USN

Internal Assessment Test II – DEC 2024

Sub:	Cloud Computing					Sub Code:	BAD 515C	Branch:	AINDS
Date:	13/12/2024	Duration:	90 minutes	Max Marks:	50	Sem	V		

Answer any FIVE Questions

			MARKS	CO
1		What are the primary security concerns faced by cloud users and how can they be mitigated?	10	CO
2		With respect to Data-Center design and Interconnection networks, discuss below topics (a) and (b):		
		(a) Cooling System of a Data-Center room	5	C
		(b) Fault Tolerance and Graceful Degradation	5	
3		Compare the benefits and the potential problems due to virtualization on public, private and hybrid clouds.	10	
4		Analyze the risks posed by foreign mapping and the solution adopted by Xoar. What is the security risk posed by XenStore?	10	
5		What are the RTO and RPO in the design of a disaster recovery scheme? Explain the role of snapshots in disaster recovery schemes.	10	
6		Analyze the main differences between programming on AWS versus Microsoft Azure.	10	HO

1..What are the primary security concerns face by cloud user and how they can be mitigated?
Ans:

Top concerns of cloud users

Cloud computing has become a cornerstone of modern technology, empowering individuals, businesses, and governments with scalable, cost-efficient, and on-demand access to computing resources. However, as cloud adoption accelerates, concerns persist that highlight vulnerabilities and challenges associated with its use.

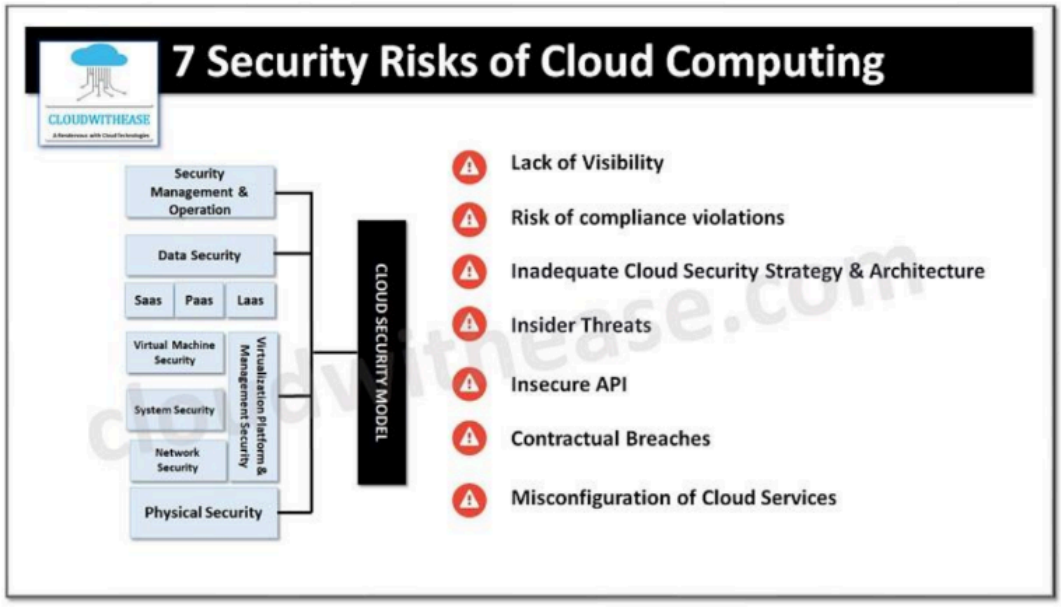

Below, we explore the top concerns of cloud users, diving into their implications and providing a comprehensive analysis.

Insider threats refer to risks posed by employees, contractors, or vendors who misuse their access to cloud resources.

- **Causes:**
  - Malicious intent (e.g., theft, sabotage).
  - Unintentional errors, such as misconfigurations or accidental data sharing.
- **Impact:**
  - Exposure of confidential data.
  - Disruption of services.
  - Financial losses due to operational downtime.
- **Mitigation Strategies:**
  - Implement a **zero-trust security model** where every access is verified.
  - Monitor and log all activities performed by users with privileged access.
  - Regularly review and restrict permissions based on roles and responsibilities.

---

c. Malware and Ransomware Attacks

Malware or ransomware introduced into the cloud can encrypt or steal data, demanding payment for its release or causing disruption.

- **Causes:**
  - Uploading infected files to cloud storage.
  - Compromised endpoints connected to the cloud environment.
- **Impact:**
  - Data unavailability or loss.
  - Business disruptions.
  - Potential financial and legal consequences.
- **Mitigation Strategies:**
  - Use advanced anti-malware and endpoint protection solutions.

For detailed answer pl visit below link:

https://drive.google.com/file/d/1reNFCFzjZwdaLRIOUHPYGpuU3tLAmzyV/view?usp=sharing

2a.Cooling system of a data center room

How does data center cooling work?

Data center cooling removes excess heat from the air and replaces it with cooler air. This is typically done in one of several ways:

- Venting hot air outside and bringing outside air in, cooling it and circulating it in the facility.
- Recycling internal air by cooling it, usually through a hot and cold aisle design to maximize cooling efficiency.
- Venting hot air outside and then drawing pre-chilled outside air into the facility to cool it down. This approach is known as *free cooling*, and it only works for facilities in colder climates.
- Cooling or heating the facility to the highest recommended temperature and replacing equipment once it fails. Using this so-called *heat cooling* or *close-coupled cooling* method can be

cheaper, as other cooling methods might cost significantly more than equipment replacement costs.

Current data center cooling systems and technologies

Air cooling and liquid cooling are two of the most popular types of data center cooling.

