

VMware vSphere basics

Contents

- Definition of hypervisors and their types
- ESXi vs ESX.
- VMware vCenter, hosts, differences in vmRefs.
- Levels of inception (VM level, ESXi level, vC level).
- Authentication, network ports, Single Sign-On.
- Managed Object Browser (MOB).
- VM files.
- VM management, VMware Tools, vC features (vMotion, Storage vMotion, Distributed Resource Scheduler, High Availability, Fault Tolerance).
- Snapshots (copy-on-write, redirect-on-write).
- VMFS and RDM.
- VM disk types
 - Thin vs thick
 - Dependent vs independent
- Virtual switches and networks.
- Permissions.
- Licensing.

Definition of a hypervisor

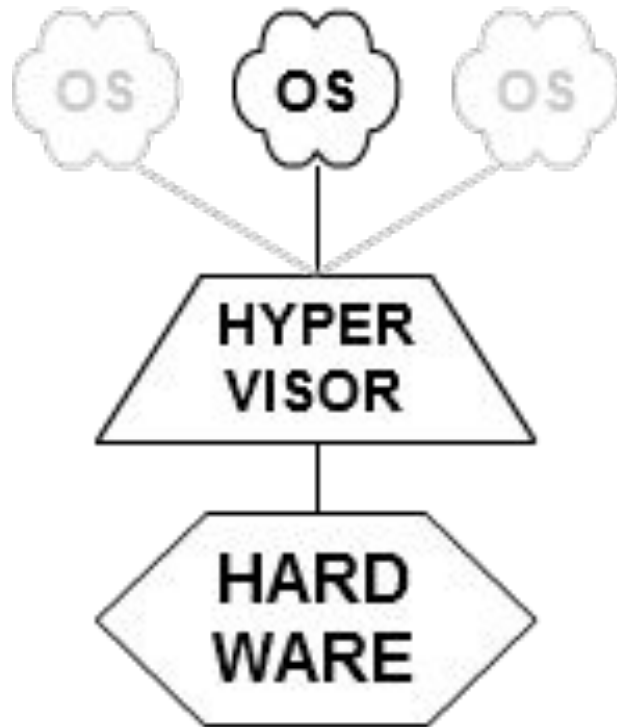
A *hypervisor* is a piece of computer software, firmware or hardware that creates and runs *virtual machines*.

A computer on which a hypervisor is running one or more virtual machines is defined as a *host machine*. Each virtual machine is called a *guest machine*. The hypervisor presents the guest operating systems with a virtual operating platform and manages the execution of the guest operating systems.

Hypervisor types

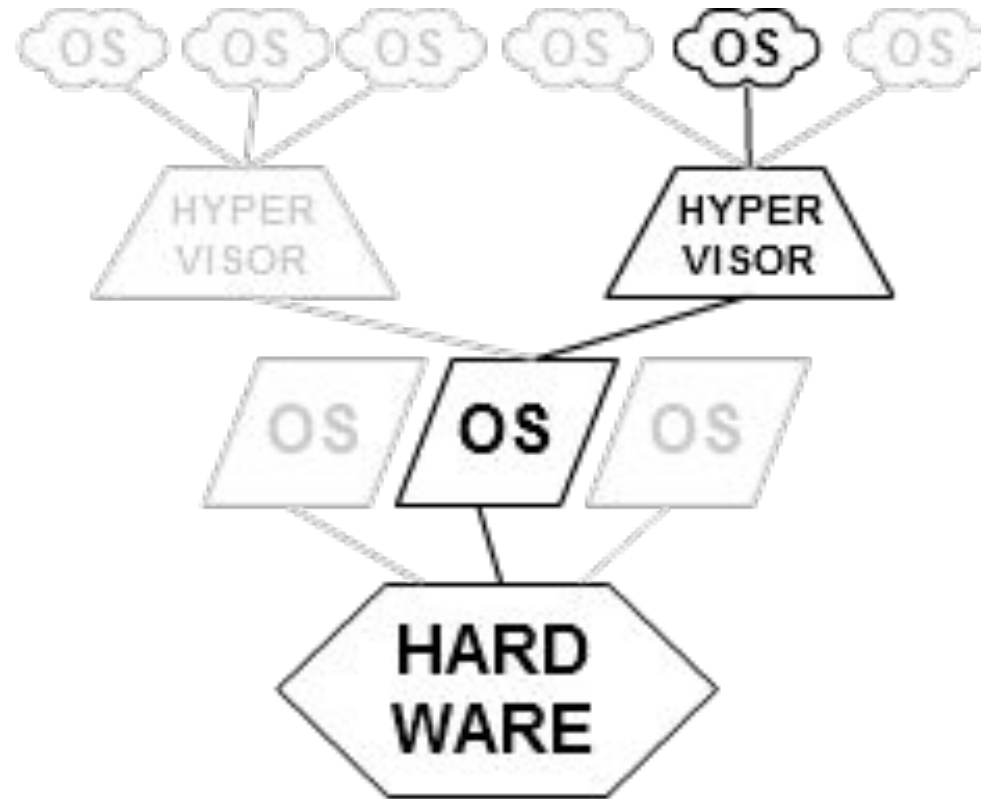
- *Type 1* (or *native*, bare metal) hypervisors run directly on the host's hardware to control the hardware and to manage guest operating systems. A guest operating-system thus runs on another level above the hypervisor.
 - VMware ESX(i), Microsoft Hyper-V
- *Type 2* (or *hosted*) hypervisors run within a conventional operating-system environment. With the hypervisor layer as a distinct second software level, guest operating-systems run at the third level above the hardware
 - VMware Workstation, VirtualBox

