Virtualization Technologies

Concepts and applications

Orange Education Program

July 2020

Mihai Huică

General concepts

Definition

- Creating a virtual version / instance (software based) of a traditionally physical entity
  - Compute
  - Storage
  - Network
  - Application
- Important component of cloud computing

Advantages of virtualization

- Efficient use of resources - cost reduction
- Easier administration
- Agility, flexibility, speed
- High availability and disaster recovery
- Reducing downtime
- Increased productivity
- Flexible charging, in accordance with the actual usage

Virtualization types

- Desktop
- Network
- Storage
- Data
- Application
- Data center
- CPU
- GPU
- OS
- Cloud

Containers

- Lighter-weight, more agile virtualization
- Packages everything needed to run a small piece of software
- Improved portability

Virtualization vs. containerization

Criteria	Virtualization	Containers
What is being virtualized?	physical hardware	OS
What is the result?	virtual machine	container (smaller, faster, portable)
What is included in the result?	virtual copy of the hardware guest OS application associated libraries and dependencies	application libraries and dependencies
Application architecture	monolithic	microservices (more granular deployment and scalability)

The hypervizor

- Placed between physical HW and Guest OS, it coordinates and runs virtual machines (guest) using physical machine (host) resources

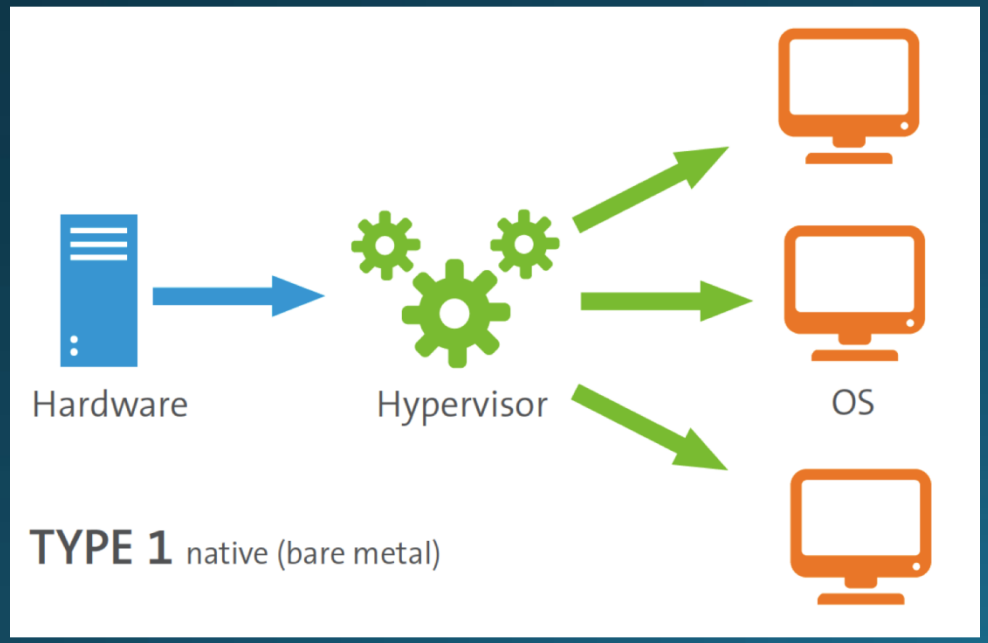

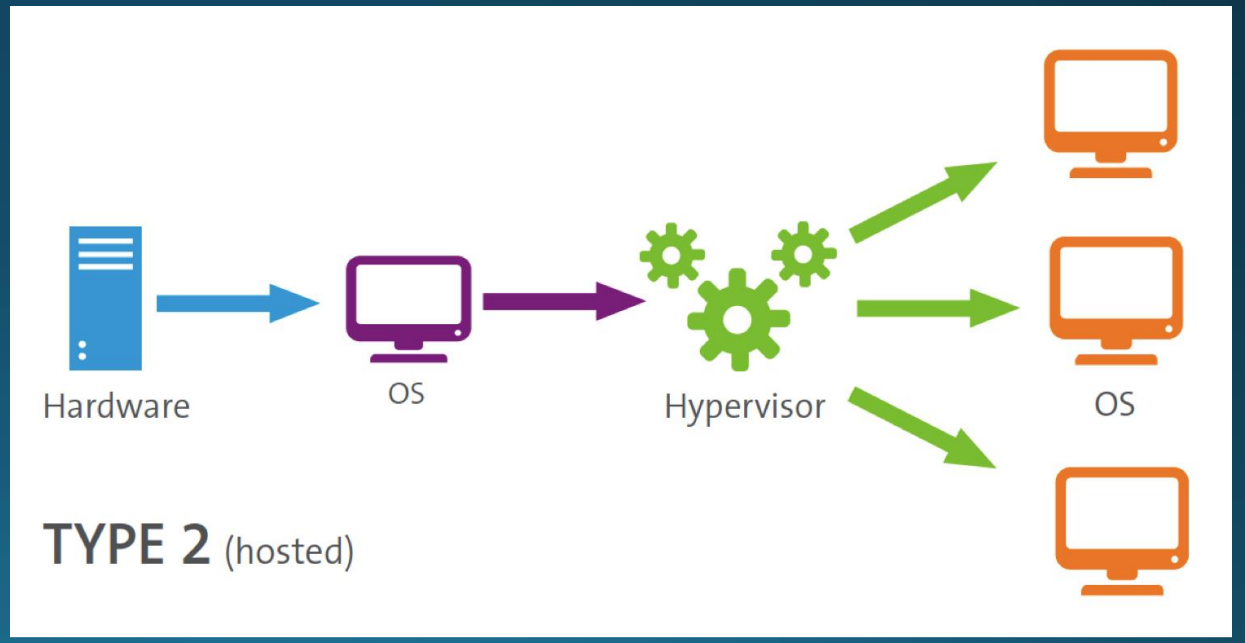

VMware vSphere

VMware ESXi

- Type 1 („bare metal”) hypervisor
  - replaces the underlying OS altogether
- Alternatives
  - Microsoft Hyper-V
  - Citrix XenServer
  - KVM

VMware Workstation

- Type 2 hypervisor
  - runs as an application on the desktop OS
- Alternatives
  - Oracle VirtualBox

Virtual machines

- Virtual environments that simulate in software-based form a physical computational environment
- Function just like physical machines
- Have an operating system and software applications
- Independent of each other

