Smart Education System

Prof. Chitrangada Chaubey[1], Nitesh Sharma[2], Akshay Unnikrishnan[3], Jaspal[4]
Assistant Professor, Department of Computer Science and Engineering
Students, Department of Computer Science and Engineering
Dronacharya Group of Institutions, Greater Noida, India

ABSTRACT

As we all know, during this covid most of us are working from home or taking sessions from our comfort zone.

Microsoft establishes a way to make this helpful pandemic. We have deployed a solution for educational institutes on how they can run their classes with ease, consistently, secure, and without spending much on their infrastructure.

We have used Microsoft Azure, M365, and other Azure Technologies like Azure Active Directory, Virtual Network, Storage Account, Workspace, and Host pool to accomplish this need.

In colleges, we have segregation between different departments like Mechanical, Electrical, Civil and Computer Science and additional years of student. We have divided the department and years and created a workspace for each segregation. The other department needs different apps and configuration to run the class smoothly. Each workspace has its virtual machine and set of applications to session smoothly.

This technology is based on a pay-as-you-go model, which reduces the total cost of ownership for college premises. Students and professors need to install remote client applications provided by Microsoft on any modern devices like windows, mac, iOS, and android just using the internet. Thus, the application offers seamless connectivity to users after authenticating from Azure Active Directory. The users can access the remote server to do their tasks.

This way, we can simplify the management.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

CHAPTER 1: INTRODUCTION

1.1 Problem Introduction

During the pandemic, all the students were away from their campus in their houses. However, online classes were held; what about the practical's for the students who don't have access to a high-end device. Sound is also a crucial part of the study. As a computer science student, I can relate to the issue of not having a high-end device.

In recent times, as we all know, the sector has been at its best, and technologies are built-in daily with the advancement in technology.

To overcome this, Microsoft introduced its cloud solution in 2019, which is known as WVD. Later, its name changed to Azure Virtual Desktop.

By using AVD, anyone can connect remotely to a virtual machine deployed over the internet. You can connect with any device (ex: mac, android, windows) which needs a minimal internet connection, and you are good to go with a high-end device. It is also based on the pay-as-you-go model, which means you must pay only for the resources you are using. In this way, it doesn't cost too much for educational institutions to set up a lab over the cloud. It is the cheap and best way to connect with students and faculties.

So, we used AVD to deploy an intelligent college lab infrastructure over the cloud.

1.2 Motivation

As the pandemic hits the world at the end of 2019 whole world goes into lockdown. So as a computer science student, I can relate to the issue that most of the students could not access their campus labs and other resources they needed to craft themselves during a bachelor's degree. To access labs from a remote location was not possible at that time, and we also faced the same.

It was the motivation for my team to develop a solution for this significant problem that was not faced by just us but by the students worldwide at that time.

1.3 Project Objective

Depending on the problems students and faculties faced during the pandemic, the interaction between the faculty and the student was not sufficient. During the recent events of pandemic and the implementation of lockdown as the measure, everyone was ceased to their homes, due to these

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

online classes were taken as an alternative and practical courses or classes which requires a high spec pc were suffering, since many students couldn't afford such PC's, we tried to solve the above problem using AVD (Azure Virtual Desktop).

By using AVD, we took the concept of the virtual machine to provide a high spec machine to run virtually even on a relatively low spec pc or system; using AVD, we can deploy several virtual machines, and since it's a pay as you go model, You can access and pay for only the resources you need, which keeps it budget-friendly and efficient, Since it would be much easier for the institution to provide their batch of students and modifying the virtual machine as you require.

1.4 Scope for the project

Remote Workers:

AVD environments make it much easier to provide access for remote workers to organizational standard desktop environments across a broad range of devices. With the virtual desktop, access to the core software systems can be controlled, and access can be granted to any remote worker at a remote site with minimal investment. Regardless of location, each team member has access to the same organizational network and resources while maintaining central entry and application controls.

Call Centers:

A nonpersistent desktop is a significant advantage over simply consuming a standard desktop from a pool of identical desktops. The typical call center is an excellent example of how this model directly supports the needs of a team of people. Each team member is only required to do a specific set of tasks, which do not require the desktop to be nonstandard. With a workforce that can flex up or down, the ability to log in to an available desktop during the person's work hours minimizes the overhead of providing individual physical desktop machines. The standard nonpersistent desktop instance can be easily patched and deployed with just the requisite software installed and across physical sites with minimal complications.

Contract Employees:

When temporary contractors join a team, they need access to some core assets and team members, but security is an important consideration. By using a virtual desktop, it's possible to control access to corporate resources while delivering the connection point for the temporary workers. Contractors can perform tasks that use organizational resources without having access to systems unrelated to the contract. This control also minimizes the investment in providing a physical endpoint device during the term contract.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

CHAPTER 2: LITERATURE SURVEY

Virtual Desktop technology is one of the trending technologies in today's market. Since the pandemic hits the world, every institution and organization ultimately needs something that could find a way to access their devices from their home remotely. There comes the AV, which saves the entire world and helps to go more online. So AVD is building as it can serve the institutions and big organizations a great value with work flexibility and productivity. By using the Virtualization technology, anyone can easily access the remote servers. Either it is a student at the university who needs to work on their projects or needs a high-end device and software/hardware, or it could be an employee who can remotely work for you without being physically available. It gives flexibility in the work environment and saves a lot of money.

Considering the history of hosted desktops, it is essential to examine the birth of virtual machines within their larger cloud computing context. Throughout the 1960s, computer development gained much ground, but the devices at the time were still only capable of completing one instruction at once. This would often mean that computational tasks had to be completed in batches, which is a time-consuming endeavor.

In 1963 the Massachusetts Institute of Technology (MIT) began developing Project MAC (Multiple Access Computer) with a $2 million Defense Advanced Research Projects Agency grant. MAC aimed to research computer processes in greater depth, particularly in artificial intelligence, computational theory, and operating systems. A portion of this research was geared towards creating computers that could complete multiple tasks at once and devices that could have more than one simultaneous user. In response to the study, International Business Machines (IBM) created the CP-67 system – the first mainframe computer that supported virtualization.

Running a user's desktop operating system did not appear until the late 90s when VMware first introduced it in 1999 with their first product, the VMware Workstation. However, the technology as we know it today did not become mainstream until VMware released a superior product repertoire in 2007. Thousands of global businesses have since utilized hosted virtual desktops.

When virtualization was first conceptualized in the mid-'60s, computers were helpful, but they had their limitations, as previously explored. They were only capable of performing one task at a time, with functions having to be queued in batches. In subsequent decades, software and hardware virtualization were invented to divide large mainframe computers into manageable entities that allowed users to connect to networks and shared sets of data simultaneously. It was believed that by dividing the mainframe and implementing virtual machines, businesses could better focus resources and increase efficiency.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

Virtualization was also intended to increase security and stability by removing dependence on a solitary device. Prior, data had to be stored on individual machines, and should that machine become compromised, either by human error or natural disaster, the data would have been lost. By storing data on a secure, shared network and granting access through virtual machines, businesses could use that data by allowing employees to access it from their separate operating systems, removing the likelihood of corruption.

The technology that allows for machine and desktop virtualization did not see significant changes between the 1960s and 90s, other than gaining access to centrally stored data from multiple devices.

When the technology became slightly more popular in the late 90's/early 2000s, it still had to be painstakingly managed on all levels, down to the code that processed actions; in 2001, the increasing migration of business files and processes online created new complex security threats which could have had lasting negative impacts. Many businesses rush to implement cloud technology as a new security tool. New developments meant files stored in the cloud were secure from external and internal threats because they could only be accessed securely with an internet connection.

Virtualization gained real traction after the Sarbanes-Oxley legislation was passed in 2004. Because of high-profile scandals, the new law introduced a strict set of management responsibilities surrounding data security. Providers, therefore, increased hosted desktop functionality as a security tool and marketed it as such.

In the present era, industry-leading companies such as Citrix and VMware are continuously releasing product revisions and updates, releasing six collectively since 2009. This represents an increasing requirement to change the technology to reflect the needs and wants of the customer base. In recent versions of hosted desktop software, such as a friendlier user interface, competitive pricing, disaster recovery technology, and bring your device integration.

CPAs) withing provides computing services, servers, focus on databases, networking, and software over the Internet whenever it is required. Cloud is a shared pool of data one can use Cloud Computing transformed databases and application software to the hastate centers. In this paper, we focus on two primary cloud computing services mainly Infrastructure as a service (IAS), and Platform as a service (PAAS) with respective service providers (AWS and Microsoft Azure) and their implementation industries

NIST defines cloud computing as a model for enabling sustainable, on-demademandwork access to a shared po of configurable com ting reso ces that an b qui ly ovisioned and free with

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

minimal management effort or service provider interaction. Cloud is network-based environment that focus on sharing computation and resources. In cloud environment several kinds of virtual machines are loaded on same physical server as infrastructure (IAAS) In This paper, some papers dealing with the security issues and their solution methodology provided by service providers is focused. The work will be focusing as follows: In the first part we describe the overview of cloud computing service delivery models and in the second part, we discussed comparative study of different service providers. The third part, deals with security analysis and the fourth we end up with the
conclusion

Three models are used to deliver cloud computing: Software or Application as a Service (SaaS), Platform as a Service (PaaS) and Infrastructure as a Service (IaaS). [10] Infrastructure-as-a-Service (IaaS): It provides the platform for computational requirements, computing resources and infrastructural utilities such as (storage and archiving, servers, and networking) to setup the needed environment for hosted applications. Example: Microsoft Azure, Amazon's EC2. Platform-as-a-Service (PaaS): It provides the platform for computational requirements, and the solution stack required by the customers to build their own applications and host their own data. Example: Google Apps, AWS Elastic Beanstalk. Software-as-a-Service (SaaS): It provides the platform for computational requirements and applications for the customers to utilize. Example: Office 365, Twitter, Myspace and Facebook, and emails access. Figure fig: Cloud shows the different cloud service delivery models and their respective domains. Generally, Infrastructure as a Service (IaaS) give users/clients more controls and flexibility and Platform as a Service (PaaS) tend to be more opinionated but has less things to maintain and support. Currently there are different cloud service providers Amazon web services, Microsoft AZURE, Google cloud and many more. These service providers allow subscribers to use various cloud computing services on demand.

In this section we are going to discuss two major cloud service providers mainly Amazon web services and Microsoft Azure based on the type of IAAS and PAAS services offered
by them. The AWS Cloud infrastructure is built around Regions and Availability Zones (AZs) [5]. It consists of many cloud services that you can use in combinations fitted to your business or organizational needs. Some of the IAAS by AWS are Amazon EC2, Amazon Elastic Block Store, Autoscaling, Elastic Load Balancing. Amazon EC2 is used for Virtual Machine hosting to provide compute capacity in the cloud.
Amazon EC2 provides you Linux and Windows based servers with cores in the range of 1 to 60+ and RAM in the range of 1.7 GB to 244 GB. Amazon EBS is used to provide storage to EC2 instances. With the help of the Autoscaling feature you can scale up or down an automatically. Amazon ELB is used to redirect traffic to healthy instance in order to handle incoming traffic [6]. PAAS offered by AWS is AWS Beanstalk, AWS Lambda [11] AWS Lambda is PAAS service offered by

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

Amazon, it is a server less compute service than runs your code irrespective of no of request as it offers automatic scaling feature. Microsoft Azure provides different IAAS services typically Azure Container service, AZURE load balancer, AZURE Autoscaling and Azure virtual machines. Where Azure container services and virtual machine services are
responsible for things like Operating systems (OS), Antivirus and load balancing. Virtual machines use virtual hard disks (VHDs) to store their operating system (OS) and data. VHDs are also used for the images you can choose from to install an OS. App services, AZURE search and Azure CDN (content
delivery network) are some of the PAAS services offered by Microsoft Azure.

Cloud computing provides a large variety of architectural configurations, such as the number of cores, amount of memory, and the number of nodes. The performance of a workload an application and its input can execute up to 20 times longer or cost 10 times more than optimal. The ready flexibility in cloud offerings has created a paradigm shift. Whereas before an application was tuned for a given cluster, in the cloud the architectural configuration is tuned for the workload. Furthermore, because the cloud has a pay-as-you-go model, each configuration (cluster size VM type) has running cost and execution time. Therefore, a workload can be optimized for least cost or shortest time which are different configurations. Choosing the right cloud configuration for an application is essential to service quality and commercial competitiveness. For instance, a bad cloud configuration can result in up to 12 times more cost for the same performance target. The saving from a proper cloud configuration is even more significant for recurring jobs [5], [9] in which similar workloads are executed repeatedly. Nonetheless, selecting the best cloud configuration, e.g., the cheapest or the fastest, is difficult due to the complexity of simultaneously achieving high accuracy, low overhead, and adaptivity for different applications and workloads.

A cloud can be defined as the software and services that run on the Internet, instead of locally on a local host system. These software and services can be accessed remotely. Example of cloud services include Netflix, Google Drive, and Microsoft OneDrive.
Also, Amazon Web Service, Microsoft Azure, IBM cloud service, Google Cloud service are among the top cloud service provider. Most Common questions that's need to answer in Cloud Computing are:

- How to evaluate and choose the right cloud solution?
- How to design applications which is optimized for the cloud?
- How to integrate public cloud applications within premise and private cloud applications?
- How to integrate different cloud solutions?
- How to setup new infrastructures so that applications running on them can easily interoperate and move to the public cloud if required? In this work, a brief survey is presented to discuss the present research work addressing all these questions and their probable solutions.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

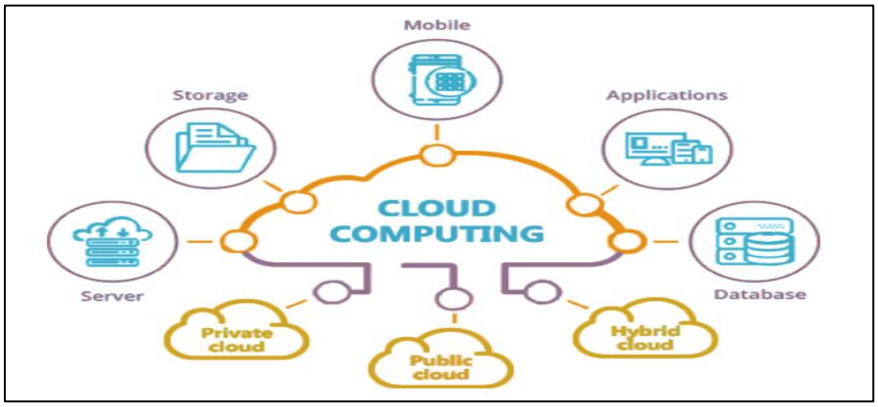

Figure 1: Cloud Computing

Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction [13]. This cloud model is composed of five essential characteristics, three service models, and four deployment models. NISTs definition identified self-service, accessibility from desktops, laptops, and mobile phones, resources that are pooled among multiple users and applications, elastic resources that can be rapidly reapportioned as needed, and measured service as the five essential characteristics of cloud computing. When these characteristics are combined, they create cloud computing infrastructure that contains both a physical layer and an abstraction layer. The physical layer consists of hardware resources that support the cloud services (i.e. servers, storage and network components). The abstraction layer consists of the software deployed across the physical layer, thereby expressing the essential characteristics of the cloud per NISTs definition

