Software-Defined Storage Concepts

From NDG In partnership with VMware IT Academy www.vmware.com/go/academy





This content will cover



- Learn key concepts related to storage in the data center
- Understand software-defined storage concepts
- Describe the architecture, technical characteristics, and benefits of vSAN
- Compare the functionalities and performance of vSAN with other traditional storage options
- Be exposed to basic information about the Hyper-Converged Infrastructure



The Data Center

- Data centers exists for the sole purpose of storing data
- Modern data centers consist of multiple layers
 - Management Layer
 - Virtual Layer
 - Physical Layer

vmware[®]

- Network Layer
- Compute Layer
- Storage Layer



The Data Center

Data Centers must provide Availability and Redundancy.

- Availability The expectation that the storage is online and running, making it accessible.
- Redundancy The duplication of critical components in a system to provide a back-up in the event of a failure in the original location.





- Abstraction In a complex system or piece of software, focusing on the most relevant details and hiding what can be ignored.
- Array Data storage that is made up of multiple storage devices and cache memory.
- Block Storage (or Block-Level Storage) Data is saved in fixed-sized volumes called 'blocks'; each block is treated as an individual storage device, has a unique identifier, and has its own file system.
- Deploy To install, test, and run hardware or software in a live environment.







- File Storage (or File-Level Storage) Data is saved in files and folders in a hierarchical system of directories and sub-directories; in order to be accessed, the storage drives must be configured with the Network File System (NFS) for Unix/Linux systems or Server Message Block (SMB) for Microsoft Windows systems.
- Logical Virtual; not physical.
- Mirror To make an exact copy of data from one storage device drive to another storage device in real-time. This serves to prevent data loss in the event of a hardware failure. This is also known as RAID 1.



- Object with vSAN, an object is a virtual machine disk (VMDK) file, a snapshot (a copy of a VMDK taken at a specific point in time), or the virtual machine home folder.
- Object Storage (or Object-Based Storage) Data is bundled together with its metadata (information such as date created, size, and author) and a unique identifier.
- Policy A set of rules about the storage requirements of virtual machines and the applications they run.





- RAID (Redundant Array of Independent Disks) Storage that is made up of multiple hard drives. The same data is stored across different disks. The method by which this data is stored is classified by RAID levels.
 - RAID 0- Striping with no fault tolerance or redundancy. Designed for speed.
 - RAID 1- Mirroring with no striping. Designed for reliability.
 - There are other RAID formats that utilize both striping and mirroring.
- Stripe To divide a piece of data into equally-sized units which are then spread across multiple storage devices; no copies of the data are made. This is often referred to as RAID 0.





Local Disk Storage



- Has a 1:1 relationship between storage device and personal device or server
- 3 types of Local Disk Storage:
 - Hard Disk Drives (HDDs) Utilizes a platter and magnets to store 1s by magnetizing and 0s by demagnetizing. Spinning platter can cause data loss in the event of abrupt shutdown. These are known for their large amount of storage at an affordable price.
 - Solid State Drives (SSDs) Utilizes transistors to store data in electrical charges. Transistors that conduct current have a value of 1 and a chain that doesn't conduct current has a value of 0.
 Faster than HDDs, but more expensive.
 - Optical Disk Drives (ODDs) Utilizes a laser to read 0s and 1s from a spinning disk. CDs, DVDs, Blu-Ray discs. Inexpensive, and highly portable.



Local Disk Storage Protocols

- A protocol is the language that local disk storage uses to communicate with a device.
- Several commonly seen protocols are:
 - Small Computer System Interface (SCSI)
 - Serial Attached SCSI (SAS)
 - Serial Advanced Technology Attachment (SATA)
 - Universal Serial Bus (USB)
 - Fibre Channel (FC)
- Servers may need an adapter to communicate with local storage. This adapter is known as a Host Bus Adapter (HBA).







......................





Network-Attached Storage

- NAS is connected to a Local Area Network (LAN).
- Each authorized user on the network can access the data on the NAS.
- Can be considered a "Personal Cloud" or "Private Cloud", where the storage is located on site rather than remotely.
- NAS will have its own IP address.
- Capable of RAID configurations.

vmware[®]

- Uses TCP/IP to send information through the network.
- File System Protocols often used are NFS for Linux/Unix, SMB for Windows, and Apple Filing Protocol (AFP) for Apple devices.



Storage Area Network



- Runs alongside a LAN and can serve several different physical locations.
- Gives access to block level storage; is capable of file level storage can be obtained via the servers' operating systems.
- Does not need to be in the same physical location as the servers. Can be off-site storage.
- Allows for easy and immediate scalability.
- Prevents network bottlenecking by running alongside LANs.
- Frees up computing resources on servers.