Air cooling

This cooling method is ideal for smaller or older data centers that combine raised floors with hot and cold aisle designs. When the computer room AC (CRAC) unit or computer room air handler (CRAH) sends out cold air, the pressure below the raised floor increases and sends the cold air into the equipment inlets. The cold air displaces the hot air, which is then returned to the CRAC or CRAH, where it's cooled and recirculated.

Hot and cold air aisles increase the efficiency of air-based cooling systems by enabling more targeted placement of intake and exhaust vents. This prevents hot and cold air mixing so the cooling CRAC or CRAH can work more efficiently.

Also, a CRAH is more efficient than a CRAC, as it draws outside air in and cools it using chilled water instead of refrigerant. A CRAC functions like a residential AC unit that uses refrigerants to cool the air. CRAC units are more

appropriate for small data center closets because they can't keep up with enterprise-level data centers.

Hot and cold aisle layouts

In this layout, server cabinets and racks are arranged in a row pattern, with each row facing the opposite direction to the one in front of it.

This equipment configuration has cold intake and hot air vents facing each other, creating alternating aisles of hot and cold air. CRAHs in each aisle either vent or pump in air so the cooling system works efficiently. Hot air is vented from the hot aisle, and cool air is pumped through the cold aisle.

Doors and walls can be added to the layout to further direct airflow. Cabinets should be as full as possible to avoid empty spaces, gaps and cable openings that can leak hot or cold air into the opposite aisle, causing the cooling system to work overtime.

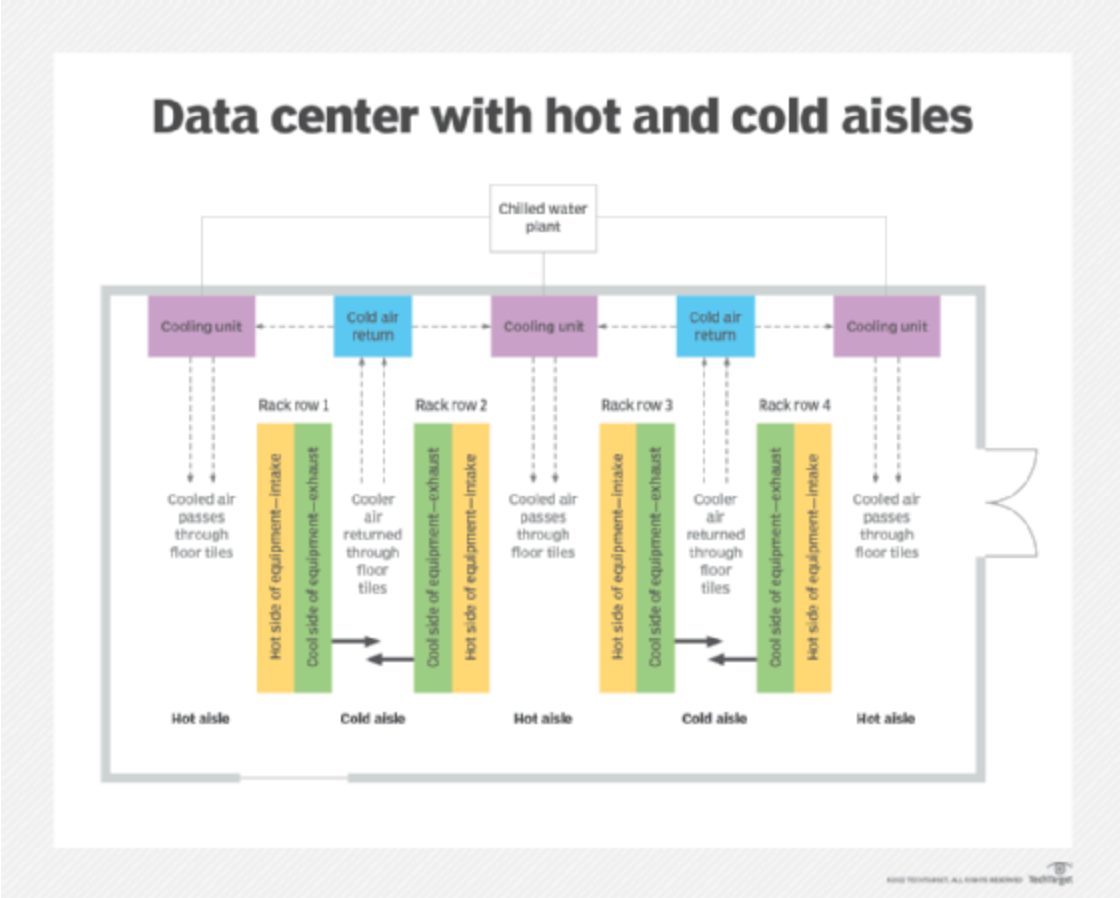

This diagram illustrates how hot and cold air circulates to maintain optimal temperature levels in the data center.

Liquid cooling

A relatively newer technology is liquid cooling. It's a more efficient and cost-effective cooling system because it can be installed on data center devices that need it the most. Liquid is more efficient than air at transferring heat away from emitting sources. It can also support greater equipment densities and items that generate higher-than-average heat, such as high-density and edge-computing data centers.

There are two main types of liquid cooling:

1. **Liquid immersion cooling.** This method places the entire electrical device into dielectric fluid in a closed system. The fluid absorbs the heat emitted by the device, turns it into vapor and condenses it, helping the device to cool down.
2. **Direct-to-chip liquid cooling.** This method uses flexible tubes to bring nonflammable dielectric fluid directly to the processing chip or motherboard component generating the most heat, such as the CPU or GPU. The fluid absorbs the heat by turning it into vapor, which carries the heat out of the equipment through the same tube.

