Cloud Security Operations

Build or Buy

- Building a DC is an expensive proposition
- <u>But advantages are:</u>
  - Location of choice
  - Service offering
  - Design elements choice

Considerations when deciding DC

- Cost of the facility vs cost of the lease over time
- Efficiency in terms of power, utility, HVAC, staffing
- Regulations or control specifications
- Velocity of growth

Location

- First aspect to be decided when choosing a DC
  - Availability of inexpensive power
  - High-speed network connectivity
  - Natural disaster zone
  - Proximity to other DCs
  - Temperature cooling challenge
  - Vendor support
  - Legal and regulatory mandates

Power Resilience

- Ensuring power requirements are adequate and available continuously is one of the key considerations in DC design
  - Ability to acquire power from multiple grids
  - No single point of failures
  - Underlying infrastructure ready to support the load of the facility India
- Brief outages should be handled by the UPS
- Generators should handle longer power outages

Communication Resilience

- Identify the current and future bandwidth needs
- Ensure network connectivity from more than one ISP (***Multi vendor Pathway communication***)
- Ensure different ISPs do not share the same upstream dependencies
- Assess connectivity paths for environmental dangers and single point of failures
- Design internal systems and networks to support redundant connectivity and resilience

Physical Security

- Vehicular traffic design to the facility
- Guest / Visitor access control
- CCTV monitoring
- Protected placement of hazardous resources away from human personnel
- Interior physical access controls
- Fire Detection and Suppression Systems
- Security controls redundancy during power interruptions

Data Centre Tiers

Tier	Description
Tier 1 DC	• Basic infrastructure required to run an IT operation • No redundancy, downtime in the event of unplanned maintenance or interruption • Uptime 99.671% **Requirements:** • UPS for line conditioning and backup • An area to house IT systems • Dedicated cooling systems • Power generator for extended power outage • Redundancy for chillers, pumps, UPS and generators
Tier 2 DC	Provides more redundancy that Tier 1 DC Uptime 99.741% **Requirements:** • UPS and generators • Chillers and cooling units • Pumps • Fuel tanks • Intended to ensure critical applications are not interrupted during planned Maintenance

Data Centre Tiers

Tier	Description
Tier 3 DC	Design is known as "Concurrently maintainable Infrastructure" Provides N+1 redundancy Can have planned maintenance activities without disruptions Uptime 99.982% **Requirements:** • Tier 2 DC requirements + • Multiple distribution paths where only a sole path is needed to serve critical operations at any given point of time
Tier 4 DC	Highest Level DC proposed by Uptime Institute Provides 2N+1 redundancy Can withstand planned or unplanned interruptions Uptime 99.995% **Requirements:** • Independent and physically isolated systems • Provides resiliency at both the component and distribution path levels • Designed around Fault tolerance for components

Hardware Specific Security

Trusted Platform Module

- Microchip installed on the motherboard that is dedicated to carrying out security functions on the system. It is only a physical component and cannot be added later
- It is referred as Cryptographic coprocessor
- It is used to form roots of trust
- Two major functions of TPM are
  - **Binding the hard disk drive**
    - Content of the Hard disk drive is encrypted, and the decryption key is stored away in the TPM chip
    - If the TPM chip fails, the encrypted content in the HDD will be rendered useless
  - **Sealing a system configuration**
    - TPM generates hash values based on the systems configuration and stores them in TPM chips
    - Only after TPM verifies the integrity of the system's configuration will it allow activation of the system

Trusted Platform Module (TPM)

- Services provided by TPM
  - Random number generators
  - Asymmetric key generation
  - Hash generators
  - Used for storage of highly secure limited data (cryptographic keys)
  - Used to form roots of trust

Hardware Security Module (HSM)

- Dedicated hardware designed to support and perform cryptographic functions
- Functions performed by HSM
  - Secure storage of cryptographic keys
  - Encryption and decryption function
  - Cryptographic based authentication
  - Generate data needed for cryptographic functions

KVM Security

- Key attributes for a secure Keyboard Video Mouse (KVM) switch are:
  - Isolated data ports
  - Tamper-resistant/evident design
  - Secure storage
  - Secure firmware
  - Physical disconnects - button on the front of KVM to allow the user to switch between connected systems
  - USB Port and device restriction

