Virtual machines, containers, and bare-metal servers deployed in OpenStack.
VMware Cloud Foundation	**vSphere Distributed Resource Scheduler (DRS):** • DRS is responsible for load balancing VMs based on available resources via vMotion. Additional virtual machines in a cluster are allocated to a pool of resources that are underutilized and hence increase the available capacity to the VM. This is more like vertical scaling of VMs by providing them additional CPU and RAM. • It can utilize Public Cloud (such as AWS) for horizontal scaling but it will expose it to public networks and therefore increasing risk.
Microsoft Azure Stack Portfolio	**Virtual Machine Scale Sets** • As the application demand increases, this pushes the scale set to automatically increase the number of VM instances and reduces them when demand decreases. **Azure Arc** • Allows to elastically scale based on the required capacity to the public cloud.

6.4 MULTI-CLOUD SUPPORT

As discussed earlier in the literature review section, Multi-cloud is a powerful tool that uses vendor diversification strategy to allow multiple benefits and it is important that a private or a hybrid cloud infrastructure supports it. When all three platforms were compared, it was noticed that Microsoft had limited options as the Azure Arc service only supports Azure-based products whereas both OpenStack and VMware are more open platforms that can communicate with any cloud provider.

The following table shows the comparison between the three vendors on multi-cloud supports:

Table 7: Benchmarking Multi-Cloud

OpenStack	**Multi-cloud Support** • OpenStack service Heat supports multi-cloud orchestration. [85]
VMware Cloud Foundation	**VMware vRealize Suite Multi-cloud management solution** • Provides comprehensive visibility across Oracle, AWS, Azure, Google Cloud, Kubernetes, and Private cloud • Helps with cost optimization across multiple cloud platforms. • It provides the ability to inspect cloud resources for security and compliance based on industry standards. • Streamlines cloud operation and control with custom policies and automation.
Azure Stack Portfolio	**Azure Arc** • Helps with governance and management by providing a consistent Azure-based Multi-cloud (public and private on-premises) management and edge platform using Azure Resource Manager.

6.5 DISTRIBUTED CLOUD

As we have seen earlier distributed cloud is about enabling geographical distribution, central management of public cloud services so to provide on-premise optimized performance, lower latency, and edge computing to the end-user or client. From the analysis carried out in this, Microsoft Azure Stack Edge is the only provider with a complete solution.

The following table shows the comparison between the three private cloud vendors:

Table 8: Benchmarking Distributed Cloud

OpenStack	• OpenInfra Foundation Edge Computing Group is working on the development of distributed cloud with edge and fog computing. There is no finalized product from the OpenStack community and the project is in the testing stage. [86]
VMware Cloud Foundation	**VMware Telco Cloud** • VMware Telco Cloud makes use of telco edge to provide edge computing, but the service is in the early stages according to VMware. [87]
Azure Stack Portfolio	**Azure Stack Edge** • Azure Stack Edge a complete solution for edge computing from Microsoft. It is a managed Hardware-as-a-service appliance that delivers services such as compute, storage, and intelligence to the edge on-premises. It supports services such as artificial AI, Hyper-V VM, FPGA Compute, storage, Kubernetes, HA, and GPU.

6.6 CONTAINERIZATION AND VIRTUALIZATION

The technology review in the literature review section on containerization and virtualization was used for comparing containers and virtual machines. It was noticed that containers have poor isolation from the OS, lower resource utilization, portability in the same type of OS, faster boot time and they utilize the same network adapter when compared with virtual machines.

The following table displays the differences that exist in these two technologies:

Table 9: Virtualization vs Containerization [16]

Feature	**VM**	**Container**
Isolation	Isolates the application from the host OS and hence provides greater security	Provides lightweight isolation for the application from the host OS
System Resources	Require greater resources since a complete OS must run for each instance so greater utilization of CPU, RAM, and storage	Less resource-intensive as only a user-mode of an OS is used and only runs required services.
Guest Compatibility	Can run any type of OS in the VM	Runs the same version of OS as the host
Deployment	Can be deployed using tools such as OpenStack, scripting, ESXI, or System center Virtual machine manager	Can be deployed using scripting, Kubernetes services, or Docker.

OS Updates/Upgrades	Each VM will need to be updated separately, and setting up a new VM may be required if a new version of OS is being installed which can be time-consuming	Upgrading and updating are done in a uniform manner using orchestration. And overall overhead is significantly reduced when we account for maintenance, patching, and updates
Startup time	Takes few minutes to boot so poor for fault tolerance	Takes milliseconds to boot so good for fault tolerance as can be rapidly recreated
Storage	Makes use of virtual hard disks	A container does not use virtual hard disks.
Load balancing	VM load balancing helps them to move to a different server during a failover	A container does not move on its own and requires an orchestrator that can automatically start or stop nodes during a failover
Network adapters	Use separate virtual adapters	Uses a single isolated virtual adapter allowing a shared firewall hence reducing the overhead

6.7 CONTAINERIZATION AND ORCHESTRATION SUPPORT

A Private or a Hybrid Cloud platform needs to support modern cloud-native applications that run on containers. For this reason, Kubernetes support is essential as it is the orchestration platform for multiple containers. During the comparison, it was found that all the platforms support container orchestration summarized in the table below:

Table 10: Container Orchestration Support

OpenStack	**Magnum** • In project Magnum, OpenStack has support for provisioning container orchestration engines such as Kubernetes. It makes use of the Heat service for orchestration of Operating System images that have orchestration engines inside them and runs the image in VM or on bare metal inside a cluster.
VMware Cloud Foundation	**VMware Tanzu** • A new service Tanzu as a part of vSphere which has transformed it for running Kubernetes container workloads directly on the hypervisor or the ESXi hosts. **VMware NSX** • The NSX Distributed Firewall has support for Kubernetes or Cloud Foundry application instances and Kubernetes Networking policy. It is well integrated with the Tanzu Kubernetes grid.
Azure Stack Portfolio	**Azure Stack Edge** • Supports the deployment of Kubernetes clusters. The local edge Compute makes use of containers that are orchestrated via Kubernetes. Container management can be carried out using Azure IoT Hub. **Azure Stack HCI** • It also has built-in support for Kubernetes container orchestration. Azure stack HCI hybrid services can scan containers for vulnerabilities and protects Azure Kubernetes Service instances. **Azure Stack Hub** • Has built-in support for Kubernetes container orchestration

7.0 BUSINESS MODEL OF GLOBALIZATION AND LOCALIZATION

7.1.1 BUSINESS MODEL FOR CLOUD COMPUTING

In this section, each business model will be investigated, and recommendations will be given based on the type of data the industry handles. The following analysis is done based on the Processing Personal Data Across Borders Guidelines provided by the Office of the Privacy Commissioner of Canada (PIPEDA) [89] [90] and information compiled in this report:

Disclaimer: It is always a good idea that an organization uses tools such as Privacy Impact Assessments (PIA) or Threat Risk Assessments (TRA) so to evaluate their risk tolerance and access their safeguards.

Table 11: Analysis of Cloud for Business Model

Business Model	Industry Type	Recommended Cloud Model(s)	Reasons
Small and Medium Business (SMB)	Industries that do not hold Personal Identifiable Information (PII) and Protected Health Information (PHI).	Public Cloud e.g. Azure, AWS, and Google Cloud	• No Capex which is crucial for small businesses • No hardware maintenance • Ease of use • Lower staff requirements • Less time to implement
	Industries that do hold Personal Identifiable Information (PII), Protected Health Information (PHI), and federally regulated financial institutions	Community Cloud via CXP e.g. Azure, AWS, and Google Cloud.	• No Capex which is crucial for small businesses • More Secure • Law requires you to protect PII: failure to comply could result in complaints and legal action • Customers expect you to be transparent • Lowers risk • Lower staff requirements

Small and Medium Enterprise (SME)	Industries that do not hold Personal Identifiable Information (PII) and Protected Health Information (PHI).	Hybrid Cloud e.g. Azure Stack HCI	• Lowers Capex • Less hardware maintenance • Ease of use • Hybrid Cloud will provide cloud bursting feature during heavy traffic • Resource intensive application can be placed on the local cloud so to reduce long term OPEX that can incur on public cloud
	Industries that do hold Personal Identifiable Information (PII), Protected Health Information (PHI), and federally regulated financial institutions	Hybrid (Private Cloud and Community Cloud via CXP) e.g. with a combination of OpenStack + AWS, VMware Cloud Foundation + AWS	• Usage of Private Cloud will lower OPEX for resource-intensive applications • Improves ROI • More Secure • Law requires you to protect PII: failure to comply could result in complaints and legal action • Customers expect you to be transparent • Government regulation can change • Lowers risk • Hybrid Cloud will provide cloud bursting feature during heavy traffic

Large Enterprise	Industries that do not hold Personal Identifiable Information (PII), Protected Health Information (PHI), and federally regulated financial institutions	Multi-Cloud (Private Cloud and multiple Public Clouds) e.g. with a combination of OpenStack + AWS + Google Cloud, VMware Cloud Foundation + AWS + Azure	• Improves reliability • Prevents Vender-lock in • Improves ROI as Multi-Cloud provides the option to run the resource-intensive application on private cloud and provides more options when looking for cheaper cloud services. • Provides cloud bursting feature during heavy traffic
	Industries that do hold Personal Identifiable Information (PII), Protected Health Information (PHI), and federally regulated financial institutions	Multi-Cloud (Private Cloud and multiple Community Clouds via CXP connection) e.g. with a combination of OpenStack + AWS + Azure, VMware Cloud Foundation + AWS + Google Cloud	• Improves reliability • Prevents Vender-lock in • Improves ROI as Multi-Cloud provides the option to run a resource-intensive application on private cloud and provides more options when looking for cheaper cloud services. • Provides cloud bursting feature during heavy traffic • Government regulation can change • Lowers risk and more secure • Customers expect you to be transparent • Law requires you to protect PII: failure to comply could result in complaints and legal action

8.0 PROJECT PLAN AND MILESTONES

The Draft submission and final submission was updated accordingly:

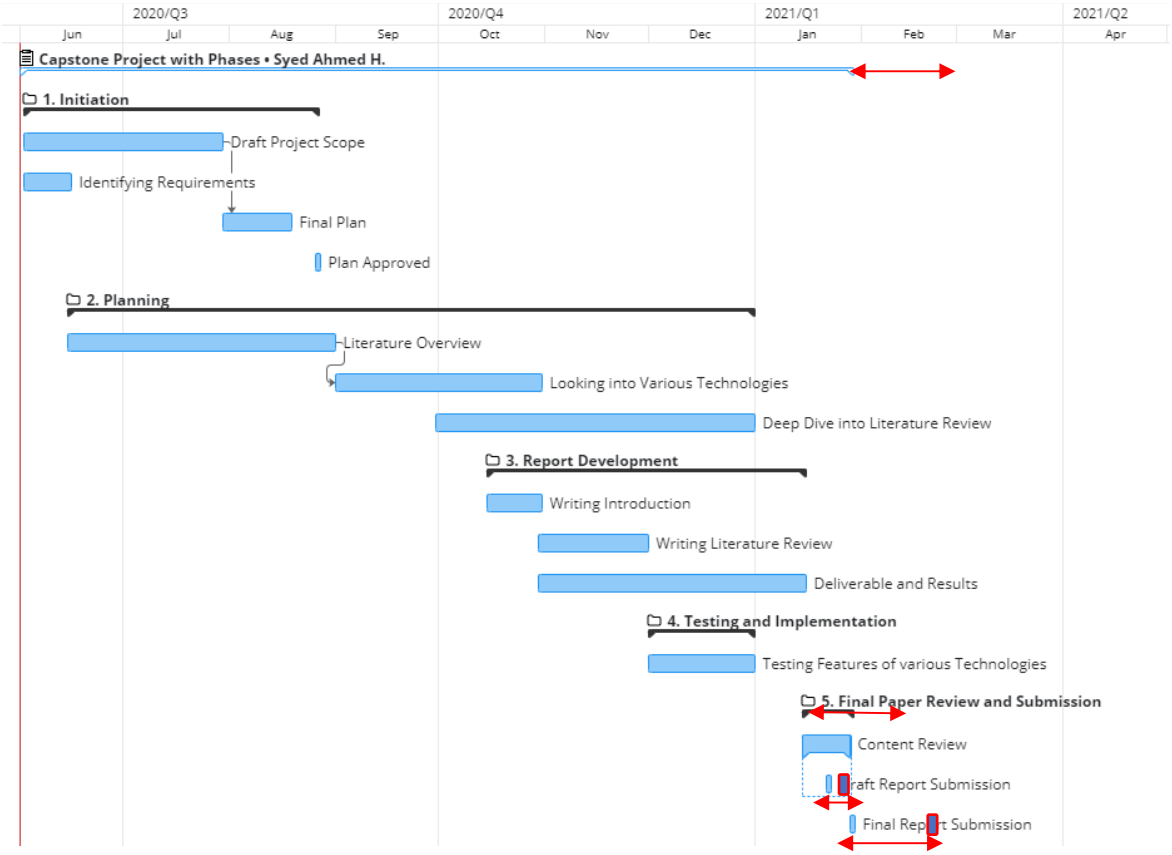

The red indicators show the updated schedule, all the other tasks were completed in time as and Testing and implementation were removed due to time constraints and lack of hardware availability.

9.0 CONCLUSION AND SUMMARY

The focus of this report was to carry out an analysis on different local and hybrid cloud providers and to benchmark features such as Scalability, Latency, Security, Distributed Computing. Another objective involved recommending the cloud model based on different business models in the industry and the implication of businesses going global. The objectives also involved a comparison between Virtualization and Containerization and analysis of Multi-Cloud technologies.

The feature availability of each local and hybrid cloud provider was extensively discussed in the literature review section. The features of virtualization and containerization were also explored in the literature review section and compared in the Benchmarking section.

Scalability was discussed in the benchmarking of local and hybrid technologies section of the report. The two scaling technologies that were considered were vertical scaling and horizontal scaling. It was found that all three vendors had the ability to scale the resources, OpenStack named horizontal scaling as Auto-scaling and Octavia. VMware Cloud Foundation made use of DRS technology. VMware used vMotion and vertical scaling technology. Horizontal scaling in VMware Cloud Foundation can only be utilized using a public cloud. Azure Stack Portfolio used technology called Virtual Machine Scale Sets which was able to perform horizontal scaling and it also supported elasticity via Azure Arc using the public cloud.

Latency in cloud computing is a very critical issue that was discussed in this report. The latency issues were extensively discussed in the literature review section for each and every vendor and later it was also used when comparing the vendors in the latency section. During the analysis, Microsoft Azure Stack Portfolio outperforms as compared to the rest of the vendors based on technologies such RDMA, VMMQ, and Azure Stack Edge. VMware also utilized RDMA with PVRDMA for reducing latency and whereas OpenStack had a Storage Performance Development Kit driver that could be programmed to run SSDs in polled mode, asynchronous and lockless- NVMe and also to run technologies such as RDMA.