Storage Area Network

- A SAN consists of 3 layers:
 - Host Layer- Made up of servers which grant access to files to users.
 - Fabric Layer- Contains the physical equipment and cabling of the SAN.
 - Storage Layer- Consists of the actual physical storage devices.
- Several protocols a SAN may utilize are:
 - Fibre Channel (FC) High speed.
 - Fibre Channel over Ethernet (FCoE) Cheaper alternative to FC.
 - Internet Small Computer Systems Interface (iSCSI) Uses IP, good for SMB.
 - ATA over Ethernet (AoE) Simplified protocol, good for economical networks.
 - InfiniBand (IB)- Very high speed, often utilized between super-computers.



Virtualized Storage

vmware[®]

- Virtualized storage gives VMs the storage they require to host their operating system and applications.
- The Virtual Disks of a VM are stored on a 'datastore'.
- A partition on the physical drive is created, forming a space called a Logical Unit Number (LUN).
- A datastore is made up of one or more volumes, which are made up of one or more LUNs.
- File Systems give names and metadata to stored data.
- VMware utilizes either VMFS or NFS file systems.





Virtual Machine File System

- Developed for use by VMware.
- Allows multiple servers to read/write at the same time.
- As a Clustered File System (CFS), it is simultaneously mounted on multiple servers.
- Allows multiple VMs to share a single file system.
- Can be linked to a single SAN LUN or span multiple SAN LUNs.
- Each VM is encapsulated into a small set of files in a single directory. This makes mirroring and recovery much simpler in the event of a disaster.





Network File System

_ NFS

vmware[®]

- IP-based file sharing protocol utilized by NAS systems.
- Uses file-level access, the storage device is controlled by the NAS.
- NFS has three components: client, server, and protocol.
 - Files appear to be local to the client machine.
- While typically used on Unix/Linux and MacOS, it is OS Independent.
- Useful for freeing up resources and transferring files between different operating systems.
- VMware vSphere supports NFS version 3 over TCP/IP.



Virtual Volumes



- Virtual Volumes (VVOL)s are an industry wide standard focused on increasing the flexibility of virtual storage.
- VVOLs allow a focus on VMs for storage management, rather than being limited to LUNs. They encapsulate virtual disks and other virtual machine files on a physical storage device without using a file system.
- A VVOL is created every time a virtual machine is created, cloned, or a snapshot is made of it.
- Unlike LUNs, VVOLs can have their size and number adjusted.
- ESXi hypervisors must use protocol endpoints to access VVOLs.





VVOLS

- VVOLs can broadly be classified into five types:
 - Config-VVOLs- Contain VMX (primary configuration file), NVRAM (file that contains the state of the virtual machine's BIOS), and log files.
 - Data-VVOLs- Contain data related to VMDKs (virtual disk drives that store the contents of the VM's storage device) and delta files (such as snapshots).
 - Mem-VVOLs- Contain data related to memory snapshots.
 - Swap-VVOLs- Contain information about swap files.
 - Other-VVOLs- A generic type of VVOL containing files relating to particular vSphere features.





Virtualized Storage Area Networks

• vSAN is included in the ESXi hypervisor.

vmware[®]

- Virtualizes the physical storage resources of ESXi hosts and pools them into a vSAN datastore.
- vSAN datastores are accessible to all hosts in the vSAN cluster.
- Virtual routing and switching reduces need for physical networking equipment, such as cabling.
- Requires at least one flash-based storage device per disk group.





Software-Defined Storage

- "Virtualized storage with a storage management interface."
- Storage Virtualization is only a piece of the SDS stack.
- SNIA states that SDS must include:
 - Automation
 - Standard Interfaces
 - Virtualized Data Path
 - Scalability
 - Transparency



• SDS allows a greatly increased amount of flexibility in storage options.





Advantages of Software-Defined Storage

- Separation of hardware and software eliminates the necessity to use proprietary products. This allows most hardware to be used and allows the software to be upgraded independently.
- Storage can be allocated specifically to certain amounts per application, increasing resource efficiency.
- SDS is capable of prioritizing where information is stored and moving it to optimal media depending on usage.
- SDS greatly reduces the complexity of storage solutions, as well as the required overhead for storage administrators.



mware[®]

Types of Software-Defined Storage

Several types of SDS exist, these may include, but are not limited to:

- Hypervisor-based A storage hypervisor that helps manage multiple pools of storage. Ex: VMware's vSAN technology.
- Hyper-Converged Infrastructure Package This option packages compute, storage, networking, and virtualization in the same hardware.

mware[®]

IT ACADEMY

• Container-based - Software-Defined Storage that is specifically built for container environments such as Docker.