Hypervisor types



TYPE 1

native
(bare metal)



TYPE 2

hosted

ESX vs ESXi

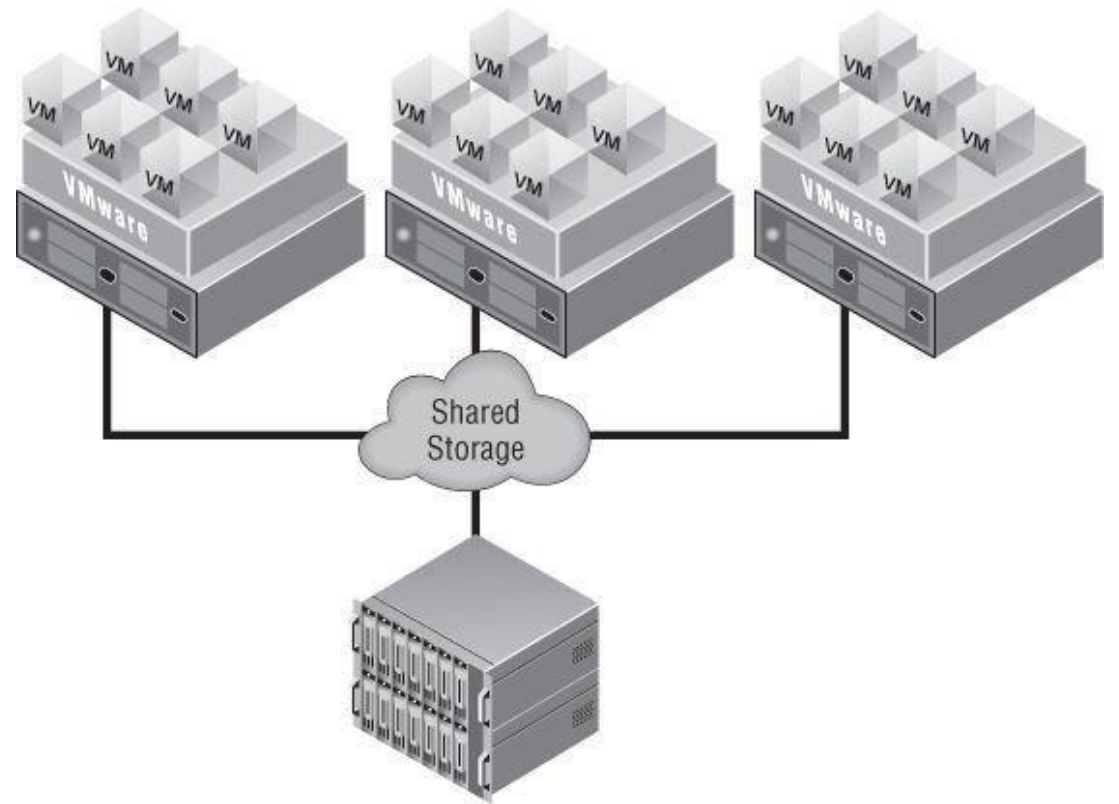
- ESX
 - Existed up until ESX 4.1 (year 2010)
 - Includes a Linux service console
 - Now considered deprecated, but still supported by Veeam and VMware
- ESXi
 - Smaller footprint
 - No service console
 - As of 2016, latest version is 6.0 (previously 5.5, 5.1, 5.0)

Hosts and VMs

Every VM runs on a certain host and is provided the following resources:

- CPU, RAM, storage, networking

One can connect to a ESX(i) host and manage it via the vSphere Client.