Analyzing cloud security was another major agenda of this report, it was covered in depth in the literature review section of this report and was later compared for each platform in the benchmarking section. During the comparison, it was found that VMware had better built-in network security when compared with the rest. VMware Cloud Foundation made use of NSX for network security which had the support of

intrusion detection system (IDS), intrusion prevention system (IPS), gateway firewall, distributed firewall, web application firewall WAF, and context-aware micro-segmentation. OpenStack had a better built-in Keystone Identity service. Whereas Azure Stack HCI did provide advanced security monitoring but at the expense of exposing the infrastructure to the public network as it requires the internet. It was noticed that both VMware Native VMkernel Cryptographic Module and Microsoft BitLocker Drive Encryption were both FIPS 140-2 compliant.

Multi-cloud technology was discussed in the literature review section of the report. It uses a combination of different cloud service providers and private cloud to deliver workload portability, prevent vendor lock-in, provide reliability, better ROI, more variety of tools, and adaptable to changing regulatory requirements. In the benchmarking section, the comparison showed that Azure Arc service only supported Azure-based products whereas both OpenStack Heat and VMware vRealize Suite Multi-cloud management solution are open platforms that are compatible with any other cloud service providers.

Distributed cloud computing was another area that was explored in this project. An extensive study was carried out that is revied in the literature review section of this report. Distributed cloud computing enabled geographical distribution, centralized management of cloud services so to provide on-premise optimized performance, reduced latency, and edge computing to the end-user at the remote site. From the comparison carried in the benchmarking section, it was found that Microsoft Azure Stack Edge was the only vendor with a complete out-of-the-box solution for the distributed cloud at this moment.

It was found that the container architecture comprises of application and some lightweight APIs and services that run on top of the operating system kernel which runs on top of the physical hardware. Whereas, in virtualization, VMs run an entire operating system and a separate kernel for each application and service on top of the physical hardware. Other significant differences were compared between the two in the benchmarking section such as isolation, system resource usage, guest compatibility, deployment methodology, operating system updates, startup time, storage, load balancing, and network adapters were also compared in Table 12: Virtualization vs Containerization in the benchmarking section. Containers had poor isolation from the operating system, lower utilization of the resource, portability was in the same type of operating system, they had faster boot time as they had lower overhead and they utilized the same network adapter in contrast with VMs. It was found that all three platforms

had built-in support for container orchestration as it is an essential requirement for modern cloud-native applications.

Finally, in the business model of globalization and localization section, each business model was matched with a specific cloud computing model based on what type of data those businesses hold and reasoning was given as to why those specific models were chosen.

10.0 FUTURE WORK

In the future, further studies have to be carried on the compliance of cloud technologies such as FIPS, Common Criteria, NIST, and ITSG-33. There is also a need to evaluate container orchestration technologies such as Kubernetes based on importance, functionality, and need for cloud-native applications in much more depth.

REFERENCES

[1] Microsoft, "azure.microsoft," Microsoft, 01 06 2020. [Online]. Available: https://azure.microsoft.com/en-ca/overview/what-is-cloud-computing/#benefits. [Accessed Monday June 2020].

[2] D. S. Linthicum, Cloud Computing and SOA Convergence in Your Enterprise: A Step-by-Step Guide, Pearson Education, 2009.

[3] P. Z. Mahmood, "Cloud Computing: Challenges, Limitations and R&D Solutions," Springer, 2014, p. 19.

[4] O. o. t. P. C. o. Canada, "Cloud computing and privacy," Office of the Privacy Commissioner of Canada, 01 08 2011. [Online]. Available: https://www.priv.gc.ca/en/privacy-topics/technology/online-privacy-tracking-cookies/online-privacy/cloud-computing/02_05_d_51_cc/. [Accessed Monday June 2020].

[5] L. M. Yousef Farhaoui, "Big Data and Smart Digital Environment," Springer, 2019, pp. 360-361.

[6] NIST, "The NIST Definition of Cloud," 11 2011. [Online]. Available: https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf. [Accessed 02 01 2021].

[7] microsoft, 24 07 2013. [Online]. Available: https://docs.microsoft.com/en-us/archive/blogs/xdot509/getting-started-with-windows-azure-part-2-what-are-cloud-services. [Accessed 010 01 2021].

[8] N. K. T. Stephen R Smoot, in *Private Cloud Computing: Consolidation, Virtualization, and Service-Oriented Infrastructure* , Elsevier, 2011, pp. 3-4.

[9] microsoft, "what-is-cloud-bursting," 01 01 2021. [Online]. Available: https://azure.microsoft.com/en-ca/overview/what-is-cloud-bursting/#:~:text=In%20cloud%20computing%2C%20cloud%20bursting,with%20peaks%20in%20IT%20demand.&text=An%20application%20can%20be%20applied,necessary%20to%20meet%20peak%20demands.. [Accessed 30 01 2021].

[10] OpenStack, "security-boundaries-and-threats," 04 02 2021. [Online]. Available: https://docs.OpenStack.org/security-guide/introduction/security-boundaries-and-threats.html. [Accessed 10 02 2021].

[11] GC, "Levels of security," 16 06 2020. [Online]. Available: https://www.tpsgc-pwgsc.gc.ca/esc-src/protection-safeguarding/niveaux-levels-eng.html . [Accessed 29 01 2021].