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

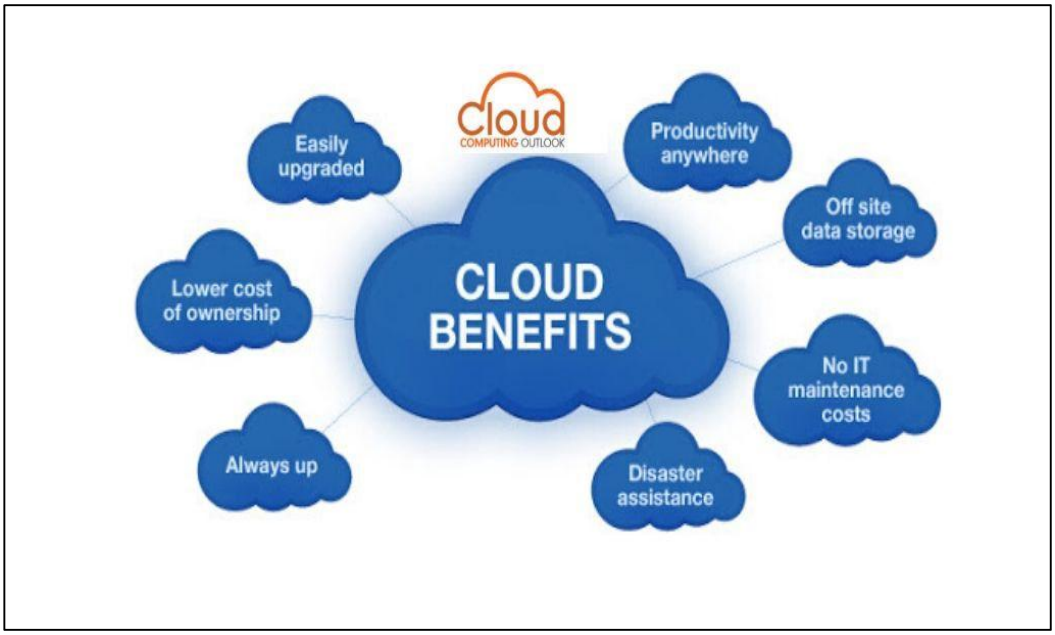

Figure 2: Benefits of Cloud Computing

Characteristics of Cloud Computing:

• On-demand self-service. A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.

• Broad network access. Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops, and workstations).

• Resource pooling. The providers computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, and network bandwidth.

• Rapid elasticity. Capabilities can be elastically provisioned and released, in some cases automatically, to Fig. 3. Cloud-Computing-Architecture scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.

• Measured service. Cloud systems automatically control and optimize resource use by leveraging a metering capability1 at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

According to Amazon, clouds enable 7 transformations of how applications are designed, built and used.

- Cloud makes distributed architectures easy
- Cloud enables users to embrace the security advantages of shared systems
- Cloud enables enterprises to move from scaling by architecture to scaling by command
- Cloud puts a supercomputer into the hands of every developer
- Cloud enables users to experiment often and fail quickly
- Cloud enables big data without big servers
- Cloud enables a mobile ecosystem for a mobile-first world

Service Models:

- Software as a Service (SaaS). The capability provided to the consumer is to use the providers applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, except for limited user specific application configuration settings.
- Platform as a Service (PaaS). The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider.3 The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.
- Infrastructure as a Service (IaaS). The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer can deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls).

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

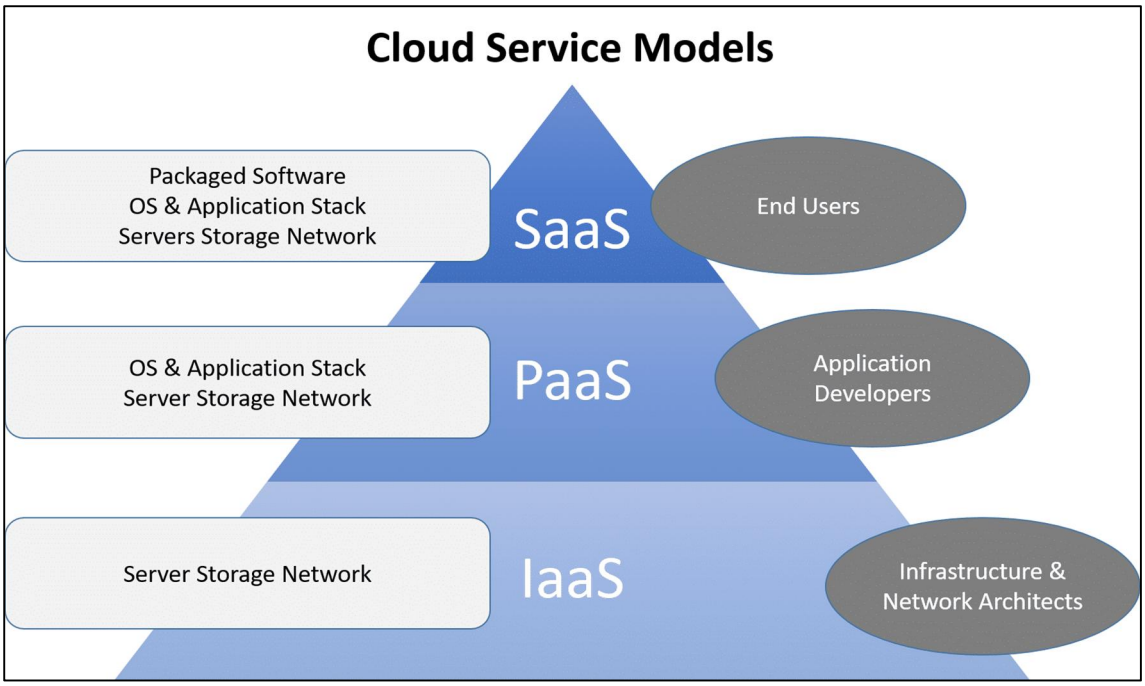

Figure 3: Cloud Computing Model

Deployment Models:

• Private cloud. The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.
• Community cloud. The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.
• Public cloud. The cloud infrastructure is provisioned for open use by the public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.
• Hybrid cloud. The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).

Challenges in Cloud Computing There are several challenges for picking the best cloud configurations for big data analytics jobs. Complex performance model:

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

The running time is affected by the number of resources in the cloud configuration in a non-linear way. For instance, a regression job on Spark (with fixed number of CPU cores) sees a diminishing return of running time at 256GB RAM. This is because the job does not benefit from more RAM beyond what it needs. Therefore, the running time only sees marginal improvements. In addition, performance under a cloud configuration is not deterministic. In cloud environments, which is shared among many tenants, stragglers can happen. [6] measured the running time of TeraSort-30GB on 22 different cloud configurations on AWS EC2 five times. In [6] they then computed the coefficient of variation (CV) of the five runs. Their results show that the median of the CV is about 10% and the 90 percentile is above 20%. This variation is not new [9]. Cost model: The cloud charges users based on the amount of time the VMs are up. Using configurations with a lot of resources could minimize the running time, but it may cost a lot more money. Thus, to minimize cost, we must find the right balance between resource prices and the running time.

Cloud Computing is hot research are since couple of years. Good number of papers have been published in this domain. Some state-of-the-artwork in cloud computing are described below. Performance Prediction: There have been several recent efforts at modeling job performance in datacenters to support SLOs or deadlines. Techniques proposed in Jockey [9] and ARIA [15] use historical traces and dynamically adjust resource allocations to meet deadlines. In Ernest we build a model with no historic information and try to minimize the amount of training data required. Bazaar [12] proposed techniques to model the network utilization of MapReduce jobs by using small subsets of data. In Ernest we capture computation and communication characteristics and use high level features that are framework independent. Projects like Tuner [16] model MapReduce jobs at very fine granularity and set optimal values for options like memory buffer sizes etc. In Ernest we use few simple features and focus on collecting training data will help us maximize their utility. Finally scheduling frameworks like Quasar [7] try to estimate the scale out and scale up factor for jobs using the progress rate of the first few tasks. Ernest on the other hand runs the entire job on small datasets and can capture how different stages of a job interact in a long pipeline. Query Optimization: Database query progress predictors [14] solve a performance prediction problem like Ernest. Database systems typically use summary statistics of the data like cardinality counts to guide this process. Further, these techniques are typically applied to a known set of relational operators. Similar ideas have also been applied to linear algebra operators [11]. In Ernest we use advanced analytics jobs where we know little about the data, or the computation being run. Recent work has also looked at providing SLAs for OLTP and OLAP workloads in the cloud and some of the observations in [6] about variation across instance types in EC2 are also known to affect database queries. Tuning and Benchmarking: Ideas related to experiment design, where we explore a space of possible inputs and choose the best inputs, have been used in other applications like server benchmarking [17]. Related techniques like Latin Hypercube Sampling have been used to efficiently explore file system design space [8]. Autotuning BLAS libraries like ATLAS [19] also solve a similar problem of exploring a state space efficiently

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

In this work, they observe that scale-out workloads share many inherent characteristics that place them into a distinct workload class from desktop, parallel, and traditional server workloads. They perform a detailed micro-architectural study of a range of scale-out workloads, finding a large mismatch between the demands of the scale-out workloads and todays predominant processor microarchitecture. They observe significant over-provisioning of the memory hierarchy and core micro-architectural resources for the scale-out workloads. Moreover, continuing the current processor trends will result in further widening the mismatch between the scale-out workloads and server processors. Conversely, they find that the characteristics of scale-out workloads can be leveraged to gain area and energy efficiency in future servers

The key contributions of this paper are:

• An experimental characterization of performance tradeoff of various VM types for realistic workloads across Amazon AWS and Microsoft Azure.

• A novel hybrid offline and online data collection and modeling framework which eliminates the O(n2) data collection overhead while providing accurate performance predictions across cloud providers.

• A detailed experimental evaluation demonstrating that PARIS accurately estimates multiple performance metrics and their variabilities (P90 values), for several real-world workloads across two major public cloud providers, thereby reducing user cost by up to 45 percent relative to strong baseline techniques.

Scout: An Experienced Guide to Find the Best Cloud Configuration

Selecting the best cloud configuration from the service provider is a challenge. Several methods have been proposed to find the best cloud configuration

These methods can be broadly classified into prediction which uses elaborate offline evaluation to generate a machine learning model that predicts the performance of workloads and search-based techniques which successively evaluate configurations looking for one that is near optimal. Prediction, as proposed in PARIS, is not reliable because of high variance in prediction results. A search-based method does not require an accurate model but can have a high evaluation cost (measured in terms of configurations evaluated). They choose the search-based method because it better tolerates prediction error and delivers effective solutions. Any search-based method has two aspects.

• Exploration: Gather more information about the search space by executing a new cloud configuration.

• Exploitation: Choose the most promising configuration based on information collected. Additional exploration incurs higher search cost, and insufficient exploration may lead to sub-optimal solutions.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

This is the exploration-exploitation dilemma appeared in many machine learning problems. For example, cherry-pick requires a good exploration strategy to characterize the search space [6]. In this paper, they argue that it is possible to trade exploration with exploitation without settling for a suboptimal configuration. The central insight of this paper is that the cost of the search for the right cloud configuration can be significantly reduced if They could learn from the historical data experiences of finding the right cloud configuration for other workloads. In this paper, they present a SCOUT, which uses historical data to find the best cloud configuration for a workload. In doing so, they enable practitioners to find a near-optimal cloud configuration with a lower search cost than state of the art. Additionally, they answer the following questions about improving the performance of the search-based method and reducing the search-cost. Their key contributions are:

• They propose a novel method, SCOUT, that finds (near) optimal solutions and solves the shortcomings of the prior work.

• They present a novel way to represent the search space, which can be used to transfer knowledge from historical measurements

• They evaluate SCOUT and other state-of-the-art methods using more than 100 workloads on three different data processing systems. and

• They make their performance data available

POPULAR CLOUD COMPUTING PLATFORM

In this section, a survey of some of the dominant cloud computing products will be discussed.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

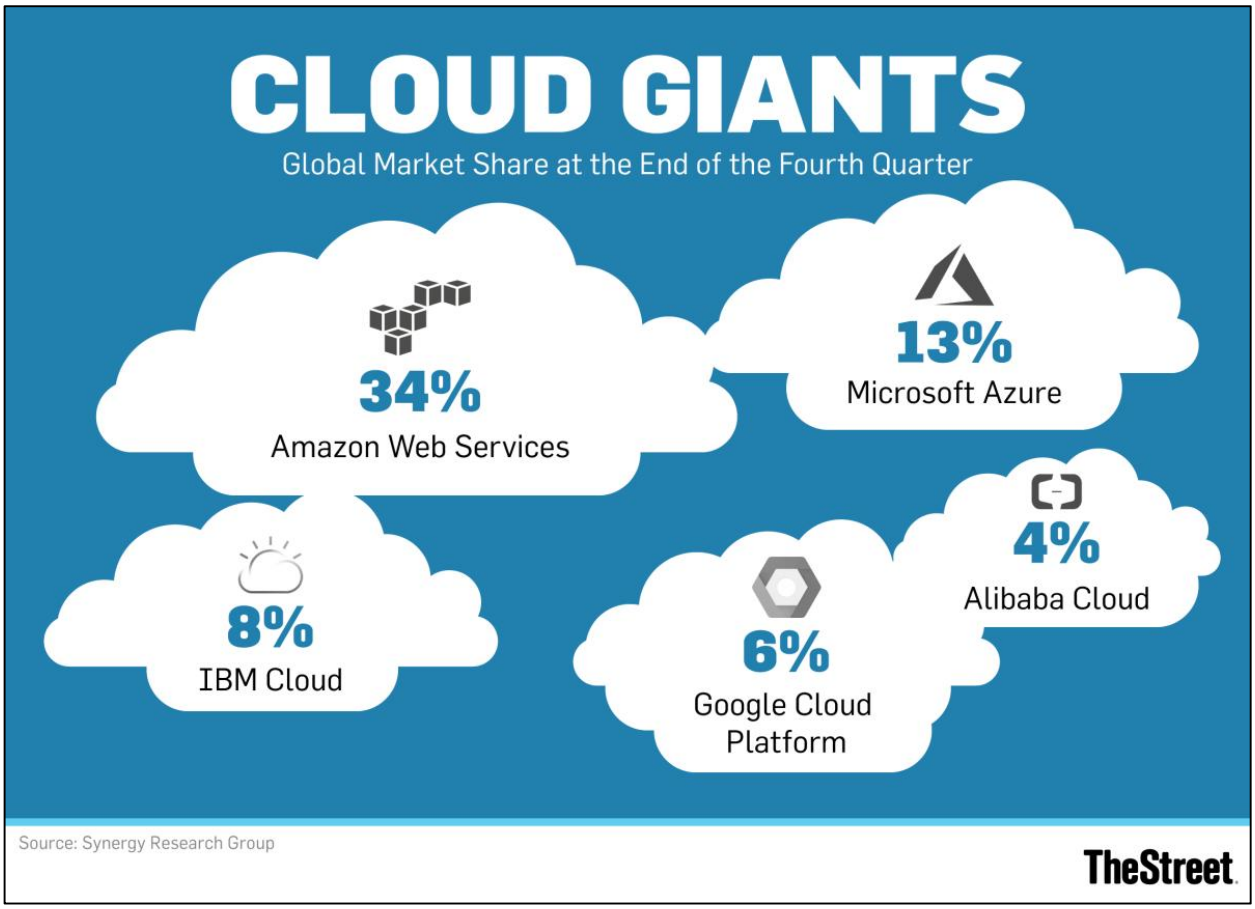

Figure 4: Cloud Computing Market Share

Amazon Web Services (AWS)

Amazon Web Services (AWS) is a set of cloud services, providing cloud-based computation, storage and other functionality that enable organizations and individuals to deploy applications and services on an on-demand basis and at commodity prices. Amazon Web Services offerings are accessible over HTTP, using REST and SOAP protocols. Amazon Elastic Compute Cloud (Amazon EC2) enables cloud users to launch and manage server instances in data centers using APIs or available tools and utilities.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

Figure 5: Amazon Web Services

EC2 instances are virtual machines running on top of the Xen virtualization engine [4]. After creating and starting an instance, users can upload software and make changes to it. When changes are finished, they can be bundled as a new machine image. An identical copy can then be launched at any time. Users have nearly full control of the entire software stack on the EC2 instances that look like hardware to them. On the other hand, this feature makes it inherently difficult for Amazon to offer automatic scaling of resources.