Storage protocols and connection type

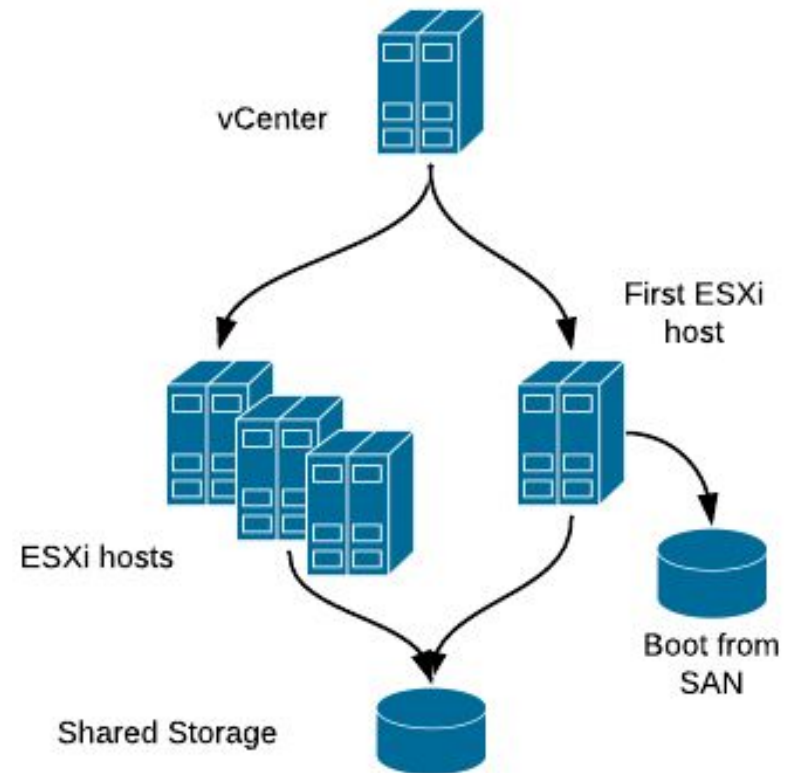
- Locally attached
 - SATA
 - SAS
 - SCSI
- Shared
 - iSCSI
 - Fibre-channel
 - Fibre-channel over Ethernet (FCoE)
 - NFS 3 / NFS 4.1

Storage devices

- SAN (Storage Area Network)
 - High-grade storage device
 - Expensive
 - Will likely support most protocols: iSCSI, FibreChannel, NFS, CIFS/SMB
 - Suitable for production (VM files)
 - Advanced features (e.g. SAN snapshots)
- NAS (Network-Attached Storage)
 - Cheaper device
 - Less protocols: CIFS/SMB, NFS, sometimes iSCSI.
 - Suitable for backups
- Deduplication appliance NAS boxes

vCenter

- vCenter serves as a single point of management with access to all of the resource, and it enables new features.
- Several **hosts** may be added to vCenter and then joined to a **cluster** and grouped into **datacenters**.



vCenter deployment options

vCenter deployment options:

- Installed on Windows Server 2008+ (physical or virtual)
 - Database options: PostgreSQL (built-in), MSSQL, Oracle
- Deployed as *vCenter Server Appliance (VCSA)*
 - Based on SUSE Linux Enterprise Server
 - Only virtual
 - Database options: PostgreSQL (built-in), Oracle

Core components:

- vCenter Server (vpxd)
- Single Sign-On
- VMware vSphere Web Client

vSphere/vCenter features

- vMotion
 - Move a VM without interruption from one host to another
- Storage vMotion
 - Move VM's files (disks) without interruption from one datastore to another
- Distributed Resource Scheduler
 - Uses (s)vMotion to migrate VMs between hosts and datastores for resource balancing
- High Availability
 - In case of host failure will restart the VMs on a different host, short downtime
- Fault Tolerance
 - Runs a perfect copy of a VM and switches over to it in case of hardware failure