[12] GC, "Government of Canada White Paper: Data Sovereignty and Public Cloud," 28 07 2020. [Online]. Available: https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/cloud-services/gc-white-paper-data-sovereignty-public-cloud.html#toc4. [Accessed 18 12 2020].

[13] C. C. f. C. Security, "ITSP.50.104 Guidance on Defence in Depth for Cloud-Based Services," 05 01 2020. [Online]. Available: https://cyber.gc.ca/en/guidance/itsp50104-guidance-defence-depth-cloud-based-services. [Accessed 21 02 2020].

[14] "distributed-cloud," IBM, 3 11 2020. [Online]. Available: https://www.ibm.com/cloud/learn/distributed-cloud . [Accessed 27 01 2021].

[15] "Mobile Services Meet Distributed Cloud: Benefits, Applications, and Challenges," 01 06 2018. [Online]. Available: https://www.researchgate.net/publication/325458691_Mobile_Services_Meet_Distributed_Cloud_Benefits_Applications_and_Challenges . [Accessed 1 2 2021].

[16] F. Z. C. Brian E. Whitaker, "Cloud Edge Computing: Beyond the Data Center," 2020. [Online]. Available: https://www.openstack.org/use-cases/edge-computing/cloud-edge-computing-beyond-the-data-center/?lang=en_US . [Accessed 01 01 2021].

[17] microsoft, "Windows and containers," 22 10 2019. [Online]. Available: https://docs.microsoft.com/en-us/virtualization/windowscontainers/about/. [Accessed 15 01 2021].

[18] Microsoft, "What is a container?," 01 01 2021. [Online]. Available: https://azure.microsoft.com/en-us/overview/what-is-a-container/ . [Accessed 21 02 2021].

[19] Microsoft, "Containers vs. virtual machines," 21 10 2019. [Online]. Available: https://docs.microsoft.com/en-us/virtualization/windowscontainers/about/containers-vs-vm. [Accessed 28 11 2020].

[20] J. Chen, "Making Containers More Isolated: An Overview of Sandboxed Container Technologies," 06 06 2019. [Online]. Available: https://unit42.paloaltonetworks.com/making-containers-more-isolated-an-overview-of-sandboxed-container-technologies/. [Accessed 20 02 2021].

[21] Kubernetes, "What is Kubernetes?," 01 02 2021. [Online]. Available: https://kubernetes.io/docs/concepts/overview/what-is-kubernetes/. [Accessed 25 02 2021].

[22] Canonical, "CIO guide to multi-cloud operations," 01 08 2018. [Online]. Available: https://pages.ubuntu.com/rs/066-EOV-335/images/CIO_MultiCloud_WP_Canonical_31.08.18.V4.pdf?_ga=2.48577702.1698745104.1612407643-294589640.1612244813&_gac=1.258568056.1612407654.CjwKCAiAsOmABhAw

EiwAEBR0ZvTsUiSJwI5rhKGJEq-xG7QF4NlasF0nJw2RRaG82-h0hu8F-ZHskh. [Accessed 25 12 2020].

[23] OpenStack, "Introduction to OpenStack," 04 02 2021. [Online]. Available: https://docs.openstack.org/security-guide/introduction/introduction-to-openstack.html. [Accessed 9 02 2021].

[24] OpenStack, "THE OPENSTACK LANDSCAPE," 01 07 2020. [Online]. Available: https://www.OpenStack.org/software/. [Accessed 12 12 2020].

[25] OpenStack, "OpenStack Services," 01 07 2020. [Online]. Available: https://www.openstack.org/software/project-navigator/openstack-components/#openstack-services. [Accessed 12 1 2020].

[26] "Nova System Architecture," 16 12 2020. [Online]. Available: https://docs.OpenStack.org/nova/pike/user/architecture.html. [Accessed 25 12 2020].

[27] OpenStack, "Data encryption," 04 02 2021. [Online]. Available: https://docs.OpenStack.org/security-guide/tenant-data/data-encryption.html#ephemeral-disk-encryption . [Accessed 10 02 2021].

[28] OpenStack, "OpenStack Glance Architecture," 21 08 2019. [Online]. Available: https://docs.OpenStack.org/glance/pike/contributor/architecture.html. [Accessed 13 12 2020].

[29] C. Ltd., "What makes up OpenStack?," 2021. [Online]. Available: https://ubuntu.com/openstack/what-is-openstack. [Accessed 01 01 2021].

[30] OpenStack, 04 02 2021. [Online]. Available: https://docs.openstack.org/security-guide/networking/architecture.html. [Accessed 10 02 2021].

[31] T. Sonia Lelii, "cloud object storage," 01 12 2017. [Online]. Available: https://searchstorage.techtarget.com/definition/cloud-object-storage. [Accessed 06 01 2021].

[32] IBM, "Block Storage," 24 Jun 2019. [Online]. Available: https://www.ibm.com/cloud/learn/block-storage#:~:text=Block%20storage%2C%20sometimes%20referred%20to%20as%20block-level%20storage%2C,they%20require%20fast%2C%20efficient%2C%20and%20reliable%20data%20transportation.. [Accessed 15 12 2020].