TLS Security

- Transport Layer Security protocol is a set of cryptographic protocols that provide encryption for data in transit
- TLS can be used both for encryption and authentication
- TLS Communication steps:
  - Client initiates a request – "ClientHello". Informs the list of cipher suites and TLS versions it supports
  - Server chooses the highest TLS Version from the list and communicates back along with its certificate
  - Client and server negotiate a session key
  - Session key is then used to encrypt all data that is shared

DHCP

- Runs over UDP
- Utilizing ports:
  - 67 – connections to server
  - 68 – connections to client
- Extension of BOOTP (protocol used for simple interaction)
- All interactions are initiated by a client
- Server only replies
- Uses client–server model

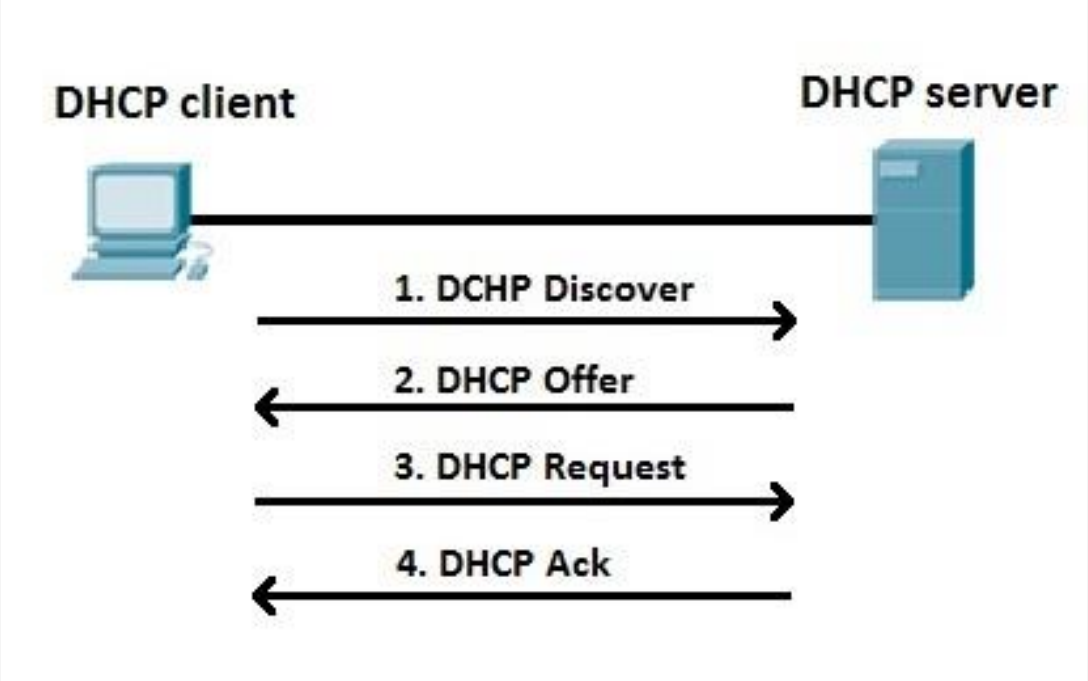

DNS

- The mechanism by which Internet software translates names to attributes such as addresses
- A globally distributed, scalable, reliable database
- Comprised of three components
    - A “name space”
    - Servers making that name space available
    - Resolvers (clients) which query the servers about the name space
- The name space is the structure of the DNS database
    - An inverted tree with the root node at the top
- Each node has a label
    - The root node has a null label, written as “”
- A domain name is the sequence of labels from a node to the root, separated by dots (“.”), read left to right
    - The name space has a maximum depth of 127 levels
    - Domain names are limited to 255 characters in length
- A node’s domain name identifies its position in the name space

DNS attacks

- **DNS cache poisoning**
  - These attacks capture and divert queries to another website unknown to users
- **Denial of service (DoS)**
  - Attempts to make a given service impossible or very hard to access.
- **Distributed denial of service (DDoS)**
  - An elaborate form of DoS that involve thousands of computers generally as part of a botnet or robot network
- **Reflected attacks**
  - Send thousands of requests with the victim's name as the source address. When recipients answer, all replies converge on the official sender, whose infrastructures are then affected