Licensing – vSphere Editions

- Hypervisor (“Free ESXi”)
 - No vCenter, no vSphere or vStorage APIs (thus unsupported by Veeam B&R).
- Essentials
 - The most basic feature set.
- Essentials Plus
 - Adds High Availability, vMotion.
- Standard
 - Adds Storage vMotion, FT, Virtual Volumes (new in 6.0).
- Enterprise
 - Adds DRS, MPIO.
- Enterprise Plus
 - Adds sDRS, dvSwitch and more.
- Further details: <http://kb.vmware.com/kb/2109507>

Authentication and ports

- ESX(i) Host:
 - root is the default user with full privileges
 - Access via:
 - VMware vSphere Client, uses HTTPS on TCP port 443
 - SSH, uses TCP port 22
 - TCP port 902 for NFC (“network file copy”) requests
- vCenter:
 - administrator@vsphere.local is the default user with full privileges
 - Single Sign-On allows using Active Directory domain user accounts for authorization, e.g. DOMAIN\user
 - Access via:
 - Recommended: vSphere Web Client e.g. <https://vcenter:9443/vsphere-client/>
 - Legacy: VMware vSphere Client, port 443

Managed Object Browser

- The Managed Object Browser (MOB) is a graphical interface that allows you to navigate the objects on a server and to invoke methods. Any changes you make through the MOB take effect on the server.
- Separate MOB's for each host and every vCenter.
- Accessed on e.g. <https://vcenter/mob/>
- Note that Veeam will use VMware Web Services SDK instead, e.g. <https://vcenter/sdk/>

VM Files

- *machine.vmx*, *machine.vmx*
 - VM configuration and extended VM configuration files
- *machine.vmdk*
- *machine-flat.vmdk*
 - Disk descriptor and disk data blocks
- *machine-000001.vmdk*
- *machine-000001-delta.vmdk*
 - Delta disk descriptor and disk data blocks, when a VM is running on snapshots
- *machine-ctk.vmdk*
 - Changed block tables used for Changed Block Tracking (binary)
- *machine.nvram*
 - Essentially VM BIOS
- *machine.vmsd*, *machine.vmsn*
 - Snapshot tree (text) and snapshot state file (binary)
- *vmware.log*
 - VM-specific events are logged here

Virtual Machine Compatibility Levels

| vSphere release | Virtual Machine Hardware Version | vSphere compatibility |
|-----------------|----------------------------------|-------------------------------|
| vSphere 4.0 | Version 7 | VMware ESX/ESXi 4.0 and later |
| vSphere 4.1 | Version 7 | VMware ESX/ESXi 4.0 and later |
| vSphere 5.0 | Version 8 | VMware ESXi 5.0 and later |
| vSphere 5.1 | Version 9 | VMware ESXi 5.1 and later |
| vSphere 5.5 | Version 10 | VMware ESXi 5.5 and later |
| vSphere 6.0 | Version 11 | VMware ESXi 6.0 and later |

In the .vmx file:
`virtualHW.version = "10"`

VMware Tools

- A set of drivers and utilities that enhance the performance of the VM's guest operating system and improve management of the VM.
 - Specific network, graphics and other drivers.
 - Reporting of hostname, IP address, options to shutdown/restart guest.
- Allows for communicating with guest via network-less VIX API
- Windows: an .msi package
- Linux: a .tar package with a Perl script
 - Recommended: [Open Virtual Machine Tools](#) (**open-vm-tools** package)

Snapshots

- Snapshot is a persistent saved state of a VM at a point in time.
- Snapshots allow for reverting the VM to that point in time and also to read data from it
- Two ways to implement a snapshot:
 - Redirect-on-write (“delta”) snapshots (used on ESXi)
 - Copy-on-write snapshots (used on NTFS shadow copies)

Redirect-on-write (“delta”) snapshot

- Used on VMware ESX(i)
- Once snapshot is created, disk changes are written to a delta disk instead. Base disk is no longer locked for read or write.
- Snapshot trees can be maintained. VMware supports up to 32 snapshots in a single branch.
- VM I/O performance suffers greatly when running on snapshots
- Snapshot deletion causes a short VM freeze (vSphere < 6.0)