EC2 provides the ability to place instances in multiple locations. EC2 locations are composed of Regions and Availability Zones. Regions consist of one or more Availability Zones, are geographically dispersed. Availability Zones are distinct locations that are engineered to be insulated from failures in other Availability Zones and provide inexpensive, low latency network connectivity to other Availability Zones in the same Region. EC2 machine images are stored in and retrieved from Amazon Simple Storage Service (Amazon S3). S3 stores data as objects that are grouped in buckets. Each object contains from 1 byte to 5 gigabytes of data. Object names are essentially URI pathnames. Buckets must be explicitly created before they can be used. A bucket can be stored in one of several Regions. Users can choose a Region to optimize latency, minimize costs, or address regulatory requirements. Amazon Virtual Private Cloud (VPC) is a secure and seamless bridge between a company's existing IT infrastructure and the AWS cloud. Amazon VPC enables enterprises to connect their existing infrastructure to a set of isolated AWS compute resources via a Virtual Private Network (VPN) connection, and to extend their existing management capabilities such as security services, firewalls, and intrusion detection systems to include their AWS resources. For cloud users, Amazon CloudWatch is a useful management tool which collects raw data from partnered AWS services such as Amazon EC2 and then processes the information into readable, near real-time metrics. The metrics about EC2 include, for example, CPU utilization, network in/out bytes, disk read/write operations, etc.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

Microsoft Windows Azure platform

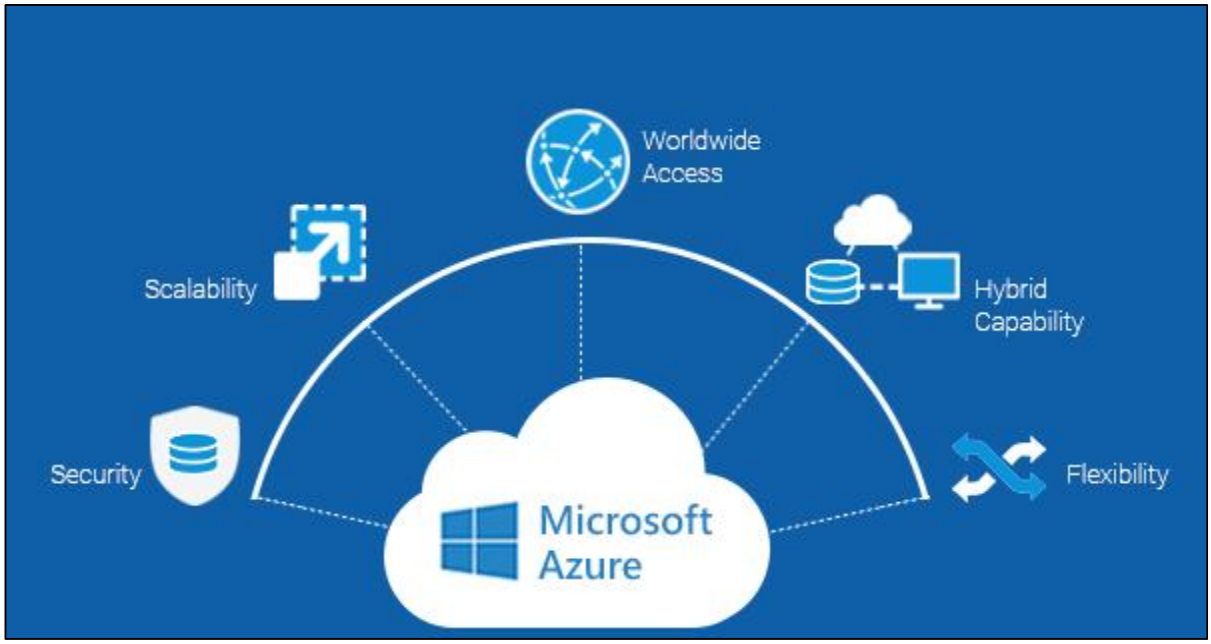

Figure 6: Microsoft Azure

Microsoft's Windows Azure platform consists of three components and each of them provides a specific set of services to cloud users. Windows Azure provides a Windows based environment for running applications and storing data on servers in data centers; SQL Azure provides data services in the cloud based on SQL Server; and .NET Services offer distributed infrastructure services to cloud-based and local applications. Windows Azure platform can be used both by applications running in the cloud and by applications running on local systems. Windows Azure also supports applications built on the .NET Framework and other ordinary languages supported in Windows systems, like C, Visual Basic, C++, and others. Windows Azure supports general-purpose programs, rather than a single class of computing. Developers can create web applications using technologies such as ASP.NET and Windows Communication Foundation (WCF), applications that run as independent background processes, or applications that combine the two. Windows Azure allows storing data in blobs, tables, and queues, all accessed in a RESTful style via HTTP or HTTPS. SQL Azure components are SQL Azure Database and Huron Data Sync. SQL Azure Database is built on Microsoft SQL Server, providing a database management system (DBMS) in the cloud. The data can be accessed using ADO.NET and other Windows data access interfaces. Users can also use on-premises software to work with this cloud-based information. Huron Data Sync synchronizes relational data across various on-premises DBMSs.

Today's researchers have access to a greater variety and volume of data than ever before. Scientific instruments and sensors can generate data around the clock day in and day out, while computer modeling and simulation programs produce rich data sets in lieu of or in conjunction with experiments. Science now includes a computational perspective that requires researchers to increasingly rely on

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

computers to aid them in gathering data, analyzing it, and fitting it to drive conclusions. In the book The Fourth Paradigm: Data Intensive Scientific Discovery, you can find many examples of the growth in tools and processes to provide more and better data to science. Although this greater access to data is a huge benefit to the advancement of scientific knowledge, the computing power required to effectively analyze this data is limited by the available technical infrastructure. For most researchers, reliance on desktop PCs and small computing clusters constrains the advancement of their research by slowing the speed at which processing can occur, increasing the cost of research because of increased processing time, and restricting the ability to share their discoveries with the larger research community when data sets become too large to easily move. Furthermore, substantial up-front financial investments and a lack of skills needed to manage an advanced computing infrastructure are significant barriers to entry preventing these researchers from transitioning to larger systems. Finally, as more publicly available scientific data sets become available, the growing volumes of this data make it impractical to move data to the desktop for analysis, but instead require moving the questions to the data. Collectively, these issues require a new computing paradigm for science—cloud computing. Advancements in cloud computing in recent years promise to remove these barriers. Microsoft has invested heavily in the development of data centers for a public cloud infrastructure, known as Microsoft Azure, which is ideally suited to serve the needs of the scientific community. Microsoft Azure provides a variety of cloud services enabling you to pick and choose the right combination to meet your needs, from setting up a community website to document and discuss research findings to performing complex data analysis in a scalable environment. Microsoft Azure has already proven successful for a variety of research projects and future enhancements promise to support research in new and exciting ways as cloud computing continues to evolve.

According to the National Institute of Standards and Technology, cloud computing "…is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction." Importantly, it dynamically scales to meet the current demand, whether the demand results from the execution of a resource intensive application by a single user or the sudden influx of multiple users requesting access to a centralized shared resource. Similarly, the cloud resources can be released once demand begins to diminish. Having a cloud infrastructure means researchers need not worry about having or hiring the skills to build, manage, and maintain a centralized and scalable environment. Instead, they can rely on continuous access to a data center managed by a third party. Rather than investing upfront to secure the necessary hardware for an infrastructure capable of supporting computational science, researchers can instead pay for access to cloud computing only as the need arises.

Access to vast arrays of managed resources is another compelling aspect of the cloud for researchers. Cloud computing platforms maintain the infrastructure and services on which applications run, such as operating systems and database services, among others. Because all hardware is abstracted by the cloud platform, there is no dependency on any specific piece of hardware. Consequently, vendors can apply patches and upgrade components without adversely impacting researchers. In addition, cloud

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

services have built-in redundancy so routine failures are handled automatically, usually with minimal or no impact to users. That means you can rely on access to the data, applications, or services you choose to relegate to the cloud on a schedule that suits you.

The architecture of a cloud service resembles that of an on-premises solution, except that the responsibility for managing individual components of the architecture differs for a cloud service. There are three models used most for cloud services:

- Infrastructure as a Service (IaaS). The cloud service provider maintains the physical or virtual machines, storage, and a networking layer, whereas you build and maintain one or more virtual machines that you load with an operating system, applications, and data. This model is stateful, which means that even when you shut down your virtual machines, their contents are saved to disk when you shut them down and are available again when you restart the machines.

- Platform as a Service (PaaS). With this model, the cloud service provider manages everything for you to support an application that you build. This model is considered best practice due to its statelessness. Application components do not persist a current state on the current node but rely on external persistent storage so that no data is lost if hardware fails.

- Software as a Service (SaaS). The cloud service provider provides everything from the hardware to the applications running on the server in this model, leaving you simply to use the application.

Microsoft Azure is Microsoft's IaaS and PaaS solution, first announced in 2008 and enhanced with additional features since then. It relies on a global network of data centers managed by Microsoft to provide a collection of services that facilitate the development, deployment, and management of scalable cloud-based applications and services. Currently, Microsoft maintains data centers in four regions in North America, two regions in Europe, and two regions in Asia, as shown in Figure 1, and has plans to expand into additional sub-regions in the future. Each region contains one or more data centers. In turn, each data center, considered state of the art design for large computing and high data volumes, holds from ten to hundreds of thousands of servers. When you create a virtual machine in the cloud or develop a cloud-based application service, you select a data center region in which to store the virtual machine or execute the code. An alternative is to use the same service in multiple regions concurrently and direct users to the region closest to them. This option also provides disaster recovery in the event of a failure or network outage in a single data center.

Google Cloud

Google Cloud is a platform for traditional web applications in Google-managed data centers. Currently, the supported programming languages are Python and Java. Web frameworks that run on the Google App Engine include Django, Cherrypie, Pylons, and web2py, as well as a custom Google-written web application framework like JSP or ASP.NET. Google handles deploying code to a cluster, monitoring,

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

failover, and launching application instances as necessary. Current APIs support features such as storing and retrieving data from a Bigtable non-relational database, making HTTP requests and caching. Developers have read only access to the filesystem on App Engine. Several papers studied the performance of big data applications on scale-out platform and clouds

All these works use performance counters to monitor the performance and behavior of applications. In, authors perform a set of comprehensive experiments to analysis the impact of memory subsystem on the performance of data intensive applications running on cloud environment. In, author uses compress sensing to improve data movement after finding the performance bottleneck using performance counters. Performance counters also can be used to trace the applications behavior to find the malicious behavior. Moreover, there are new approaches to improve the performance of modern computing systems such as hardware acceleration, and cloud computing.

Figure 7: Google Cloud Platform

Cloud computing has recently emerged as a compelling paradigm for managing and delivering services over the Internet. The rise of cloud computing is rapidly changing the landscape of information technology, and ultimately turning the long-held promise of utility computing into a reality. However, despite the significant benefits offered by cloud computing, the current technologies are not matured enough to realize its full potential. Many key challenges in this domain, including automatic resource provisioning, power management and security management, are only starting to receive attention from the research community. Therefore, I believe there is still tremendous opportunity for researchers to make groundbreaking contributions in this field and bring significant impact to their development in the industry. In this paper, I have surveyed the state-of-the-art of cloud computing, covering its essential concepts, architectural designs, prominent characteristics, key technologies as well as research directions. As the development of cloud computing technology is still at an early stage, I hope this work

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

will provide a better understanding of the design challenges of cloud computing and pave the way for further research in this area.

CHAPTER 3: METHODOLOGY AND IMPLEMENTATION

To deploy and establish an infrastructure over the cloud firstly you must get a subscription of whichever cloud provider you are using in this case we are using Microsoft Azure.

3.1 Azure Subscription

Follow these steps to retrieve the ID for a subscription in the Azure portal.