Importance of energy efficiency in data center cooling

Cooling systems should be part of a data center's overall energy-efficiency strategy. According to the latest stats, these systems typically consume a significant amount of the facility's energy consumption -- up to 33%. And with most facilities spending 50% of their operating expenses on energy bills, cooling is a good place to focus efficiency improvements.

Ensuring a state of good repair of the facility infrastructure, such as HVAC and power systems, is a good first start. Next, operators can look at the IT hardware they use to confirm it's always working optimally. Replacement and sunsetting processes can help by introducing more modern, efficient technologies as needed.

Exploring new cooling technologies is another way to manage energy efficiency. New and evolved technologies, such as free cooling and liquid cooling systems, can greatly reduce cooling needs and increase energy consumption efficiency across the facility.

Future data center cooling systems and technologies

Although liquid cooling is still relatively new, other data center cooling technologies are on the horizon, such as geothermal cooling methods, smart technologies that use AI and machine learning to better monitor and manage cooling, and evaporative cooling.

Using nature to cool data centers

Here are some ways data centers can use nature to cool their facilities:

- Geothermal cooling uses the near-constant temperature of the Earth below surface level to provide cooling. It's a centuries-old idea, once used to keep food cold, adapted to our modern era. In data centers, geothermal cooling uses a closed-loop pipe system with water or another coolant that runs through vertical wells underground, filled with a heat-transferring liquid. Iron Mountain's western Pennsylvania data center, Verne Global in Iceland and Green Mountain in Norway use geothermal cooling for their data centers.
- Evaporative cooling, or swamp cooling, takes advantage of the drop in temperature that occurs when water is exposed to moving air and

begins to vaporize and change to a gas. A fan draws warm data center air through a water- or coolant-moistened pad, and as the liquid evaporates, the air is chilled and pushed back into the data center. It can cost a fraction of an air-cooled HVAC system and works best in low-humidity climates.

- Solar cooling converts heat from the sun into cooling that can be used in data center air cooling systems. The system collects solar power and uses a thermally driven cooling process to decrease the air temperature in a building. This is useful in areas with a lot of sunlight or for data centers looking to supplement their current cooling with a more environmentally friendly method.
- Kyoto Cooling is an enhancement of the free cooling method that uses a thermal wheel to control hot and cold airflows across the data center. Internal hot air is vented to the outside as the wheel rotates; the outside air then cools the wheel and the air that is drawn back into the facility. It uses between 75% to 92% less power to run than other CRAH systems, reduces carbon dioxide emissions and eliminates the need for water in the cooling system. The technology is used by United Airlines' data center outside of Chicago and HP in its data center outside of Toronto.

Using smart technology for data center cooling

Because many of the newer data center cooling technologies require significant investment by facility owners, smart technology has become popular. Data center smart assistants, AI and machine learning technologies

can monitor facilities more efficiently and make real-time adjustments to ensure optimal temperatures and humidity levels. Google, for example, uses smart temperature controls to reduce heat output and cooling usage. The company has also used its DeepMind AI product to reduce cooling energy use by 40% in 18 months.

Data center cooling robots can move within the facility, monitoring temperatures and humidity levels in specific server cabinets. One challenge with monitoring temperatures in cabinets manually is that conditions change as soon as they're opened. Companies such as OneNeck IT Solutions have developed a robot sensor probe that fits into standard cabinets. The robot moves up and down a belt-driven rail inside the cabinet to collect temperature data for each rack. It then transmits the data via Bluetooth to connected devices so data center pros can create a full heat map of the cabinet.

https://aws.amazon.com/what-is/service-oriented-architecture/

2b. Fault tolerance and graceful degradation

Fault tolerance is a process that enables an operating system to respond to a failure in hardware or software. This fault-tolerance definition refers to the system's ability to continue operating despite failures or malfunctions.

An operating system that offers a solid definition for faults cannot be disrupted by a single point of failure. It ensures business continuity and the high availability of crucial applications and systems regardless of any failures.

How Does Fault Tolerance Work?

Fault tolerance can be built into a system to remove the risk of it having a single point of failure. To do so, the system must have no single component that, if it were to stop working effectively, would result in the entire system failing.

Fault tolerance is reliant on aspects like load balancing and failover, which remove the risk of a single point of failure. It will typically be part of the operating system's interface, which enables programmers to check the performance of data throughout a transaction.

A fault-tolerance process follows two core models:

Normal functioning

This describes a situation when a fault-tolerant system encounters a fault but continues to function as usual. This means the system sees no change in performance metrics like throughput or response time.

Graceful degradation

Other types of fault-tolerant systems will go through graceful degradation of performance when certain faults occur. That means the impact the fault has on the system's performance is proportionate to the fault severity. In other words, a small fault will only have a small impact on the system's performance rather than causing the entire system to fail or have major performance issues.