Redirect-on-write (“delta”) snapshot

| | | | |
|---|---|---|---|
| 1 | 0 | 1 | 1 |
|---|---|---|---|

vmname-flat.vmdk

Disk data before snapshot

| | | | |
|---|---|---|---|
| 1 | 0 | 1 | 1 |
|---|---|---|---|

vmname-flat.vmdk

Snapshot is created

| | | | |
|--|--|--|--|
| | | | |
|--|--|--|--|

vmname-flat-00001.vmdk

| | | | |
|---|---|---|---|
| 1 | 0 | 1 | 1 |
|---|---|---|---|

vmname-flat.vmdk

A ‘0’ is written to block 3 by OS, but all changes are now saved to the snapshot disk instead

| | | | |
|--|--|---|--|
| | | 0 | |
|--|--|---|--|

vmname-flat-00001.vmdk

VMFS

- VMware VMFS (Virtual Machine File System) is a cluster file system.
- It is used to store virtual machine disk images, including snapshots. Multiple servers can read/write the same file system simultaneously, while individual virtual machine files are locked.
- Current version: VMFS 5 (block size always 1 MB)
 - up to 62 TB file size (and thus VM disk limit) on vSphere 5.5 and newer
 - up to 2 TB file size on vSphere 5.1
- For VMFS 3, the block size limits maximum file size:
 - 1 MB block > 256 GB max file size
 - 2 MB block > 512 GB max file size
 - 4 MB block > 1 TB max file size
 - 8 MB block > 2 TB max file size

Raw Device Mapping (RDM)

- Allows for connecting LUNs attached to the host directly to a VM, bypassing the VMFS layer
- For example a physical NTFS-formatted volume can be attached to a Windows VM
- Physical RDM, aka “pass-through”
 - (almost) All SCSI commands are sent to the device, unfiltered
 - Snapshots are **not** supported
- Virtual RDM
 - Read and write SCSI commands are translated by the ESXi host
 - Snapshots **are** supported

Disk provisioning types

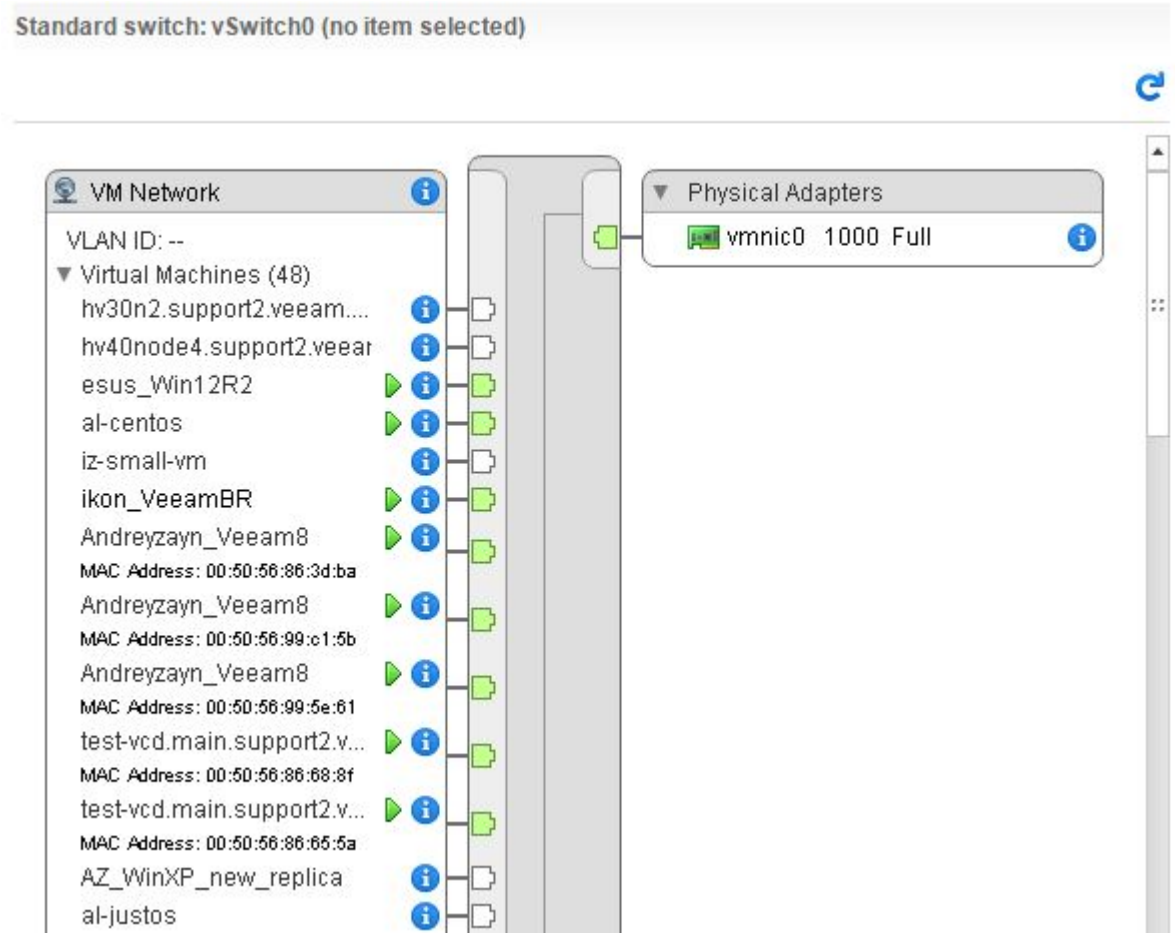
- Thin provisioning
 - Only the blocks that have been written to are actually allocated, such disk slowly grows in size on the datastore
- Thick lazy-zeroed
 - The space is immediately allocated for such disk
- Thick eager-zeroed
 - The space is immediately allocated and also zeroed

