ELECTIVE

5.4.1CLOUDCOMPUTING LTP 3-3

RATIONALE

Thiscourseoffersagoodunderstandingof cloudcomputingconceptsandchallengesfacedin implementation of cloud computing.

LEARNINGOUTCOMES

Afterundergoingthesubject,thestudentswouldbeableto:

- Explaincoreconceptsofcloudcomputingparadigm.
- ExplainvariousServiceModels
- ExplainvariousDeploymentModels.
- DescribeSLAmanagementinCloudComputing
- Explainandapplytheconceptof virtualization.
- Describetheschedulingoftasksincloud.
- Illustratethefundamentalconceptsofcloudstorage.
- Describevarioussecurityissuesinthecloud.
- Makeuseofcloud.

DETAILEDCONTENTS

1. Introduction(6Periods)

EvolutionofCloudComputing,CloudComputingOverview,Characteristics,Applications,B

enefits, Challenges.

2. ServiceandDeploymentModels(6Periods)

- CloudComputingServiceModels:InfrastructureasaService,PlatformasaService, Software as a Service;
- CloudComputingDeploymentModels:Private Cloud;PublicCloud,CommunityCloud, Hybrid Cloud, Major Cloud Service providers.

3. ServiceLevelAgreement(SLA)Management(4Periods)

OverviewofSLA,TypesofSLA,SLALifeCycle,SLAManagementProcess.

4. VirtualizationConcepts(8Periods)

OverviewofVirtualization,TypesofVirtualization,BenefitsofVirtualization,Hypervisors.

5. CloudSecurity(6Periods)

InfrastructureSecurity,DataSecurity&PrivacyIssues,LegalIssuesinCloudComputing.

6. CloudStorage(6Periods)

Overview;StorageasaService,BenefitsandChallenges,StorageAreaNetworks(SANs).

7. SchedulinginCloud(12Periods)

OverviewofSchedulingproblem,Differenttypesofscheduling,Schedulingforindependenta

nd dependent tasks, Static vs. Dynamic scheduling.

LISTOFPRACTICALS

1. IntroductiontoCloudVendors:Amazon,Microsoft,IBM.
2. SettingupVirtualizationusingVirtualbox/VMWareHypervisor
3. IntroductiontoOwnCloud
4. InstallationandconfigurationofOwnCloudsoftwareforSaaS
5. AccessingMicrosoftAZUREcloud-services
6. CloudSimulationSoftwareIntroduction:CloudSim

INSTRUCTIONALSTRATEGY

Inadditiontocassroom teaching,the

teachershoulddemonstratethepracticalusageofcloudusing real cloud services.

MEANSOFASSESSMENT

- AssignmentsandQuiz/classtests,mid-termandend-termwrittentests
- ActuallaboratoryandpracticalworkandViva-Voce

RECOMMENDEDBOOKS

1. RajkumarBuyya,JamesBroberg,AndrzejGoscinski(Editors):CloudComputing:Pri

nciplesand Paradigms, Wiley, 2011

2. KumarSaurabh,CloudComputing,Wiley,2012.

3. BarrieSosinsky:CloudComputingBible,Wiley,2011.

4. JudithHurwitz,RobinBloor,MarciaKaufman,FernHalper:CloudComputingforDummi es,Wiley, 2010

5. e-books/e-tools/relevantsoftwaretobeusedasrecommendedbyAICTE/HSBTE/NITTTR.

WebsitesforReference:

http://swayam.gov.in

Unit 1
Introduction

Evolution of Cloud Computing:- Cloud computing is all about renting computing services. This idea first came in the1950s.In making cloud computing what it is today, five technologies played a vital role. These are distributed systems and its peripherals, virtualization, web 2.0, service orientation, and utility computing.

I. Distributed Systems and its peripherals
II. Virtualization
III. Web2.0
IV. Service Orientation
V. Utility Computing

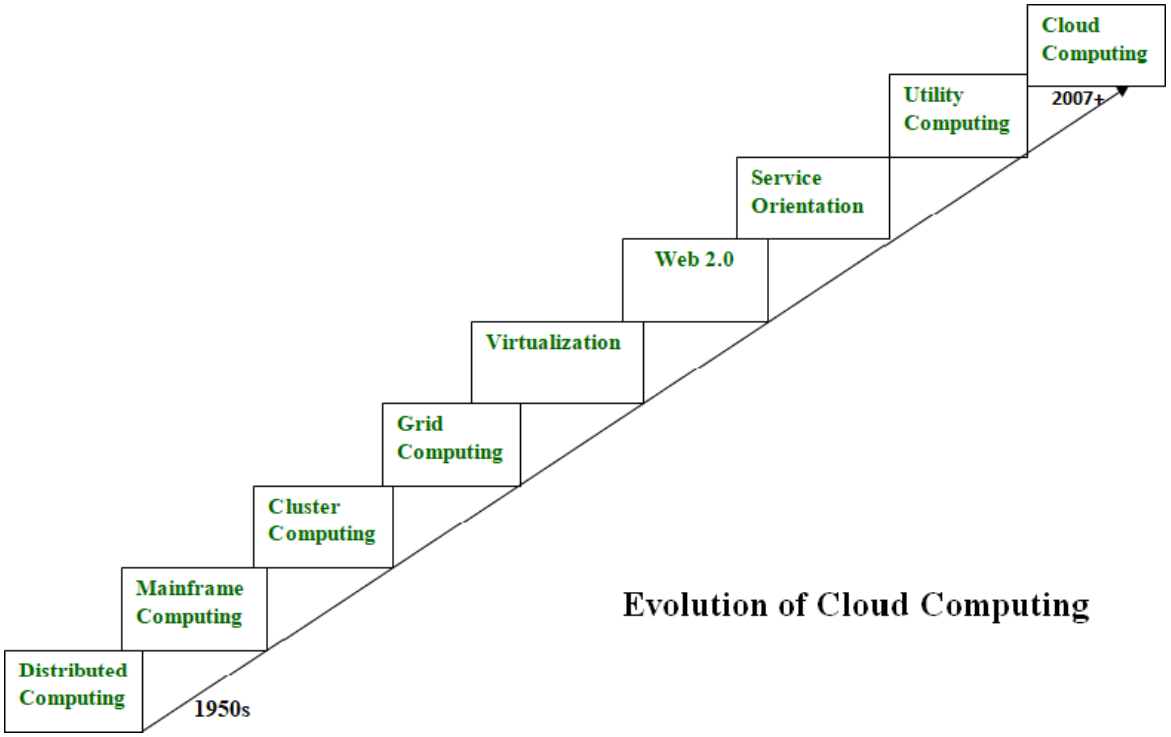

Evolution of Cloud Computing

1) Distributed Systems:
It is a composition of multiple independent systems but all of them are depicted as a single entity to the users. The purpose of distributed systems is to share resources and also use them effectively and efficiently. Distributed systems possess characteristics such as scalability, concurrency, continuous availability, heterogeneity, and independence in failures. But the main problem with this system was that all the systems were required to be present at the same geographical location. Thus to solve this problem, distributed

Computing led to three more types of computing and they were- Mainframe computing, cluster computing, and grid computing.

2) Main frame computing:
   Main frames which first came into existence in1951are highly powerful and reliable computing machines. These are responsible for handling large data such as massive input-output operations. Even today these are used for bulk processing tasks such as online transactions etc. These systems have almost no downtime with high fault tolerance. After distributed computing, these increased the processing capabilities of the system. But these were very expensive. To reduce this cost, cluster computing came as an alternative to mainframe technology.
3) Cluster computing:
   In 1980s, cluster computing came as an alternative to mainframe computing. Each machine in the cluster was connected to each other by a network with high bandwidth. These were way cheaper than those mainframe systems. These were equally capable of high computations.
   Also, new nodes could easily be added to the cluster if it was required. Thus, the problem of the cost was solved to some extent but the problem related to geographical restrictions still pertained. To solve this, the concept of grid computing was introduced.

4) Grid computing:
   In 1990s, the concept of grid computing was introduced. It means that differentsystemswereplacedatentirelydifferentgeographicallocations and these all were connected via the internet. These systems belonged to different organizations and thus the grid consisted of heterogeneous nodes. Although it solved some problems but new problems emerged as the distance between the nodes increased. The main problem which was encountered was the low availability of high bandwidth connectivity and with it other network associated issues. Thus cloud computing is often referred to as “Successor of grid computing”.

5) Virtualization:
   It was introduced nearly40 years back. It refers to the process of creating a virtual layer over the hardware which allows the user to run multiple instances simultaneously on the hardware. It is a key technology used in cloud computing. It is the base on which major cloud computing services such as Amazon EC2, VMware v Cloud, etc work on. Hardware virtualization is still one of the most common types of virtualization.

6) Web2.0:
It is the interface through which the cloud computing services interact with the clients. It is because ofWeb 2.0 that we have interactive and dynamic web pages. It also increases flexibility among web pages.
Popular examples ofweb2.0include Google Maps, Facebook, Twitter, etc. Needless to say, social media is possible because of this technology only. In gained major popularity in 2004.

7) Serviceorientation:
Itactsasareference modelforcloudcomputing.Itsupports low-cost, flexible, and evolvable applications. Two important concepts were introduced in this computing model. These were Quality of Service (QoS) which also includes the SLA (Service Level Agreement) and Software as a Service (SaaS).

8) Utilitycomputing:
It is a computing model that defines service provisioning techniques for services such as computeservicesalongwithothermajorservicessuchas storage, infrastructure, etc which are provisioned on a pay-per-use basis.

Thus,theabovetechnologiescontributedtothemakingofcloudcomputing.

Introduction to Cloud Computing

Cloud Computing is the delivery of computing services such as servers, storage, databases,networking,software,analytics,intelligence,andmore,overthe Cloud (Internet).