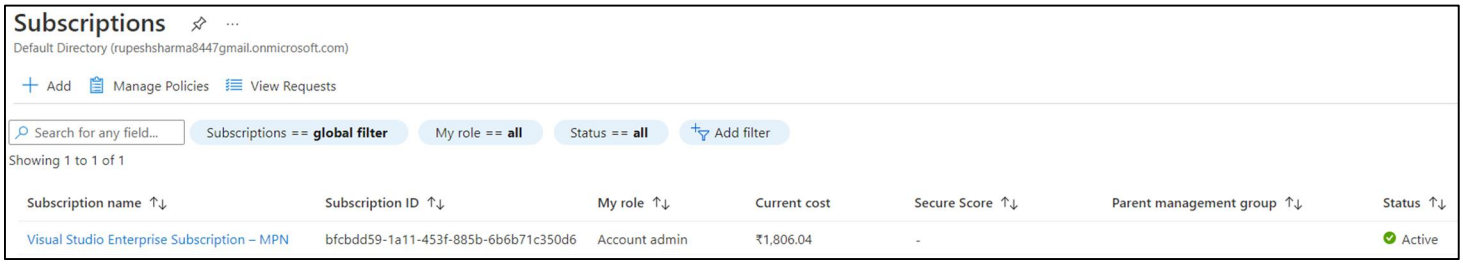

Figure 8: Azure Subscription

- Log in to the Azure Portal.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

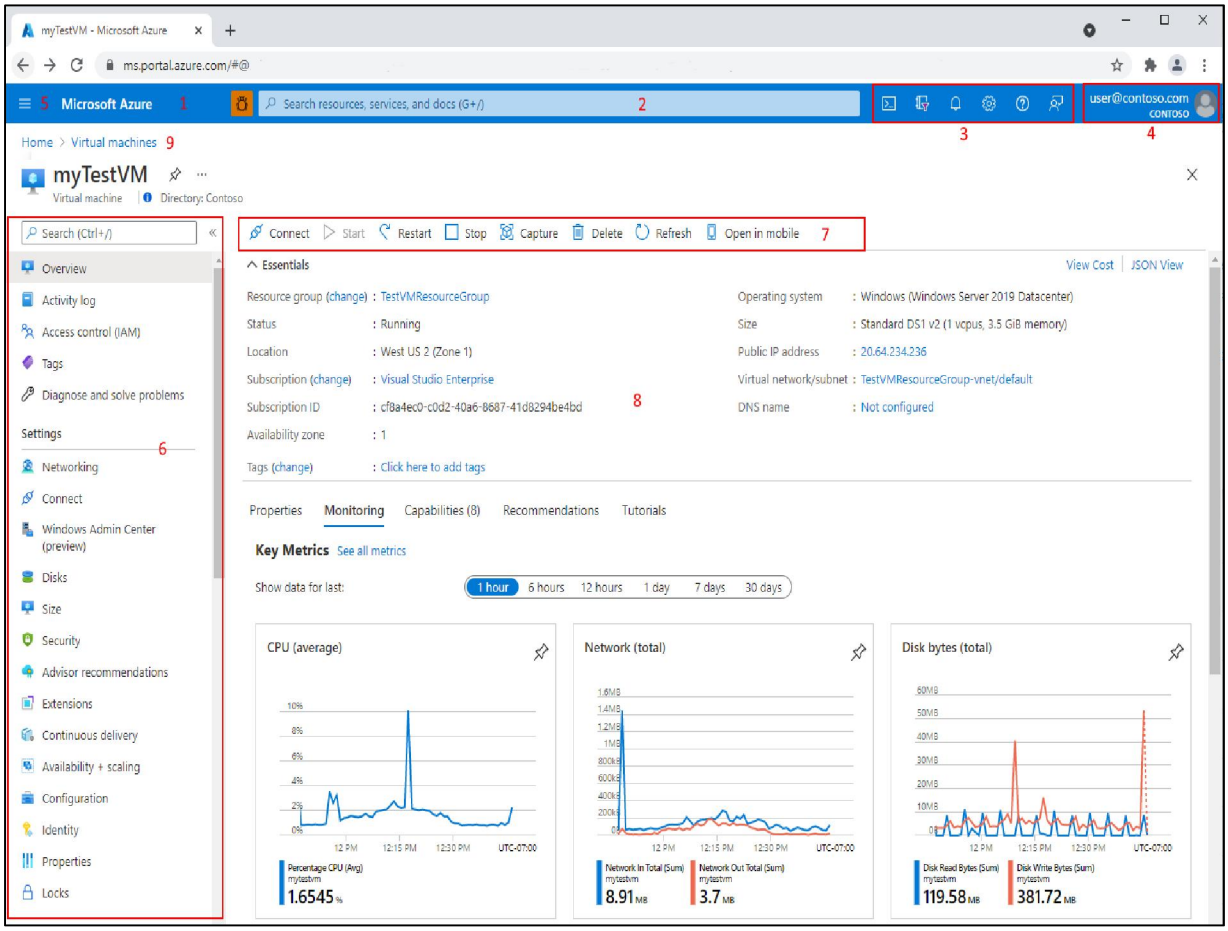

Figure 9: Azure Portal Description

KEY DESCRIPTION
1. Page header. Appears at the top of every portal page and holds global elements.
2. Global search. Use the search bar to quickly find a specific resource, a service, or documentation.
3. Global controls. Like all global elements, these features persist across the portal and include Cloud Shell, subscription filter, notifications, portal settings, help and support, and send us feedback.
4. Your account. View information about your account, switch directories, sign out, or sign in with a different account.
5. Azure portal menu. This global element can help you to navigate between services. Sometimes referred to as the sidebar. (Items 9 and 10 in this list appear in this menu.)
6. Resource menu. Many services include a resource menu to help you manage the service. You may see this

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

element referred to as the left pane. Here, you'll see commands that are contextual to your current focus.
7. Command bar. These controls are contextual to your current focus.
8. Working pane. Displays details about the resource that is currently in focus.
9. Breadcrumb. You can use the breadcrumb links to move back a level in your workflow.

- Under the Azure services heading, select **Subscriptions**. If you don't see Subscriptions here, use the search box to find it.

- Find the Subscription ID for the subscription shown in the second column. If no subscriptions appear, or you don't see the right one, you may need to switch directories to show the subscriptions from a different Azure AD tenant.

3.2 Active Directory Domain Service (ADDS)

Azure Active Directory Domain Services (Azure AD DS) provides managed domain services such as domain join, group policy, lightweight directory access protocol (LDAP), and Kerberos/NTLM authentication. You use these domain services without the need to deploy, manage, and patch domain controllers (DCs) in the cloud.

An Azure AD DS managed domain lets you run legacy applications in the cloud that can't use modern authentication methods, or where you don't want directory lookups to always go back to an on-premises AD DS environment. You can lift and shift those legacy applications from your on-premises environment into a managed domain, without needing to manage the AD DS environment in the cloud.

Azure AD DS integrates with your existing Azure AD tenant. This integration lets users sign into services and applications connected to the managed domain using their existing credentials. You can also use existing groups and user accounts to secure access to resources. These features provide a smoother lift-and-shift of on-premises resources to Azure.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

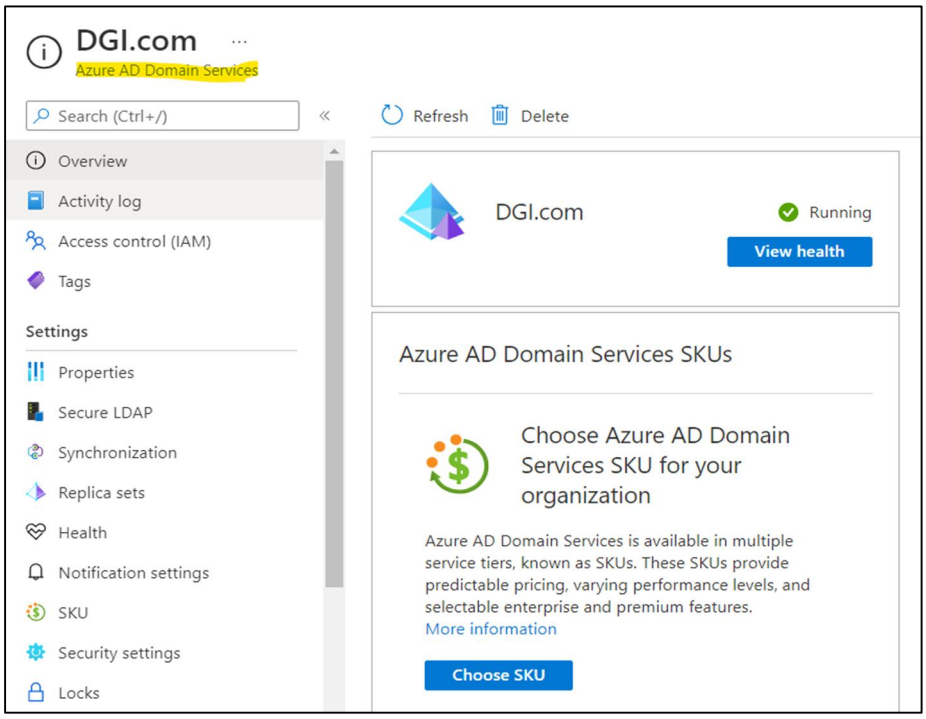

Figure 10: Active Directory Domain Services (DGI.com)

When you create an Azure AD DS managed domain, you define a unique namespace. This namespace is the domain name, such as “DGI.com”.

3.3 Virtual Network (VNET)

A virtual network is a network where all devices, servers, virtual machines, and data centers that are connected are done so through software and wireless technology. This allows the reach of the network to be expanded as far as it needs to for peak efficiency, in addition to numerous other benefits.

Azure Virtual Network (Vnet) is the fundamental building block for your private network in Azure. Vnet enables many types of Azure resources, such as Azure Virtual Machines (VM), to securely communicate with each other, the internet, and on-premises networks. Vnet is like a traditional network that you'd operate in your own data center but brings with it additional benefits of Azure's infrastructure such as scale, availability, and isolation.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

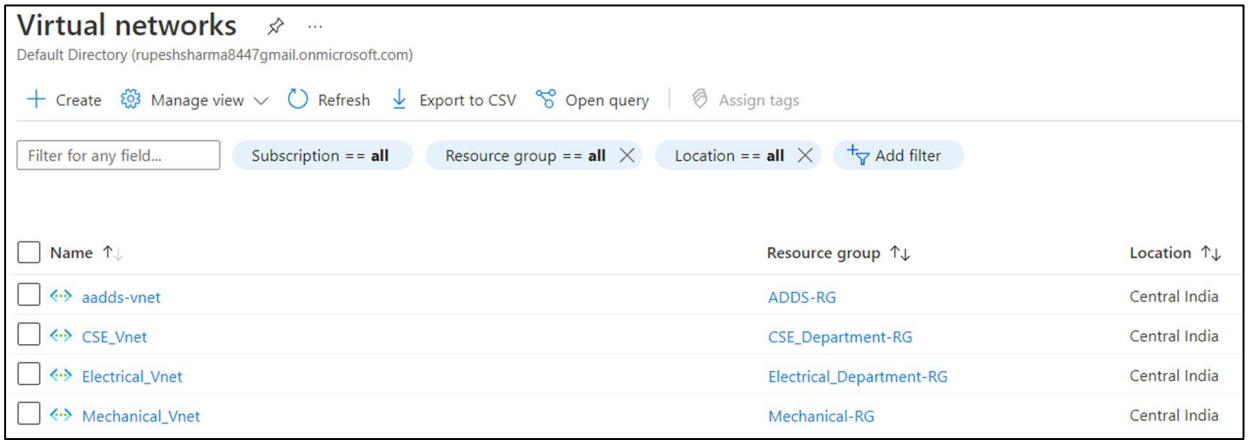

Figure 11: Virtual Networks

Virtual Network Components

Following are the virtual network components:

- Subnets
- Network Security Groups

What are Subnets?

Each Virtual Network can be divided into sub parts, these sub parts are called subnets.

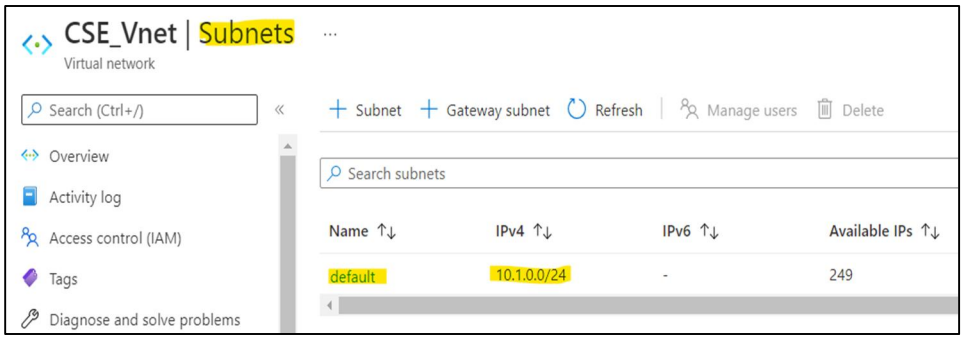

Figure 12: Subnet

A subnet can further be divided into:

- Private Subnet – A network in which there is no internet access.
- Public Subnet – A network in which there is internet access

Let's look at an example and understand how Virtual Networks are used:

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

3.4 Network Security Groups

This is where you do all your connection settings, like which ports to open, by default all are closed. Don't get scared, this blog will guide you through all the settings, and all of them are very easy to configure.

But first, let me show you how the final architecture for a Virtual Network looks like:

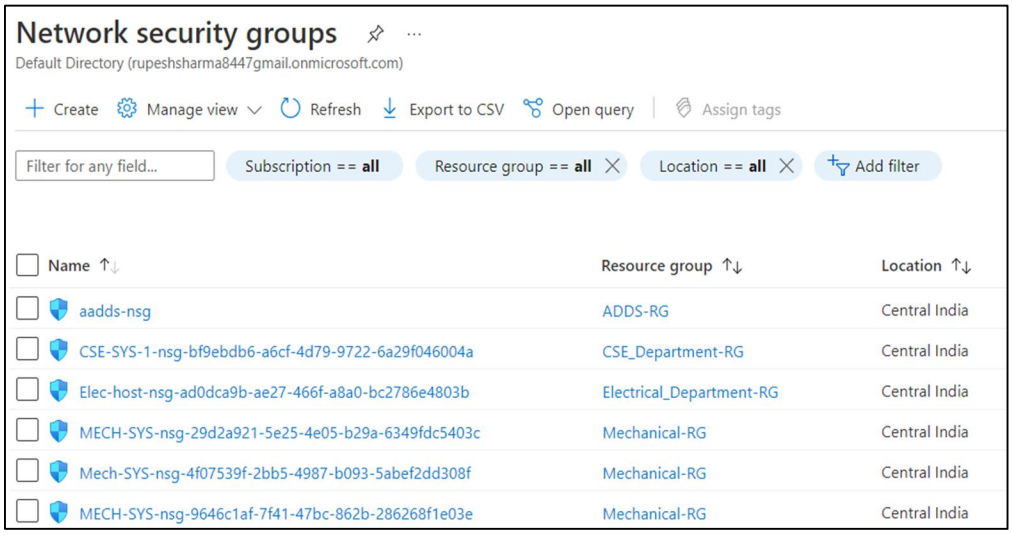

Figure 13: Network Security Groups

This is how Virtual Network works:

- First you create a virtual network.
- Then, in this virtual network you create subnets.
- You associate each subnet with the respective Virtual Machines or Cloud Instances.
- Attach the relevant Network Security Group to each subnet.
- Configure the properties in the NSGs and you are set!

3.5 Azure Virtual Desktop

Azure Virtual Desktop (AVD) is an Azure service that, combined with appropriate licenses, services, and resources, delivers a complete virtualized multi-user Windows 10 (or a single-user Windows 7) experience together with Office 365 Pro Plus. AVD includes centralized management and monitoring; system administrators can quickly deploy and manage desktops, apps, and Windows servers in the Azure Cloud.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

AVD helps businesses seamlessly scale their virtualization requirements while benefitting from the top-of-the-line security features on Azure along with the cost benefits of its subscription-based model. With AVD, users can enjoy a richer virtualization experience for accessing hosted applications when compared to the existing Windows Server-based Remote Desktop Services (RDS) platform that leverages Microsoft Partner community support for similar solutions.