Dependent vs independent disks

- Dependent
 - The most common type, support snapshots
- Independent persistent
 - Doesn't support snapshots, changes are always written to the disk
- Independent non-persistent
 - Changes are always written to a delta disk file which is discarded once the VM is powered off
 - Useful for testing and lab setups

Networking

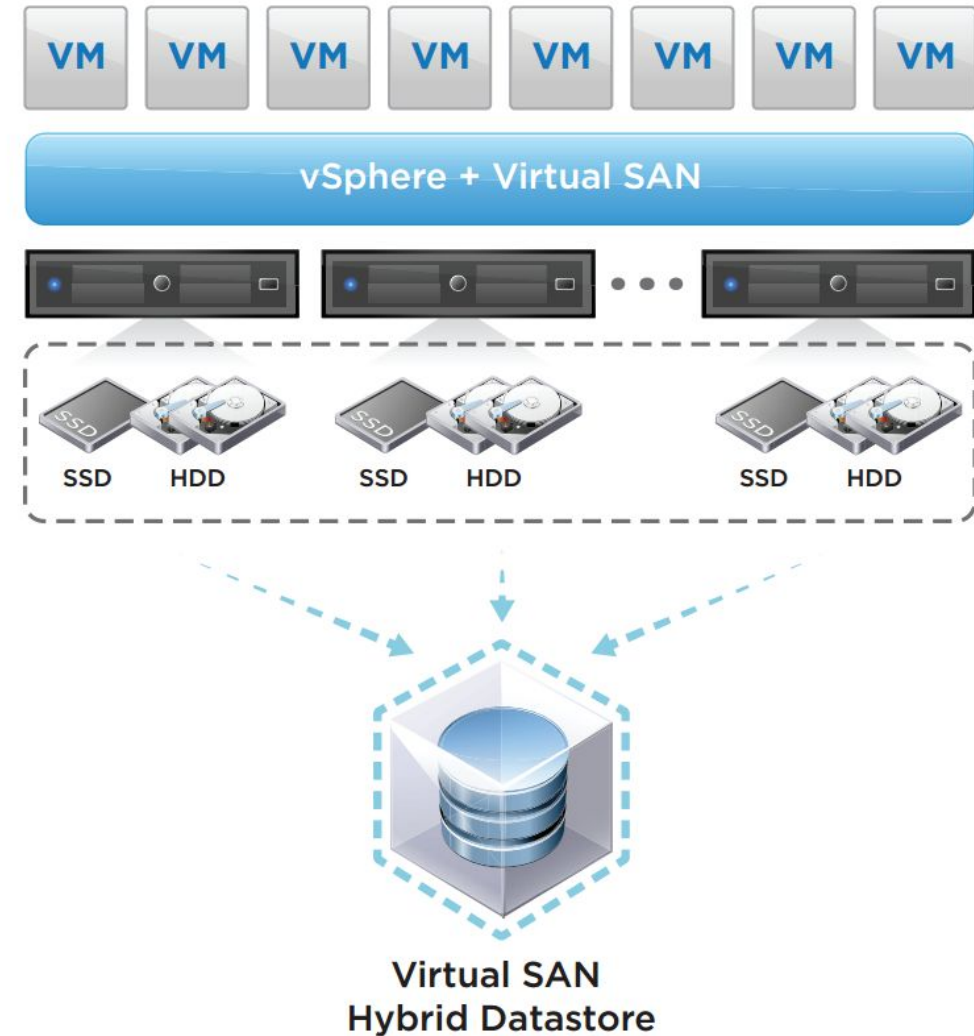