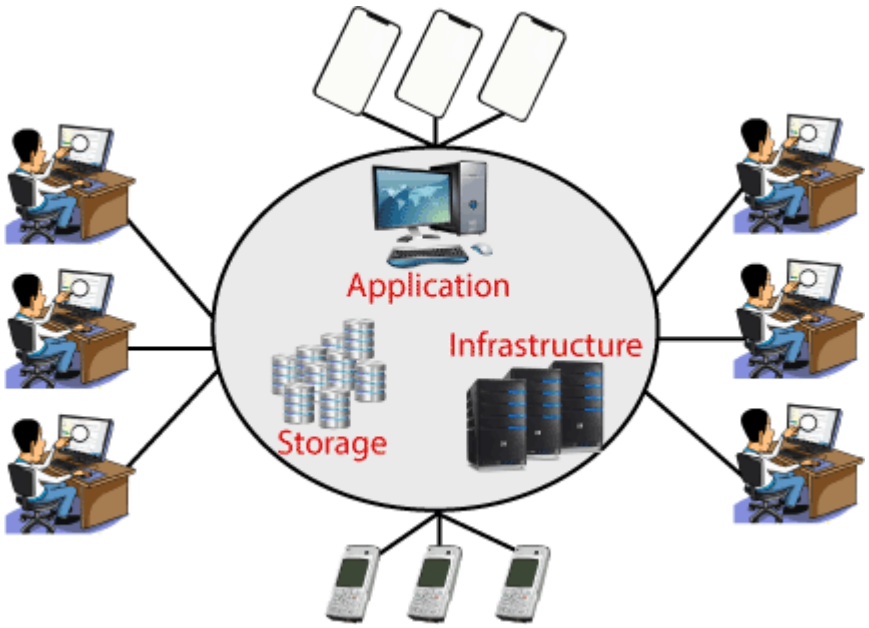

Cloud Computing provides an alternative to the on-premises datacentre. With an on-premisesdatacentre,wehavetomanageeverything, such as purchasingandinstalling hardware, virtualization, installing the operating system, and any other required applications,settingupthenetwork,configuringthe firewall,andsettingup storagefor data. After doing all the set-up, we become responsible for maintaining it through its entire lifecycle.

But if we choose Cloud Computing, a cloud vendor is responsible for the hardware purchaseandmaintenance.Theyalsoprovideawidevarietyof softwareandplatformas a service. We can take any required services on rent. The cloud computing services will be charged based on usage.

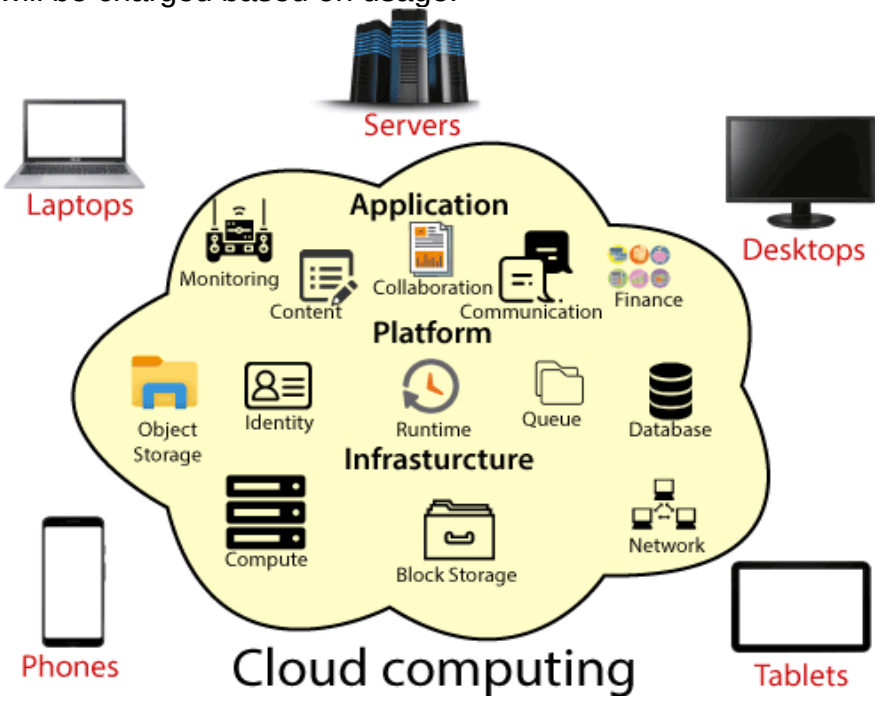

Thecloudenvironmentprovidesaneasilyaccessibleonlineportalthatmakeshandyfor the user to manage the compute, storage, network, and application resources. Some cloud service providers are in the following figure.

Main Characteristics of Cloud Computing

ThefivemajorcharacteristicsoftheCloudComputingareasfollows:

1. On-demandself-service-Theserviceofthecloudisavailableroundthe clock and provides computing capabilities on-demand of the user automatically.
2. Broadnetworkaccess -Userscanaccesstheservicesviadifferent modes as the heterogeneous thin and thick client platforms.
3. Resourcepooling-Thefeatureofmulti-tenancywhereusersare assigned resources dynamically, based on demands.
4. Rapidelasticity– Theservice is flexibleandcanbescaledupordownto suit the business requirements. Resources and programs can be used based on requirement and the user is billed only for the usage.
5. Measuredservice –Usage meteringisavailableand youpayonlyfor what you use. You need not pay for any infrastructure that youdo not use.

Advantagesofcloudcomputing

- Cost:Itreducesthehugecapitalcostsofbuyinghardwareandsoftware.
- Speed:Resourcescanbeaccessedinminutes,typicallywithina few clicks.
- Scalability:Wecanincreaseordecreasetherequirementofresources according to the business requirements.
- Productivity: While using cloud computing, we put less operational effort.Wedonotneedtoapplypatching,aswellasno needto maintain hardware and software. So, in this way, the IT team can be more productive and focus on achieving business goals.
- Reliability:Backupandrecoveryofdataarelessexpensiveand veryfast for business continuity.
- Security:Manycloudvendorsofferabroad setofpolicies,technologies, and controls that strengthen our data security.

Unit 2: Cloud Computing Service Models:

Various Cloud Computing Service model are as:

1) InfrastructureasaService(IaaS)
2) PlatformasaService(PaaS)
3) SoftwareasaService(SaaS)

TypesofCloudServices

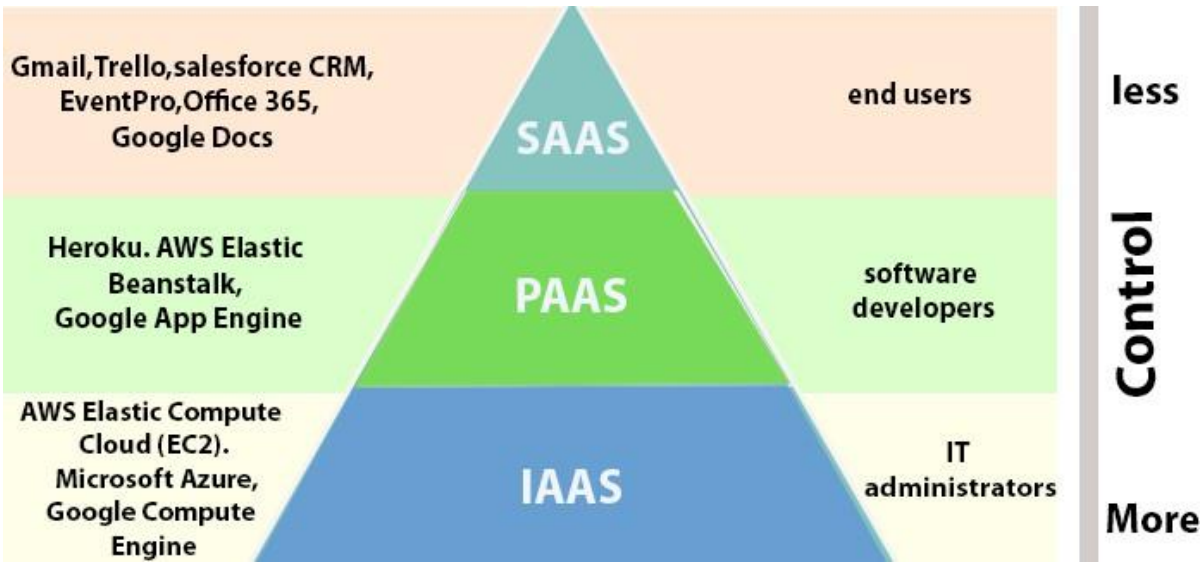

1. Infrastructure as a Service (IaaS): In IaaS, we can rent IT infrastructures like servers and virtual machines (VMs), storage, networks, operating systems from a cloud service vendor. We can create VM runningWindows or Linuxand installanythingwe wantonit. Using IaaS, we don't need to care about the hardware or virtualizationsoftware, butotherthanthat,wedohaveto manageeverythingelse.UsingIaaS,we

getmaximumflexibility,butstill,weneedtoputmoreeffortinto maintenance.

2. Platform as a Service (PaaS): This service provides an on-demand environmentfordeveloping,testing,delivering,andmanagingsoftware applications. The developer is responsible for the application, and the PaaS vendor provides the ability to deploy and run it. Using PaaS, the flexibilitygets reduce, butthe management ofthe environment is taken care of by the cloud vendors.
3. Software as a Service (SaaS): It provides a centrally hosted and managed software services to the end-users. It delivers software over the internet,on-demand,andtypicallyonasubscriptionbasis.E.g.,Microsoft One Drive, Dropbox, WordPress, Office 365, and Amazon Kindle. SaaS is used to minimize the operational cost to the maximum extent.

CloudComputingDeploymentModels

The cloud services can be deployed in different methods. The deploymentmodel is based on the service model, organizational structure, location, user base, and so on. The four most commonly used deployment models are as follows:

1. PublicCloud
2. PrivateCloud
3. CommunityCloud
4. HybridCloud

1.PublicCloud

In this model, the infrastructure is accessible to the public and it is owned by a vendor, who offers the services of the cloud to the users. The cloud vendor shares the cloud resources withthe end users. The resource pool is huge and the services are shared bylots ofusers. The services ofthis cloud modelcanbe free or available for nominal charges. Google uses a public cloud deploymentmodel. With this model, users need not purchase any infrastructure but can use that of the vendor. A drawback of the public cloud model is that it poses a security threat. If you have very confidential data running in your network, it is not safe to use the public cloud model.

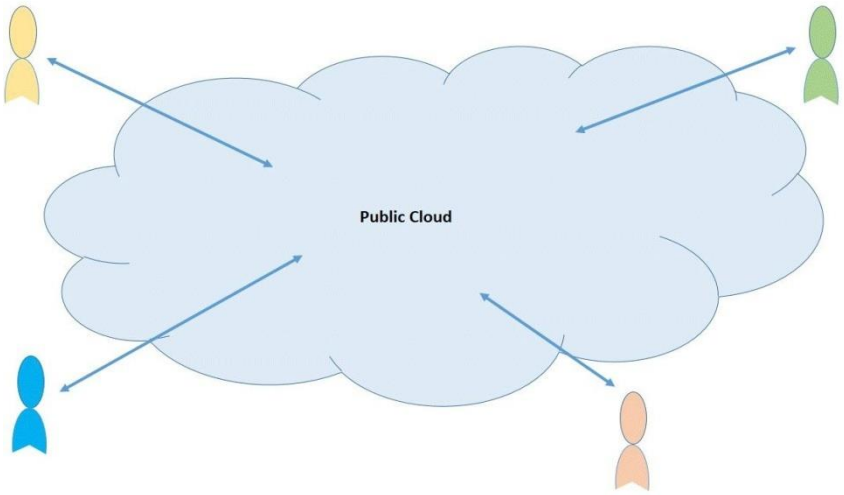

Fig2.2.1:PublicCloud

1. PrivateCloud

As the name suggests, this would be a privatelyowned cloud. Here, the user or organization owns the cloud and only the user or employees of the company have access to the cloud, thereby makingdata and transactions secure. There is more control over resources when compared to the Public Cloud model. The Privatecloud modelusestheVirtualizationsolutionandthedatacentersbelong to the company. The major advantage of this model is the security and the control that the users have over the resources and application. However, the drawbackisthatmore financialinvestmentisrequiredandtheofferingis notas big scale as that of a public cloud model.