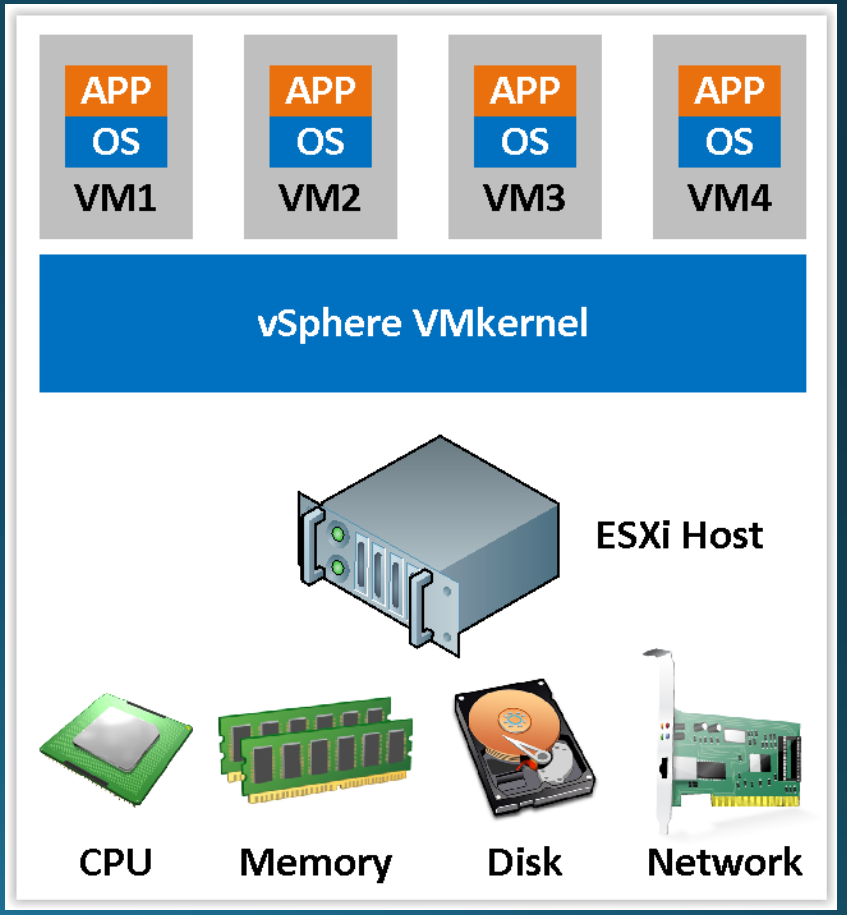

Physical resources sharing

- Each VM shares the resources of one physical computer
- Physical resources are sliced and allocated to VMs
  - CPU
  - Memory
  - Disk
  - Network

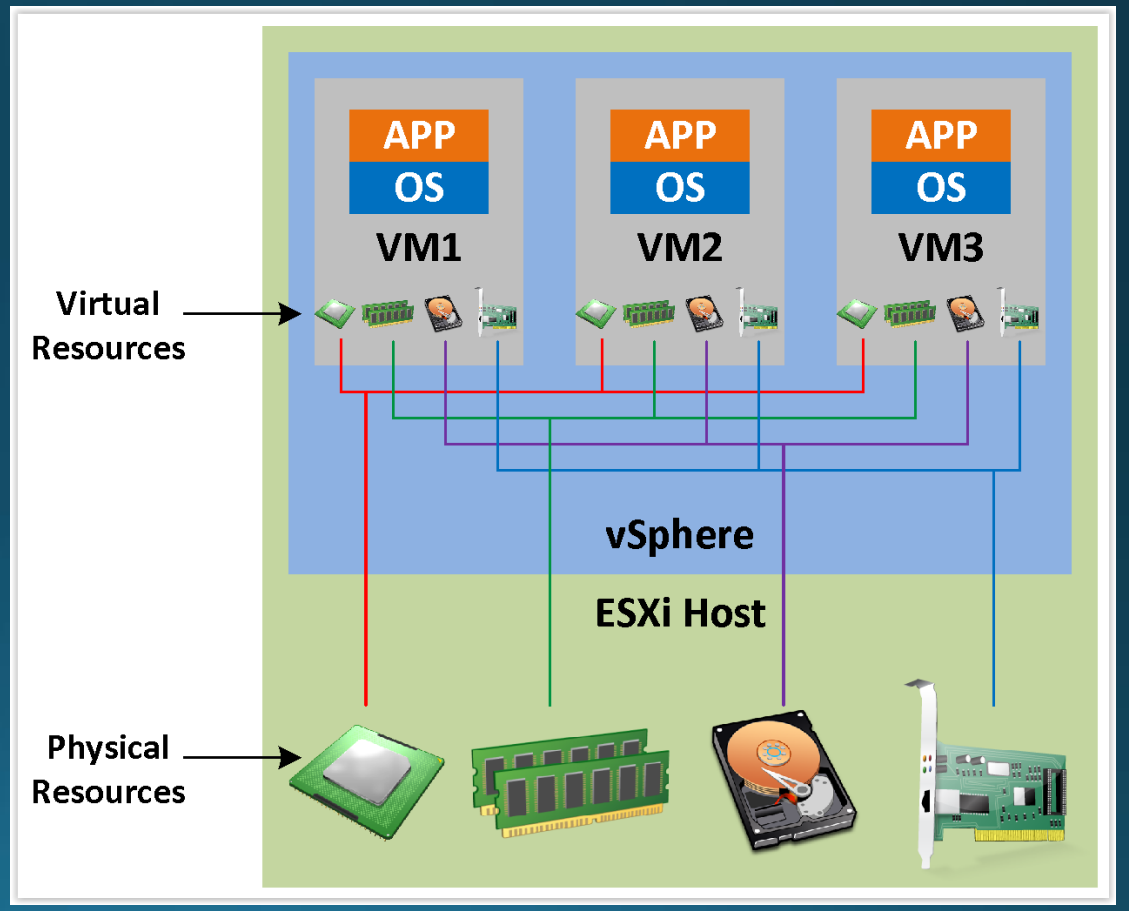

Storage

- Datastores
  - logical storage containers
  - uniform model for storing virtual machine files
- Types
  - VMFS
  - NFS
  - vSAN
  - vSphere Virtual Volumes