3. Benefits and the potential problems due to virtualization on public, private and hybrid cloud

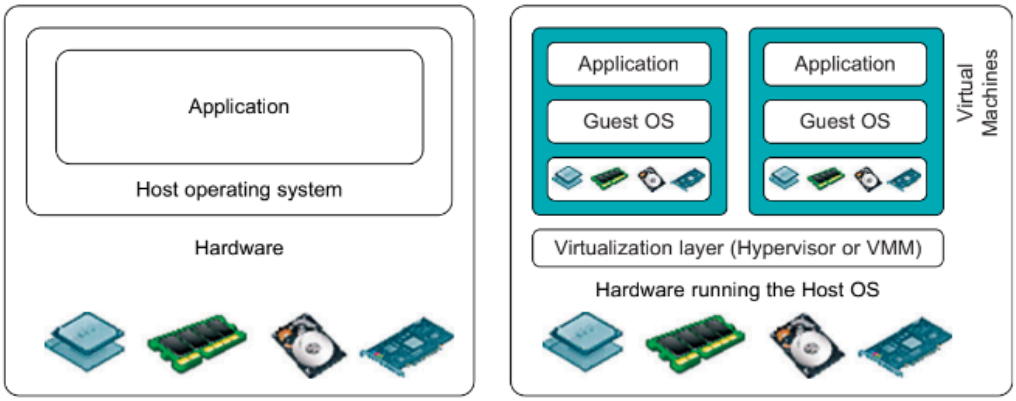

FIGURE 3.1

The architecture of a computer system before and after virtualization, where VMM stands for virtual machine monitor.

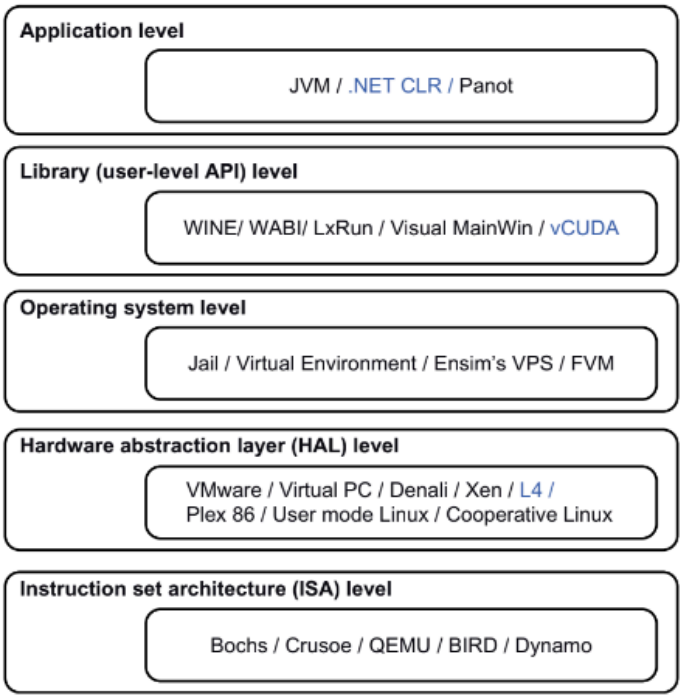

FIGURE 3.2

Virtualization ranging from hardware to applications in five abstraction levels.

Public Cloud

A Public Cloud is Cloud Computing in which the infrastructure and services are owned and operated by a third-party provider and made available to the

public over the internet. The public can access and use shared resources, such as servers, storage, and applications and the main thing is you pay for what you used. . Examples of public cloud providers – are Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP)

Advantages

- **Cost Efficient**: In the public cloud, we have to pay for what we used. So it is more cost-efficient than maintaining the physical servers or their own infrastructure.
- **Automatic Software Updates**: In the public cloud, there are automatic software updates. we don't have to update the software manually.
- **Accessibility**: Public clouds allow users to access their resources and applications from anywhere in the world. We just need an internet connection to access it.

Disadvantages

- **Security and Privacy Concerns:** Public clouds can be vulnerable to data breaches, cyber attacks, and other security risks. Since data is stored on servers owned by a third-party provider, there is always a risk that confidential or sensitive data may be exposed or compromised.

- **Limited Control:** With public cloud services, users have limited control over the infrastructure and resources used to run their applications. This can make it difficult to customize the environment to meet specific requirements.
- **Reliance on Internet Connectivity:** Public cloud services require a reliable and stable internet connection to access the resources and applications hosted in the cloud. If the internet connection is slow or unstable, it can affect the performance and availability of the services.
- **Service Downtime:** Public cloud providers may experience service downtime due to hardware failures, software issues, or maintenance activities. This can result in temporary loss of access to applications and data.
- **Compliance and Regulatory Issues:** Public cloud services may not meet certain compliance or regulatory requirements, such as those related to data privacy or security. This can create legal or contractual issues for businesses that are subject to these requirements.
- **Cost Overruns:** Public cloud services are typically billed on a pay-per-use basis, which can result in unexpected cost overruns if usage exceeds anticipated levels. Additionally, the cost of using

public cloud services may increase over time, as providers adjust their pricing models or add new features and services.

Private Cloud

A Private Cloud is a cloud computing environment in which the infrastructure and services are owned and operated by a single organization, for example, a company or government, and it is accessed by only authorized users within that organization. Private Cloud organizations have their own data center. private cloud provides a higher level of security. Examples – HPE, Dell, VMware, etc.

Advantages

- **Security Status**: Private clouds provide a higher level of security. as the organization has full control over the cloud service. They can customize the servers to manage their security.
- **Customization of Service**: Private clouds allow organizations to customize the infrastructure and services to meet their specific requirements. and also can customize the security.
- **Privacy**: Private clouds provide increased privacy as the organization(company or government) has more control over who has access to their data and resources.

Disadvantages

- **Higher Cost:** Private clouds require dedicated hardware, software, and networking infrastructure, which can be expensive to acquire and maintain. This can make it challenging for smaller businesses or organizations with limited budgets to implement a private cloud.
- **Limited Scalability:** Private clouds are designed to serve a specific organization, which means that they may not be as scalable as public cloud services. This can make it difficult to quickly add or remove resources in response to changes in demand.
- **Technical Complexity:** Setting up and managing a private cloud infrastructure requires technical expertise and specialized skills. This can be a challenge for organizations that lack in-house IT resources or expertise.
- **Security Risks:** Private clouds are typically considered more secure than public clouds since they are operated within an organization's own infrastructure. However, they can still be vulnerable to security risks such as data breaches or cyber attacks.
- **Lack of Standardization:** Private clouds are often built using proprietary hardware and software, which can make it challenging to integrate with other cloud services or migrate to a different cloud provider in the future.
- **Maintenance and Upgrades:** Maintaining and upgrading a private cloud infrastructure can be time-consuming and resource-intensive.