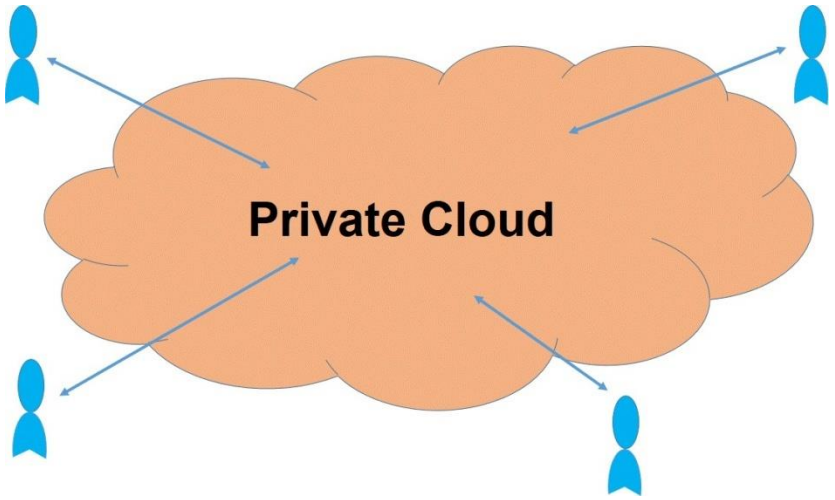

Fig2.2.2:PrivateCloud

1. Community Cloud

Inthe CommunityCloud model, the infrastructure is owned jointlybydifferent organizations. The organizations may have a similar set of requirements, policies, and customer base. So, they can combine the offerings and make the customer base even bigger. Duplication of same or similar applications and resources are avoided. This model helps reduce the costs, which would otherwisebehigheriftheorganizationdeploysthePrivateCloudmodel.This is again a classification of the Private Cloud, as it is available to only a certain group of users.

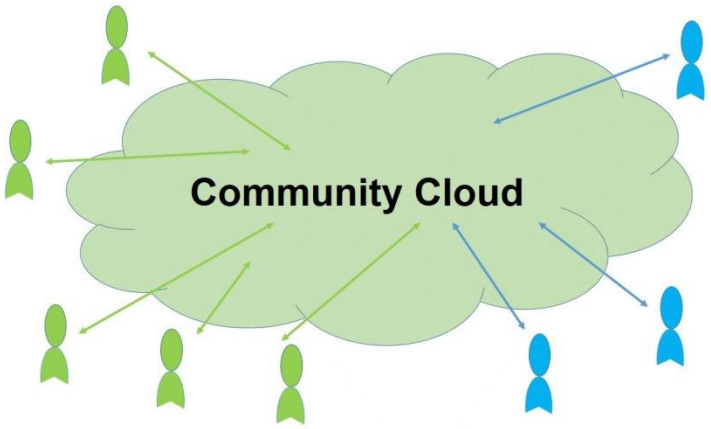

Fig2.2.3:Community Cloud

2. HybridCloud

The Hybrid Cloud deployment model comprises of two or more clouds. Thiscan be a combination of the other three cloud types – public, private, or community. The hybrid deployment is complex compared to the other three owingtotheexecutionand managementtasksinvolved. Anexamplescenarioof this model can be where an organization is on the private cloud but there are load spikes which the private cloud cannot handle. For this the organization dependsonthepubliccloudtosupportthe load.Theshift fromtheprivatetothe public cloud and back will be seamless to the end user

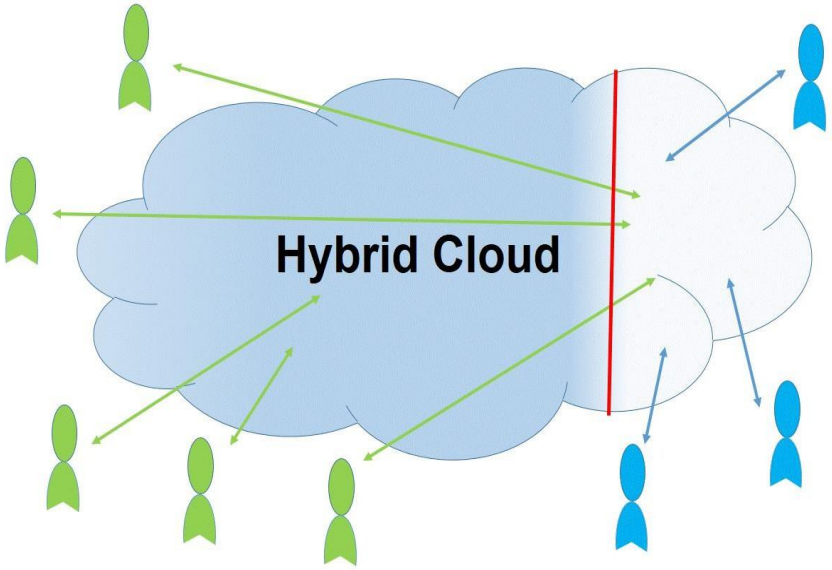

Fig2.2.4:HybridCloud

Major Cloud Service Providers

- Kamatera
- Softchoice
- Prolifics
- ScienceSoft
- phoenixNAP
- Cloudways
- pCloud
- AmazonWebServices
- MicrosoftAzure
- GoogleCloudPlatform
- Adobe
- VMware
- IBM Cloud
- Rackspace
- RedHat
- Salesforce
- OracleCloud
- SAP
- VerizonCloud
- Navisite
- Dropbox
- Egnyte

The followingtable summarizesthetop3keyplayersandtheirofferings inthecloud computing world:

	AWS	Azure	GoogleCloud
Company	AWS Inc.	Microsoft	Google
Launchyear	2006	2010	2008
Geographical Regions	25	54	21
Availability Zones	78	140(countries)	61
Keyofferings	Compute, storage, database, analytics, networking, machine learning,andAI,mobile, developer tools, IoT, security, enterprise applications,blockchain.	Compute, storage, mobile, data management, messaging,mediaservices, CDN, machine learning and AI, developer tools, security, blockchain, functions,IoT.	Compute, storage, databases,networking, big data, cloud AI, management tools, Identity and security, IoT, API platform
Compliance Certificates	46	90	
Annual Revenue	$33 billion	$35 billion	$8billion

Unit 3 :- Cloud SLA (cloud service-level agreement)

A cloud SLA (cloud service-levelagreement)is anagreement betweena cloudservice providerand a customer that ensures a minimum level of service is maintained. Itguarantees levelsof reliability,availabilityandresponsiveness to systemsandapplications,whilealsospecifyingwhowillgovernwhenthereisa service interruption.

Acloudinfrastructurecanspangeographies,networksandsystemsthatare both physical and virtual. While the exact metrics of a cloud SLA can vary by service provider, the areas covered are uniform: volume and quality of work -- including precision and accuracy -- speed, responsiveness and efficiency. The document aims to establish a mutual understanding of the services, prioritized areas, responsibilities, guarantees and warranties provided by the service provider.

Metrics and responsibilities among the parties involved in cloud configurations are clearlyoutlined, suchasthespecificamountofresponsetime forreportingoraddressing system failures.

Financialpenaltiesaprovider mustpayforfailingtoliveupto theguaranteed termsare also included. These penalties are often in the form of credits for service time.

Whattolookforina cloudSLA

Service-level agreements have become more important as organizations move theirsystems,applicationsanddatatothecloud.AcloudSLAensuresthatcloud providers meet certain enterprise-levelrequirements and provide customers with a clearly defined set of deliverables.

Thedefined levelofservicesshould be specificand measureable ineacharea. Thisallowsthe qualityofservice(QoS)tobebenchmarkedand,ifstipulatedby the agreement, rewarded or penalized accordingly.

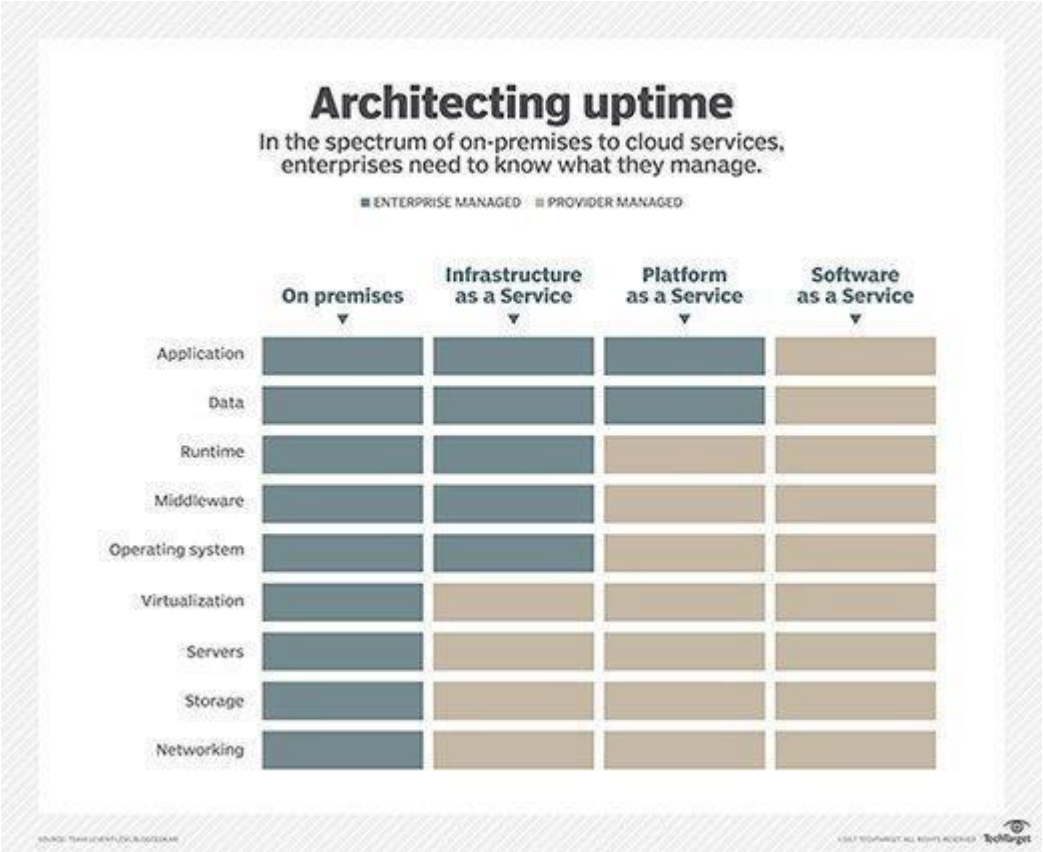

AnSLA willcommonlyusetechnicaldefinitionsthat quantifythelevelofservice, such as mean time between failures (MTBF) or mean time to repair (MTTR), which specifies a target or minimum value for service-level performance.