[33] OpenStack, "Cinder System Architecture," 17 10 2019. [Online]. Available: https://docs.OpenStack.org/cinder/latest/contributor/architecture.html. [Accessed 16 11 2020].

[34] OpenStack, "Components," 15 12 2020. [Online]. Available: https://docs.OpenStack.org/swift/pike/admin/objectstorage-components.html. [Accessed 27 12 2020].

[35] O'REILLY, "Chapter 1. Fundamental Keystone Topics," 2021. [Online]. Available: https://www.oreilly.com/library/view/identity-authentication-and/9781491941249/ch01.html. [Accessed 20 1 2021].

[36] OpenStack, "Keystone Architecture," 14 04 2020. [Online]. Available: https://docs.OpenStack.org/keystone/latest/getting-started/architecture.html. [Accessed 15 11 2020].

[37] R. HAT, "OpenShift on OpenStack: Introduction to availability zones," 29 09 2018. [Online]. Available: https://www.redhat.com/en/blog/openshift-openstack-introduction-availability-zones. [Accessed 01 01 2021].

[38] OpenStack, "Theory of Auto-Scaling," 07 06 2019. [Online]. Available: https://docs.OpenStack.org/auto-scaling-sig/latest/theory-of-auto-scaling.html . [Accessed 28 11 2020].

[39] OpenStack, "Horizon Basics," 29 06 2017. [Online]. Available: https://docs.OpenStack.org/horizon/latest/contributor/intro.html#contributor-intro. [Accessed 26 11 2020].

[40] VMWARE, "Hyperconverged Infrastructure," 2021. [Online]. Available: https://www.vmware.com/products/hyper-converged-infrastructure.html. [Accessed 05 01 2021].

[41] Vmware, "VCF," 10 08 2020. [Online]. Available: https://docs.vmware.com/en/VMware-Cloud-Foundation/4.1/vcf-41-introducing/GUID-7EBCC024-9604-4064-90A1-4851A78F7641.html. [Accessed 29 11 2020].

[42] VMware, "SDDC Manager," 13 11 2017. [Online]. Available: https://docs.vmware.com/en/VMware-Cloud-Foundation/2.1.3/com.vmware.vcf.ovdeploy.doc_213/GUID-F16F5CA4-ABF1-4282-974D-7CBB96028964.html . [Accessed 17 12 2020].

[43] vJenner, "Horizon 7 on VMware Cloud Foundation 2.3," 06 03 2018. [Online]. Available: http://www.vjenner.com/2018/03/horizon-7-on-vmware-cloud-foundation-2-3/ . [Accessed 25 12 2020].

[44] G. Blake, "Introducing VMware Cloud Builder, Automated Deployment of VMware Validated Designs," 23 01 2019. [Online]. Available: https://blogs.vmware.com/cloud-foundation/2019/01/23/introducing-vmware-cloud-builder-automated-deployment-of-vmware-validated-designs/#:~:text=%20VMware%20Cloud%20Builder%20has%20been%20developed%20to,Design%20for%20Management%20and%20Workload%20Consolid. [Accessed 28 11 2020].

[45] VMware, "VMware vSphere Documentation," 2020. [Online]. Available: https://docs.vmware.com/en/VMware-vSphere/index.html . [Accessed 19 01 2021].

[46] VMware, "Architecture of VMware ESXi," 15 10 2008. [Online]. Available: https://www.vmware.com/techpapers/2007/architecture-of-vmware-esxi-1009.html#:~:text=VMware%20ESXi%20is%20the%20next-generation%20hypervisor%2C%20providing%20a,offering%20improved%20security%2C%20increased%20reliability%2C%20and%20simplified%20management.. [Accessed 18 11 2020].

[47] onekobo, "Explain ESXi and vCenter Server Architecture," 15 03 2011. [Online]. Available: http://www.onekobo.com/articles/vsphere/obj01/obj1-5/obj1-5b.html. [Accessed 09 12 2020].

[48] VMware, "VMware vSphere," 10 01 2020. [Online]. Available: https://docs.vmware.com/en/VMware-Cloud-Foundation/4.1/vcf-41-introducing/GUID-92386FF4-4D6B-4670-BA5A-3B20A40F9950.html. [Accessed 29 11 2020].

[49] VMware, "vSphere High Avability," 2021. [Online]. Available: https://www.vmware.com/products/vsphere/high-availability.html#:~:text=VMware%20vSphere%20High%20Availability%20delivers%20the%20availability%20required,operating%20system%20outages%20within%20your%20virtualized%20IT%20environment . [Accessed 10 02 2021].

[50] VMware, "VMware VMotion," 2007. [Online]. Available: https://www.vmware.com/pdf/vmotion_datasheet.pdf . [Accessed 18 12 2020].

[51] R. Castagna, "Storage vMotion," 04 2019. [Online]. Available: https://searchvmware.techtarget.com/definition/Storage-vMotion#:~:text=Storage%20vMotion%20is%20a%20component%20of%20VMware%20vSphere,users.%20This%20migration%20occurs%20while%20maintaining%20data%20integrity. . [Accessed 19 12 2020].