This can be a challenge for organizations that need to focus on other core business activities.

Hybrid Cloud

A hybrid cloud is a combination of both public and private cloud environments that allows organizations to take advantage of the benefits of both types of clouds. It manages traffic levels during peak usage periods It can provide greater flexibility, scalability, and cost-effectiveness than using a single cloud environment. Examples – IBM, DataCore Software, Rackspace, Threat Stack, Infinidat, etc.

Advantages

- **Flexibility**: Hybrid cloud stores its data (also sensitive) in a private cloud server. While public server provides Flexibility and Scalability.
- **Scalability**: Hybrid cloud Enables organizations to move workloads back and forth between their private and public clouds depending on their needs.
- **Security**: Hybrid cloud controls over highly sensitive data. and it provides high-level security. Also, it takes advantage of the public cloud's cost savings.

Disadvantages

- **Complexity:** Hybrid clouds are complex to set up and manage since they require integration between different cloud environments. This can require specialized technical expertise and resources.
- **Cost:** Hybrid clouds can be more expensive to implement and manage than either public or private clouds alone, due to the need for additional hardware, software, and networking infrastructure.
- **Security Risks:** Hybrid clouds are vulnerable to security risks such as data breaches or cyber attacks, particularly when there is a lack of standardization and consistency between the different cloud environments.
- **Data Governance:** Managing data across different cloud environments can be challenging, particularly when it comes to ensuring compliance with regulations such as GDPR or HIPAA.
- **Network Latency:** Hybrid clouds rely on communication between different cloud environments, which can result in network latency and performance issues.
- **Integration Challenges:** Integrating different cloud environments can be challenging, particularly when it comes to ensuring compatibility between different applications and services.
- **Vendor Lock-In:** Hybrid clouds may require organizations to work with multiple cloud providers, which can result in vendor lock-in and limit the ability to switch providers in the future.

4 XOAR

XOAR- Breaking the monolithic design of TCB

- Xoar is a version of Xen designed to boost system security; based on micro-kernel design principles. The design goals are:
  - Maintain the functionality provided by Xen.
  - Ensure transparency with existing management and VM interfaces.
  - Tight control of privileges, each component should only have the privileges required by its function.
  - Minimize the interfaces of all components to reduce the possibility that a component can be used by an attacker.
  - Eliminate sharing. Make sharing explicit whenever it cannot be eliminated to allow meaningful logging and auditing.
  - Reduce the opportunity of an attack targeting a system component by limiting the time window when the component runs.
- The security model of Xoar assumes that threats come from:
  - A guest VM attempting to violate data integrity or confidentiality of another guest VM on the same platform, or to exploit the code of the guest.
  - Bugs in the initialization code of the management virtual machine.

Xoar system components

- Permanent components -> XenStore-State maintains all information regarding the state of the system.
- Components used to boot the system; they self-destruct before any user VM is started. They discover the hardware configuration of the server including the
  - PCI drivers and then boot the system:
  - PCIBack - virtualizes access to PCI bus configuration.
  - Bootstrapper - coordinates booting of the system.
- Components restarted on each request:
  - XenStore-Logic.
  - Toolstack - handles VM management requests, e.g., it requests the Builder to create a new guest VM in response to a user request.
  - Builder - initiates user VMs.
- Components restarted on a timer; the two components export physical storage device drivers and the physical network driver to a guest VM.
  - Blk-Back - exports physical storage device drivers using udev rules.
  - NetBack - exports the physical network driver.

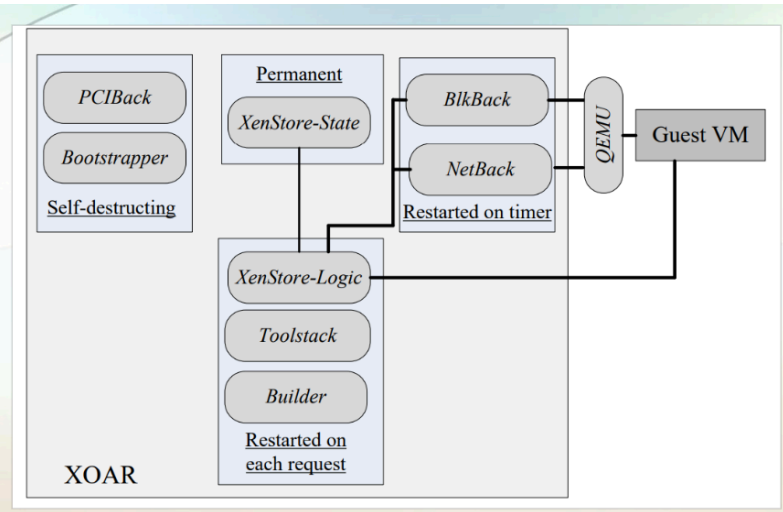

Xoar has nine classes of components of four types: permanent, self-destructing, restarted upon request, and restarted on timer. A guest VM is started using the by the Builder using the Toolstack; it is controlled by the XenStore-Logic. The devices used by the guest VM are emulated by the Qemu component. Qemu is responsible for device emulation

https://docs.google.com/presentation/d/1zypVMp6JQ2RA0qXCgpDyWFy3oraLIRb-/edit?usp=sharing&ouid=100088343219679463685&rtpof=true&sd=true

5. Disaster recovery- RTO and RPO

Determining how to protect and recover an application can often be easier than determining how quickly your business needs that application recovered. Establishing the correct recovery objective targets at an application level is a critical part of business continuity planning, though. This blog is intended to help customers as they establish or reevaluate recovery targets, build a recovery plan, and determine how AWS services fit within that plan.