Software-Defined Storage Model

- VMware's SDS model prioritizes the application, allowing the storage to fit the needs of the application, rather than the application conforming to the limitations of storage.
- This means that in addition to simplifying the storage process, reducing costs, and increasing flexibility, the SDS model also removes waste; applications get exactly what features from storage that they require in the lines of performance, capacity, protection, and so on.





vmware[®]

Virtual Data Plane



- In virtualization, "plane" refers to a specific level or layer in an organization's IT structure.
- The Virtual Data Plane stores data for later retrieval and applies data services before passing it to the Policy-Driven Control Plane.
- The Virtual Data Plane also applies the policy selected to the objects in the Virtual Datastore.
- The Virtual Data Plane is delivered through vSAN x86 for hyper-converged storage, or through vSphere virtual volumes for external storage SAN and NAS.



Policy-driven Control Plane



- Serves as the bridge between applications and infrastructure, creating standardized management and automation.
- Ensures that legal compliance is met through the life-cycle of the application
- Grants "perspectives" which are visible to different roles in a typical IT environment (application administrator, storage administrator, data protection administrator, business operations, etc.). This is more efficient than each role relying on its on set of proprietary management tools.
- The control plane is programmed through public APIs.



Storage Policy-based Management



- VMware vSphere SPBM automates the provisioning and monitoring of services based on the policies set to them.
- Can allocate storage based on need, re-optimizing as need changes.
- The default storage policy is compatible with any vSAN datastore in the vCenter server.
- Policies can be applied to VMs or individual disks.
- It is highly recommended that you do NOT edit the settings of the default storage policy.
- Instead, clone the default storage policy and use it as a template.



Virtual Data Services

- Data services applied by the Virtual Data Plane may include:
 - Compression- Reducing the size of files to conserve space.
 - Replication- The continuous copying of data as a backup safety feature.
 - Caching- Temporarily storing frequently used information for quick access.
 - Snapshots- Images of a system taken to conserve a specific point in time.
 - Deduplication- Removing redundant information to conserve space.
 - Encryption- Encoding data so that it remains secure to only those authorized.
- Data services are applied on a per-VM basis, allowing you to customize and change services as need arises.
- The Control plane manages resource allocation for storage services.



Hyper-Converged Infrastructure

- Compute, Storage, Networking, and Management are integrated to run as software on the hypervisor.
- Run on non-proprietary servers with common management tools.
- The common way to achieve this is to run third-party storage software in the VM that sits on top of the hypervisor. This comes at a cost of resources and performance.
- VMware implements storage software into the hypervisor itself, causing convergence inside the hypervisor rather than on top of it. This increases performance and resource efficiency.





Benefits of an HCI Model

- Virtualization with a hypervisor, combined with hyper-converged storage, a single set of management tools, and a wide compatibility with various hardware.
 Hyper-Convergence offers several benefits:
 - Fewer resources will be consumed, particularly when using storage software converged inside the hypervisor.
 - Increased performance and reduced latency.
 - Simpler hardware setup with single management interface.
 - Reduced cost via increased storage efficiency and fewer hardware purchases.
 - Improved security via software-based security, often built into modern HCI.
 - Industry-leading hypervisor makes 99.999% availability possible.
 - Improved scalability and flexibility.





Software Stack Components

- VMware's HCI is made up of three industry-leading solutions:
- Virtualization Virtualization Virtualization Virtualization Virtualization Virtualization Virtualization Virtualization
- VMware vCenter Server- A unified server management software that provides a centralized platform for controlling your VMware vSphere environments.
- VMware vSphere- The world's leading server virtualization software and the heart of a modern software-defined data center (SDDC). This software helps users run, manage, connect and secure their applications in a common operating environment across clouds. Advanced security features integrated into the hypervisor and powered by machine learning provide better protection against and response times for security incidents.
- VMware vSAN- The only vSphere embedded, flash-optimized storage for virtual machines and containers. It joins all storage devices in a vSphere cluster into a shared data pool.
 vSAN-powered HCI lowers storage costs by approximately 40% or more compared to traditional server and storage architectures.



Storage Policies Management

vmware[®]

- Virtual Machine Storage Policies are sets of rules that define how the vSAN stores files for the VM.
- Storage policies contain data placement rules and data service rules.
- Storage policies can be applied during any phase of a VM's cycle.
- When a VM is cloned or migrated, a new storage policy can be applied, or it can carry over the original.
- During application, the SPBM will list which datastores are compatible with the current policy.



