

Episode 3.01

Cloud Storage Types

Network Attached Storage (NAS)

- NAS servers connected to the network
- Provide access to storage using varying protocols
 - SMB
 - Network File System (NFS)
 - HTTP/HTTPS
 - File Transfer Protocol (FTP)
 - No encryption
 - SFTP
 - SSH tunnel
 - FTPS
 - SSL or TLS for encryption

Network Attached Storage (NAS)

- Generally provides file-based storage
- Cloud providers often provide this service

Direct Attached Storage (DAS)

- Storage attached directly to computers (virtual or physical)
 - Generally SATA/SCSI
 - Could be USB/FireWire/Thunderbolt
- Provided with SSDs (solid-state drives) and HDs (hard disks)
- Cloud providers offer both types

Storage Area Network (SAN)

- A network providing shared storage
- Common protocols include Fibre Channel and iSCSI
- Locally, required a hardware host-bus adapter (HBA)
 - If not using IP-SAN
- Cloud providers may offer this, but it is typically abstracted
- Provides block-level storage

Object-Based Storage

- Object-based storage allows the storage of data as objects
 - Usually no file hierarchy
 - Not block-level
- Examples of object storage
 - S3 Buckets (AWS)
 - Azure Blobs (Azure)
 - Buckets (GCP)

Quick Review

- Network Attached Storage (NAS) is available in some public cloud services and is commonly used in small private clouds
- Direct Attached Storage (DAS) is available for compute in all public cloud services and is used in private clouds as well
- Storage Area Network (SAN) is used in private clouds and is available in public clouds, but usually hidden from view

Episode 3.02

Provisioning Storage

Storage Provisioning Planning

- What data requires storage?
- What is the data format and structure?
- How much data?
- How sensitive is the data?
- What policies relate to storing the data?
- What regulations relate to storing the data?

Thick Storage Provisioning

- Allocates virtual storage at the time of the request
- May cost more
- Potentially provides improved performance
- Mostly used with virtual DAS

Thin Storage Provisioning

- Allocates storage as needed
- May cost less
- May not perform as well
- Mostly used for file/object storage

Encryption

- Data encryption can be at rest or in transit
- At-rest encryption is used on storage media
 - Protects against unauthorized local access
 - Protects against data access with media theft
- In-transit encryption is used during data transfer
 - Protects against eavesdropping
 - Protects against man-in-the-middle attacks

Tokenization

- Allows sensitive data to be stored in more secure solution
 - Token is stored in place of the data
 - Token is used to retrieve the actual data from secure storage
- Ex:
 - Personally Identifiable Information (PII)
 - Protected Health Information (PHI)

Tokenization

- Store a token in place of the social security number
 - Retrieve the social security number from the tokenization system when required
- Often used for payment processing

Quick Review

- Thick storage provisioning allocates all storage at the time of request
- Thin storage provisioning allocates storage as needed
- At rest encryption is used when data is stored
- Tokenization stores a token in place of the data and the actual data is stored in more secure storage

Episode 3.03

Storage Protection Capabilities

High Availability

- Redundancy
- Replication
- Cloning and mirroring
- Clustering
- Load Balancing
- Failover zones

Redundancy

- Redundant Array of Inexpensive/Independent Disks (RAID)
 - RAID redundancy levels/factor
 - RAID 0 – striping (no redundancy)
 - RAID 1 – mirroring
 - RAID 5 – striping with parity

Redundancy

- Nested RAID
 - Left number equals physical, right number equals logical
 - RAID 1+0 (RAID 10)
 - RAID 0+1 (RAID 01)

Replication

- Data replication can result in duplicated data stored in multiple locations
 - Regional/Same-Region Replication (SRR)
 - Replication within a region
 - Multiregional/Cross-Region Replication (CRR)
 - Replication across regions

Replication

- Synchronous replication occurs at the time of data modification
 - Immediate
- Asynchronous replication occurs later
 - Scheduled
 - Opportunistic

Cloning and Mirroring

- Cloning creates a point-in-time exact copy of a source
 - Less intensive ongoing
 - Cannot recover to the current time

Cloning and Mirroring

- Mirroring creates a continually updated copy of a source
 - More intensive ongoing
 - Can recover to the current time