AtypicalcomputeandcloudSLAarticulatespreciselevelsofservice,aswellas the recourse or compensation the user is entitled to should the provider fail to deliver the service as described. Another area to consider carefully is service availability, which specifies the maximum amount of time a read request can take; how many retries are allowed; and so on.

The SLA should also define compensationfor users if the specifications aren't met.Acloudstorageserviceproviderusuallyoffersatieredservicecredit plan that gives userscredits based on the discrepancybetween SLA specifications and the actual service levels delivered.

Mostpubliccloudstorageservicesprovidedetailsoftheservicelevelsthatusers can expect on their websites, and these will likely be the same for all users. However,anenterpriseestablishingaservicewithaprivatecloudstorage

providermaybeableto negotiateamorecustomizeddeal.Inthiscase,thecloud SLA might include specifications for retention policies, the number of copies that will be retained, storage locations and so on.

Cloud service-level agreements may be more detailed to cover governance, securityspecifications,compliance,andperformanceand uptimestatistics.They should address security and encryption practices for data privacy, disaster recovery expectations, data location, as well as data access and portability.

Data protectionprocesses, such as backup and disaster recovery, should also be addressed. Theagreementshouldoutlinethe responsibilitiesofeachparty, theacceptable performance parameters, a description of the applications and services covered under the agreement, procedures for monitoring service levels, and a schedule for the remediation of outages.

Examine theramifications ofthecloudSLAbeforesigning.Forexample,99.9%uptime, a common stipulation, translates to nine hours of outage per year. For some mission-critical data, that maynot be adequate. You should also check to see how terms are defined.

SLAsthatscale

Most SLAs are negotiated to meet the needs of the customer at the time of signing, butmanybusinesseschangedramaticallyinsizeovertime.Asolidcloud service-levelagreementoutlines intervalsforreviewinga contractso thatitmeets the changing needs of an organization.

Some vendors even build in notification workflows that indicate when a cloud service-levelagreementisclosetobeingbreachedso newnegotiationscanbe initiated based on the changes in scale. When entering any cloud SLA negotiation, it's important to protect the business by clarifying uptimes.A good SLA protects both the customer and supplier from missed expectations.

Finally,thecloudSLAshouldincludeanexitstrategythatoutlinesthe expectations of the provider to ensure a smooth transition.

Service Level Agreements

https://www.youtube.com/watch?v=9Btb72NZwg4

ServiceLevelAgreementsintheCloud:Whocares?

ServiceLevelAgreements (SLAs)in theCloud. Therehavebeen manyarticleswrittenon thetopic,but stillthereisconfusionabout the importanceofSLAs. Most people requirea blueprint for architects and contractors to start building a new home and similarlywould expect a new car to come with a warranty. An SLA serves as both the blueprint and warranty for cloud computing.

There isanarticlewrittenforEducauseQuarterlybyThomasJ.Trapplercalled"IfIt'sinthe Cloud, Get it on Paper: Cloud Computing Contract Issues". In his paper he recommends that the contract:

- Codifies the specific parameters and minimum levels required for each elementoftheservice, aswellasremediesforfailuretomeet thoserequirements.
- Affirmsyourinstitution'sownershipof itsdatastoredontheservice provider's system, and specifies your rights to get it back.
- Detailsthesysteminfrastructureandsecuritystandardstobemaintainedby the service provider, along with your rights to audit their compliance.
- Specifiesyourrightsandcosttocontinueanddiscontinueusingtheservice.

UsingTrappler'sthoughts,let'sdelve intowhya SLAis importanttoensuringthecloud meets the requirements of the enterprise.

In order to survive in today's world, one must be able to expect the unexpected as there are always new, unanticipated challenges. The only way to consistently overcome these challengesistocreatea stronginitialsetofgroundrules,andplan for exceptionsfromthe start. Challenges can come from many fronts, such as networks, security, storage, processing power, database/software availability or even legislation or regulatorychanges. As cloud customers, we operate in an environment that can spans geographies, networks, and systems. It only makes sense to agree on the desired service level for your customers and measure the real results. It only makes sense to set out a plan for when things go badly, so that a minimum level of service is maintained. Businesses depend on computing systems to survive.

In some sense, the SLA sets expectations for both parties and acts as the roadmap for changeinthecloud service–both expectedchangesandsurprises.JustasanyITproject would have a roadmap with clearlydefined deliverables, an SLA is equallycritical for working with cloud infrastructure. That raises the next question in the journey: what should be in the SLA?

Inorder toconsistentlydevelop aneffectiveSLA,alist ofimportantcriterianeeds tobe established. Let's start with an initial list:

- Availability(e.g.99.99%duringworkdays,99.9%fornights/weekends)
- Performance(e.g.maximumresponsetimes)

- Security/privacyofthedata(e.g.encryptingallstoredandtransmitteddata)
- DisasterRecoveryexpectations(e.g.worsecaserecoverycommitment)
- Locationofthedata(e.g.consistentwithlocallegislation)
- Accesstothedata(e.g.dataretrievablefromproviderinreadableformat)
- Portabilityofthedata(e.g.abilitytomove datatoadifferentprovider)
- Processtoidentifyproblemsandresolutionexpectations(e.g.callcenter)
- ChangeManagementprocess(e.g.changes–updatesornewservices)
- Disputemediationprocess(e.g.escalationprocess,consequences)
- ExitStrategywithexpectationsontheprovidertoensuresmoothtransition

With a core setof criteria established, the next step is to evaluate the criticalityof the cloud service and associated data. Nearly any computing system can be made extremely reliable, but the costs may be too high. Not every system needs the samedegreeofreliabilityasNASAdesignedforthespace shuttles, andfewcould afford the costs.

Forexample,providinga read-onlycatalogue forcustomersis fairlysimple.While the catalogue maybe veryhigh value, it is fairlyeasy to restore from backup with minimal customer impact. However, if the same service has an online shopping with financial transactions and customer data, then the risk level and also importance to the business just increased. The nature of the service is integral to determining the right SLA.

Foreverynewcloud service,anSLAassessment processshouldbedone.TheSLAis a living agreement though and as services change, the SLA should be reassessed.

The SLA should act as a guide for handling potential problems. We need to look atthe SLAas a toolforprotecting the stabilityof the service, protecting the assetsof thecompanyandminimizingthe expense should drasticactionsberequired. Asan example, changing serviceproviders andundoingthecontractsinplace,shouldbe a last resort; it's a veryexpensive and painful solution. Nonetheless, it needs to be covered in the SLA so that both parties can disengage a lawsuit.

LessonsthathavebeenlearnedwithrespecttoSLAs:

1. Read yourcloudprovider'sSLAverycarefully –Thealmostfour-dayAmazon outage of April 2011 did not breach Amazon's EC2 SLA, which as a FAQexplains, "guarantees 99.95% availability of theservice within a Region over atrailing365period."Since ithas beentheEBSandRDS servicesratherthan EC2 itself that has failed (and all the failures have been restricted to Availability Zones within a single Region), the SLA has not been breached, legally speaking.
2. Get technical staff involved to validate SLAs against common outagescenarios–Anothersetofoutagescomealongtodelightthosewhofollow

Microsoft'sjourneytocloudnirvana.OnAugust7it wasa poweroutagethat affected their Dublin datacenter and affected service to European users of Business Productivity Online Services (BPOS).

3. Havecontingencyplansinplacetosupportworsecasescenarios –When we re-designed for the cloud this Amazon failure was exactly the sort of issue that we wanted to be resilient to. Our architecture avoids using EBS as our maindatastorageservice,andtheSimpleDB,S3andCassandraservicesthat we do depend upon were not affected by the outage.

Bottom line, the SLA is your contractwith the service provider and sets expectations for the relationship. It needs to be written to protect your cloud service(s) according to the levelofrisk youare preparedtoaccept. ThegoalistohaveanSLAwhich both thecloud consumer and providercan understand and agree to, including an exitstrategy. The SLA should be looked at as the document that establishes the partnership between the parties and is used to mitigate any problems.

Hopefully you can now agree thata SLAis required fora Cloud service and isforthe benefitofboththeconsumerandtheprovider.In thelong run,itwillsavebothparties money and drive satisfaction for not only the parties directly involved, but more importantly, the end-users.

In should be noted, at the recent December member meeting, the Cloud StandardsCustomerCouncillaunchedanewWorkingGrouptofocusonthedevelopmentofa Service Level Agreement Cookbook.

What is hypervisor?

Unit 4 :- Virtualization

What is Virtualization?

Virtualization is technology that you can use to create virtual representations of servers, storage, networks, and other physical machines. Virtual software mimics the functions of physical hardware to run multiple virtual machines simultaneously on a single physical machine. Businesses use virtualization to use their hardware resources efficiently and get greater returns from their investment. It also powers cloud computing services that help organizations manage infrastructure more efficiently.

Why is virtualization important?

By using virtualization, you can interact with any hardware resource with greater flexibility. Physical servers consume electricity, take up storage space, and need maintenance. You are often limited by physical proximity and network design if you want to access them. Virtualization removes all these limitations by abstracting physical hardware functionality into software. You can manage, maintain, and use your hardware infrastructure like an application on the web.

Virtualization example

Consider a company that needs servers for three functions:

1. Store business email securely
2. Run a customer-facing application
3. Run internal business applications

Each of these functions has different configuration requirements:

- The email application requires more storage capacity and a Windows operating system.
- The customer-facing application requires a Linux operating system and high processing power to handle large volumes of website traffic.
- The internal business application requires iOS and more internal memory (RAM).

To meet these requirements, the company sets up three different dedicated physical servers for each application. The company must make a high initial investment and perform ongoing maintenance and upgrades for one machine at a time. The company also cannot optimize its computing capacity. It pays 100% of the servers' maintenance costs but uses only a fraction of their storage and processing capacities.

What are the benefits of virtualization?

Virtualization provides several benefits to any organization:

Efficient resource use

Virtualization improves hardware resources used in your data center. For example, instead of running one server on one computer system, you can create a virtual server pool on the same computer system by using and returning servers to the pool as required. Having fewer underlying physical servers frees up space in your data center and saves money on electricity, generators, and cooling appliances.