Application Programming Interfaces

• APIs allow applications to speak to one another.

mware[®]

- APIs can serve as software intermediaries between the user interface and the server database or website.
- APIs are only accessible by developers; they are not user-facing.
- APIs give developers access to assets to develop new software without starting from scratch.
- Public APIs are considered open and are shared outside of the owner-organizations.
- Private APIs are restricted to use only within the owner-organization.





Hyper-Converged Storage vSAN

- There are two types of vSAN clusters:
 - "All Flash" vSAN clusters are made up entirely of SSDs and PCI-E storage devices. These are extremely high performance.
 - "Hybrid" vSAN clusters combine server-attached flash devices for caching purposes and magnetic drives for storage. These are more cost effective.
- Combines all the storage from ESXi hosts into a single pool of storage. It then allocates this storage to VMs based on their policies.
- vSAN is an enterprise-class storage solution for any virtualized application that allows seamless integration with vSphere and the entire VMware stack.



© Network Development Group reserved for use with NDG.tech/vmware content

- Ease of use. vSAN provies step by step guidance on how to create a vSAN cluster, in addition to scaling up with new drives or scaling out with new hosts at a moment's notice without disruption.
- vSAN integration into the ESXi hypervisor simplifies management and removes the need for dedicated hardware and complicated networking.
- vSAN is designed to utilize the newest developments in flash technology to maximize performance. This couples with the ability to use industry standard hardware, rather than proprietary hardware.







- Deduplication and Compression both help to reduce the amount of storage required and aid in getting the most out of your storage solutions.
- The VMware Update Manager (VUM) brings increased efficiency to the update process by centralizing all updates in a single location and scanning for issues post-update.
- Storage Policy-based Management allows VMs to get precisely what they need out of your storage hardware, no more and no less. This increases storage efficiency.
- vSAN works extremely well with both APIs and SDKs.





mware[®]

- vSAN encryption is the industry's first native HCI encryption solution. This can be enabled or disabled easily, and does not require self-encrypting drives.
- vSAN offers replication, continuously copying data from one server to another to minimize disruption in the event of a failure.
- The Snapshot feature vSAN offers allows you to save the state of a VM at a specific point in time. This is useful in a wide variety of situations, not limited to testing and developing.
- Cloning a VM creates a copy with its own MAC address and ID. Any changes made to the clone will not affect the original VM.



 vSAN contains a Quality of Service feature that can throttle the amount of Input/Output Operations per Second (IOPS). This prevents one VM from consuming all available resources and ensures that all VMs can access the resources they need.







Cache Layer and Capacity Layer

- vSAN architecture utilizes two layers:
 - Cache Layer- Used for read caching and write buffering, this is for "hot" data. 0
 - Capacity Layer- Used for long term storage, this is for "cold" data. 0
- The Cache Layer must always consist of a flash device, such as an SSD.
- The Capacity Layer may contain all flash devices in an "All Flash" format, or one or more magnetic devices in a "Hybrid" format.
- vSAN organizes disks into disk groups. A disk group will contain 1 drive on the Cache Layer and 1-7 devices on the Capacity Layer. A vSAN host can contain up to 5 disk groups.

SSD

FLASH





Object and Component Layout

- Virtual Machines contain five types of objects:
 - VM Home- Stores all virtual machine configuration files.
 - VMDK- This is the Virtual Machine's disk file.



- VM Swap- Reduces the amount of memory the host must reserve for VM operations. Created when the VM powers on.
- Snapshot- A copy of the VM's state at a specific point in time.
- Memory- A backup of the VM's memory stored on the host file system.
- An object is created for each virtual disk.
- These objects are logical and are mirrored across the cluster.



Object and Component Layout

- A vSAN cluster can store and manage tens of thousands of objects.
- Objects are striped into components. A single component can not exceed 255GB. If an object is larger, it is striped into multiple components.
- Policies can also affect the number of components and object is striped into.
- Objects and components can reside on different hosts.
- Components make up the 'leaves' of the object tree.







Summary of vSAN Benefits

- Ease of configuration, provisioning, and management. Due to being embedded in the hypervisor, configuration and installation are performed with a few clicks in the vSphere Client.
- Flexibility and Agility both in meeting changes in demand (both with application needs and scaling needs).
- Savings accrued from being able to use industry standard x86 hardware rather than proprietary, custom hardware.
- Higher performance, lower latency reduced resource consumption and need, increased security, streamlined operations, and potential 99.999% availability all granted from using HCI.