As a quick refresher, RTO stands for Recovery Time Objective and is a measure of how quickly after an outage an application must be available again. RPO, or Recovery Point Objective, refers to how much data loss your application can tolerate. Another way to think about RPO is how old can the data be when this application is recovered? With both RTO and RPO, the targets are measured in hours, minutes, or seconds, with lower numbers representing less downtime or less data loss. Within the context of a Business Continuity Plan, applications having similar RTO targets are grouped together in Tiers, with Tier 0 having the lowest RTO.

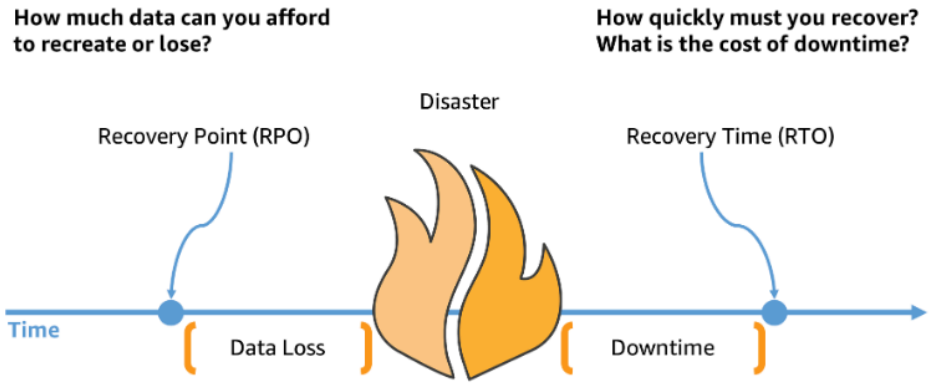

Image1:Data loss is measured from most recent backup to the point of disaster. Downtime is measured from the point of disaster until fully recovered and available for service.

https://aws.amazon.com/blogs/mt/establishing-rpo-and-rto-targets-for-cloud-applications/

6. Programming on AWS Versus Programming on MS Azure

Difference between AWS and Azure: AWS vs Azure

The following are the differences between Microsoft Azure and Amazon Web Services:

Azure	AWS

Azure was launched in 2010	AWS was launched in 2006
In the Cloud, For computation, virtual machines are used.	In AWS, For computation, Elastic Compute Cloud is used.
Azure uses blocks to store.	While it uses Simple Storage Service to store.
Azure is a virtual network.	While AWS is a virtual private cloud.
Azure Cloud spans 140 availability zones. (as of Feb 2023).	AWS cloud spans 61 availability zones. (as of Feb 2023).
SQL databases, MySQL, Cosmos DB, etc., are used in Azure for databases.	In AWS for database, RDS and DynamoDB are used.

The pricing model offered by Microsoft is less flexible.	The pricing model offered by AWS is more flexible.
There are four levels of certification in Azure.	AWS has six levels of specialty certifications.
Microsoft Azure has a 22% market share.	Amazon Web Services has a 33% market share
Some famous clients of Azure are: Nike, Dell, Starbucks, etc	Some famous clients of AWS area include: Netflix, Adobe, Spotify, etc

Key Differences Between AWS and Azure

The following are the key differences between AWS And Azure:

1. Market Share and Reach

- **AWS:** AWS cloud providers have with the largest share and longest history in the cloud industry. It has extensive global reach because of having more regions and availability zones.

- **Azure:** Azure is the second largest share marketer after AWS. It is growing rapidly with establishing strong integration with Microsoft services and enterprise solutions.

2. Service Offerings

- **AWS:** It extensively comes with a wide range of AWS services with a broader selection of computing, storage, database, and machine learning. It is a more mature and diversified service portfolio.
- **Azure:** Azure comes with comprehensive service particularly supporting strongly in hybrid cloud and enterprise services. It has excellent integration with Microsoft products such as (Windows Server, Active Directory, and Office 365).

3. Pricing Models

- **AWS:** It provides flexible pricing with options like On-demand, Reserved Instances, and Spot Instances. It comes with complex pricing structures that might be difficult without detailed analysis.
- **Azure:** It comes with competitive pricing with similar options to AWS including Pay-As-You-Go, Reserved Instances, and Spot pricing. It often offers cost benefits for existing customers of Microsoft through discounts and Credits.

4. Hybrid Cloud and On-premises Integration

- **AWS:** The AWS Outposts service will support hybrid cloud solutions. it focuses mainly on cloud-native approaches.
- **Azure:** Azure strongly emphasizes the hybrid cloud with services like Azure Arc and Azure Stack. It provides seamless integration with on-premises Microsoft environments.

5. Open Source and DevOps

- **AWS:** AWS comes with strong support for a wide range of open-source tools and applications. The extensive services that AWS provides for DevOps include AWS Codepipeline, CodeBuild, CodeDeploy, and CodeCommit.
- **Azure:** Azure facilitates excellent support for open-source technologies through having various partnerships with various open-source communities.

Similarities Between AWS and Azure

The following are the similarities between AWS and Azure:

1. **Core Services:** Both AWS and Azure Cloud Providers offers the essential services such as Computing, storage, databases and networking.