Automated IT management

Now that physical computers are virtual, you can manage them by using software tools. Administrators create deployment and configuration programs to define virtual machine templates. You can duplicate your infrastructure repeatedly and consistently and avoid error-prone manual configurations.

Faster disaster recovery

When events such as natural disasters or cyberattacks negatively affect business operations, regaining access to IT infrastructure and replacing or fixing a physical server can take hours or even days. By contrast, the process takes minutes with virtualized environments. This prompt response significantly improves resiliency and facilitates business continuity so that operations can continue as scheduled.

What are the different types of virtualization?

You can use virtualization technology to get the functions of many different types of physical infrastructure and all the benefits of a virtualized environment. You can go beyond virtual machines to create a collection of virtual resources in your virtual environment.

Server virtualization

Server virtualization is a process that partitions a physical server into multiple virtual servers. It is an efficient and cost-effective way to use server resources and deploy IT services in an organization. Without server virtualization, physical servers use only a small amount of their processing capacities, which leave devices idle.

Storage virtualization

Storage virtualization combines the functions of physical storage devices such as network attached storage (NAS) and storage area network (SAN). You can pool the storage hardware in your data center, even if it is from different vendors or of different types. Storage virtualization uses all your physical data storage and creates a large unit of virtual storage that you can assign and control by using management software. IT administrators can streamline storage activities, such as archiving, backup, and recovery, because they can combine multiple network storage devices virtually into a single storage device.

Network virtualization

Any computer network has hardware elements such as switches, routers, and firewalls. An organization with offices in multiple geographic locations can have several different network technologies working together to create its enterprise network. Network virtualization is a process that combines all of these network resources to centralize administrative tasks. Administrators can adjust and control these elements virtually without touching the physical components, which greatly simplifies network management.

The following are two approaches to network virtualization.

Software-defined networking

Software-defined networking (SDN) controls traffic routing by taking over routing management from data routing in the physical environment. For example, you can program your system to prioritize your video call traffic over application traffic to ensure consistent call quality in all online meetings.

Network function virtualization

Network function virtualization technology combines the functions of network appliances, such as firewalls, load balancers, and traffic analyzers that work together, to improve network performance.

Data virtualization

Modern organizations collect data from several sources and store it in different formats. They might also store data in different places, such as in a cloud infrastructure and an on-premises data center. Data virtualization creates a software layer between this data and the applications that need it. Data virtualization tools process an application's data request and return results in a suitable format. Thus, organizations use data virtualization solutions to increase flexibility for data integration and support cross-functional data analysis.

Application virtualization

Application virtualization pulls out the functions of applications to run on operating systems other than the operating systems for which they were designed. For example, users can run a Microsoft Windows application on a Linux machine without changing the machine configuration. To achieve application virtualization, follow these practices:

- Application streaming – Users stream the application from a remote server, so it runs only on the end user's device when needed.
- Server-based application virtualization – Users can access the remote application from their browser or client interface without installing it.
- Local application virtualization – The application code is shipped with its own environment to run on all operating systems without changes.

Desktop virtualization

Most organizations have nontechnical staff that use desktop operating systems to run common business applications. For instance, you might have the following staff:

- A customer service team that requires a desktop computer with Windows 10 and customer-relationship management software
- A marketing team that requires Windows Vista for sales applications

You can use desktop virtualization to run these different desktop operating systems on virtual machines, which your teams can access remotely. This type of virtualization makes desktop management efficient and secure, saving money on desktop hardware. The following are types of desktop virtualization.

Hypervisor :- The hypervisor is a keyto enable virtualization. In its simpler form, the hypervisor is specialized firmwareorsoftware,orboth,installedonsinglehardware thatwouldallow you to host several virtual machines. It allows physical hardware to be shared across several virtual machines. A computer on which hypervisor runs one or more virtual machines is called a host machine. The virtual machine is called a guest machine.

Basically,thehypervisorallowsthephysicalhostmachineto run variousguestmachines. This helps in achieving maximum benefits from computing resources such as memory, network bandwidth, and CPU cycles.

Advantages of Hypervisors

- Though virtual machines operate on the same physical hardware, they are separatedfromeachother. Thisalsodepictsthatifone virtualmachineundergoes a crash, error, or a malware attack, it doesn't affect the other virtual machines.
- Another benefit is that virtual machines are very mobile as theydon't depend on theunderlyinghardware.Sincetheyarenotlinkedtophysicalhardware,switching between local or remote virtualized servers gets a lot easier as compared to traditional applications.

Types of Hyper visor In Cloud Computing

Therearetwomaintypesofhypervisorincloudcomputing.

Type I Hypervisor

A type I hypervisor operates directlyon the host's hardware to monitor hardware and guest virtual machines, and it's referred to as the bare metal. Usually, theydon't require theinstallationof softwareaheadoftime. Instead, youcan installrightontothehardware. This type of hypervisor tends to be powerful and requires a great deal of expertise to function it well. In addition, Type I hypervisor are more complex and have certain hardwarerequirementstorunadequately.Dueto this, it is mostlychosenbyIToperations and data center computing.

Examplesoftype Ihypervisorsinclude Xen,OracleVMServerforSPARC,Oracle VM Server for x86, Microsoft Hyper-V and VMware’s ESX/ESXi.

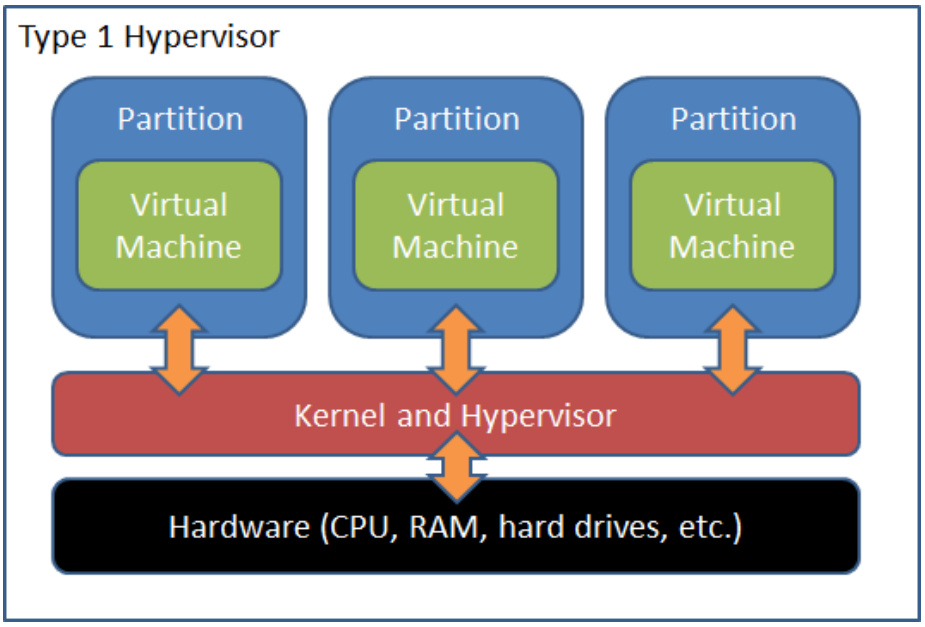

Type II Hypervisor

It's also called a hosted hypervisor because it is usually installed onto an existing operating system. Theyare not much capable to run more complex virtual tasks. People use it for basic development, testing, and emulation. If there is anysecurityflaw found inside thehostOS,itcanpotentiallycompromiseallofvirtualmachinesrunning. Thisis whytype IIhypervisorscannotbeusedfordata centercomputing. Theyaredesignedfor end-user systems where security is a lesser concern. For instance, developers could use type II Hypervisor to launch virtual machines in order to test software product before their release.

AfewexamplesareVirtualbox,VMwareworkstation,fusion.

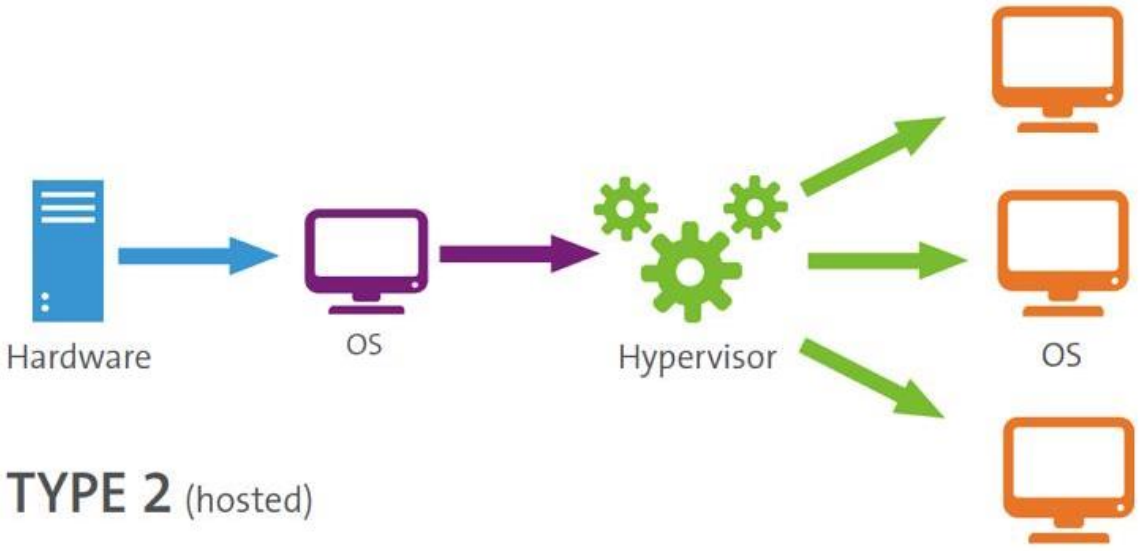

Conclusion

When you achieve virtualization, it brings a consolidation of multiple resources. This tends to reduce costs and improves manageability. In addition to it, a hypervisor can manage increased workloads. In a situation when a specific hardware node gets overheated, youcaneasilyswitch those virtualmachinesonto someotherphysicalnodes. Virtualization also delivers other benefits of security, debugging and support.

What makes virtualization possible are hypervisors.

Server virtualizationallows different operating systems running separate applicationsononeserverwhilestillusingthesamephysicalresources.These virtual machines make it possible for a system and network administrators to have a dedicated machine for every service they need to run.

Notonlydoesthisreducethenumberof physicalserversrequired, butitalso saves time when trying to pinpoint issues.