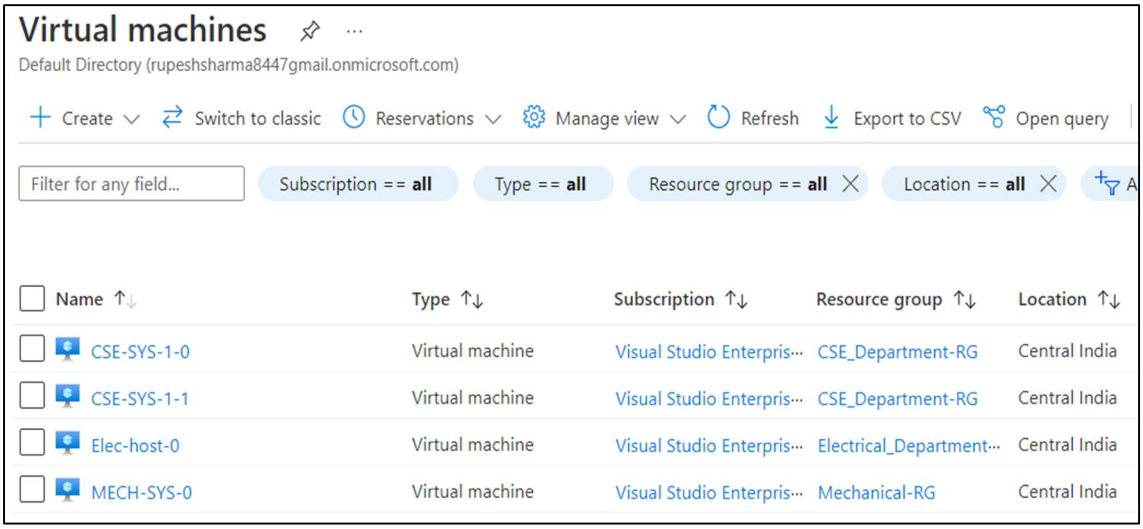

Figure 14: Azure Virtual Desktop (AVD)

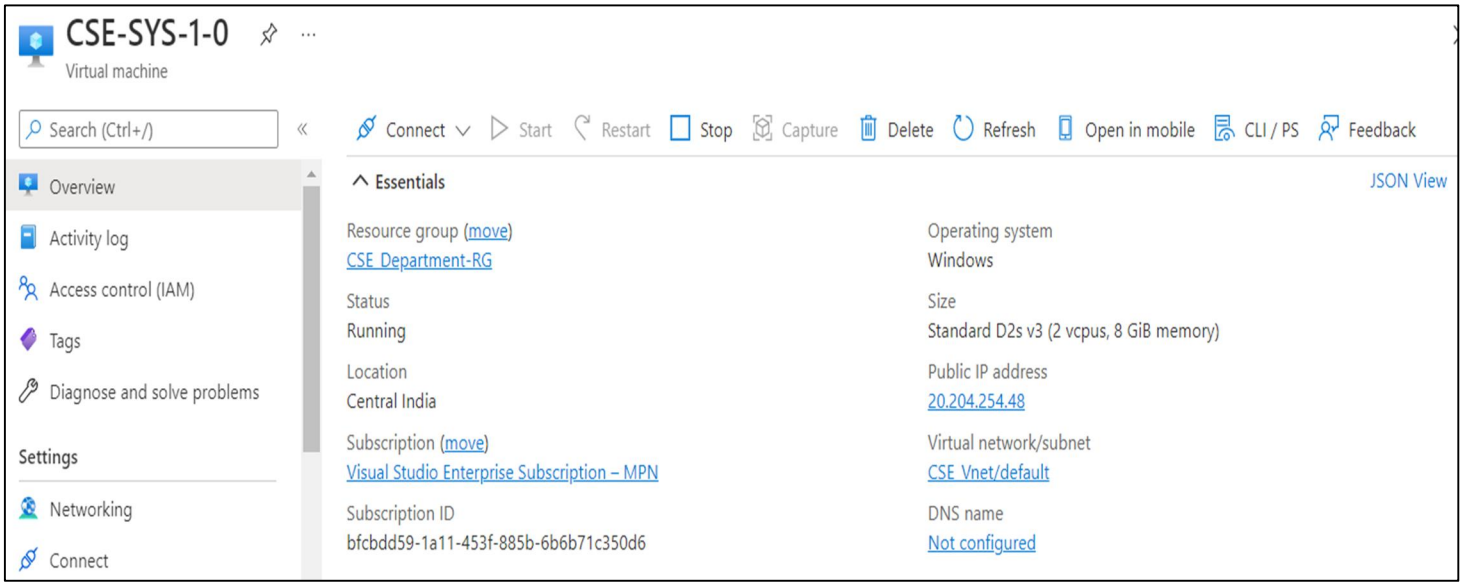

Figure 15: Computer Science Department AVD

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

3.5.1 Ideal WVD Use Scenarios

Businesses of varying sizes (10-1000+) with a significant percentage of mobile staff working from various locations would find Azure Windows Virtual Desktop an ideal solution. Here are some scenarios that benefit from WVD.

- Security and Regulation Compliance Requirements
- Financial Services, Healthcare, and Government sectors with their regulatory policies and rigorous security norms will not have to worry about data privacy, security, and compliance with WVD.
- Standardizing Operations
- When businesses undergo mergers and acquisitions or must collaborate with contractor businesses, employees can be provided with WVD for standardizing operations.
- Enabling Workforce Flexibility
- Employees with BYOD and mobile staff, call center workforce, and branch employees working from various locations can work with a unified solution despite the differences in the devices they might be using to access the organization's resources.
- Specialized Workload Cases
- Engineering & Designing companies, businesses using legacy applications in the IT sector for software development and testing often require unique solutions for deploying their unique workloads. Azure WVD is the perfect solution for managing such scenarios.
- Windows Virtual Desktop Advantages
- WVD offers serval advantages over traditional virtualization solutions.

3.5.2 Windows Virtual Desktop Advantages

WVD offers serval advantages over traditional virtualization solutions.
Decreases costs

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

By moving their desktops and applications to Azure with WVD, businesses can efficiently manage their cost constraints for running virtual desktops with Windows OS such as Windows 10 and 7.
Windows 10 Virtual Desktop reduces costs in the cloud by allowing session-based desktops on Azure Virtual Machines (VMs) which results in improved utilization of costly resources. By simplifying licensing for the use of Windows 10 on the WVD environment, and not requiring additional Client Access Licenses (CALs) for access, as a Server OS would require, the Microsoft licensing costs come down.

Companies that prefer Windows 7 on a desktop can continue getting free updates and support for Windows 7 for another three years, which expired in January 2020.
Improves flexibility

Moving Office IT infrastructure to the Cloud simplifies operations for both IT staff and the end-users.

End-users can access their desktops and applications from various devices and browsers. Office 365 applications have now been optimized to work with WVD, and the end-user experience is smoother.

After moving to the Cloud, IT staff do not spend as much time managing physical machines and local networks. All apps, data, and resources can move to the Cloud.
Desktop apps are easier to manage as WVD allows Windows 10 or Windows 7 OS. IT staff can generate persistent and personal desktops from desktop images.
In WVD, the user profiles and apps are stored in separate containers, and this can improve flexibility and scalability.

Superior Office 365 Compatibility
Other advantages of moving to Azure WVD are the availability of an infinite variety of services and resources in the Cloud and better integration with Office 365 products. Office 365 products, themselves on the Microsoft Cloud, will see improved performance in WVD environments.

3.5.3 Windows Virtual Desktop Implementation Requirements

While Azure Cloud takes away the headaches of dealing with physical machines, cabling, and network equipment; deploying Azure WVD requires at least mid-level system administration and networking skills, and familiarity with Azure services and resources.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

Remember, when your resources are in the Cloud, they need to be managed in a manner that is different from on-premises resources. Unless you have such resources in-house, it is advisable to find a dependable Azure Cloud Service Provider to collaborate with you.

3.5.4 Azure Windows Virtual Desktop Components

When you are buying an Azure WVD subscription from Apps4Rent, here is what you get as a part of the plan.

- Subscription to Windows 10 Enterprise for each WVD user
- Azure Active Directory (AAD) tenant
- Active Directory Domain Services (AD DS) deployment
- Azure subscription
- File Server

Subscription to Windows 10 Enterprise

Windows Virtual Desktop Management Service and Windows 10 desktop OS are licensed via a subscription to Windows 10 Enterprise. Unlike a standalone solution when you would need an Office 365 or Windows 10 Enterprise license, you get a bundled solution with Apps4Rent. You can use the same per-user subscription license that comes with your plan.

Azure Active Directory (AAD) tenant

For deploying and managing Windows Virtual Desktop and assigning to users to desktops/apps, an Azure Active Directory (AAD) tenant is required. This is included in our plan.
The tenant gets a unique domain name called “Directory” or “Account.” Apps4Rent will provide and manage the tenant for you so that you can start working on WVD without worrying about its administration.

Azure Subscription

Azure subscription is required for creating and running Windows Virtual Desktop session host virtual machines (VMs). The subscription includes WVD Management Service, Windows 10 VMs, and infrastructure. The plans include the cost of all these components.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

File Server

WVD uses Logix profile management technology to enhance, enable, and simplify pooled (non-persistent) Windows computing environments.

The user profiles (encased in virtual hard disk files) are stored in a file server free from Windows 10 session host virtual machines. When a user is allocated a non-persistent/pooled desktop, the profiles (including the Search cache) remain available independent of the virtual desktop machine the user logs into. With Apps4Rent plans, user profile management is taken care of by Apps4Rent.

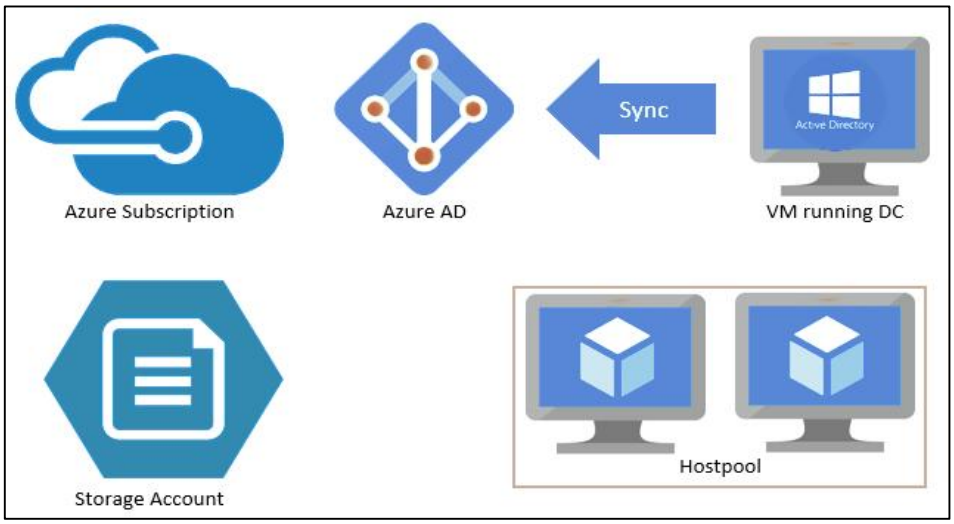

Figure 16: AVD Working Process

3.6 Azure Active Directory (AAD)

Azure Active Directory (Azure AD) is a cloud-based identity and access management service. This service helps your employees access external resources, such as Microsoft 365, the Azure portal, and thousands of other SaaS applications. Azure Active Directory also helps them access internal resources like apps on your corporate intranet network, along with any cloud apps developed for your own organization. For more information about creating a tenant for your organization, see QuickStart.

To learn the differences between Active Directory and Azure Active Directory, see Compare Active Directory to Azure Active Directory. You can also refer to Microsoft Cloud for Enterprise Architects Series posters to better understand the core identity services in Azure like Azure AD and Microsoft-365.

Azure AD is intended for:

IT admins: As an IT admin, use Azure AD to control access to your apps and your app resources, based on your business requirements. For example, you can use Azure AD to require multi-factor

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

authentication when accessing important organizational resources. You can also use Azure AD to automate user provisioning between your existing Windows Server AD and your cloud apps, including Microsoft 365. Finally, Azure AD gives you powerful tools to automatically help protect user identities and credentials and to meet your access governance requirements. To get started, sign up for a free 30-day Azure Active Directory Premium trial.

App developers: As an app developer, you can use Azure AD as a standards-based approach for adding single sign-on (SSO) to your app, allowing it to work with a user's pre-existing credentials. Azure AD also provides APIs that can help you build personalized app experiences using existing organizational data. To get started, sign up for a free 30-day Azure Active Directory Premium trial. For more information, you can also see Azure Active Directory for developers.

3.7 Assign WVD To Users

To assign a personal desktop in the Azure portal:

- Sign into the Azure portal at https://portal.azure.com.
- Enter Azure Virtual Desktop into the search bar.
- Under Services, select Azure Virtual Desktop.
- At the Azure Virtual Desktop page, go the menu on the left side of the window and select Host pools.
- Select the host pool you want to modify user assignment for.
- Next, go to the menu on the left side of the window and select Session hosts.
- Select the checkbox next to the session host you want to reassign to a different user, select the ellipses at the end of the row, and then
- Select Assign to a different user. You can also select Assignment>Assign to a different user.
- Select the user you want to assign the session host to from the list of available users.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

- When you are done, click on Select.