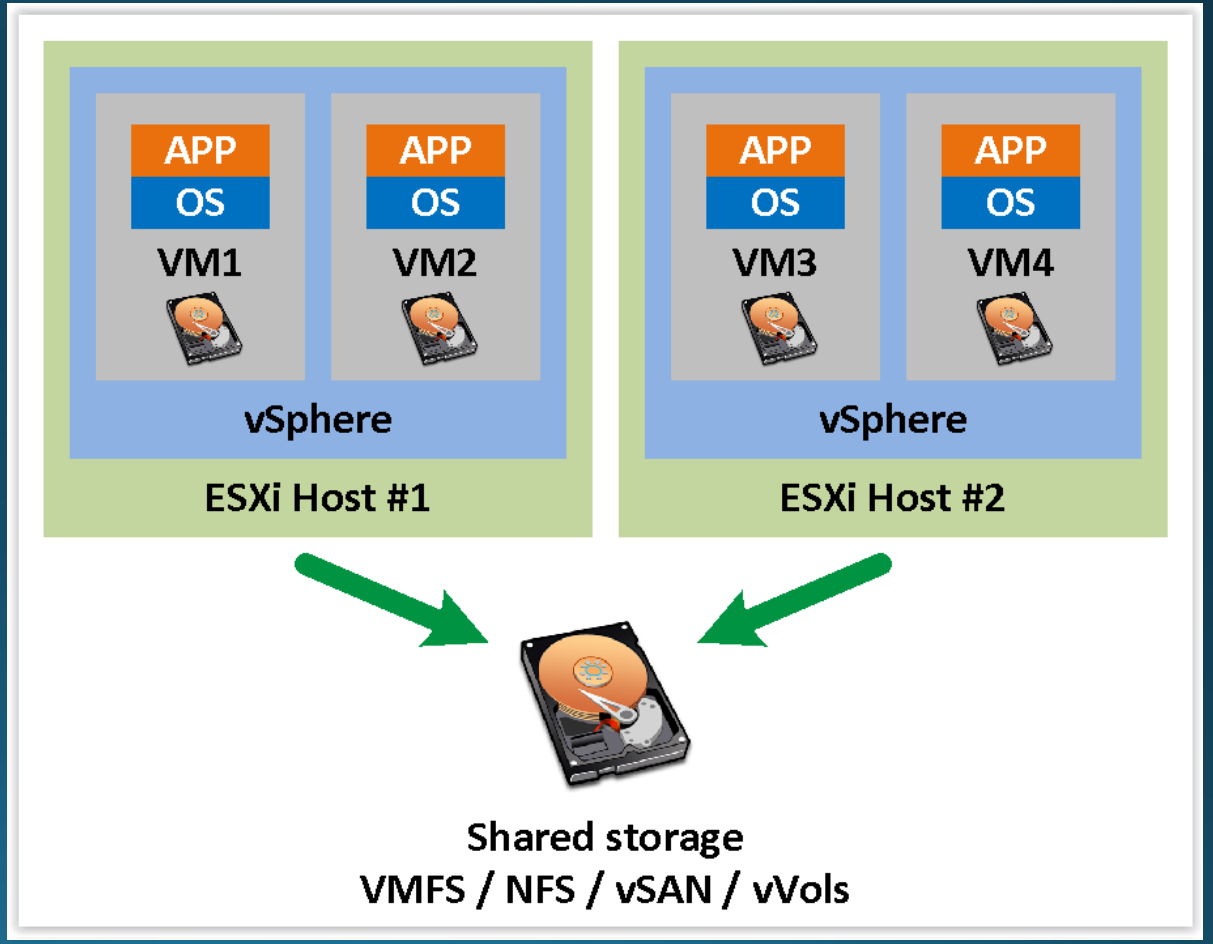

Storage – Datastores

- Protocols
  - DaS
  - FC
  - FCoE
  - iSCSI
  - NAS

	VM Virtual Disks						
Datastore type	VMFS				NFS	vSAN	vSphere Virtual Volumes
Transport	Direct Attached	FC	FCoE	iSCSI	Ethernet	Direct Attached	FC / Ethernet
Backing	LUN				File System	vSAN Cluster	Storage Container

Storage – provisioning models

- Thick Provision Lazy-Zeroed
- Thick Provision Eager-Zeroed
- Thin Provision
- Impact
  - virtual disk layout
  - zeroing allocated file blocks

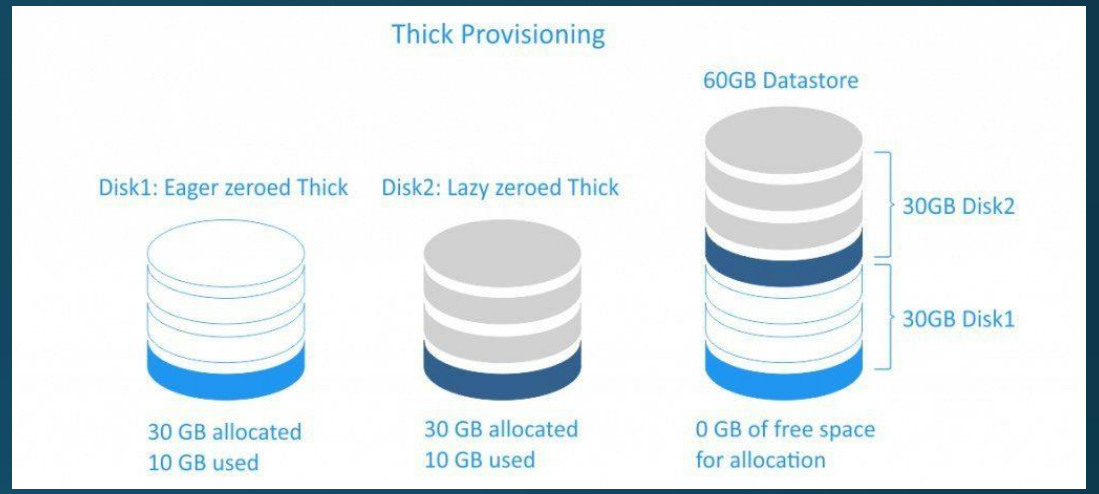

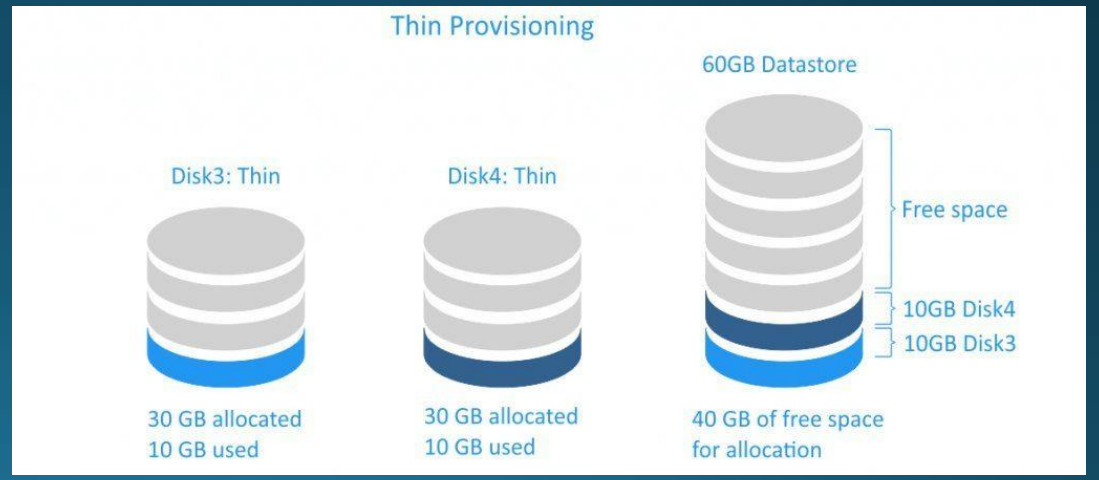

Storage – vSAN

- Hypervisor-converged, software defined storage for virtual environments
- Physical support
  - hybrid storage
  - all-flash storage