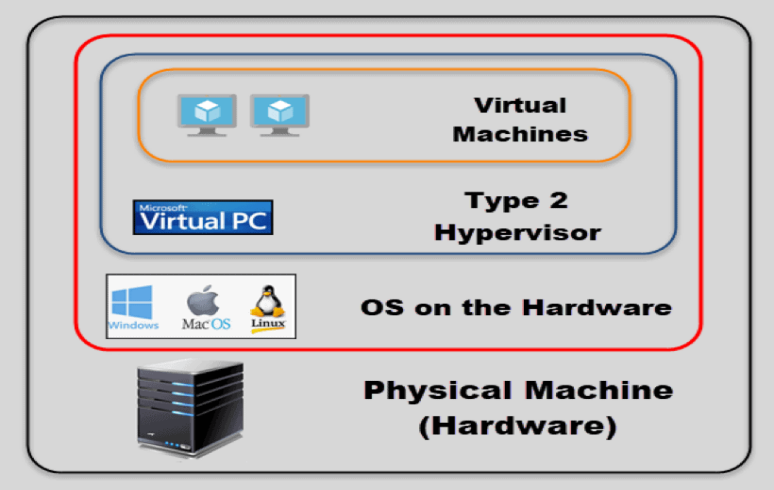

Hypervisor Type2 Performance

Hostedhypervisorsessentiallyalsoactasmanagementconsolesforvirtual machines, you can perform any task using the built-in functionalities.

There is no need to install separate software on another machine to create and maintainyourvirtualenvironment.Yousimplyinstallandruna type2 hypervisor as you would any other application within your OS. With it, you can create snapshots or clone your virtual machines, import or export appliances, etc.

Hereisoneexampleofatype2hypervisorinterface(VirtualBoxbyOracle):

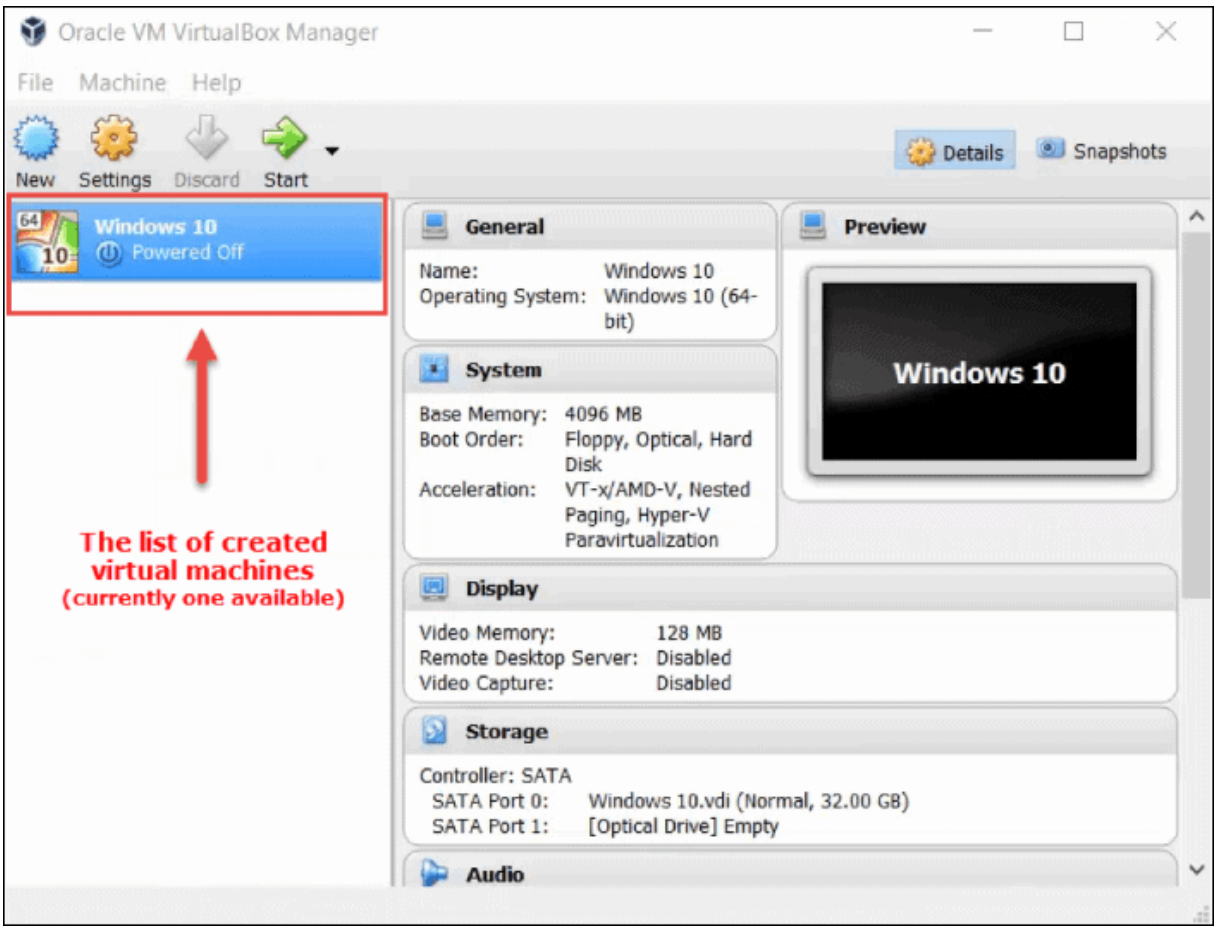

Youdo needto becarefulwhenallocatingactualresourceswiththistypeof hypervisor.

Bare-metalhypervisorscandynamicallyallocateavailableresourcesdepending on the current needs of a particularVM.A type 2 hypervisor occupies whatever you allocate to a virtual machine.

Whenyouassign8GB ofRAMto aVM,thatamountwillbe taken upevenifthe VM is using only a fractionof it. If the host machine has 32GB of RAM and you create three VMs with 8GB each, you are left with 8GB of RAM to keep the physical machine running. Creating another VM with 8GB of ram would bring downyoursystem.Thisiscriticaltokeepinmind,soastoavoidover-allocating resources and crashing the host machine.

Type2 hypervisorsareconvenientfortestingnewsoftwareandresearch projects.

It is possible to use one physical machine to run multiple instances with different operating systems to test how an application behaves in each environment or to createaspecific networkenvironment.Youonlyneedtomakesurethatthereare

enoughphysicalresourcesto keepboththe hostandthevirtualmachines running.

Unit 5 Cloud Security

What is cloud security?

Cloud security refers to the cybersecurity policies, best practices, controls, and technologies used to secure applications, data, and infrastructure in cloud environments. In particular, cloud security works to provide storage and network protection against internal and external threats, access management, data governance and compliance, and disaster recovery.

Cloud computing has become the technology of choice for companies looking to gain the agility and flexibility needed to accelerate innovation and meet the expectations of today's modern consumers. But migrating to more dynamic cloud environments requires new approaches to security to ensure that data remains secure across online infrastructure, applications, and platforms.

Cloud security risks and challenges

Cloud suffers from similar security risks that you might encounter in traditional environments, such as insider threats, data breaches and data loss, phishing, malware, DDoS attacks, and vulnerable APIs.

However, most organizations will likely face specific cloud security challenges, including:

Lack of visibility

Cloud-based resources run on infrastructure that is located outside your corporate network and owned by a third party. As a result, traditional network visibility tools are not suitable for cloud environments, making it difficult for you to gain oversight into all your cloud assets, how they are being accessed, and who has access to them.

Misconfigurations

Misconfigured cloud security settings are one of the leading causes of data breaches in cloud environments. Cloud-based services are made to enable easy access and data sharing, but many organizations may not have a full understanding of how to secure cloud infrastructure. This can lead to misconfigurations, such as leaving default passwords in place, failing to activate data encryption, or mismanaging permission controls.

Access management

Cloud deployments can be accessed directly using the public internet, which enables convenient access from any location or device. At the same time, it also means that attackers can more easily gain authorized resources with compromised credentials or improper access control.

Dynamic workloads

Cloud resources can be provisioned and dynamically scaled up or down based on your workload needs. However, many legacy security tools are unable to enforce policies in flexible environments with constantly changing and ephemeral workloads that can be added or removed in a matter of seconds.

Compliance

The cloud adds another layer of regulatory and internal compliance requirements that you can violate even if you don't experience a security breach. Managing compliance in the cloud is an overwhelming and continuous process. Unlike an on-premises data center where you have complete control over your data and how it is accessed, it is much harder for companies to consistently identify all cloud assets and controls, map them to relevant requirements, and properly document everything.

Benefits of cloud security

Although cloud security has often been framed as a barrier to cloud adoption, the reality is that cloud is no more or less secure than on-premises security. In fact, cloud computing security offers many advantages for businesses that can improve your overall security posture.

The top cloud providers have secure-by-design infrastructure and layered security that is built directly into the platform and its services, including everything from zero-trust network architecture to identity and access management to multi-factor authentication, encryption, and continuous logging and monitoring. Plus, the cloud helps you to automate and manage security at an enormous scale.

Unit 5 Cloud Storage

Cloud storage is an integral part of the cloud computing sphere, and it's arguably what thrust the cloud into the mainstream when Dropbox first launched. Keep reading to find out all of the cloud storage advantages (and disadvantages).

Cloud Storage Services & Cloud Computing: How It Works

Before we go into all the benefits of using cloud storage, let's quickly go over how it works. When you upload a file to a cloud storage service, it gets saved on a remote computer called a "server."

You can access these servers from anywhere in the world via the internet. This means that your file is no longer tied to your device, and you can access it from any device, anywhere in the world.

Everytime you access an online service, you're using the cloud.

Obviously this brings with it a host of advantages (as well as some minor **disadvantages), which we'll go over in detail in the list below.**

Advantages of Cloud Services

Here are some of the biggest advantages of using cloud storage.

1. Backup Your Data to the Cloud

If you've ever had a hard drive die on you, you're familiar with the pandemonium it can cause, especially if you're an at-home worker. Not only do you have to get it replaced, but your data might be gone for good.

Some data might be salvageable via a very pricey restoration service, although those aren't always capable of recovering your data, and usually at least some of it will be gone for good. Or an even worse scenario: you lose your phone or laptop. Now there's no way of getting your data back at all.

3. Remotely Update and Sync Your Files

If you make a change to a file that's on your computer and want to update the file on another computer, you'll have to manually copy the file over to that device. On the other hand, updating a cloud file means the file now looks the same to everyone accessing it on every device.

Plus, most cloud services have a feature called "synchronization," or just "sync." To use sync, you usually need to download a cloud app called a "sync client" and log in. If you have the sync client installed on multiple devices, it will sync your files instantly across every device, which means you don't have to manually upload or download anything.

4. Share Files Easily

One of the best things about cloud storage is just how easy file sharing is. If your files are already in the cloud, you don't need to upload them to a file transfer service or even send an email. You can just find the file in your cloud account and hit "share." Usually you'll get a link that you can paste wherever you want, and the recipient will be able to download or view your file.