DNS attacks

- **Reflective amplification DoS:**
    - The same technique as reflected attacks is used, except that the difference in weight between the answer and question amplifies the extent of the attack. A variant can exploit the protective measures in place, which need time to decode the long replies; this may slow down query resolution
- **Cybersquatting**
    - Involves registering a domain name with the deliberate intent of undermining and profiting from a third party's rights or in some way harming that third party.
- **"Name-jacking" or theft**
    - Appropriating the domain name (updating the holder's field and/or contacts) or taking control by technical means to divert traffic, such as by modifying the name servers hosting the site

DNSSEC

- It is set of extensions focused on providing Integrity of DNS
- Provides cryptographic authentication of DNS data using digital signatures
- It provides proof of origin and makes cache poisoning / spoofing attacks more difficult
- It does not provide confidentiality

Software-Defined Perimeter (SDP)

- Operating Model
- SDP controllers are created and are connected to an authentication service to enforce access control
- “Accepting” SDP hosts authenticate to the SDP controllers. These hosts do not accept any new connection by default
- “Initiating” SDP hosts connect to SDP controllers for authentication and can request access to resources of an accepting host.
- The SDP controllers makes authorization decisions and provides details to both the ”accepting” and “initiating” SDP hosts
- A mutual VPN is established between the “Initiating” and “Accepting” host and the user is able to interact with the resource.

Virtualization Concepts

Virtualization Concepts

- **Distributed Resource Scheduling (DRS)**
  - A utility that balances computing workloads with available resources in a virtualized environment
  - If the workload on one or more virtual machines drastically changes, DRS redistributes the virtual machines among the physical servers
  - If the overall workload decreases, some of the physical servers can be temporarily powered-down and the workload consolidated
  - DRS intelligently allocates resources and can be configured to automatically take care of workload Migration
  - DRS enables optimal workload distribution on virtual machines based on business needs and changing priorities.
  - DRS migrates VMs based on the availability and utilization of CPU and memory resources

Virtualization Concepts

- **Dynamic Optimization**
  - Dynamic Optimization (DO) is a feature that initiates live migration of VMs that are on a cluster, to improve load balancing among cluster nodes and to correct any placement constraint violations
  - Dynamic Optimization settings can be configured for the CPU, memory, disk I/O, and network I/O
  - It relies on real-time data and defined goals to determine configuration and resource changes

Virtualization Concepts

- **Maintenance Mode**
  - When DRS maintenance mode is invoked, DRS is used to evacuate all virtual machines from one host to another, without incurring any downtime.
  - This is especially useful for performing maintenance, updating an ESXi host, installing additional memory, or upgrading firmware, and so on
  - In order to use maintenance mode in the Virtualized architecture, a cluster needs to be created and DRS must be enabled

Virtualization Concepts

- **Containerization**
  - Containerization places an application, along with all the libraries and components it needs
  - It does not virtualize an entire operating system, just the application environment alone
  - Keeping containers secures, requires the images to be secured
  - Security of container registers, signing containers, managing secrets and validating signatures are important

Virtualization Concepts

- **Ephemeral Computing**
  - Leverages the ability to quickly standup virtual machines
  - Allows for efficient horizontal scaling

Hypervisor Security

- **Common Security mandates are:**
  - Restricting access to superuser accounts
  - Requiring MFA
  - Using logging and alerting
  - Limit access to authorized users only
  - Encrypting VMs
  - Using secure boot for the underlying hardware
  - Performing regular audits of configurations and systems

Virtualization Management tool best practices

- Vendor recommended hardening installations to be followed
- Redundancy ~ HA and duplicate architecture
- Scheduled downtime and maintenance
- Isolated network and robust access control
- Configuration and Change management
- Logging and Monitoring

Storage Operations

Storage Clusters

- Storage devices are grouped in clusters to provide for
  - Performance
  - Flexibility and
  - Reliability
- Two Cluster Architectures:
  - Tightly coupled
  - Loosely coupled