[52] VMware, "What is Live Migration of Virtual Machines and how does it work?," 2021. [Online]. Available: https://www.vmware.com/products/vsphere/vmotion.html . [Accessed 09 01 2021].

[53] T. Contributor, "VMware DRS (Distributed Resource Scheduler)," 05 2019. [Online]. Available: https://searchvmware.techtarget.com/definition/VMware-DRS. [Accessed 27 12 2020].

[54] VMware, "Enhanced Application Performance and Availability," 2021. [Online]. Available: https://www.vmware.com/products/vsphere/enhanced-app-performance.html . [Accessed 06 01 2021].

[55] M. Potheri, "VMware Cloud Foundation as an enabler for GPU as a service," 18 06 2020. [Online]. Available: https://blogs.vmware.com/apps/2020/06/vmware-cloud-foundation-as-an-enabler-for-gpu-as-a-service-part-1-of-3.html . [Accessed 17 01 2021].

[56] VMware, "VMware vSAN," 31 05 2019. [Online]. Available: https://docs.vmware.com/en/VMware-Cloud-Foundation/3.0/com.vmware.vcf.ovdeploy.doc_30/GUID-E6C1B2F8-F6CB-4017-A4D1-53103FF9AB84.html. [Accessed 17 11 2020].

[57] VMware, "vSAN technology overview," 2021. [Online]. Available: https://core.vmware.com/resource/vsan-technology-overview#section1 . [Accessed 09 01 2021].

[58] Supermicro, "Hyper-converged Storage Systems," 2017. [Online]. Available: https://www.supermicro.com/solutions/datasheet_vSAN.pdf. [Accessed 13 01 2021].

[59] VMware, "VMware NSX Data Center," 2020. [Online]. Available: https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/nsx/vmware-nsx-solution-brief.pdf . [Accessed 22 01 2021].

[60] VMware, "VMware NSX Data Center," 2020. [Online]. Available: https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/nsx/vmware-nsx-datasheet.pdf. [Accessed 26 01 2021].

[61] V. Bhandari, "Security that's Designed for the Modern Data Center," 25 02 2020. [Online]. Available: https://blogs.vmware.com/networkvirtualization/2020/02/security-for-modern-data-center.html/. [Accessed 28 01 2021].

[62] VMware, "VMware vRealize Suite," 2021. [Online]. Available: https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/datasheet/vmware-vrealize-suite-datasheet.pdf. [Accessed 07 02 2021].

[63] VMware, "Essential Cloud Management with vRealize Suite and vCloud Suite," 2021. [Online]. Available: https://www.vmware.com/asean/products/vrealize-suite-vcloud-suite.html. [Accessed 07 02 2021].

[64] Microsoft, "Expanding the Azure Stack portfolio to run hybrid applications across the cloud, datacenters, and the edge," 05 11 2019. [Online]. Available: https://azure.microsoft.com/en-us/blog/expanding-the-azure-stack-portfolio-to-run-hybrid-applications-across-the-cloud-datacenter-and-the-edge/ . [Accessed 08 11 2020].

[65] Microsoft, "Azure Stack HCI Overview white paper," 21 07 2020. [Online]. Available: https://azure.microsoft.com/en-us/resources/azure-stack-hci-overview-white-paper/ . [Accessed 09 01 2021].

[66] Microsoft, "What is Azure Stack Edge?," 2021. [Online]. Available: https://docs.microsoft.com/en-us/learn/modules/intro-to-azure-stack/5-azure-stack-edge. [Accessed 01 02 2021].

[67] Microsoft, "Azure Site Recovery," 2021. [Online]. Available: https://azure.microsoft.com/en-us/services/site-recovery/. [Accessed 29 01 2021].

[68] Microsoft, "About Site Recovery," 17 03 2020. [Online]. Available: https://docs.microsoft.com/en-us/azure/site-recovery/site-recovery-overview. [Accessed 19 01 2021].

[69] Microsoft, "Deploy a Cloud Witness for a Failover Cluster," 18 01 2019. [Online]. Available: https://docs.microsoft.com/en-us/windows-server/failover-clustering/deploy-cloud-witness. [Accessed 05 01 2021].

[70] Microsoft, "What is the Azure Backup service?," 24 04 2019. [Online]. Available: https://docs.microsoft.com/en-us/azure/backup/backup-overview . [Accessed 18 01 2021].

[71] Microsoft, "Update Management overview," 22 01 2021. [Online]. Available: https://docs.microsoft.com/en-us/azure/automation/update-management/overview. [Accessed 29 01 2021].

[72] n. Microsoft, "Top 10 Networking Features in Windows Server 2019: #3 Azure Network Adapter," 14 02 2019. [Online]. Available: https://techcommunity.microsoft.com/t5/networking-blog/top-10-networking-features-in-windows-server-2019-3-azure/ba-p/339780. [Accessed 23 01 2021].

[73] Microsoft, "Azure Defender," 2021. [Online]. Available: https://azure.microsoft.com/en-us/services/azure-defender/#features. [Accessed 12 02 2021].

[74] Microsoft, "Azure Stack HCI software," 10 2019. [Online]. Available: https://azure.microsoft.com/mediahandler/files/resourcefiles/azure-stack-hci-software-architecture/190126-MSFT-AzureStackHCI%20BoM-SoftwareArchitectureWP-Final.pdf. [Accessed 11 01 2021].