Quick Review

- Redundancy is often a part of RAID storage solutions
- Synchronous replication occurs at the time of data modification
- Asynchronous replication occurs on a schedule or during time so low utilization
- Cloning creates a point-in-time copy of source storage

Episode 3.04

Storage Features

Compression

- Used to conserve storage space by rewriting data with compression algorithm
- Ex: .zip, .tar, .rar, etc

Compression Demo

- ASCII bits for the letter “e” are 01100101
 - Store the “e” as 0001 and save 4 bits
 - Do the same for the top 16 used characters and savings are 50 percent for those characters
- Real compression algorithms are much more complex

Compression

- Works best on non-binary, text-based data
 - .txt, .docx, .pptx
- Works least on compressed formats
 - .jpg, .mp4

Deduplication

- Process used to remove duplicate file copies
 - Evaluate files for duplication
 - Remove duplicate files
 - Create pointers to remaining file

Obfuscation

- Like lightweight encryption
- Modifying data according to an algorithm such that it does not appear meaningful on first viewing
- Prevents automatic searching of information

Obfuscation Demo

- Letter replacement – change the letter e to q and the letter q to s and the letter s to e
 - Obfuscated:
 - This is example text that is quite hard to read
 - Real:
 - This is example text that is quite hard to read

Obfuscation Demo

- Number replacement – change the number 7 to 4 and the number 4 to 9 and the number 9 to 7
 - Obfuscated:
 - 619-555-4971
 - Real:
 - 614-555-7491

IOPS

- Input/output operations per second (IOPS)
 - A measure of the speed of disk operations
- Read
- Write

Quick Review

- Compression allows for the storage of more data without requiring more physical drive space
- Deduplication ensures that data is stored only once within the entire cloud storage solution
- Obfuscation stores data in a format that limits searchability

Episode 3.05

Access Protocols

SMB

- Server Message Block (SMB)
 - Microsoft file access protocol
 - Latest version is SMB 3.1.1 (Windows 10 and Server 2016)
 - Supported by many cloud providers
- A client/server request/response protocol

NFS

- Network File System (NFS)
 - An IETF standard for network file access (RFC 7862)
 - Latest version is 4.2 (2016)
 - Popular in Unix/Linux environments
 - Supported by many cloud providers
- Based on Remote Procedure Calls (RPCs)

Application-Level Access Protocols

- Hypertext Transfer Protocol (HTTP)
 - Used to access files across Internet technologies (web-based)
 - HTTPS adds SSL/TLS encryption
- File Transfer Protocol (FTP)
 - Used to access files across Internet technologies
 - SFTP/FTPS add encryption
 - SFTP uses SSH with FTP
 - FTPS uses SSL/TLS with FTP

Private Cloud SAN Protocols

- Fibre Channel (FC)
- Fibre Channel over Ethernet (FCoE)
- Internet SCSI (iSCSI)
- Internet Fibre Channel Protocol (iFCP)

Zoning

- Zoning is used in SANs
- Provides access to storage by a specific device
- Ensures only that device (or authorized devices) can access the storage
- Zoning or zones may also refer to an area in which the cloud operates

Quick Review

- Server Message Block (SMB) is used mostly with Windows, but is available for other OSes as well
- Network File System (NFS) is supported by all major operating systems
- HTTP and FTP are both commonly used application layer protocols for file transfer
- Zoning is often used in SANs to provide access to specific devices

Episode 3.06

Storage Management Demo

Management Methods

- Web-based interface
 - Point-and-click
 - Provides visual cues and help along the way
- Command-line interface (CLI)
 - Faster
 - More complex
 - Can be automated easily

Storage Tiers and Classes

- Tiers or classes define the capacity of storage
 - Space available
 - Performance (IOPS)
 - Availability
- Each cloud vendor has unique terminology for tiers/classes

DEMO

- Storage classes in GCP
- Storage classes in Azure

Overcommitting

- Overcommitting is “getting more than you need”
 - Used to allow for growth
 - Used to allow for unexpected demand spikes
- Two forms of overcommitment
 - Space available
 - Performance

Quick Review

- Cloud storage solutions typically offer both web-based and CLI-based management
- Tiers or classes define the capacity of storage (performance, space, and availability)
- Overcommitment is used to provide for growth and unexpected spikes in demand