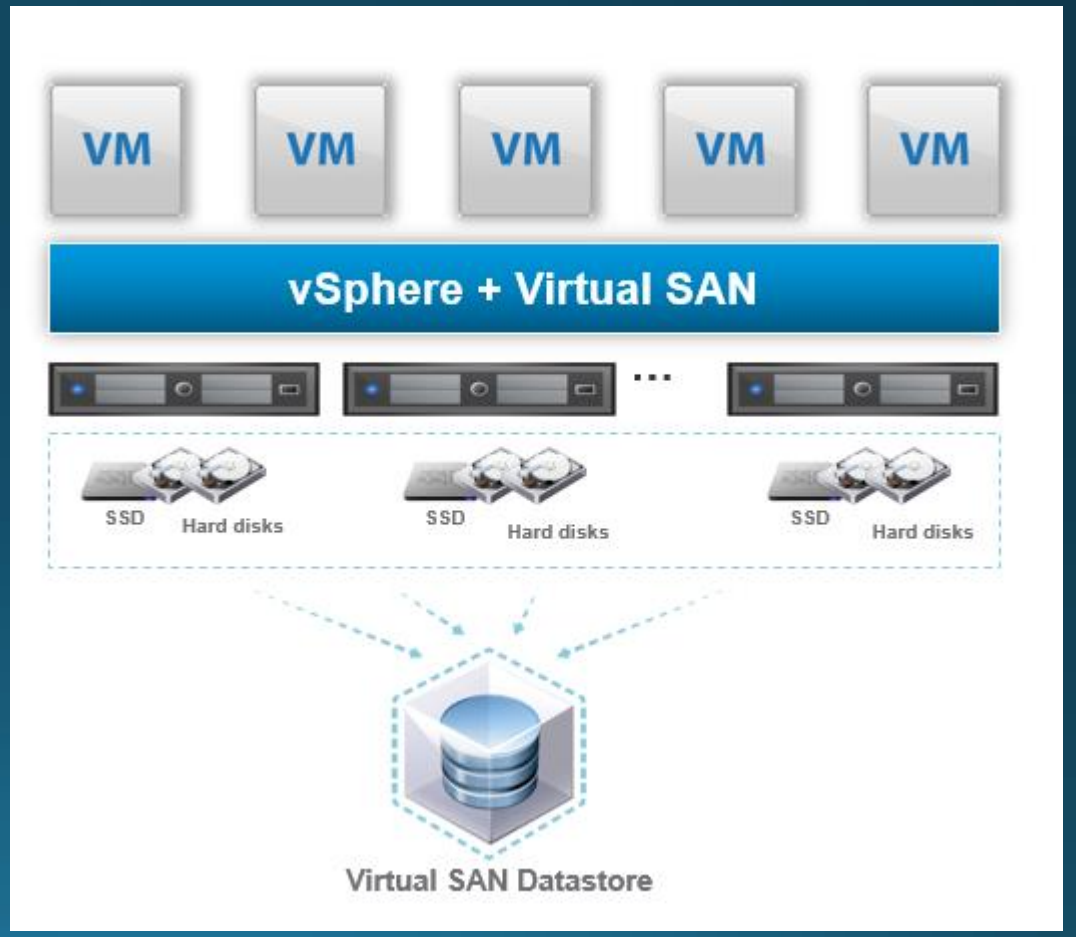

Network – Virtual switches

- Enable VMs to communicate with other machines
  - on same host
  - on different host (external)
- VLANs
- SR-IOV passthrough
- DirectPath I/O

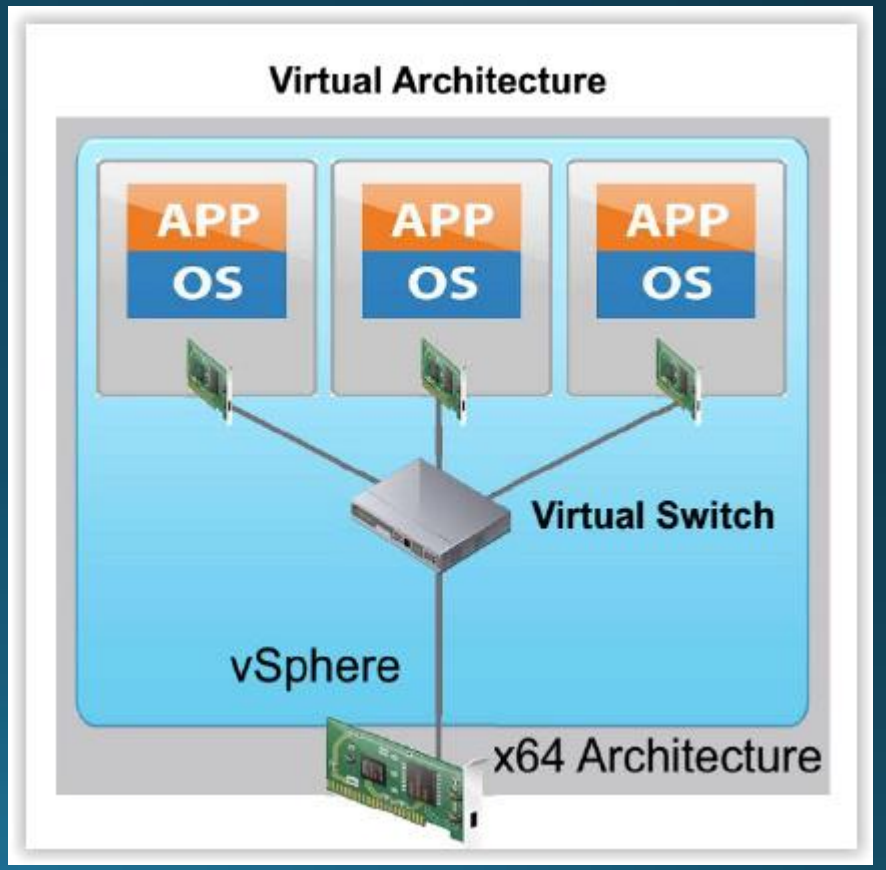

Network – standard switch

- Connection types
  - Virtual machine port groups
  - VMkernel ports
  - Uplink ports