Storage cluster Architecture

- **Tightly Coupled Architecture**
  - All storage devices are directly connected to the physical device backplane
  - Each component of the cluster is aware of the other devices
  - All components subscribe to the same policies and rulesets
  - All devices must need to be from the same vendor
  - They are confined to more restrictive design parameters
  - This architecture enhances performance due to division of data into deterministic blocks

Storage cluster Architecture

- **Loosely Coupled Architecture**
  - Each component of the cluster is independent of the others
  - Nodes can be added using any off-the-shelf parts
  - Uses file-level storage system
  - New nodes can be added as needed
  - Nodes are only logically connected and do not share physical framework
  - Performance does not necessarily scaleup
  - This architecture provides flexibility

Data Resiliency

- Two ways of providing Data protection in Cloud storage clusters:
  - RAID
  - Data Dispersion

Data Resiliency – RAID

- Redundant Array of Independent Disks is a technology used for redundancy and/or performance improvement

RAID	Activity	Name
0	Data **stripped** over serval drives, no redundancy or parity. If one volume fails entire volume may become unusable, primarily used for **performance improvement**	**Stripping**
1	Data is **written to two drives** at the same time. If one fails, the other drive has the same data	**Mirroring**
2	Data **stripping** over all drives at **the bit level,** Parity data is created with Hamming code. **Not** used **in Production** today	**Hamming code Parity**
3	Data stripping over all drives, **parity code in in one drive**. If one drive fails it can be reconstructed using parity code	**Byte-level parity**
4	Same as RAID 3, parity is created at block level instead of byte level	**Block-level parity**
5	Data is **written in disk sector** units to all drives. **Parity is also written to all drives**. Ensure no single point of failure	**Interleave parity**
6	Similar to RAID 5 with **added Fault tolerance**, second set of parity added to all drives	**Double parity**
10	Data is **mirrored and stripped** simultaneously across several drives and can support multiple drive failure	**Stripping and Mirroring**

Data Resiliency – Data Dispersion

- Concept of separating data into unrecognizable "slices" that are distributed via network connections to storage nodes locally or across the world.
- Transforms data into slices by using equations such that a subset of the slices can be used to re-create the original data
- Dispersed storage systems are well-suited for storing unstructured data like digital media of all types
- Dispersal is not optimized for transaction-oriented primary storage for databases and similar high IOP workloads
- Uses **Erasure Codes** to create redundancy for transferring and storing data
    - An Erasure Code is a Forward Error Correction (FEC) code that transforms a message of k symbols into a longer message with n symbols such that the original message can be recovered from a subset of the n symbols (k symbols)

IT Service Management Functions

IT Service Management Functions

- There are 3 Important IT Service Management Functions
  - **Service-level Management:**
    - Ensures the IT organization is fulfilling the commitments to internal and external customers
  - **Availability Management:**
    - Improves the resiliency of the IT Services to meet the customer needs
  - **Capacity Management:**
    - Ensures IT resources are sufficient to meet the current and future business needs.

Physical and Environmental Protection

- Technical Committee 9.9 of the American Society for Heating, Refrigiration and Air-conditioning Engineers (ASHRAE) has created target metrics for performance monitoring of Environmental projection in DC
- Recommendations are:
  - Temperature: 18 C to 27 C
  - Humidity: dew point of -9C to 15C, relative humidity at 60%
- High humidity will cause corrosion
- Low humidity will cause static electricity

Device Maintenance Concepts

- When a device is put into maintenance mode, the follows tasks should be completed:
  - All operational instances are removed from the system before it enters into Maintenance mode
  - Prevent all new logins
  - Begin enhanced logging

Updates

- **Due care** – adhering to vendor specifications for device updates
- **Due diligence** – adherence to documented vendor instructions

Updates Process

- Move the machines to maintenance mode
- Apply the update and annotate the asset inventory
- Verify the update coverage in all machines
- Validate the intended modifications post update
- Return to normal operations