Episode 3.07

Storage Security Lab

Authentication

- User
 - Validating user identity
 - Password-based
 - Token-based
 - Certificate-based
- Host
 - Validating device identity
 - Password-based
 - Certificate-based
- Multi-factor authentication

Authorization

- Determining access/action rights
- Implemented through access control lists (ACLs)
- May grant or deny access
- Generally processed top-to-bottom or by priority number
- First match usually wins

Hands-On

- Show access control lists in AWS

Quick Review

- Users and hosts (devices/services) should be authenticated to access storage
- Authorization should be configured appropriately so entities have the access required and no more
- Most cloud providers use ACLs for authorization

Episode 3.08

Disaster Recovery Capabilities Demo

Recovery Metrics

- Recovery time objective (RTO)
 - Defined by SLA
 - How long can the network be down?
- Recovery point objective (RPO)
 - What point in the past must I recover to?
- SLAs for disaster recovery
 - Specify both RPO and RTO

Quick Review

- The Recovery Point Objective (RPO) is the point-in-time to which you wish to recover
- The Recovery Time Objective (RTO) is the length of time required to recovery
- All major cloud providers offer disaster recovery solutions

Episode 3.09

Disaster Recover Considerations

Additional Considerations

- Corporate guidelines
 - Might have policies in place for traditional network disaster recovery
 - How to implement in the cloud?
- Bandwidth or ISP limitations
 - What is your Internet connection speed?
 - Consider alternate route to the Internet
- Site mirroring and replication

Additional Considerations

- Site mirroring and replication
 - Have a site ready to take over if primary site fails
- File transfer and archiving
 - Make sure to keep offline backups

Additional Considerations

- Third-party sites
 - Other sites you integrate with
 - Have disaster recovery plan for third-party sites
- Techniques and processes used in recovery
 - Documented disaster recovery plan

Quick Review

- Bandwidth limitations should be considered related to disaster recovery as the impact recovery times
- Data can be archived to third-party sites and the integration of such archival should be considered
- The techniques and processes used in recovery should be documented with step-by-step procedures

Episode 3.10

Business Continuity Plan Lab

Business Continuity Plan Components

- Alternate sites
- Continuity of operations
 - Keep the business going
- Connectivity
 - Stay connected to the Internet
- Edge sites
 - Where the action happens
 - IOT: location where the IOT sensors/actuators are located

Business Continuity Plan Components

- Equipment
 - Duplicate copies of equipment
- Availability
- Partners and third parties
 - Other partners
 - Maintain communication
- SLAs for BCP and HA

Hands-On

- Creating a second VPC in AWS

Quick Review

- Alternate sites are commonly used in BCPs and, in the cloud, alternate sites can also be used
- Internet connectivity redundancy is a key consideration in cloud computing
- VPCs can be duplicated or mirrored in AWS for business continuity and other cloud providers offer similar functionality

Episode 3.11

Additional Storage Considerations

Thin/Thick Provisioning

- Thin provisioning allocates resources as needed
- Thick provisioning allocates all resources at the time of creation

User Quotas


- Implements constraints on space consumption
- Must be carefully implemented and monitored
- Can prevent storage of unauthorized data

Hyperconverged

- An integrated systems including storage, compute, and networking in a software-defined platform
 - Also known as Hyperconverged infrastructure (HCI)
- Focus of HCI includes:
 - Software defined storage
 - Virtualization
 - Integrated management

Software-Defined Storage



- SDS abstracts the management and provisioning of storage resources
 - SDS characteristics include:
 - Abstraction
 - Automation
 - Scalability
 - Policy-based management
 - Hardware agnostic
- 

NVME-oF

- Non-volatile memory express over fabrics
 - A protocol that extends NVME over a network
 - Improved performance over SATA or SAS

Tiers

- Flash
- Hybrid
- Spinning disks
- Long-term storage

File Integrity Monitoring (FIM)

- A security process for monitoring files and directories/folders to detect unauthorized access or changes
- FIM components include:
 - Baseline
 - Monitoring
 - Alerts and Notifications
 - Reporting and Analysis

Episode 3.12

RAID Storage