2. **Global Reach:** The extensive global data center networks ensure high availability and redundancy.
3. **Security:** Strong security features, compliance certificates and encryption are prioritized by both the clouds.
4. **Hybrid Cloud**: Both the clouds support the hybrid cloud architectures facilitating integration with on-premises environments.
5. **Pricing Models:** Both uses flexible pricing models like pay-as-you-go, reserved instances and spot instances.

Features and Services Of Microsoft Azure

The following are the features and services of Microsoft Azure:

1. **Compute:** Virtual machines, containers and serverless computing (Azure Functions) services are provided for scalable processing power.
2. **Storage:** Blob Storage, file storage, and disk storage solutions are provided for diverse types of data needs.
3. **Databases:** Managed databases including SQL Database, Cosmos DB, and MySQL for various data requirements.
4. **AI & Machine Learning**: Azure Machine Learning, Cognitive Services, and AI tools for advanced analytics and AI applications.
5. **DevOps:** Azure DevOps, Github Integration and CI/CD pipelines for streamlined development workflows.

Advantages of Microsoft Azure

The following are the advantages of Microsoft Azure

1. **Scalability:** It provides easily scalable resources for meeting the demand.
2. **Global Reach:** It provides extensive network of data centers to worldwide.
3. **Integration:** It facilitates seamless integration with Microsoft products and services.
4. **Security:** It provides the strong security features and compliance certification to the applications and projects.

Disadvantages of Azure

The following are the disadvantages of Azure:

1. **Complexity:** It can be complex to manage and configure the resources.
2. **Cost**: It provides potentially high costs without proper management.
3. **Services Outages**: The occasional service outages can impact on the availability.
4. **Learning Curve:** It is Potential for high costs without proper management.

Features and Services of AWS

The following are the features and services of AWS:

1. **Compute**: It comes up with providing the services such as EC2 Instances, Lambda (Serverless Computing), and Elastic Beanstalk for scalable processing power.
2. **Storage:** S3 (Object storage), EBS (block storage), and Glacier (archival storage) are provided for storing diverse types of data supporting their needs.
3. **Databases:** Managed databases include Amazon RDS, DynamoDB, and Amazon Aurora for various data requirements.
4. **AI and Machine Learning**: SageMaker, Rekognition, and AI services are used for advanced analytics and AI applications.

Advantages of AWS

The following are the advantages of AWS:

1. **Scalability:** It provides highly scalable infrastructure to meet varying demands.
2. **Global Reach:** AWS extensively provides a network of data centers worldwide to facilitate global reach.
3. **Service Variety:** It provides a wide range of services and tools for different use cases.
4. **Security:** It provides strong security features and compliance certification to the applications and projects.

Disadvantages of AWS

The following are the disadvantages of AWS:

1. **Cost:** It potential for high costs, especially without proper cost management.
2. **Complexity**: It will be complex to navigate and manage due to vast number of services.
3. **Services Outages**: The occasional service outages can impact on the availability.
4. **Learning Curve:** It is Potential for high costs without proper management.

What is the Best Choice Between Azure and AWS?

The choosing between Azure and AWS depends on specific needs and existing infrastructure:

Azure

- **Best For**: Enterprises are heavily investing in microsoft products such as on Windows Servers, SQL servers and Office 365.
- **Integration**: It seamlessly integrate with microsoft services and tools.
- **Hybrid Solutions:** It provide strong hybrid cloud capabilities with Azure Arc and Azure stack.

AWS

- **Best For**: Startups, and companies needing a wide range of cloud services and flexibility.
- **Service Variety:** It extensively provides a range of services and rapid innovation.
- **Market Leader**: Long standing leader in the cloud market with a mature ecosystem.

In conclusion that Azure cloud is ideal for prioritizing the integration with Microsoft products and hybrid solutions. On the otherhand AWS is preferred for its service variety, flexibility and extensive global infrastructure.

Azure and AWS for Multicloud Solutions

The following are the some of the Azure and AWS's Multicloud solutions:

Azure for Multicloud Solutions

1. **Azure Arc**: Manages resources across multiple clouds and on-premises environments.
2. **Integration**: Seamless integration with other cloud platforms and on-premises data centers.
3. **Hybrid Capabilities**: Strong support for hybrid cloud setups with Azure Stack.

AWS for Multicloud Solutions

1. **AWS Outposts**: Extends AWS services to on-premises locations.

2. **Interoperability**: Tools like AWS CloudFormation and AWS Transit Gateway for managing multi cloud environments.
3. **Third-Party Support**: Extensive support for third-party tools and services to facilitate multicloud strategies.

Security Compliance and Support of AWS and Azure

The following are the security compliance and supports that are provided by AWS and Azure:

1. Security Features and Certifications

- **AWS:** Offers AWS Shield, IAM, and advanced threat detection.
- **Azure:** Provides Azure Security Center and Azure Active Directory.

Certifications: Both are compliant with ISO 27001, HIPAA, and GDPR.

2. Compliance with Industry Standards

- **AWS:** SOC 1/2/3, PCI DSS, FedRAMP certifications.
- **Azure:** SOC, PCI, various global government standards.

3. Customer Support and SLAs

- **AWS:** Developer, Business, and Enterprise support plans with 24/7 access and strong SLAs.
- **Azure:** Developer, Standard, and Professional Direct support plans, also 24/7 support and high SLAs.

4. Case Studies and Testimonials

- **AWS:** Success stories from startups, enterprises, and public sectors.
- **Azure:** Testimonials from major corporations, government bodies, and non-profits.