Change and Configuration Management

- **Baselines:**
  - Configuration and Change management begins with defining the baseline, which is a way of setting the desired standard state
  - Baseline is a general-purpose map of the network and systems, based on the required functionality and security
  - Security controls should be included in the baseline
  - While creating a baseline all stakeholders should be consulted
  - Baseline should be an excellent reflection of the risk appetite of the organization
  - Baseline should suit the largest population of systems in the organization
    - There can also be multiple baselines customized to specific groups or departments
  - If there are multiple exceptions to a defined baseline, the baseline parameters should be changed

Change and Configuration Management Policy

- The CM Policy should include:
  - Composition of the CM Board (CMB)
  - The process in detail
  - Exception management
  - Assignment of CM tasks
  - Procedure for addressing deviations, upon detection
  - Enforcement and responsibilities
- CMB should be composed of all stakeholders in the organization
- CMB shall meet often enough to ensure there are no delays

Release and Deployment Management

- Process responsible for arranging all elements to successfully, repeatably and verifiably deploy new software versions
- It includes Planning, scheduling and deploying new software, it encompasses all environments ~ Dev, QA/Testing and staging.
- Once the software moves into Production, it enters active maintenance phase

Business Continuity and DR

BC / DR

- Business continuity is focused on maintaining critical business operations during a disaster
- Disaster recovery is focused on the resumption of operations after an interruption
- Prioritizing health and safety is paramount in any BC/DR planning and efforts

BC / DR Plan

- The BC / DR plan should include
  - Critical Asset Inventory
  - Disaster Criteria
  - Disaster declaration process
  - Essential Points of Contact
  - Detailed Actions, Tasks and Activities

BC / DR Toolkit

- Toolkit that holds all necessary documentation and tools to conduct a BC/DR response action
- It should be secure, durable and compact
- It can be virtual or physical container
- The kit should have duplicate in at least one additional location

BC / DR Toolkit

- The kit should contain the following:
  - Current copy of the BC / DR plan
  - Emergency and backup communication equipment
  - Network and Infra diagrams and architecture
  - Copies of all software
  - Emergency contact information

BC / DR Terminology

- **RTO – Recovery Time Objective**
  - The maximum time within which a business process must be restored to an acceptable service level after a disaster
  - RTO value should be lesser than MTD
  - RTO deals with getting the infrastructure and systems back up and running
- **MTD represents the time after which the business cannot recover**
- **Work Recovery Time (WRT)**
  - Remainder of the overall MTD value after RTO has passed
  - WRT deals with restoring data, testing process, and then making the production process live

BC / DR Terminology

- **RPO – Recovery Point Objective**
  - It is the acceptable amount of data loss measured in time.
  - Represents the earliest point in time to which data must be recovered
  - The higher the value of data, the lower the RPO value
- The actual RTO, MTD, RPO values are derived from the Business impact assessment (BIA)

- **RSL – Recovery Service Level**
  - The proportion of the service, expressed in %, that is necessary for continued operations during a disaster

BCP Testing

- BCP maintenance should be incorporated into change management procedure
- Tests and DR drills should be conducted atleast once a year
- The first exercise should not include all employees rather a small representative sample of the organization
- People conducting the drills should expect to encounter problems and mistakes

BCP Testing

Checklist Test

- Copies of BCP/DR plan distributed to the different departments for review
- This ensures nothing is taken for granted or omitted
- Planning team integrates all changes to the master plan
- It is also called desktop or table top test

Structured walk-through

- Representatives from each department come together and go over the plan
- The group reviews the objective, scope, assumptions of the plan
- The group walks-through different scenarios of the plan from beginning to end to make sure nothing is left out

Simulation Test

- This test takes a lot of planning and resources
- All employees participating in operational and support functions come together to practice a specific scenario
- It raises the awareness level of the people involved
- The drill shall include only those materials that will be available in an actual disaster.
- The test continues upto the point where physical migration to new facility gets initiated

Parallel Test

- Some systems are moved to alternate site and processing takes place
- The results are compared with the regular processing done at original site
- Ensures specific systems can function adequately at alternate site during disaster

Full-Interruption Test

- Most intrusive to regular operations
- The original site is shut down and processing takes place at the alternate site
- Recovery team fulfills its obligations in preparing the systems and environments for the alternate site
- All processing is done at alternate site
- It should be performed only after all other tests are completed satisfactorily
- Senior mgmt. approval is needed before performing this test

All the best