5. Remote Work Made Easy

Cloud technology is essentially what makes remote work possible in the first place. Without remote access to your work files, you're stuck moving back and forth between the office, which defeats the purpose of remote working.

However, since the cloud leaves all your cloud files at your disposal wherever you are, it's easy to pick up where you left off when you get back home. Just open your cloud account, and all the files you worked on in the office are right there. Head over to

The cloud lets you access your work files remotely to work from home.

6. Keep Your Files Encrypted

A lot of people are scared of the internet and think their files are safer just staying offline. However, that's not the case. If your files are only stored locally on your device, it makes them susceptible to hackers who can get into your computer and hold your data for ransom. Not to mention all the various hardware issues that could cause you to lose your data permanently.

Keeping your files in the cloud is simply the smarter thing to do. Luckily most cloud services offer encryption for your files. Encryption is a process that scrambles your files into a string of unreadable code, which is only readable using an encryption key. One flaw is that if you store your files with an untrustworthy operator, the operator can decrypt and access your files.

7. Storage for a Lifetime

When you use an online service, making a long-term commitment always pays off. That's especially true with something you don't want to change frequently, like cloud storage.

Thankfully several cloud solutions offer lifetime plans. A lifetime plan lets you make a one-time purchase for a certain amount of storage and use that storage forever. The cost savings are usually greater than monthly or even yearly plans.

Disadvantages of Cloud Storage

Of course, as with everything in life, cloud storage comes with some negatives, though they are minor (or expected) ones.

1. You Need an Internet Connection

To use the cloud, you need to be connected to the internet. Sure you can sync your files to access them even when you're offline, but actually accessing the cloud absolutely requires an internet connection.

2. Extra Storage Space Comes at a Cost

This isn't really surprising, as very few online services are fully free. Although most cloud storage services offer some amount of free storage, like MEGA's enormous 20GB free plan, that's about as much as you'll ever get for free. Luckily, some services offer a

wealth of cloud storage features at a low price, like Sync.com, which tops our list of the best cloud storage deals.

3. Security and Privacy Concerns With Some Providers

People often fall into the trap of a cloud service offering hundreds of gigabytes of storage for free. Then they mysteriously start getting spammed with fishy-looking emails (or should we say "phishy"?), or their data just vanishes from the service. It happens more often than you'd think and puts people off even trustworthy cloud storage solutions.

What is a Storage Area Network (SAN)?

A Storage Area Network (SAN) is a dedicated, independent high-speed network that interconnects and delivers shared pools of storage devices to multiple servers. Each server can access shared storage as if it were a drive directly attached to the server. A SAN is typically assembled with cabling, host bus adapters, and SAN switches attached to storage arrays and servers. Each switch and storage system on the SAN must be interconnected.

Advantages of a storage area network

- Fast
- Easy to scale
- Fault tolerant
- Better hard disk utilization

One of the biggest advantages of using a storage area network is that SANs generally provide better performance than competing technologies, such as NAS. A NAS system acts as a NAS device, with storage traffic being routed across a LAN. These NAS systems are inexpensive -- with some consisting of only a single hard drive -- and easy to configure. However, NAS devices tend to be slow because they are constrained by network bandwidth limitations. Even so, they remain a popular choice -- especially in smaller environments -- for use as file servers or as a general-purpose data storage platform. Some organizations also use NAS devices for data protection.

SANs are also designed to be easy to scale. An organization can add capacity simply by adding switches and hard disks.

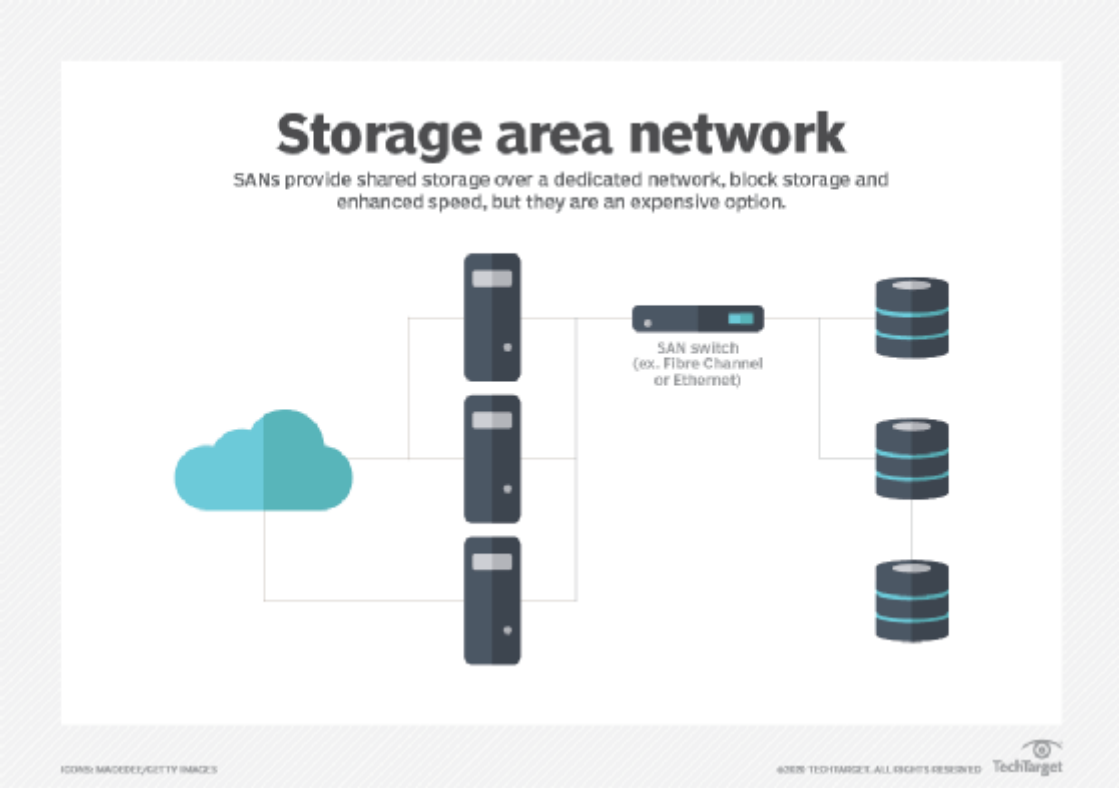

A storage area network is a dedicated high-speed network that connects pools of storage devices to multiple servers.

Another advantage to using a SAN is that it helps organizations to use their raw storage more efficiently. Consider that when DAS is used, each server has its own disks, spanned by a file system. Because it's almost impossible to estimate the exact amount of storage space that a server will need, there is almost always some amount of capacity that will go unused. Conversely, SANs treat storage space as a pool of resources that can be allocated on an as-needed basis. This helps avoid wasted capacity.

Disadvantages of a storage area network

- Expensive
- Complex

SANs have two main disadvantages. First, SANs are expensive. Redundant, high-performance hardware is costly by its very nature. However, it isn't just the acquisition costs that make SANs expensive. There are also ongoing maintenance and management costs to consider.

The other main disadvantage to operating a SAN is its complexity. Deploying, managing and maintaining a SAN requires special knowledge. As such, organizations that choose to deploy SANs must provide the IT staff with additional training or hire storage administrators.

Cloud access security broker (CASB) defined

A cloud access security broker, often abbreviated (CASB), is a security policy enforcement point positioned between enterprise users and cloud service providers. CASBs can combine multiple different security policies, from authentication and credential mapping to encryption, malware detection, and more, offering flexible enterprise solutions that help ensure cloud app security across authorized and unauthorized applications, and managed and unmanaged devices.

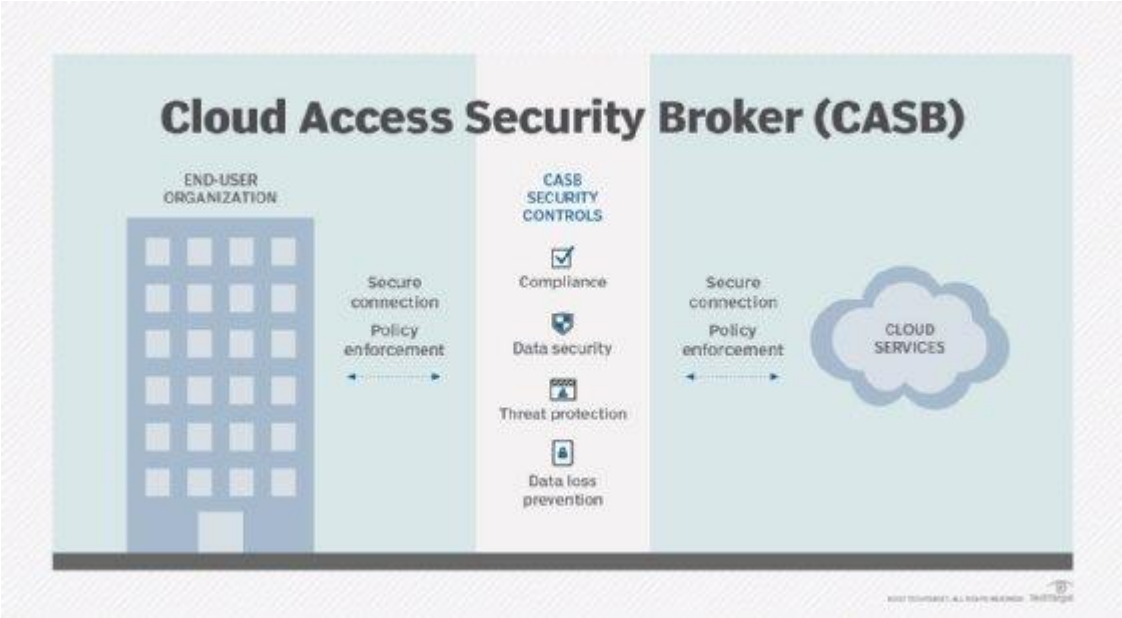

Main Cloud Security Issues and Threats

Almost every organization has adopted cloud computing to varying degrees within their business. However, with this adoption of the cloud comes the need to ensure that the organization's cloud security strategy is capable of protecting against the top threats to cloud security.

Misconfiguration

Misconfigurations of cloud security settings are a leading cause of cloud data breaches. Many organizations' cloud security posture management strategies are inadequate for protecting their cloud-based infrastructure.

Several factors contribute to this. Cloud infrastructure is designed to be easily usable and to enable easy data sharing, making it difficult for organizations to ensure that data is only accessible to authorized parties. Also, organizations using cloud-based infrastructure also do not have complete visibility and control over their infrastructure, meaning that they need to rely upon security controls provided by their cloud service provider (CSP) to configure and secure their cloud deployments. Since many organizations are unfamiliar with securing cloud infrastructure and often have multi-cloud deployments – each with a different array of vendor-provided security controls

Unauthorized Access

Unlike an organization's on-premises infrastructure, their cloud-based deployments are outside the network perimeter and directly accessible from the public Internet. While this is an asset for the accessibility of this infrastructure to employees and customers, it also makes it easier for an attacker to gain unauthorized access to an organization's cloud-based resources. Improperly-configured security or compromised credentials can enable an attacker to gain direct access, potentially without an organization's knowledge.