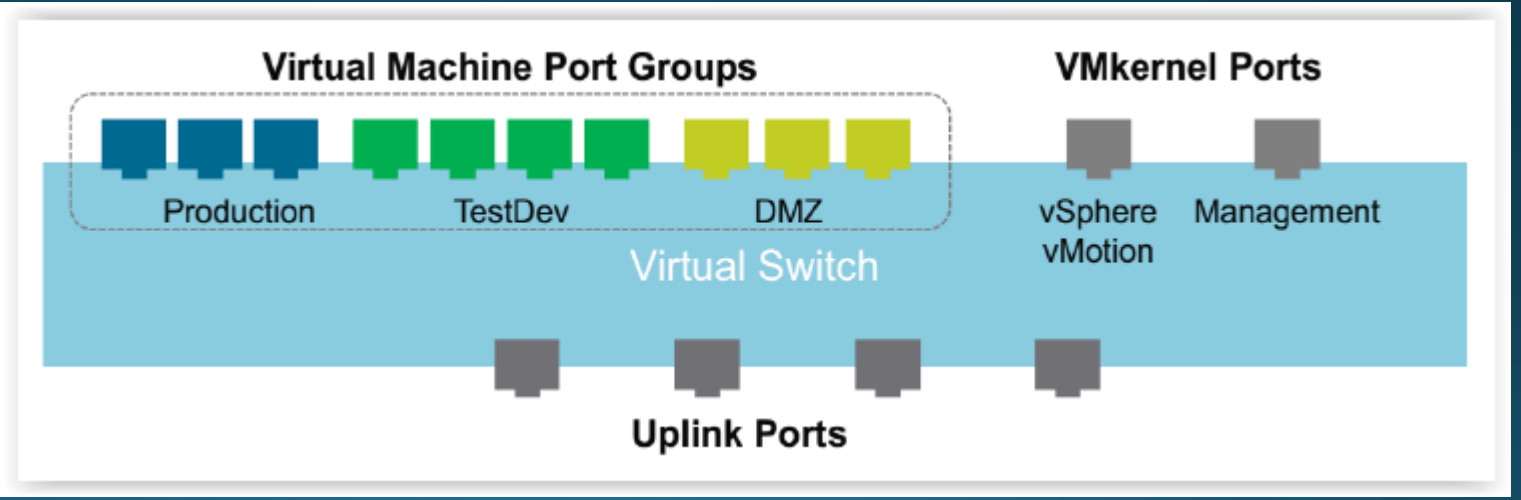

Source: VMware vSphere 6.7 Install, Configure, Manage

Network – distributed switch

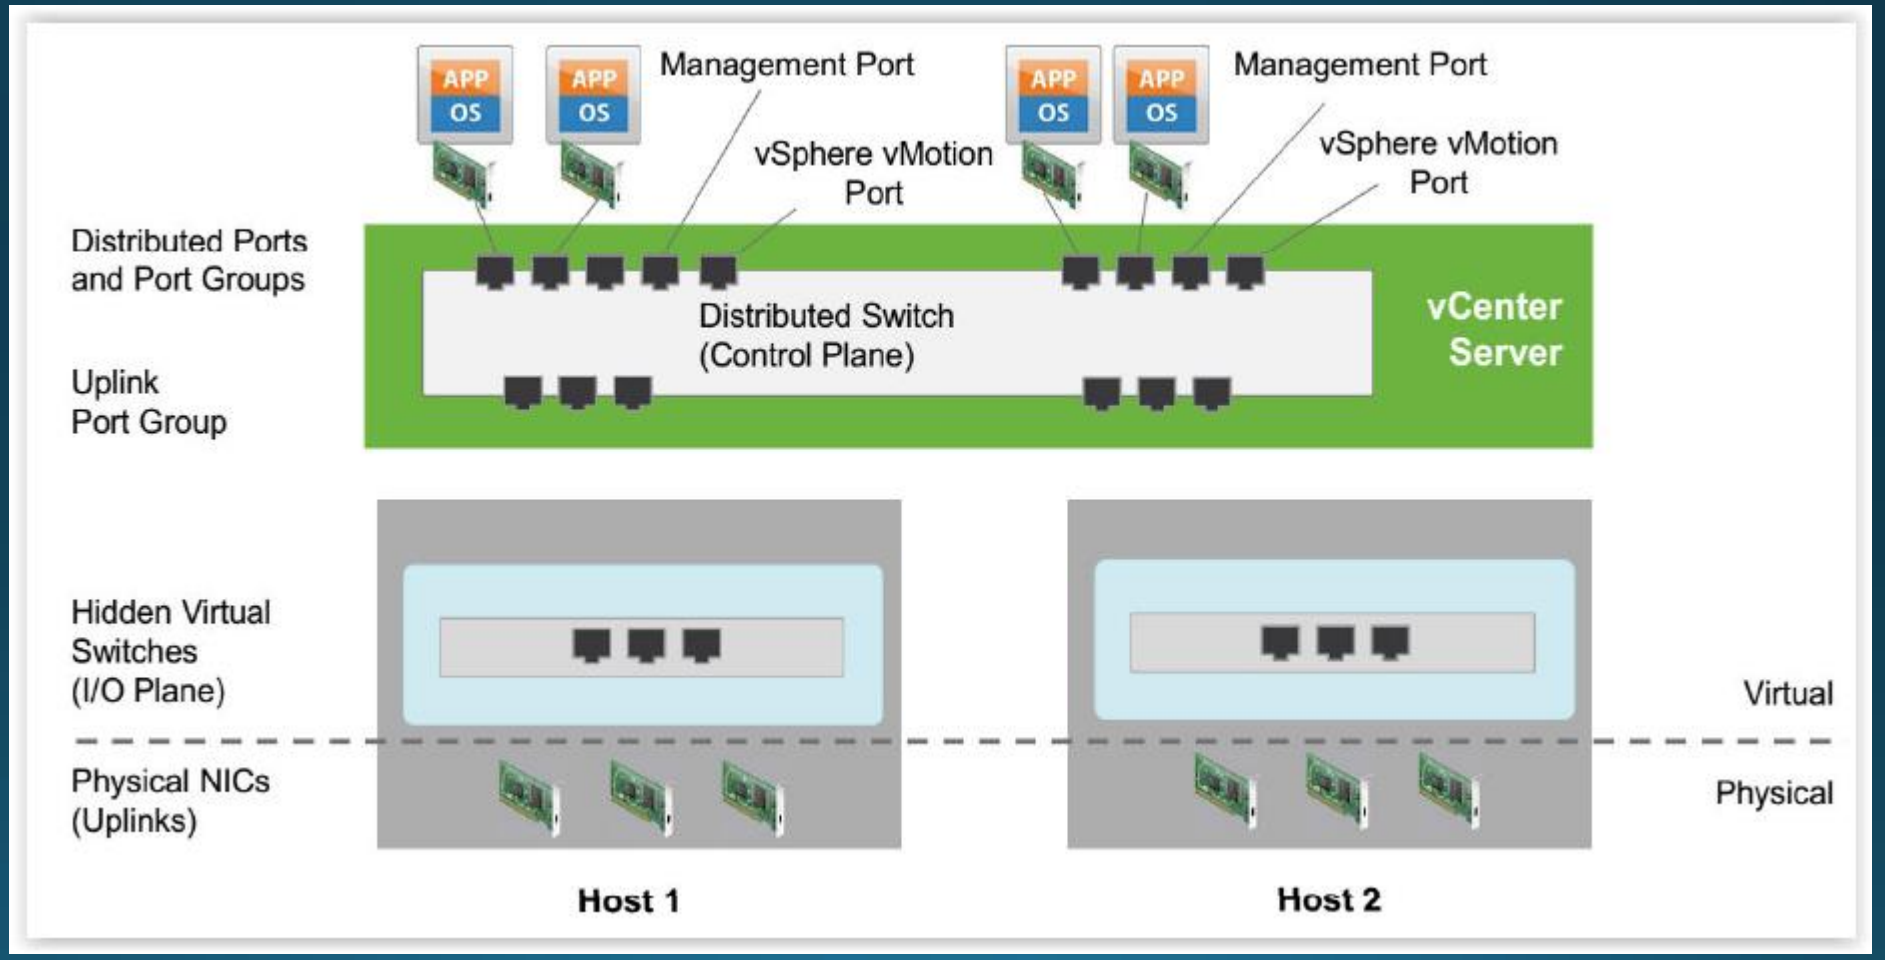

Source: VMware vSphere 6.7 Install, Configure, Manage

Network – deployment example

- Connectivity
- Traffic type
- Multipathing
- Isolation

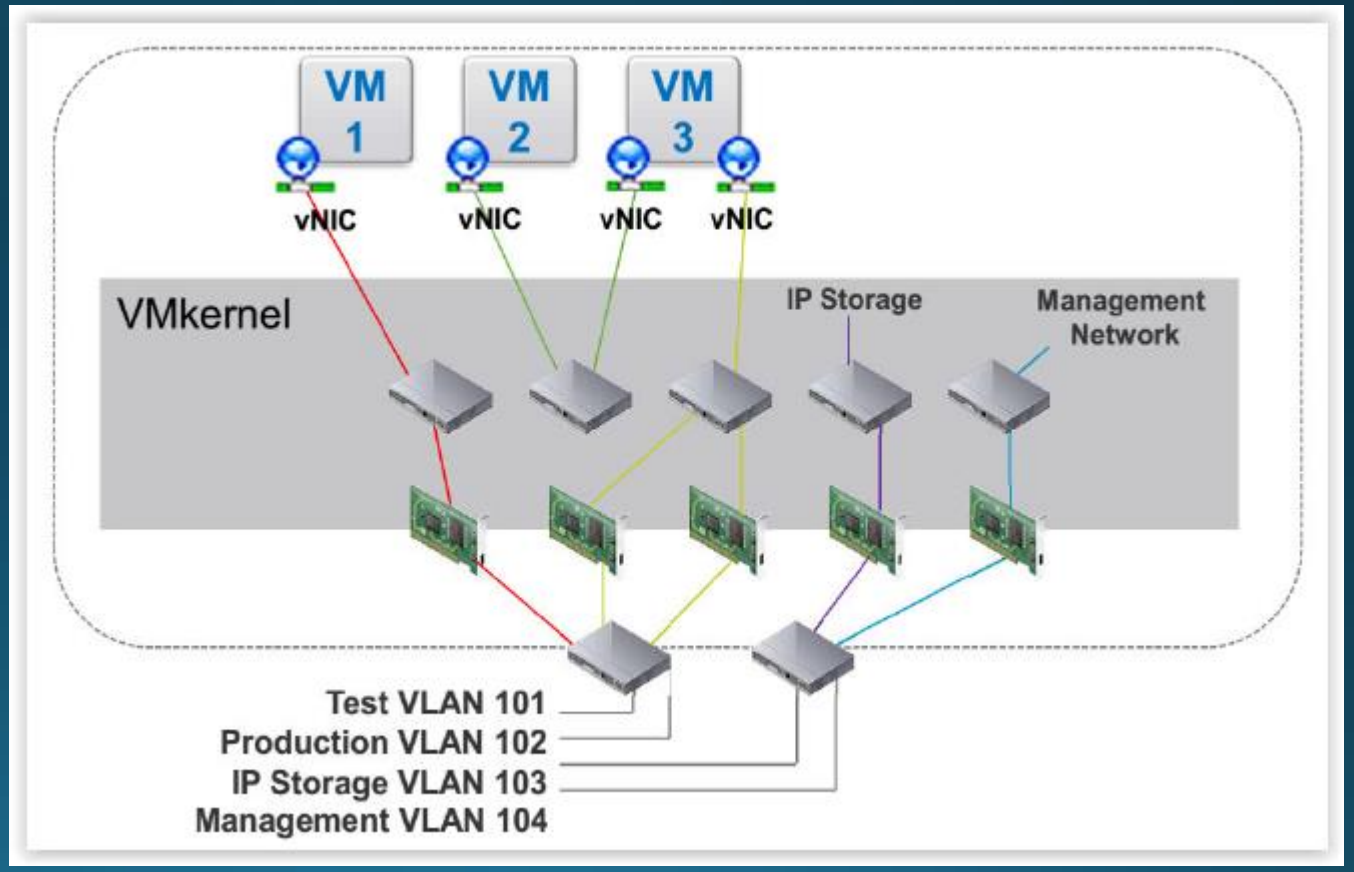

Source: VMware vSphere 6.7 Install, Configure, Manage

Network – switches comparison

- Virtual switches
  - standard
  - distributed

Feature	Standard Switch	Distributed Switch
Layer 2 switch	✓	✓
VLAN segmentation	✓	✓
IPv6 support	✓	✓
802.1Q tagging	✓	✓
NIC teaming	✓	✓
Outbound traffic shaping	✓	✓
Inbound traffic shaping		✓
VM network port block		✓
Private VLANs		✓
Load-based teaming		✓
Data center-level management		✓
vSphere vMotion migration over a network		✓
Per-port policy settings		✓
Port state monitoring		✓
NetFlow		✓