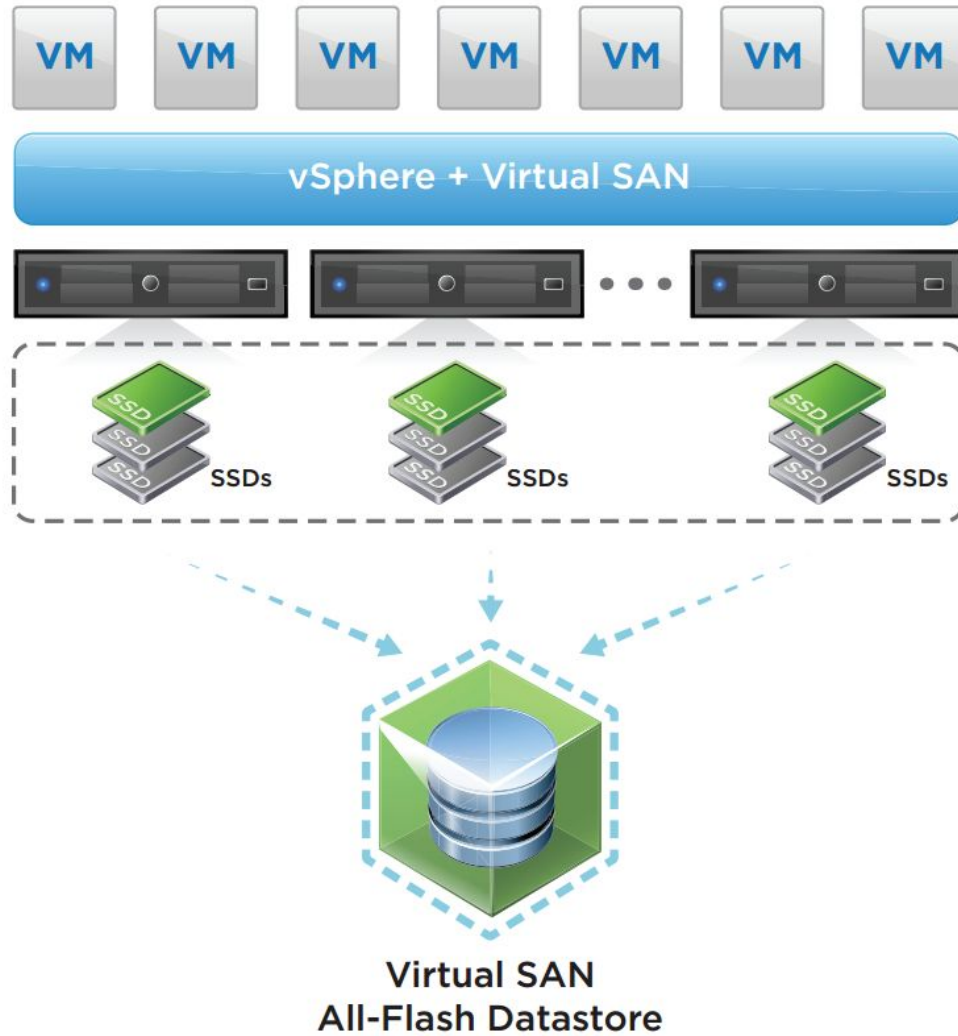
- vmnic# physical cards on the ESXi hosts.
- vSwitch (separate vSwitches per host).
- Port groups on vSwitches:
 - Virtual machine port groups
 - VM communication (locally or w/ physical network)
 - Kernel port groups
 - Host management, vMotion, FT, iSCSI traffic.
- Also dvSwitch - clustered switch for advanced setups