Insecure Interfaces/APIs

CSPs often provide a number of application programming interfaces (APIs) and interfaces for their customers. In general, these interfaces are well-documented in an attempt to make them easily-usable for a CSP's customers.

However, this creates potential issues if a customer has not properly secured the interfaces for their cloud-based infrastructure. The documentation designed for the customer can also be used by a cybercriminal to identify and exploit potential methods for accessing and exfiltrating sensitive data from an organization's cloud environment.

Hijacking of Accounts

Many people have extremely weak password security, including password reuse and the use of weak passwords. This problem exacerbates the impact of phishing attacks and data breaches since it enables a single stolen password to be used on multiple different accounts.

Account hijacking is one of the more serious cloud security issues as organizations are increasingly reliant on cloud-based infrastructure and applications for core business functions. An attacker with an employee's credentials can access sensitive data or functionality, and compromised customer credentials give full control over their online account. Additionally, in the cloud, organizations often lack the ability to identify and respond to these threats as effectively as for on-premises infrastructure.

Lack of Visibility

An organization’s cloud-based resources are located outside of the corporate network and run on infrastructure that the company does not own. As a result, many traditional tools for achieving network visibility are not effective for cloud environments, and some organizations lack cloud-focused security tools. This can limit an organization’s ability to monitor their cloud-based resources and protect them against attack.

External Sharing of Data

The cloud is designed to make data sharing easy. Many clouds provide the option to explicitly invite a collaborator via email or to share a link that enables anyone with the URL to access the shared resource.

Malicious Insiders

Insider threats are a major security issue for any organization. A malicious insider already has authorized access to an organization’s network and some of the sensitive resources that it contains. Attempts to gain this level of access are what reveals most attackers to their target, making it hard for an unprepared organization to detect a malicious insider.

On the cloud, detection of a malicious insider is even more difficult. With cloud deployments, companies lack control over their underlying infrastructure, making many traditional security solutions less effective.

Cyberattacks

Cybercrime is a business, and cybercriminals select their targets based upon the expected profitability of their attacks. Cloud-based infrastructure is directly accessible from the public Internet, is often improperly secured, and contains a great deal of sensitive and valuable data. Additionally, the cloud is used by many different companies, meaning that a successful attack can likely be repeated many times with a high probability of success. As a result, organizations’ cloud deployments are a common target of cyberattacks.

Denial of Service Attacks

The cloud is essential to many organizations’ ability to do business. They use the cloud to store business-critical data and to run important internal and customer-facing applications.

This means that a successful Denial of Service (DoS) attack against cloud infrastructure is likely to have a major impact on a number of different companies. As a result, DoS attacks where the attacker demands a ransom to stop the attack pose a significant threat to an organization’s cloud-based resources.

Unit 7 :- Scheduling in cloud

What is a Cloud Scheduler?

Cloud schedulers automate anything from cloud infrastructure to data pipelines. Each cloud service provider offers a scheduler with a unique feature set, but they're all part of a broader group of cloud automation solutions.

A cloud scheduler may have some limitations, but it can be a huge time saver as part of any cloud investment. With a need to automate a lengthy list of online services, cloud service providers (CSP), including AWS, Google Cloud Platform (GCP), and Microsoft Azure, each developed their own platform-centric automation tools. These cloud-native schedulers are finely tuned to automate their own ecosystems.

Types of Scheduling

FCFS Scheduling?

The full form of FCFS Scheduling is **First Come First Serve Scheduling**.

FCFS Scheduling algorithm automatically executes the queued processes and requests in the order of their arrival. It allocates the job that first arrived in the queue to the CPU, then allocates the second one, and so on. FCFS is the simplest and easiest CPU scheduling algorithm, managed with a FIFO queue. FIFO stands for First In First Out. The FCFS scheduling algorithm places the arriving processes/jobs at the very end of the queue. So, the processes that request the CPU first get the allocation from the CPU first. As any process enters the FIFO queue, its Process Control Block (PCB) gets linked with the queue's tail. As the CPU becomes free, the process at the very beginning gets assigned to it. Even if the CPU starts working on a longer job, many shorter ones have to wait after it. The FCFS scheduling algorithm works in most of the batches of operating systems.

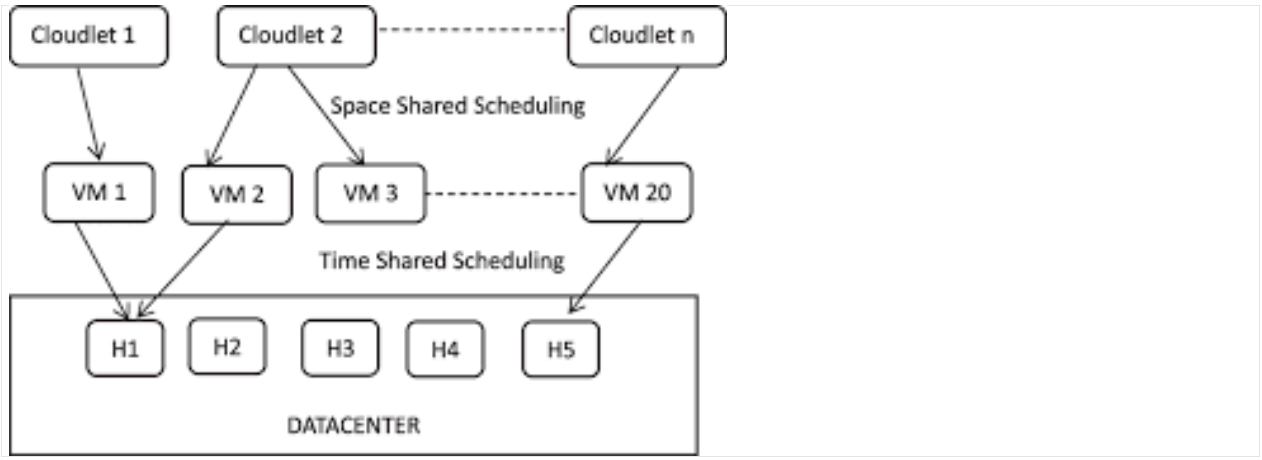

What is Priority Scheduling?

Priority Scheduling is a method of scheduling processes that is based on priority. In this algorithm, the scheduler selects the tasks to work as per the priority.

The processes with higher priority should be carried out first, whereas jobs with equal priorities are carried out on a round-robin or FCFS basis. Priority depends upon memory requirements, time requirements

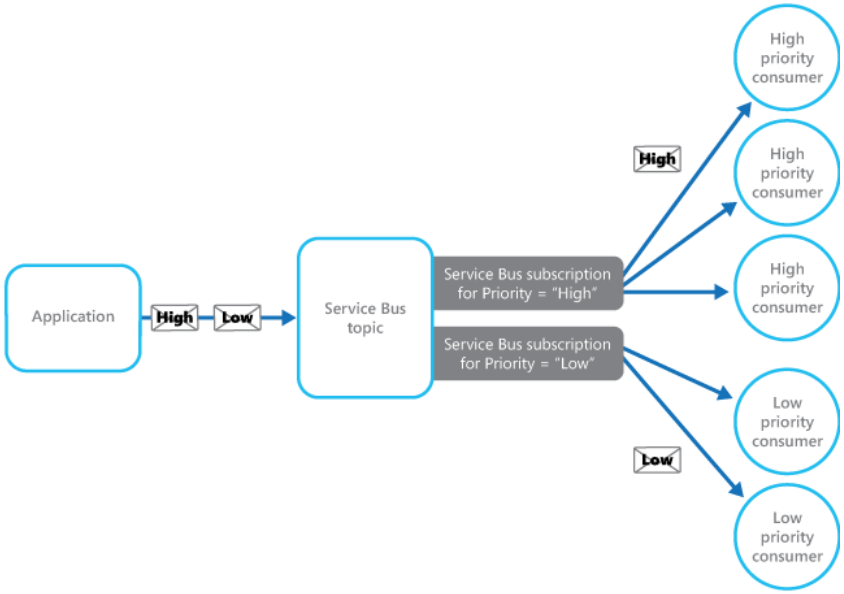

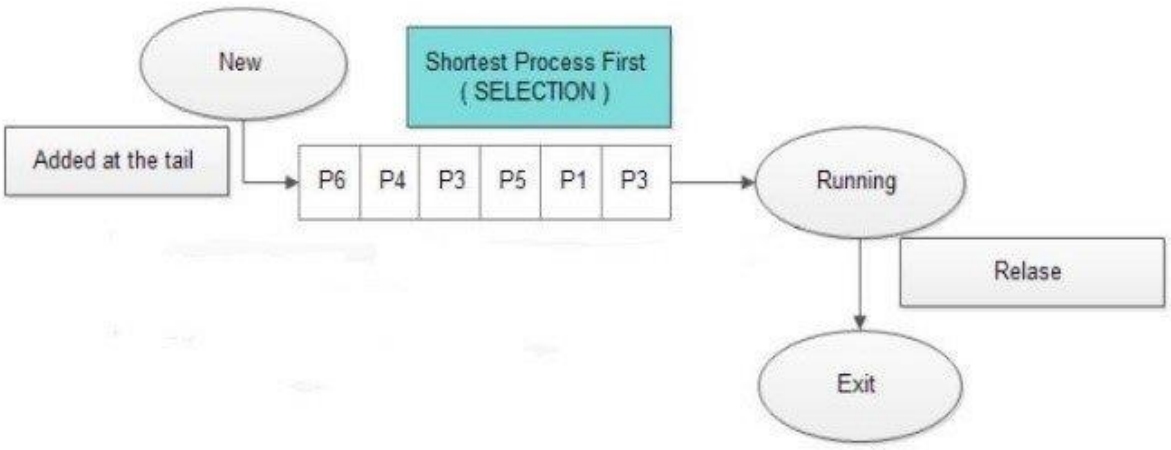

Difference between static and dynamic scheduling in a distributed system:

- In the static scheduling method, the assignment of tasks is performed offline, that is, before the real-time task is officially scheduled for execution on the processor, the task's allocation and scheduling time on the processor are scheduled, and the task is executed in accordance with the pre-scheduling plan after the task has officially started executing. Its advantage is that it can schedule the deployment in advance to reduce the overhead in the task scheduling process.
- In the dynamic scheduling method, the task assignment and schedulability testing are performed online while the system is running. Because, in a real-time system, many tasks are not on the processor in a periodic manner, especially non-periodic tasks arrive at the system randomly and are dynamically scheduled for the execution.