Port mirroring		✓

Source: VMware vSphere 6.7 Install, Configure, Manage

Network – switch policies

- Available network policies:
  - Security
  - Traffic shaping
  - NIC teaming and failover
- Policies definitions:
  - Switch level (default policies for all the ports on the switch)
  - Port group level (effective policies, overriding default policies set at the switch level)

Network Load-Balancing (1/3)

- Originating Virtual Port ID
  - traffic mapped to specific pNIC, determined by the ID of the virtual port to which this VM is connected

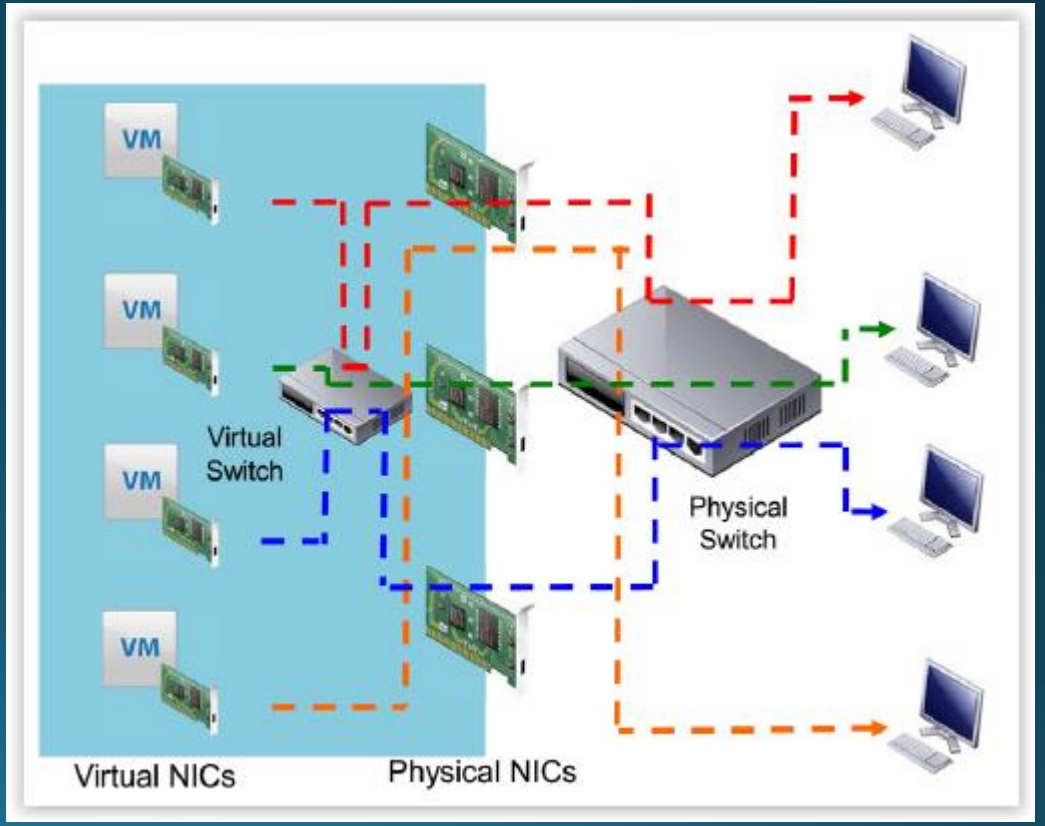

Source: VMware vSphere 6.7 Install, Configure, Manage

Network Load-Balancing (2/3)

- Source MAC Hash
  - traffic mapped to specific pNIC, based on the virtual NIC's MAC address

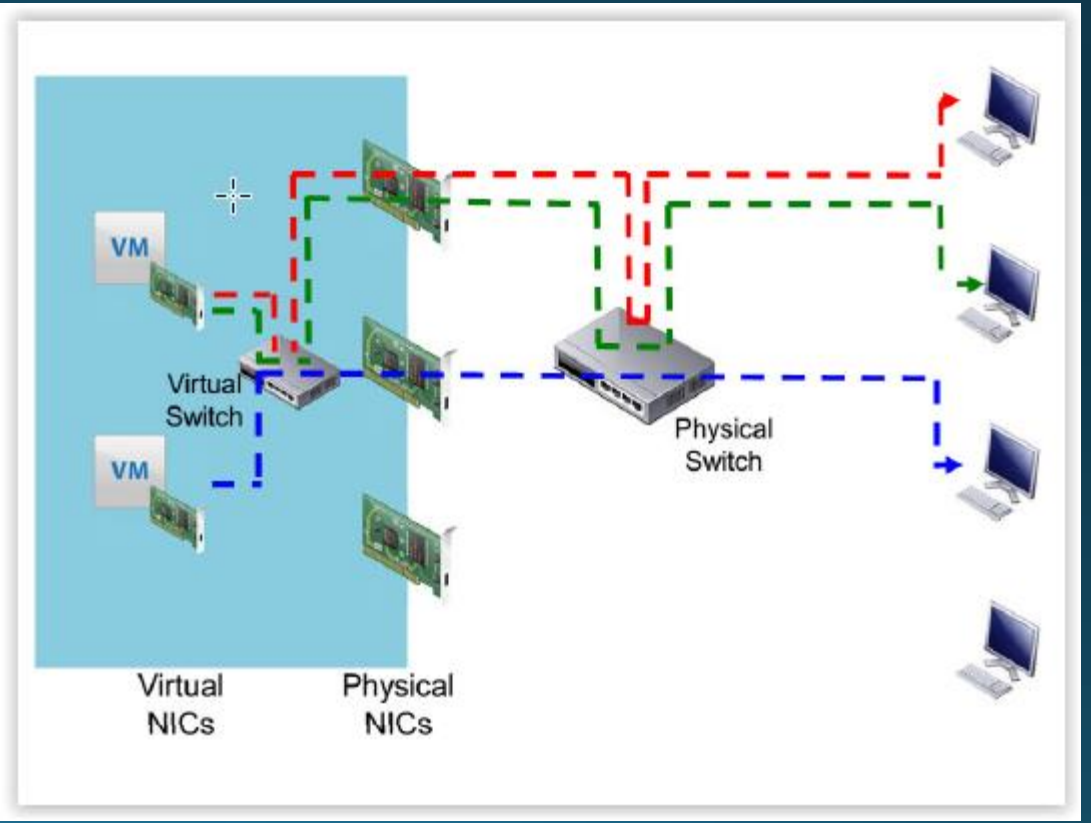

Source: VMware vSphere 6.7 Install, Configure, Manage

Network Load-Balancing (3/3)