[75] Microsoft, "Host network requirements for Azure Stack HCI," 25 11 2020. [Online]. Available: https://docs.microsoft.com/en-us/azure-stack/hci/concepts/host-network-requirements. [Accessed 15 01 2021].

[76] Microsoft, "Faster Live Migration with RDMA in Windows Server 2012 R2," 30 12 2013. [Online]. Available: https://docs.microsoft.com/en-us/archive/blogs/virtual_pc_guy/faster-live-migration-with-rdma-in-windows-server-2012-r2 . [Accessed 16 01 2021].

[77] Microsoft, "What is Azure Stack Hub?," 2021. [Online]. Available: https://docs.microsoft.com/en-us/learn/modules/intro-to-azure-stack/3-azure-stack-hub. [Accessed 02 01 2021].

[78] Microsoft, "Azure Stack Hub overview," 08 01 2021. [Online]. Available: https://docs.microsoft.com/en-us/azure-stack/operator/azure-stack-overview?view=azs-2008&viewFallbackFrom=azs-2008osoft.com%2Fen-us%2Flearn%2Fmodules%2Fintro-to-azure-stack%2F3-azure-stack-hub. [Accessed 22 01 2021].

[79] Microsoft, "What is Azure Queue Storage?," 18 03 2020. [Online]. Available: https://docs.microsoft.com/en-us/azure/storage/queues/storage-queues-introduction. [Accessed 13 01 2021].

[80] Microsoft, "What is Azure Table storage ?," 07 01 2021. [Online]. Available: https://docs.microsoft.com/en-us/azure/storage/tables/table-storage-overview . [Accessed 26 01 2021].

[81] Microsoft, "About Azure Key Vault," 01 10 2020. [Online]. Available: https://docs.microsoft.com/en-us/azure/key-vault/general/overview . [Accessed 27 01 2021].

[82] Microsoft, "Azure App Service and Azure Functions on Azure Stack Hub overview," 05 05 2020. [Online]. Available: https://docs.microsoft.com/en-us/azure-stack/operator/azure-stack-app-service-overview?view=azs-2008 . [Accessed 24 01 2021].

[83] "Use SQL databases on Azure Stack Hub," 19 08 2020. [Online]. Available: https://docs.microsoft.com/en-us/azure-stack/operator/azure-stack-sql-resource-provider?view=azs-2008 . [Accessed 29 01 2021].

[84] Microsoft, "Azure Stack Hub infrastructure security controls," 10 06 2019. [Online]. Available: https://docs.microsoft.com/en-us/azure-stack/operator/azure-stack-security-foundations?view=azs-2008 . [Accessed 12 01 2021].

[85] Microsoft, "Data at rest encryption in Azure Stack Hub," 04 03 2020. [Online]. Available: https://docs.microsoft.com/en-us/azure-stack/operator/azure-stack-security-bitlocker?view=azs-2008 . [Accessed 03 01 2021].

[86] "Azure Stack Hub infrastructure security controls," 10 06 2019. [Online]. Available: https://docs.microsoft.com/en-us/azure-stack/operator/azure-stack-security-foundations?view=azs-2008. [Accessed 29 01 2021].

[87] Microsoft, "Azure Stack Hub infrastructure security controls," 10 06 2019. [Online]. Available: https://docs.microsoft.com/en-us/azure-stack/operator/azure-stack-security-foundations?view=azs-2008 . [Accessed 06 01 2021].

[88] Microsoft, "Azure Arc overview," 23 09 2020. [Online]. Available: https://docs.microsoft.com/en-us/azure/azure-arc/overview . [Accessed 27 01 2021].

[89] OpenStack, "Multi-Clouds support," 22 01 2020. [Online]. Available: https://docs.OpenStack.org/heat/latest/template_guide/multi-clouds.html . [Accessed 23 01 2021].

[90] F. Z. C. Brian E. Whitaker, "Cloud Edge Computing: Beyond the Data Center," 2020. [Online]. Available: https://www.openstack.org/use-cases/edge-computing/cloud-edge-computing-beyond-the-data-center/?lang=en_US. [Accessed 27 01 2021].

[91] VMware, "What is an Edge Cloud Platform?," 2021. [Online]. Available: https://telco.vmware.com/solutions/edge.html . [Accessed 27 01 2021].

[92] O. o. P. C. o. Canada, "Cloud computing for small and medium-sized enterprises," 14 05 2012. [Online]. Available: https://www.priv.gc.ca/en/privacy-topics/technology/online-privacy-tracking-cookies/online-privacy/cloud-computing/gd_cc_201206/. [Accessed 01 02 2021].

[93] O. o. t. P. C. o. Canada, "Processing Personal," 2009. [Online]. Available: https://www.priv.gc.ca/media/1992/gl_dab_090127_e.pdf. [Accessed 01 01 2021].

[94] K. Downie, "Azure Stack HCI solution overview," 10 02 2021. [Online]. Available: https://docs.microsoft.com/en-us/azure-stack/hci/overview. [Accessed 11 02 2021].

End