Permissions

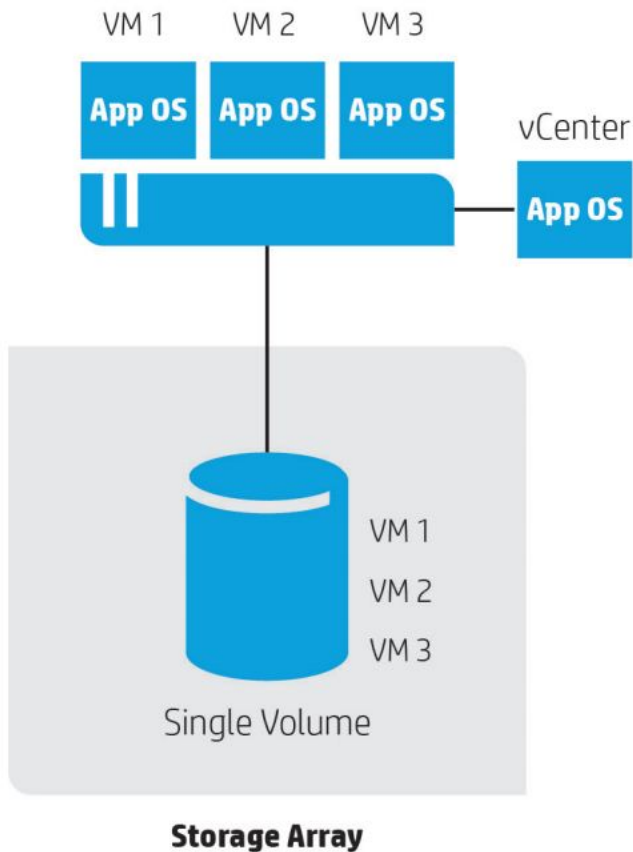
- Role — a set of granular permissions for each **type** of object handled by vSphere.
- Roles are assigned per **user** or per **group** for a single object (e.g. VM, datastore), or a container (folder, resource pool, cluster, etc.).
- “Propagate to children” option applies the role permissions also to every child object of a container.

Virtual SAN (VSAN) (new in 5.5)

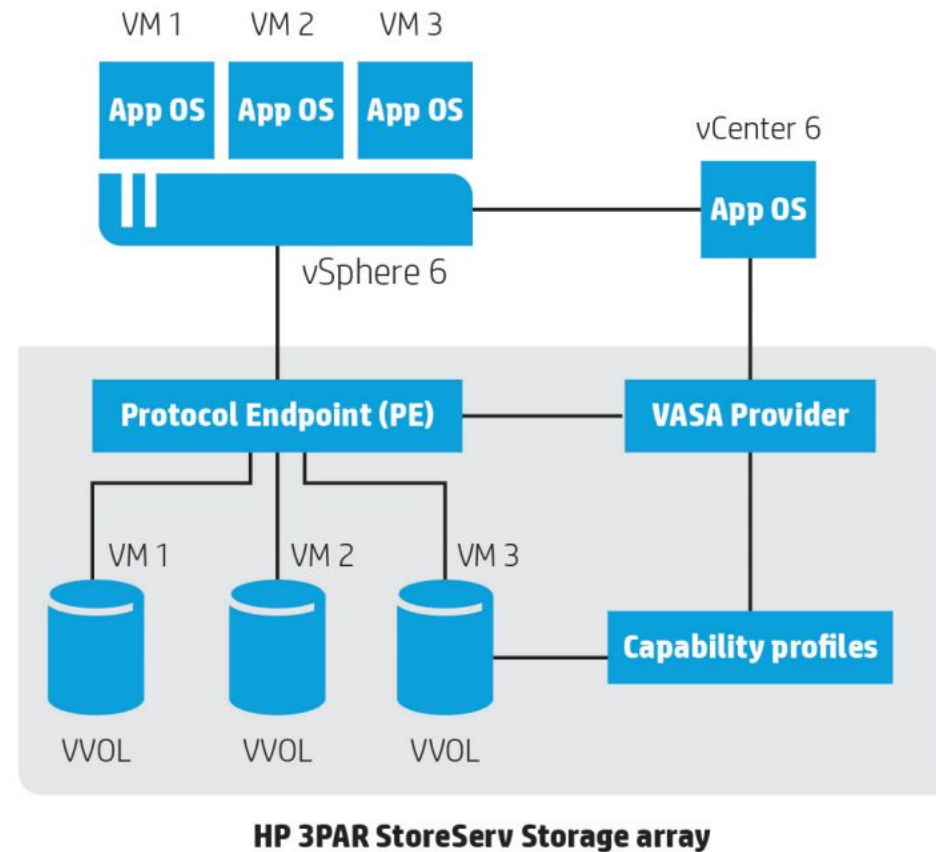


Virtual Volumes (VVOL) (new in 6.0)

Traditional VMFS Datastore (LUN-based)



VMware VVOLs on HP 3PAR StoreServ Storage array



Virtual Volumes (VVOL)