- Source and Destination IP Hash
  - a pNIC is selected for each outbound packet, based on its source and destination IP addresses

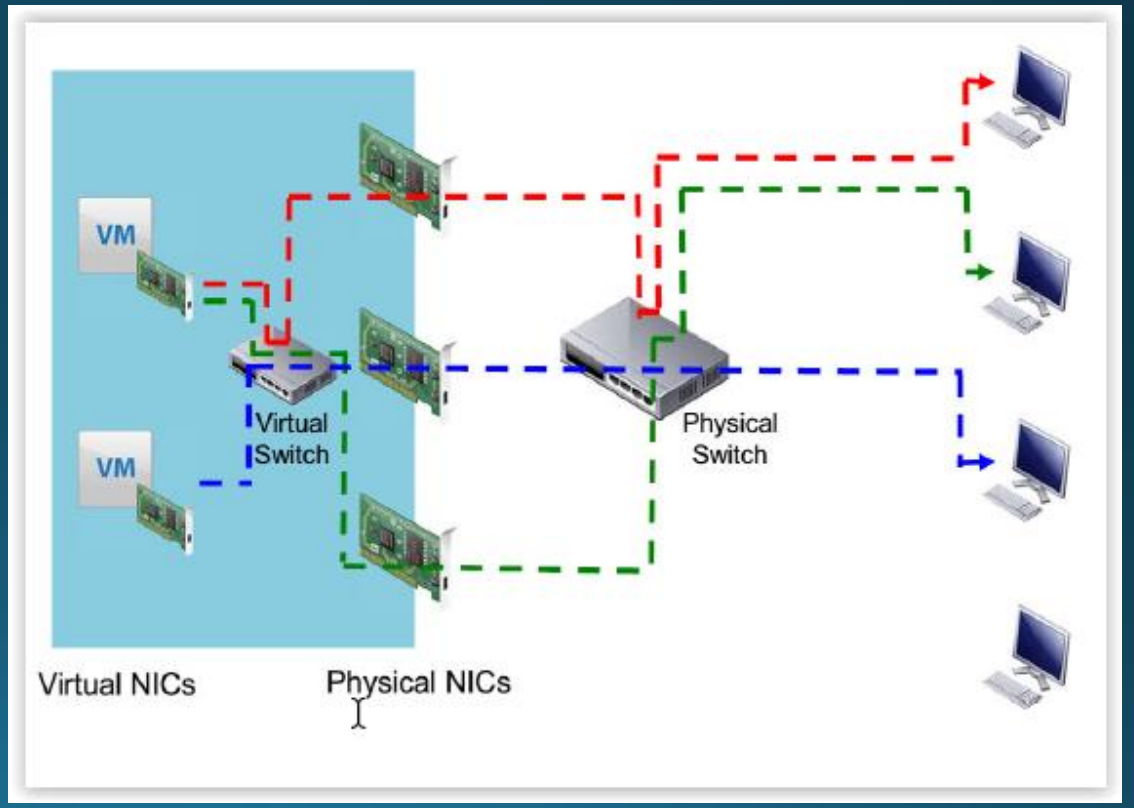

Source: VMware vSphere 6.7 Install, Configure, Manage

Backup and snapshots

- Backup
  - VM-level
- Snapshot
  - VM state in time
  - snaphots are not backups