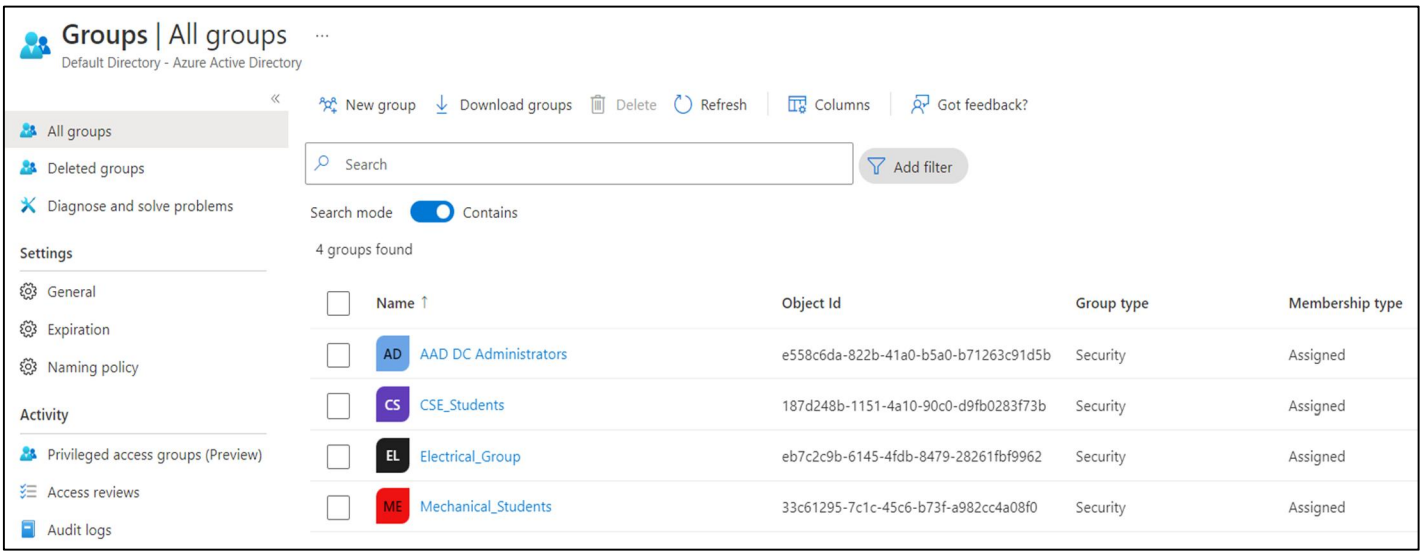

Figure 17: Groups of Students

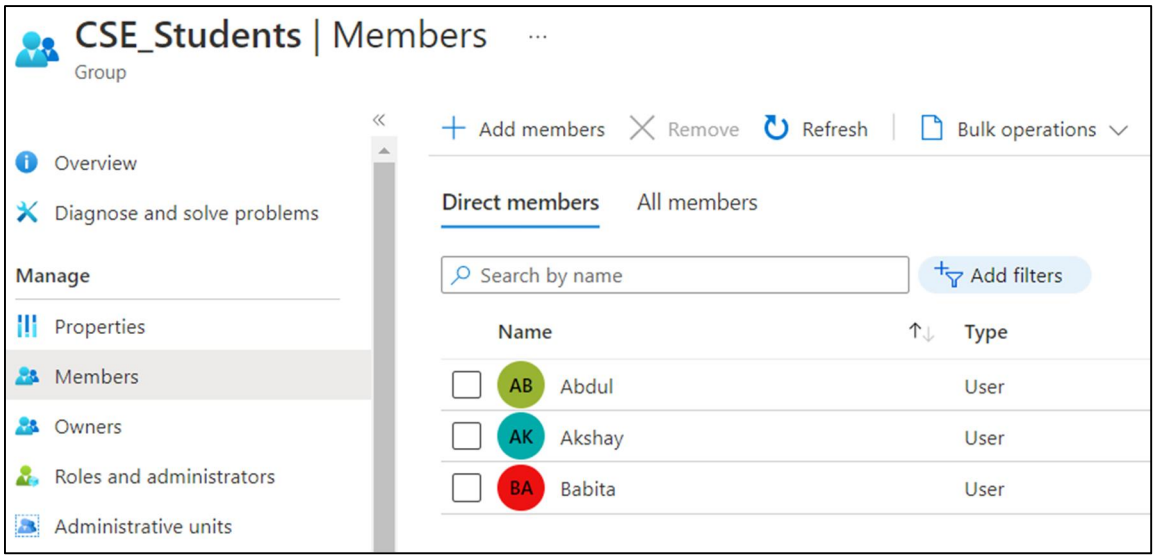

Figure 18: CSE Department Users

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

Chapter 4: Backup and File Share

4.1 Azure Backup Solution

- **On-premises** - Back up files, folders, system state using the Microsoft Azure Recovery Services (MARS) agent. Or use the DPM or Azure Backup Server (MABS) agent to protect on-premises VMs (Hyper-V and VMware) and other on-premises workloads

- **Azure VMs** - Back up entire Windows/Linux VMs (using backup extensions) or back up files, folders, and system state using the MARS agent.

- **Azure Managed Disks** - Back up Azure Managed Disks

- **Azure Files shares** - Back up Azure File shares to a storage account

- **SQL Server in Azure VMs** - Back up SQL Server databases running on Azure VMs

- **SAP HANA databases in Azure VMs** - Backup SAP HANA databases running on Azure VMs
- **Azure Database for PostgreSQL servers** - Back up Azure PostgreSQL databases and retain the backups for up to 10 years

- **Azure Blobs** - Overview of operational backup for Azure Blobs

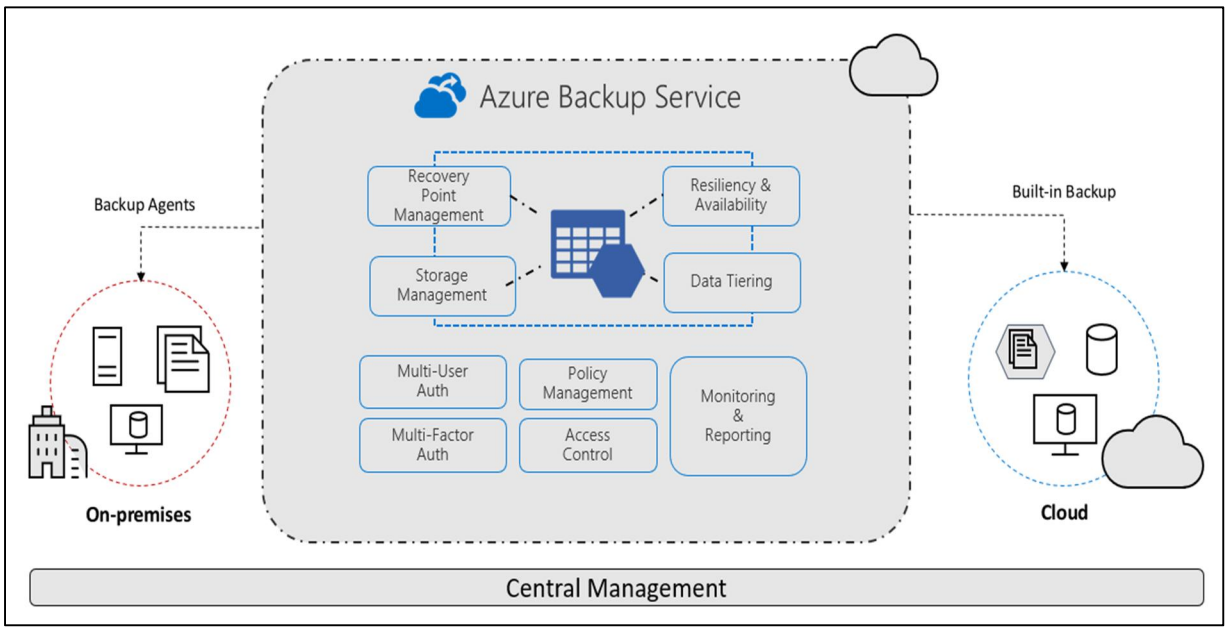

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

Figure 19: Azure Backup Services

Why use Azure Backup?

Azure Backup delivers these key benefits:

- Offload on-premises backup: Azure Backup offers a simple solution for backing up your on-premises resources to the cloud. Get short and long-term backup without the need to deploy complex on-premises backup solutions.

- Back up Azure IaaS VMs: Azure Backup provides independent and isolated backups to guard against accidental destruction of original data. Backups are stored in a Recovery Services vault with built-in management of recovery points. Configuration and scalability are simple, backups are optimized, and you can easily restore as needed.

- Scale easily - Azure Backup uses the underlying power and unlimited scale of the Azure cloud to deliver high-availability with no maintenance or monitoring overhead.

- Get unlimited data transfer: Azure Backup doesn't limit the amount of inbound or outbound data you transfer, or charge for the data that's transferred.

- Outbound data refers to data transferred from a Recovery Services vault during a restore operation.

- If you perform an offline initial backup using the Azure Import/Export service to import large amounts of data, there's a cost associated with inbound data. Learn more.

- Keep data secure: Azure Backup provides solutions for securing data in transit and at rest.

- Centralized monitoring and management: Azure Backup provides built-in monitoring and alerting capabilities in a Recovery Services vault. These capabilities are available without any additional management infrastructure. You can also increase the scale of your monitoring and reporting by using Azure Monitor.

- Get app-consistent backups: An application-consistent backup means a recovery point has all required data to restore the backup copy. Azure Backup provides application-consistent backups, which ensure additional fixes aren't required to restore the data. Restoring application-consistent data reduces the restoration time, allowing you to quickly return to a running state.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

- Retain short and long-term data: You can use Recovery Services vaults for short-term and long-term data retention.

- Automatic storage management - Hybrid environments often require heterogeneous storage - some on-premises and some in the cloud. With Azure Backup, there's no cost for using on-premises storage devices. Azure Backup automatically allocates and manages backup storage, and it uses a pay-as-you-use model. So, you only pay for the storage you consume. Learn more about pricing.

- Multiple storage options - Azure Backup offers three types of replications to keep your storage/data highly available.

- Locally redundant storage (LRS) replicates your data three times (it creates three copies of your data) in a storage scale unit in a datacenter. All copies of the data exist within the same region. LRS is a low-cost option for protecting your data from local hardware failures.

- Geo-redundant storage (GRS) is the default and recommended replication option. GRS replicates your data to a secondary region (hundreds of miles away from the primary location of the source data). GRS costs more than LRS, but GRS provides a higher level of durability for your data, even if there's a regional outage.

- Zone-redundant storage (ZRS) replicates your data in availability zones, guaranteeing data residency and resiliency in the same region. ZRS has no downtime. So your critical workloads that require data residency, and must have no downtime, can be backed up in ZRS.

Azure Backup comprehensively protects your data assets in Azure through a simple, secure, and cost-effective solution that requires zero-infrastructure. It's Azure's built-in data protection solution for a wide range of workloads. It helps protect your mission critical workloads running in the cloud, and ensures your backups are always available and managed at scale across your entire backup estate.

Azure Backup enables data protection for various workloads (on-premises and cloud). It's a secure and reliable built-in data protection mechanism in Azure. It can seamlessly scale its protection across multiple workloads without any management overhead for you. There are multiple automation channels as well to enable this (via PowerShell, CLI, Azure Resource Manager templates, and REST APIs.)

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

- Scalable, durable, and secure storage: Azure Backup uses reliable Blob storage with in-built security and high availability features. You can choose LRS, GRS, or RA-GRS storages for your backup data.

- Native workload integration: Azure Backup provides native integration with Azure Workloads (VMs, SAP HANA, SQL in Azure VMs and even Azure Files) without requiring you to manage automation or infrastructure to deploy agents, write new scripts or provision storage.

Data plane:

- Automated storage management: Azure Backup automates provisioning and managing storage accounts for the backup data to ensure it scales as the backup data grows.

- Malicious delete protection: Protect against any accidental and malicious attempts for deleting your backups via soft delete of backups. The deleted backup data is stored for 14 days free of charge and allows it to be recovered from this state.
- Secure encrypted backups: Azure Backup ensures your backup data is stored in a secure manner, leveraging built-in security capabilities of the Azure platform, such as Azure role-based access control (Azure RBAC) and Encryption.

- Backup data lifecycle management: Azure Backup automatically cleans up older backup data to comply with the retention policies. You can also tier your data from operational storage to vault storage.

- Protected critical operations: Multi-user authorization (MUA) for Azure Backup allows you to add an additional layer of protection to critical operations on your Recovery Services vaults.

Management plane:

- Access control: Vaults (Recovery Services and Backup vaults) provide the management capabilities and are accessible via the Azure portal, Backup Center, Vault dashboards, SDK, CLI, and even REST APIs. It's also an Azure role-based access control (Azure RBAC)

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

boundary, providing you the option to restrict access to backups only to authorized Backup Admins.

- Policy management: Azure Backup Policies within each vault define when the backups should be triggered and the duration they need to be retained. You can also manage these policies and apply them across multiple items.

- Monitoring and Reporting: Azure Backup integrates with Log Analytics and provides the ability to see reports via Workbooks as well.

- Snapshot management: Azure Backup takes snapshots for some Azure native workloads (VMs and Azure Files), manages these snapshots, and allows fast restores from them. This option drastically reduces the time to recover your data to the original storage.

Vault considerations:

Azure Backup uses vaults (Recovery Services and Backup vaults) to orchestrate, manage backups, and store backed-up data. Effective vault design helps organizations establish a structure to organize and manage the backup assets in Azure to support your business priorities. Consider the following guidelines when creating a vault.

Single or multiple vaults:

To use a single vault or multiple vaults to organize and manage your backup, see the following guidelines:

- Protect resources across multiple regions globally: If your organization has global operations across North America, Europe, and Asia, and your resources are deployed in East-US, UK West, and East Asia. One of the requirements of Azure Backup is that the vaults are required to be present in the same region as the resource to be backed-up. Therefore, you should create three separate vaults for each region to protect your resources.

- Protect resources across various Business Units and Departments: Consider that your business operations are divided into three separate Business Units (BU), and each business unit has its own set of departments (five departments - Finance, Sales, HR, R & D, and Marketing). Your business needs may require each department to manage and access their own backups and restores; also, enable them to track their individual usage and cost

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

expense. For such scenarios, we recommend you create one vault for each department in a BU. This way, you'll have 15 Vaults across your organization.

- Protect different workloads: If you plan to protect different types of workloads (such as 150 VMs, 40 SQL databases, and 70 PostgreSQL databases), then we recommend you create separate vaults for each type of workload (for this example, you need to create three vaults for each workload - VMs, SQL databases, and PostgreSQL databases). This helps you to separate access boundaries for the users by allowing you to grant access (using Azure role-based access control – Azure RBAC) to the relevant stakeholders.

- Protect resources running in multiple environments: If your operations require you to work on multiple environments, such as production, non-production, and developer, then we recommend you create separate vaults for each.

- Protect large number (1000+) of Azure VMs: Consider that you have 1500 VMs to back up. Azure Backup allows only 1000 Azure VMs to be backed-up in one vault. For this scenario, you can create two different vaults and distribute the resources as 1000 and 500 VMs to respective vaults or in any combination considering the upper limit.

- Protect large number (2000+) of diverse workloads: While managing your backups at scale, you'll protect the Azure VMs, along with other workloads, such as SQL and SAP HANA database running on those Azure VMs. For example, you've 1300 Azure VMs and 2500 SQL databases to protect. The vault limits allow you to back up 2000 workloads (with a restriction of 1000 VMs) in each vault. Therefore, mathematically you can back up 2000 workloads in one vault (1000 VMs + 1000 SQL databases) and rest 1800 workloads in a separate vault (300 VMs + 1500 SQL databases).

- However, this type of segregation isn't recommended as you won't be able to define access boundaries and the workloads won't be isolated from each other. So, to distribute the workloads correctly, create four vaults. Two vaults to back up the VMs (1000 VMs + 300 VMs) and the other two vaults to back up the SQL databases (2000 databases + 500 databases).

- You can manage them with:
- Backup center allows you to have a single pane to manage all Backup tasks. Learn more here.

- If you need consistent policy across vaults, then you can use Azure Policy to propagate backup policy across multiple vaults. You can write a custom Azure Policy definition that

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

uses the 'dewpoint exists' effect to propagate a backup policy across multiple vaults. You can also assign this Azure Policy definition to a particular scope (subscription or RG), so that it deploys a 'backup policy' resource to all Recovery Services vaults in the scope of the Azure Policy assignment. The settings of the backup policy (such as backup frequency, retention, and so on) should be specified by the user as parameters in the Azure Policy assignment.