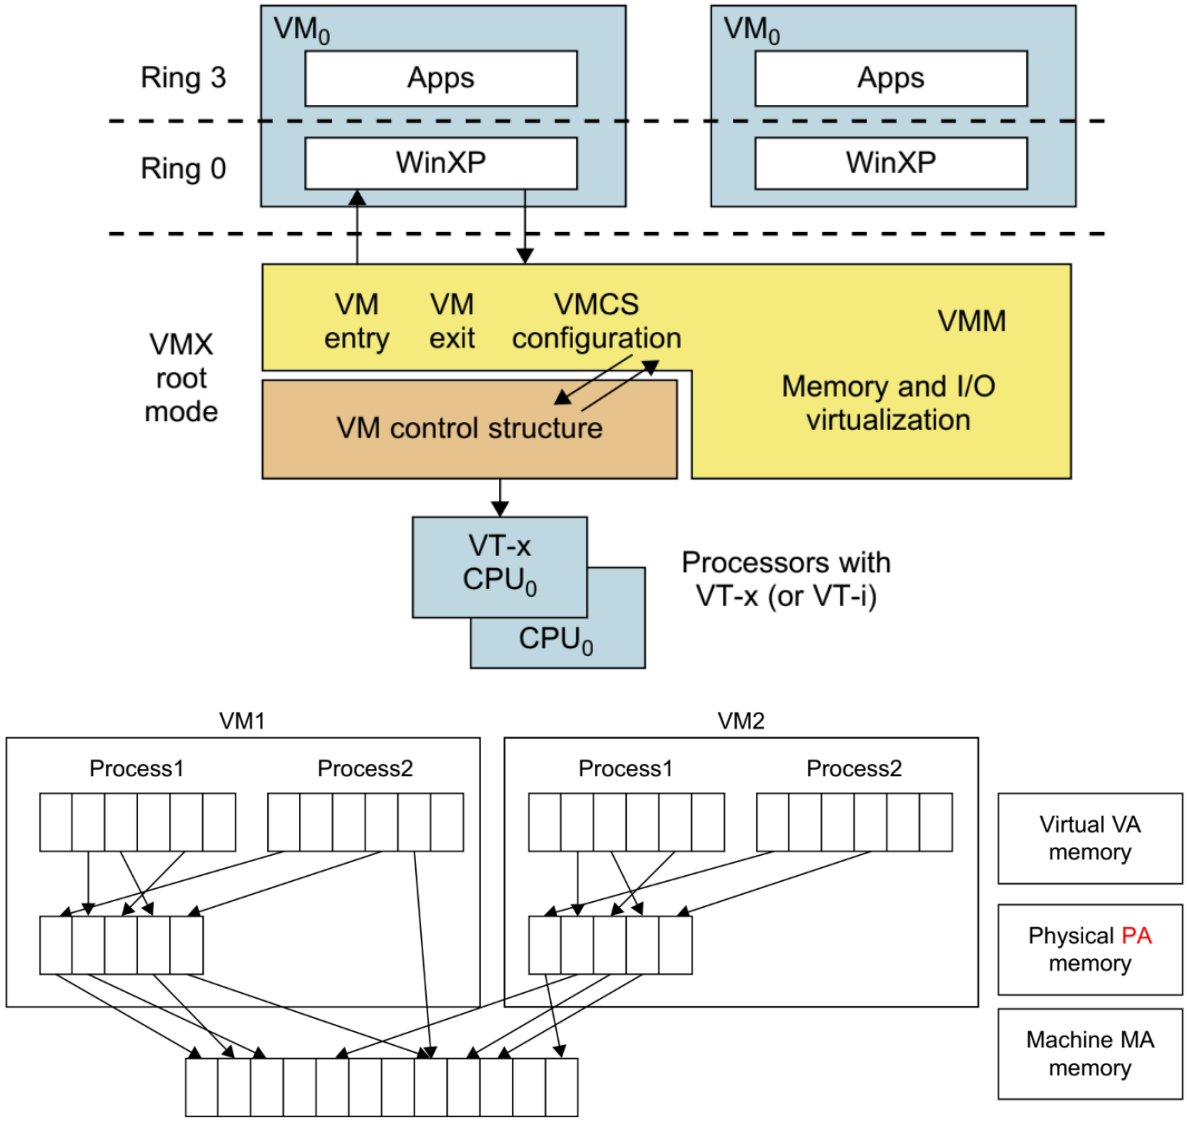

FIGURE 3.12

Two-level memory mapping procedure.

(*Courtesy of R. Rblig, et al. [68]*)

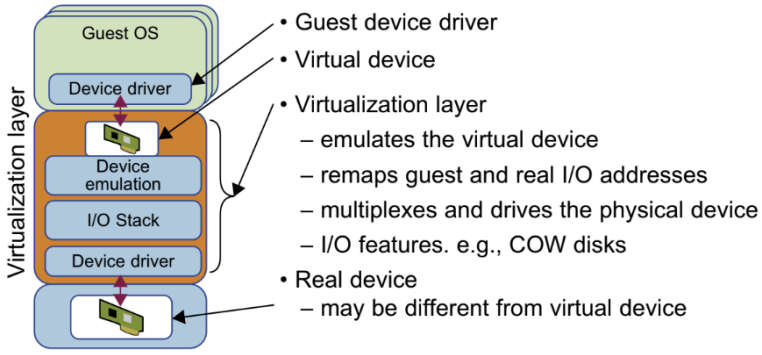

FIGURE 3.14

Device emulation for I/O virtualization implemented inside the middle layer that maps real I/O devices into the virtual devices for the guest device driver to use.

(*Courtesy of V. Chadha, et al. [10] and Y. Dong, et al. [15]*)