- As your organizational footprint grows, you might want to move workloads across subscriptions for the following reasons: align by backup policy, consolidate vaults, trade-off on lower redundancy to save on cost (move from GRS to LRS). Azure Backup supports moving a Recovery Services vault across Azure subscriptions, or to another resource group within the same subscription

Protection of backup data from unintentional deletes with soft delete:

You may encounter scenarios where you've mission-critical backup data in a vault, and it gets deleted accidentally or erroneously. Also, a malicious actor may delete your production backup items. It's often costly and time-intensive to rebuild those resources and can even cause crucial data loss. Azure Backup provides safeguard against accidental and malicious deletion with the Soft-Delete feature by allowing you to recover those resources after they are deleted.
With soft delete, if a user deletes the backup (of a VM, SQL Server database, Azure file share, SAP HANA database), the backup data is retained for 14 additional days, allowing the recovery of that backup item with no data loss. The additional 14 days retention of backup data in the soft delete state doesn't incur any cost.

Multi-User Authorization (MUA):

Any administrator that has the privileged access to your backup data has the potential to cause irreparable damage to the system. A rogue admin can delete all your business-critical data or even turn off all the security measures that may leave your system vulnerable to cyber-attacks.
Azure Backup provides you with the Multi-User Authorization (MUA) feature to protect you from such rouge administrator attacks. Multi-user authorization helps protect against a rogue administrator performing destructive operations (that is, disabling soft delete), by ensuring that every privileged/destructive operation is done only after getting approval from a security administrator.

Ransomware Protection:

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

- Direct access to Azure Backup data to encrypt by malicious actor is ruled out, as all operations on backup data can only be performed through Recovery-Services vault or Backup Vault, which can be secured by Azure role-based access control (Azure RBAC) and MUA.
- By enabling soft delete on backup data (which is enabled by default) will hold deleted data for 14 days (at free of cost). Disabling soft delete can be protected using MUA.
- Use longer retention (weeks, months, years) to ensure clean backups (not encrypted by ransomware) don't expire prematurely, and there're strategies in place for early detection and mitigation of such attacks on source data.

Internet connectivity:

- Azure VM backup: All the required communication and data transfer between storage and Azure Backup service happens within the Azure network without needing to access your virtual network. So backup of Azure VMs placed inside secured networks don't require you to allow access to any IPs or FQDNs.

- SAP HANA databases on Azure VM, SQL Server databases on Azure VM: Requires connectivity to the Azure Backup service, Azure Storage, and Azure Active Directory. This can be achieved by using private endpoints or by allowing access to the required public IP addresses or FQDNs. Not allowing proper connectivity to the required Azure services may lead to failure in operations like database discovery, configuring backup, performing backups, and restoring data. For complete network guidance while using NSG tags, Azure firewall, and HTTP Proxy, refer to these SQL and SAP HANA articles.

- Hybrid: The MARS (Microsoft Azure Recovery Services) agent requires network access for all critical operations - install, configure, backup, and restore. The MARS agent can connect to the Azure Backup service over Azure ExpressRoute by using public peering (available for old circuits) and Microsoft peering, using private endpoints or via proxy/firewall with appropriate access controls.

Azure Backup cost considerations:

The Azure Backup service offers the flexibility to effectively manage your costs; also, meet your BCDR (business continuity and disaster recovery) business requirement. Consider the following guidelines:

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

- Use the pricing calculator to evaluate and optimize cost by adjusting various levers. Learn more

- Optimize backup policy,

- Optimize schedule and retention settings based on workload archetypes (such as mission-critical, non-critical).

- Optimize retention settings for Instant Restore.

- Choose the right backup type to meet requirements, while taking supported backup types (full, incremental, log, differential) by the workload in Azure Backup.
- Reduce the backup storage cost with Selectively backup disks: Exclude disk (preview feature) provides an efficient and cost-effective choice to selectively back up critical data. For example, you can back up only one disk when you don't want to back up all disks attached to a VM. This is also useful when you have multiple backup solutions. For example, to back up your databases or data with a workload backup solution (SQL Server database in Azure VM backup), use Azure VM level backup for selected disks.

- Speed up your Restores and minimize RTO using the Instant Restore feature: Azure Backup takes snapshots of Azure VMs and stores them along with the disks to boost recovery point creation and to speed up restore operations. This is called Instant Restore. This feature allows a restore operation from these snapshots by cutting down the restore times. It reduces the time needed to transform and copy data back from the vault. Therefore, it'll incur storage costs for the snapshots taken during this period. Learn more about Azure Backup Instant Recovery capability.

- Choose correct replication type: Azure Backup vault's Storage Replication type is set to Geo-redundant (GRS), by default. This option can't be changed after you start protecting items. Geo-redundant storage (GRS) provides a higher level of data durability than locally redundant storage (LRS), allows an opt-in to use Cross Region Restore, and costs more. Review the trade-offs between lower costs and higher data durability and choose the best option for your scenario. Learn more

- Use Archive Tier for Long-Term Retention (LTR) and save costs: Consider the scenario where you've older backup data that you rarely access but is required to be stored for a long period (for example, 99 years) for compliance reasons. Storing such huge data in a Standard Tier is costly and isn't economical. To help you optimize your storage costs, Azure

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

Backup provides you with Archive Tier, which is an access tier especially designed for Long-Term Retention (LTR) of the backup data.

- If you're protecting both the workload running inside a VM and the VM itself, ensure if this dual protection is needed.

Alerts:

In a scenario where your backup/restore job failed due to some unknown issue. To assign an engineer to debug it, you would want to be notified about the failure as soon as possible. There could also be a scenario where someone maliciously performs a destructive operation, such as deleting backup items or turning off soft-delete, and you would require an alert message for such incident.
You can configure such critical alerts and route them to any preferred notification channel (email, ITSM, webhook, runbook, and so on). Azure Backup integrates with multiple Azure services to meet different alerting and notification requirements:

- Azure Monitor Logs (Log Analytics): You can configure your vaults to send data to a Log Analytics workspace, write custom queries on the workspace, and configure alerts to be generated based on the query output. You can view the query results in tables and charts; also, export them to Power BI or Grafana. (Log Analytics is also a key component of the reporting/auditing capability described in the later sections).
- Azure Monitor Alerts: For certain default scenarios, such as backup failure, restore failure, backup data deletion, and so on, Azure Backup sends alerts by default that are surfaced using Azure Monitor, without the need for a user to set up a Log Analytics workspace.

- Azure Backup provides an in-built alert notification mechanism via e-mail for failures, warnings, and critical operations. You can specify individual email addresses or distribution lists to be notified when an alert is generated. You can also choose whether to get notified for each individual alert or to group them in an hourly digest and then get notified.

- These alerts are defined by the service and provide support for limited scenarios - backup/restore failures, Stop protection with retain data/Stop protection with delete data, and so on. Learn more here.

- If a destructive operation such as stop protection with delete data is performed, an alert is raised and an email is sent to subscription owners, admins, and co-admins even if notifications are not configured for the Recovery Services vault.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

- Certain workloads can generate high frequency of failures (for example, SQL Server every 15 minutes). To prevent getting overwhelmed with alerts raised for each failure occurrence, the alerts are consolidated. Learn more here.

- The in-built alerts can't be customized and are restricted to emails defined in the Azure portal.

- If you need to create custom alerts (for example, alerts of successful jobs) then use Log Analytics. In Azure Monitor, you can create your own alerts in a Log Analytics workspace. Hybrid workloads (DPM/MABS) can also send data to LA and use LA to provide common alerts across workloads supported by Azure Backup.

- You can also get notifications through built-in Recovery Services vault activity logs. However, it supports limited scenarios and isn't suitable for operations such as scheduled backup, which aligns better with resource logs than with activity logs. To learn more about these limitations and how you can use Log Analytics workspace for monitoring and alerting at scale for all your workloads that are protected by Azure Backup, refer to this article.

How does Azure Backup work?

You can back up machines and data by using several methods:

- Back up on-premises machines:

- You can back up on-premises Windows machines directly to Azure by using the Azure Backup Microsoft Azure Recovery Services (MARS) agent. Linux machines aren't supported.

- You can back up on-premises machines to a backup server - either System Center Data Protection Manager (DPM) or Microsoft Azure Backup Server (MABS). You can then back up the backup server to a Recovery Services vault in Azure.

- Back up Azure VMs:

- You can back up Azure VMs directly. Azure Backup installs a backup extension to the Azure VM agent that's running on the VM. This extension backs up the entire VM.

- You can back up specific files and folders on the Azure VM by running the MARS agent.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

- You can back up Azure VMs to the MABS that's running in Azure, and you can then back up the MABS to a Recovery Services vault.

Where is data backed up?

Azure Backup stores backed-up data in vaults - Recovery Services vaults and Backup vaults. A vault is an online-storage entity in Azure that's used to hold data, such as backup copies, recovery points, and backup policies.
Vaults have the following features:

- Vaults make it easy to organize your backup data, while minimizing management overhead.
- You can monitor backed-up items in a vault, including Azure VMs and on-premises machines.
- You can manage vault access with Azure role-based access control (Azure RBAC).
- You specify how data in the vault is replicated for redundancy:
- Locally redundant storage (LRS): To protect your data against server rack and drive failures, you can use LRS. LRS replicates your data three times within a single data center in the primary region. LRS provides at least 99.999999999% (11 nines) durability of objects over a given year. Learn more
- Geo-redundant storage (GRS): To protect against region-wide outages, you can use GRS. GRS replicates your data to a secondary region. Learn more.
- Zone-redundant storage (ZRS): replicates your data in availability zones, guaranteeing data residency and resiliency in the same region. Learn more
- By default, Recovery Services vaults use GRS.
- Recovery Services vaults have the following additional features:
- In each Azure subscription, you can create up to 500 vaults.

Backup policy essentials:

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

- A backup policy is created per vault.
- A backup policy can be created for the backup of following workloads: Azure VMs, SQL in Azure VMs, SAP HANA in Azure VMs and Azure file shares. The policy for files and folder backup using the MARS agent is specified in the MARS console.
- Azure File Share
- A policy can be assigned to many resources. An Azure VM backup policy can be used to protect many Azure VMs.
- A policy consists of two components
- Schedule: When to take the backup
- Retention: For how long each backup should be retained.
- Schedule can be defined as "daily" or "weekly" with a specific point of time.
- Retention can be defined for "daily", "weekly", "monthly", "yearly" backup points.
- "weekly" refers to a backup on a certain day of the week
- "monthly" refers a backup on a certain day of the month
- "yearly" refers to a backup on a certain day of the year
- Retention for "monthly", "yearly" backup points is referred to as Long Term Retention (LTR)
- When a vault is created, a "DefaultPolicy" is also created and can be used to back up resources.
- Any changes made to the retention period of a backup policy will be applied retroactively to all the older recovery points aside from the new ones.

Architecture: Built-in Azure VM Backup:

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

- For Azure VMs that are selected for backup, Azure Backup starts a backup job according to the backup schedule you specify.
- During the first backup, a backup extension is installed on the VM if the VM is running.
- For Windows VMs, the Snapshot extension is installed.
- For Linux VMs, the VM Snapshot Linux extension is installed.
- For Windows VMs that are running, Backup coordinates with Windows Volume Shadow Copy Service (VSS) to take an app-consistent snapshot of the VM.
- By default, Backup takes full VSS backups.
- If Backup can't take an app-consistent snapshot, then it takes a file-consistent snapshot of the underlying storage (because no application writes occur while the VM is stopped).
- For Linux VMs, Backup takes a file-consistent backup. For app-consistent snapshots, you need to manually customize pre/post scripts.
- After Backup takes the snapshot, it transfers the data to the vault.
- The backup is optimized by backing up each VM disk in parallel.
- For each disk that's being backed up, Azure Backup reads the blocks on the disk and identifies and transfers only the data blocks that changed (the delta) since the previous backup.
- Snapshot data might not be immediately copied to the vault. It might take some hours at peak times. Total backup time for a VM will be less than 24 hours for daily backup policies.
- Changes made to a Windows VM after Azure Backup is enabled on it are:
- Microsoft Visual C++ 2013 Redistributable(x64) - 12.0.40660 is installed in the VM
- Startup type of Volume Shadow Copy service (VSS) changed to automatic from manual
- IaaSVmProvider Windows service is added

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

- When the data transfer is complete, the snapshot is removed, and a recovery point is created.

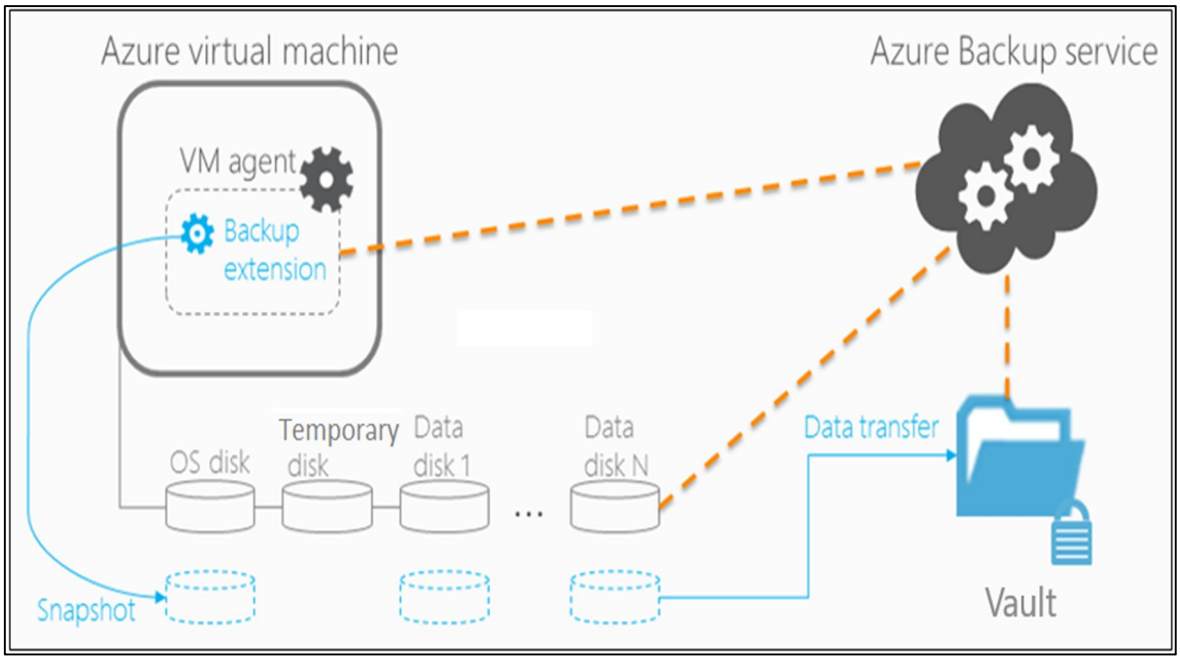

Figure 20: Azure Backup Services

4.2 Azure File Share

To create an Azure file share, you need to answer three questions about how you will use it:

- What are the performance requirements for your Azure file share? Azure Files offers standard file shares which are hosted on hard disk-based (HDD-based) hardware, and premium file shares, which are hosted on solid-state disk-based (SSD-based) hardware.

- What are your redundancy requirements for your Azure file share? Standard file shares offer locally redundant (LRS), zone redundant (ZRS), geo-redundant (GRS), or geo-zone redundant (GZRS) storage, however the large file share feature is only supported on locally redundant and zone redundant file shares. Premium file shares do not support any form of geo-redundancy.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

- Premium file shares are available with locally redundancy and zone redundancy in a subset of regions. To find out if premium file shares are currently available in your region, see the products available by region page for Azure. For information about regions that support ZRS, see Azure Storage redundancy.

- In local and zone redundant storage accounts, Azure file shares can span up to 100 TiB. However, in geo- and geo-zone redundant storage accounts, Azure file shares can span only up to 5 TiB.

Create a storage account:

Azure file shares are deployed into *storage accounts*, which are top-level objects that represent a shared pool of storage. This pool of storage can be used to deploy multiple file shares.

Azure supports multiple types of storage accounts for different storage scenarios customers may have, but there are two main types of storage accounts for Azure Files. Which storage account type you need to create depends on whether you want to create a standard file share or a premium file share:

- General purpose version 2 (GPv2) storage accounts: GPv2 storage accounts allow you to deploy Azure file shares on standard/hard disk-based (HDD-based) hardware. In addition to storing Azure file shares, GPv2 storage accounts can store other storage resources such as blob containers, queues, or tables. File shares can be deployed into the transaction optimized (default), hot, or cool tiers.

- FileStorage storage accounts: FileStorage storage accounts allow you to deploy Azure file shares on premium/solid-state disk-based (SSD-based) hardware. FileStorage accounts can only be used to store Azure file shares; no other storage resources (blob containers, queues, tables, etc.) can be deployed in a FileStorage account.

The other basics fields are independent from the choice of storage account:

- Storage account name: The name of the storage account resource to be created. This name must be globally unique. The storage account name will be used as the server name when you mount an Azure file share via SMB. Storage account names must be between 3 and 24 characters in length and may contain numbers and lowercase letters only.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

- Location: The region for the storage account to be deployed into. This can be the region associated with the resource group, or any other available region.

- Replication: Although this is labeled replication, this field means redundancy; this is the desired redundancy level: locally redundancy (LRS), zone redundancy (ZRS), geo-redundancy (GRS), and geo-zone-redundancy (GZRS). This drop-down list also contains read-access geo-redundancy (RA-GRS) and read-access geo-zone redundancy (RA-GZRS), which do not apply to Azure file shares; any file share created in a storage account with this selected will be either geo-redundant or geo-zone-redundant, respectively.

Create a file share:

Once you've created your storage account, you can create your file share. This process is mostly the same regardless of whether you're using a premium file share or a standard file share. You should consider the following differences:

Standard file shares may be deployed into one of the standard tiers: transaction optimized (default), hot, or cool. This is a per file share tier that is not affected by the blob access tier of the storage account (this property only relates to Azure Blob storage - it does not relate to Azure Files at all). You can change the tier of the share at any time after it has been deployed. Premium file shares cannot be directly converted to any standard tier.

The quota property means something slightly different between premium and standard file shares:

- For standard file shares, it's an upper boundary of the Azure file share, beyond which end-users cannot go. If a quota is not specified, standard file shares can span up to 100 TiB (or 5 TiB if the large file shares property is not set for a storage account). If you did not create your storage account with large file shares enabled, see Enable large files shares on an existing account for how to enable 100 TiB file shares.
- For premium file shares, quota means provisioned size. The provisioned size is the amount that you will be billed for, regardless of actual usage. The IOPS and throughput available on a premium file share is based on the provisioned size. For more information on how to plan for a premium file share, see provisioning premium file shares.

If you just created your storage account, you can navigate to it from the deployment screen by selecting **Go to resource**. Once in the storage account, select the **File shares** in the table of contents for the storage account.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

In the file share listing, you should see any file shares you have previously created in this storage account; an empty table if no file shares have been created yet. Select **+ File share** to create a new file share.

The new file share blade should appear on the screen. Complete the fields in the new file share blade to create a file share:

- Name: the name of the file share to be created.
- Quota: the quota of the file share for standard file shares; the provisioned size of the file share for premium file shares. For standard file shares, the quota will also determine what performance you receive.
- Tiers: the selected tier for a file share. This field is only available in a general purpose (GPv2) storage account. You can choose transaction optimized, hot, or cool. The share's tier can be changed at any time. We recommend picking the hottest tier possible during a migration, to minimize transaction expenses, and then switching to a lower tier if desired after the migration is complete.

Select **Create** to finishing creating the new share.

File shares deployed in general purpose v2 (GPv2) storage account can be in the transaction optimized, hot, or cool tiers. You can change the tier of the Azure file share at any time, subject to transaction costs as described above
There are several other storage account types you may come across in the Azure portal,

PowerShell, or CLI. Two storage account types, BlockBlobStorage and BlobStorage storage accounts, cannot contain Azure file shares. The other two storage account types you may see are general purpose version 1 (GPv1) and classic storage accounts, both of which can contain Azure file shares. Although GPv1 and classic storage accounts may contain Azure file shares, most new features of Azure Files are available only in GPv2 and FileStorage storage accounts.

On the main storage account page, select **File shares** select the tile labeled **File shares** (you can also navigate to **File shares** via the table of contents for the storage account).

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

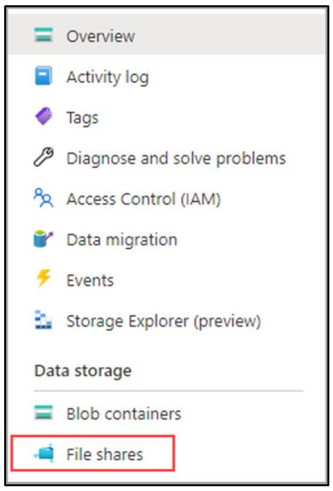

Figure 21: File Share dialogue Box

In the table list of file shares, select the file share for which you would like to change the tier. On the file share overview page, select **Change tier** from the menu.

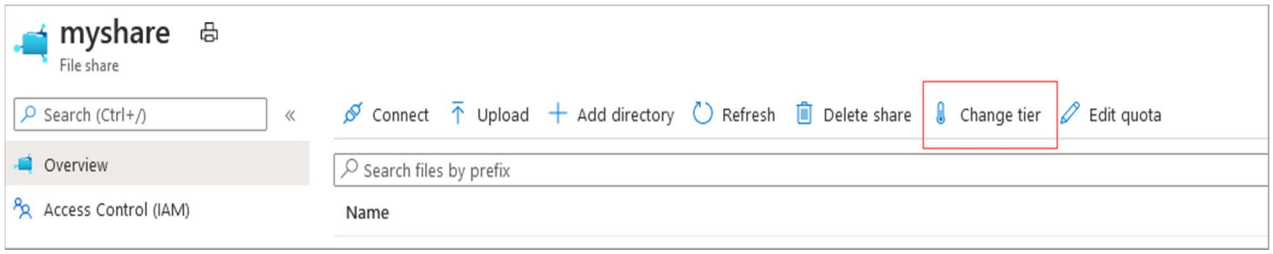

Figure 22: File Share Tier Change

Storage tiers:

Azure Files offers four different tiers of storage, premium, transaction optimized, hot, and cool to allow you to tailor your shares to the performance and price requirements of your scenario:

- Premium: Premium file shares are backed by solid-state drives (SSDs) and provide consistent high performance and low latency, within single-digit milliseconds for most IO operations, for IO-intensive workloads. Premium file shares are suitable for a wide variety of workloads like databases, web site hosting, and development environments. Premium file shares can be used with both Server Message Block (SMB) and Network File System (NFS) protocols.

- Transaction optimized: Transaction optimized file shares enable transaction heavy workloads that don't need the latency offered by premium file shares. Transaction optimized file shares are offered on the standard storage hardware backed by hard disk drives (HDDs). Transaction optimized has historically been called "standard", however this refers to the storage media type rather than the tier itself (the hot and cool are also "standard" tiers, because they are on standard storage hardware).

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

- Hot: Hot file shares offer storage optimized for general purpose file sharing scenarios such as team shares. Hot file shares are offered on the standard storage hardware backed by HDDs.

- Cool: Cool file shares offer cost-efficient storage optimized for online archive storage scenarios. Cool file shares are offered on the standard storage hardware backed by HDDs.

- Premium file shares are deployed in the FileStorage storage account kind and are only available in a provisioned billing model. For more information on the provisioned billing model for premium file shares, see Understanding provisioning for premium file shares. Standard file shares, including transaction optimized, hot, and cool file shares, are deployed in the general-purpose version 2 (GPv2) storage account kind, and are available through pay as you go billing.

- When selecting a storage tier for your workload, consider your performance and usage requirements. If your workload requires single-digit latency, or you are using SSD storage media on-premises, the premium tier is probably the best fit. If low latency isn't as much of a concern, for example with team shares mounted on-premises from Azure or cached on-premises using Azure File Sync, standard storage may be a better fit from a cost perspective.

- Once you've created a file share in a storage account, you cannot move it to tiers exclusive to different storage account kinds. For example, to move a transaction optimized file share to the premium tier, you must create a new file share in a FileStorage storage account and copy the data from your original share to a new file share in the FileStorage account. We recommend using AzCopy to copy data between Azure file shares, but you may also use tools like robocopy on Windows or rsync for macOS and Linux.

- File shares deployed within GPv2 storage accounts can be moved between the standard tiers (transaction optimized, hot, and cool) without creating a new storage account and migrating data, but you will incur transaction costs when you change your tier. When you move a share from a hotter tier to a cooler tier, you will incur the cooler tier's write transaction charge for each file in the share. Moving a file share from a cooler tier to a hotter tier will incur the cool tier's read transaction charge for each file in the share.

On the resulting dialog, select the desired tier: transaction optimized, hot, or cool.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

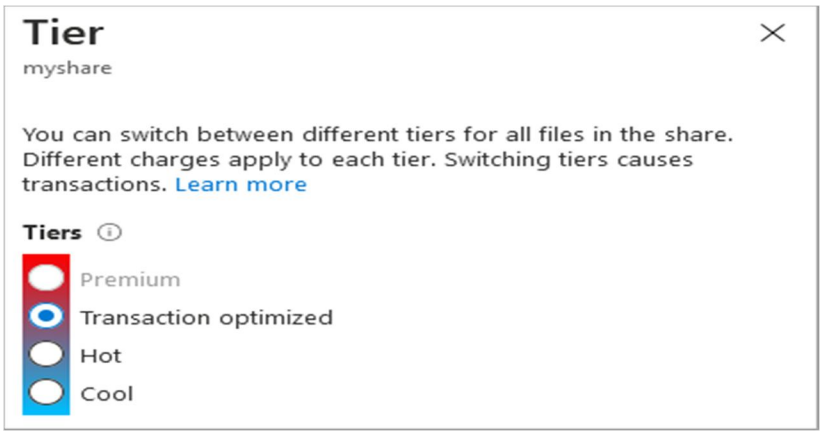

Figure 23: Tier Switch Dialog Box

If you enable large file shares on an existing storage account, you must expand existing file shares in that storage account to take advantage of the increased capacity and scale.

File share total cost of ownership checklist:

If you are migrating to Azure Files from on-premises or comparing Azure Files to other cloud storage solutions, you should consider the following factors to ensure a fair, apples-to-apples comparison:

- How do you pay for storage, IOPS, and bandwidth? With Azure Files, the billing model you use depends on whether you are deploying premium or standard file shares. Most cloud solutions have models that align with the principles of either provisioned storage, such as price determinism and simplicity, or pay-as-you-go storage, which can optimize costs by only charging you for what you actually use. Of particular interest for provisioned models are minimum provisioned share size, the provisioning unit, and the ability to increase and decrease provisioning.

- Are there any methods to optimize storage costs? With Azure Files, you can use capacity reservations to achieve an up to 36% discount on storage. Other solutions may employ storage efficiency strategies like deduplication or compression to optionally optimize storage, but remember, these storage optimization strategies often have non-monetary costs, such as reducing performance. Azure Files capacity reservations have no side effects on performance.

- How do you achieve storage resiliency and redundancy? With Azure Files, storage resiliency and redundancy are baked into the product offering. All tiers and redundancy levels ensure that data is highly available and at least three copies of your data are accessible. When

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

considering other file storage options, consider whether storage resiliency and redundancy is built-in or something you must assemble yourself.

- What do you need to manage? With Azure Files, the basic unit of management is a storage account. Other solutions may require additional management, such as operating system updates or virtual resource management (VMs, disks, network IP addresses, etc.).

- What are the costs of value-added products, like backup, security, etc.? Azure Files supports integrations with multiple first- and third-party value-added services. Value-added services such as Azure Backup, Azure File Sync, and Azure Defender provide backup, replication and caching, and security functionality for Azure Files. Value-added solutions, whether on-premises or in the cloud, have their own licensing and product costs, but are often considered part of the total cost of ownership for file storage.

Chapter 5: Conclusion and Future Scopes

5.1 Conclusion

As we discussed earlier During the pandemic, all the students were away from their campus in their houses. However, online classes were held; what about the practical's for the students who do not have access to a high-end device. Practical is also a crucial part of the study. As a computer science student, I can relate to the issue of not having a high-end device.

We had successfully deployed a Smart Education System for the students and faculties over the cloud by using Microsoft Azure Virtual Desktop.

By using AVD, we took the concept of the virtual machine to provide a high spec machine to run virtually even on a relatively low spec pc or system; using AVD, we can deploy several virtual machines, and since it's a pay as you go model, You can access and pay for only the resources you

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

need, which keeps it budget-friendly and efficient, Since it would be much easier for the institution to provide their batch of students and modifying the virtual machine as you require.

5.2 Future Scope

Access your desktop and applications from virtually anywhere

Set up Azure Virtual Desktop (formerly Windows Virtual Desktop) in minutes to enable secure remote work. Provide the familiarity and compatibility of Windows 11 with the new scalable multi-session experience for your end-users and save costs by using existing eligible Windows licenses. Manage your end-to-end Azure Virtual Desktop deployment alongside other Azure services within the Azure portal.

Deliver Windows 11 desktops and applications virtually anywhere
Provide employees the best virtualized experience with the only solution fully optimized for Windows 11 and Microsoft 365.

Built-in intelligent security
Help keep your applications and data secure and compliant with security capabilities that can proactively detect threats and take remedial action.

Deploy and scale in minutes
Simplify deployment and management of your infrastructure and scale quickly based on your business needs.

Reduce cost using existing licenses
Use existing eligible licences to reduce costs with a modern cloud-based virtual desktop infrastructure (VDI) and pay only for what you use.

IJARSCT
ISSN (Online) 2581-9429
International Journal of Advanced Research in Science, Communication and Technology (IJARSCT)
Volume 2, Issue 6, May 2022
Impact Factor: 6.252

Chapter 6: References

- https://docs.microsoft.com/en-us/learn/certifications/exams/az-900
- https://docs.microsoft.com/en-us/learn/certifications/exams/az-104
- https://docs.microsoft.com/en-us/learn/certifications/exams/az-140
- https://docs.microsoft.com/en-us/azure/
- https://docs.microsoft.com/en-us/learn/modules/intro-to-azure-virtual-machines/
- https://azure.microsoft.com/en-in/features/azure-portal/
- https://www.udemy.com/course/wvd-az-140-practice
- https://www.edureka.co/blog/azure-portal/