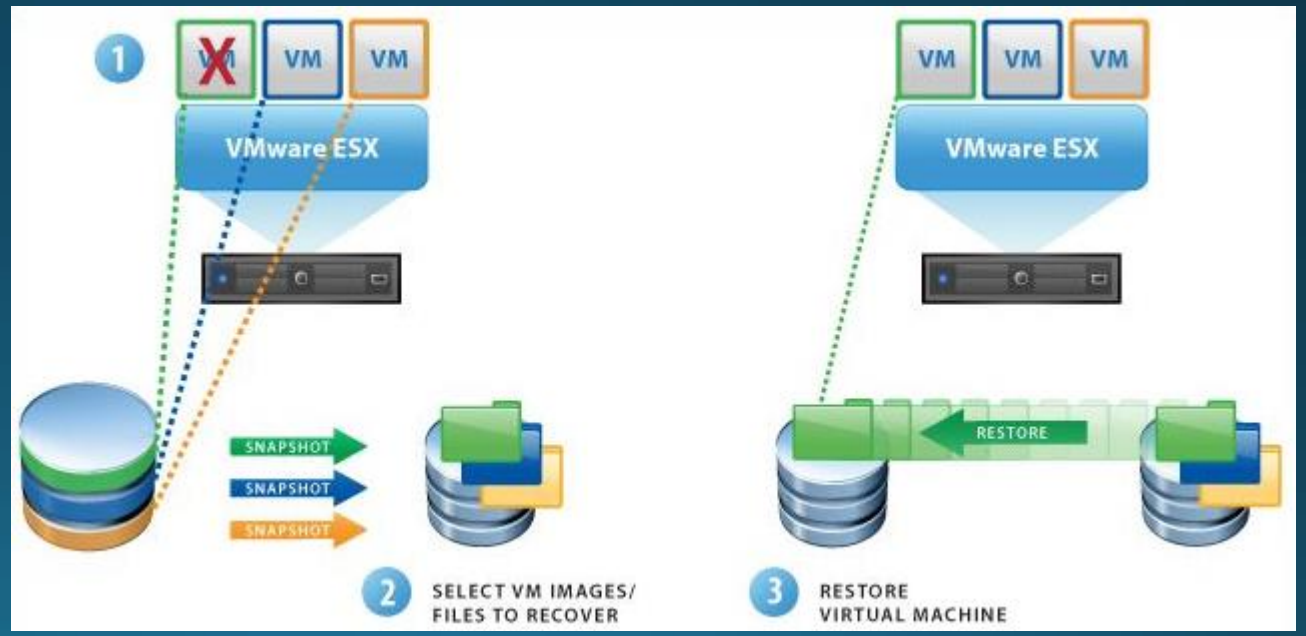

Source: https://petri.com/vmware-data-recovery-backup-and-restore

Virtual Machine Snapshots

- Snaphot data
  - settings
  - disk
  - memory
- Tree-like

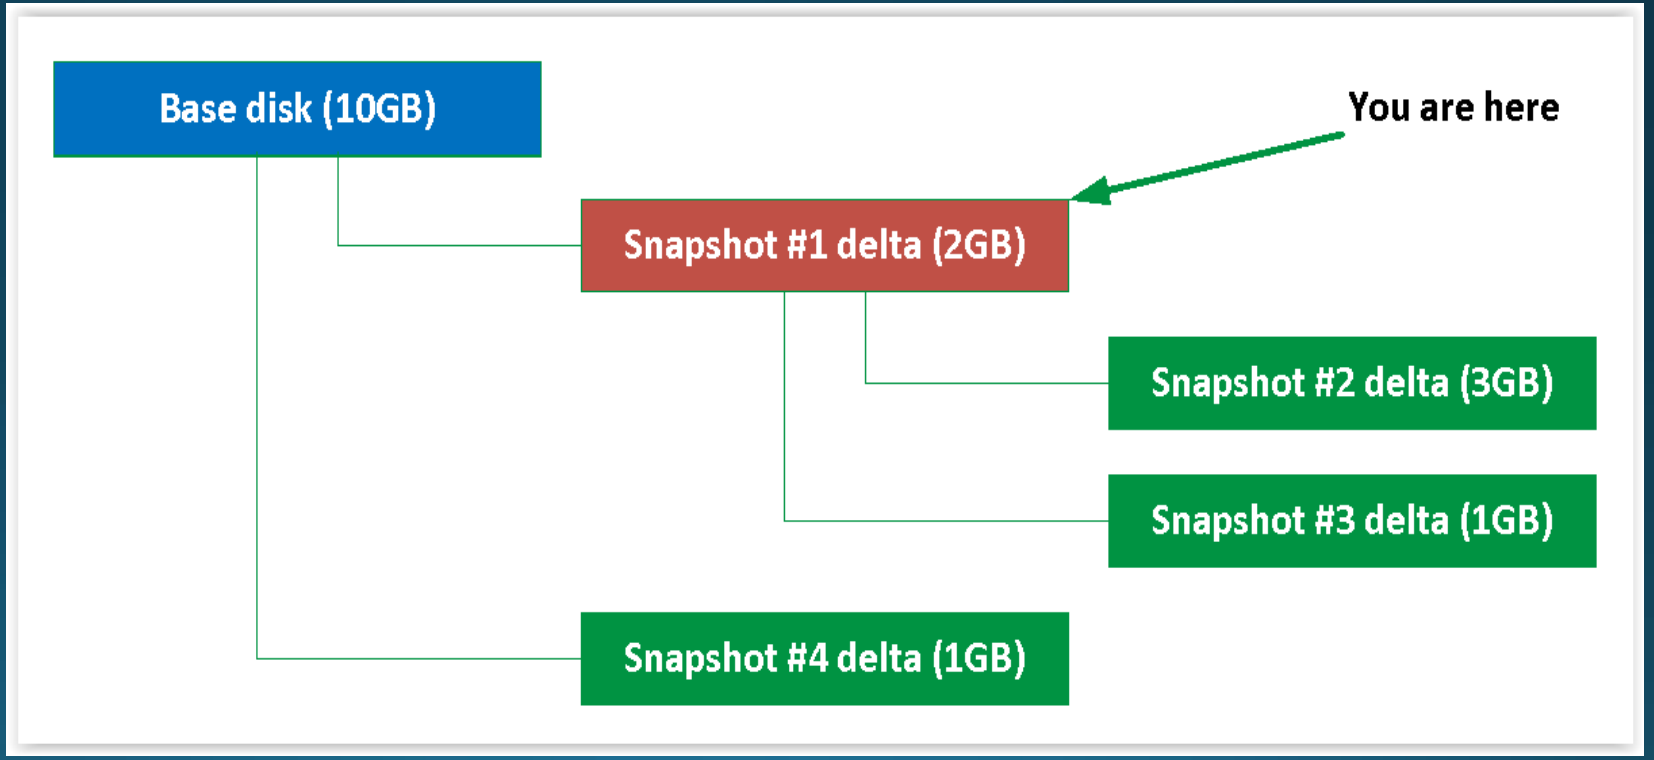

Resource management

- Shares
- Limits
- Reservations

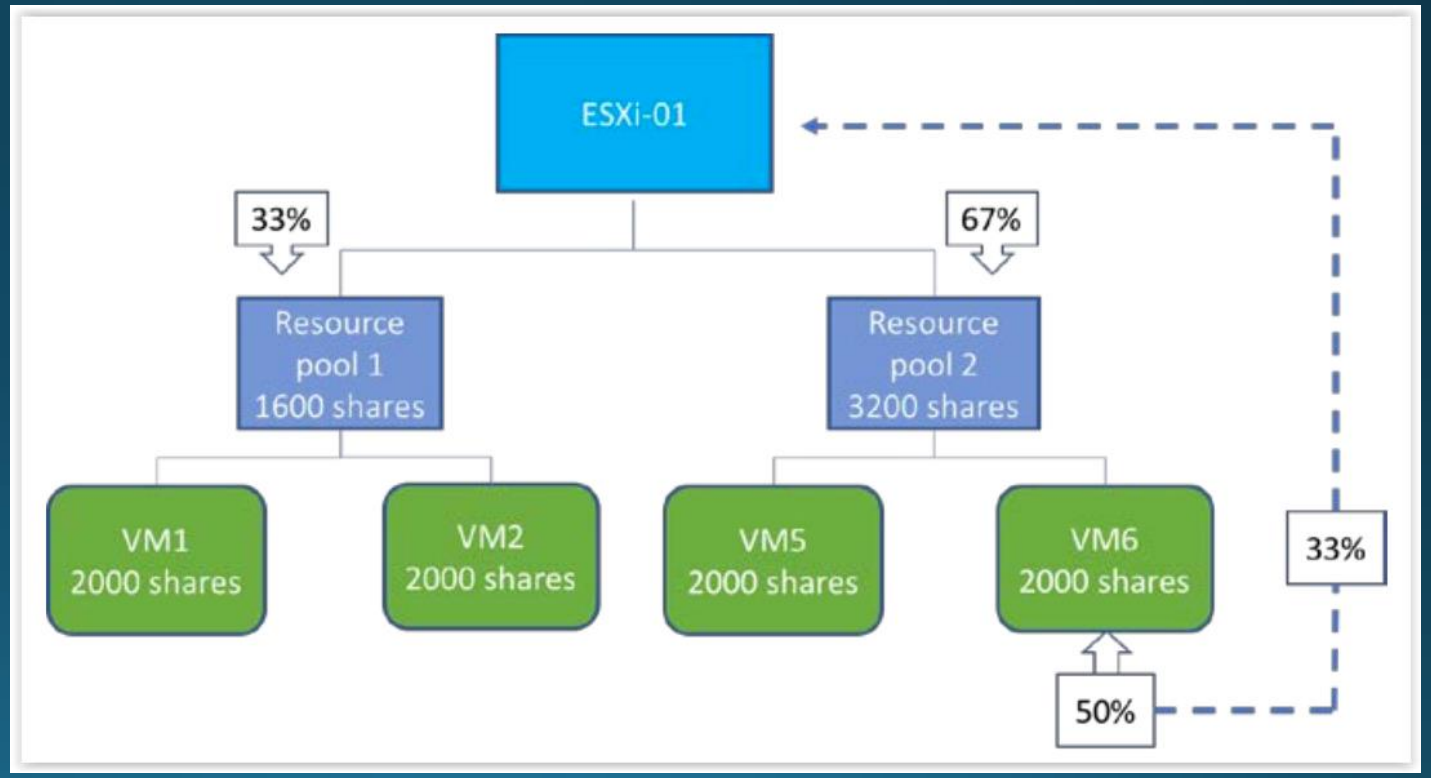

Source: https://www.ntpro.nl/blog/archives/3283-DRS-Cluster-Management-with-Reservation-and-Shares.html

vSphere HA

- Restarts VMs on host failure
- Host selection (DRS)
- Datastore heartbeat

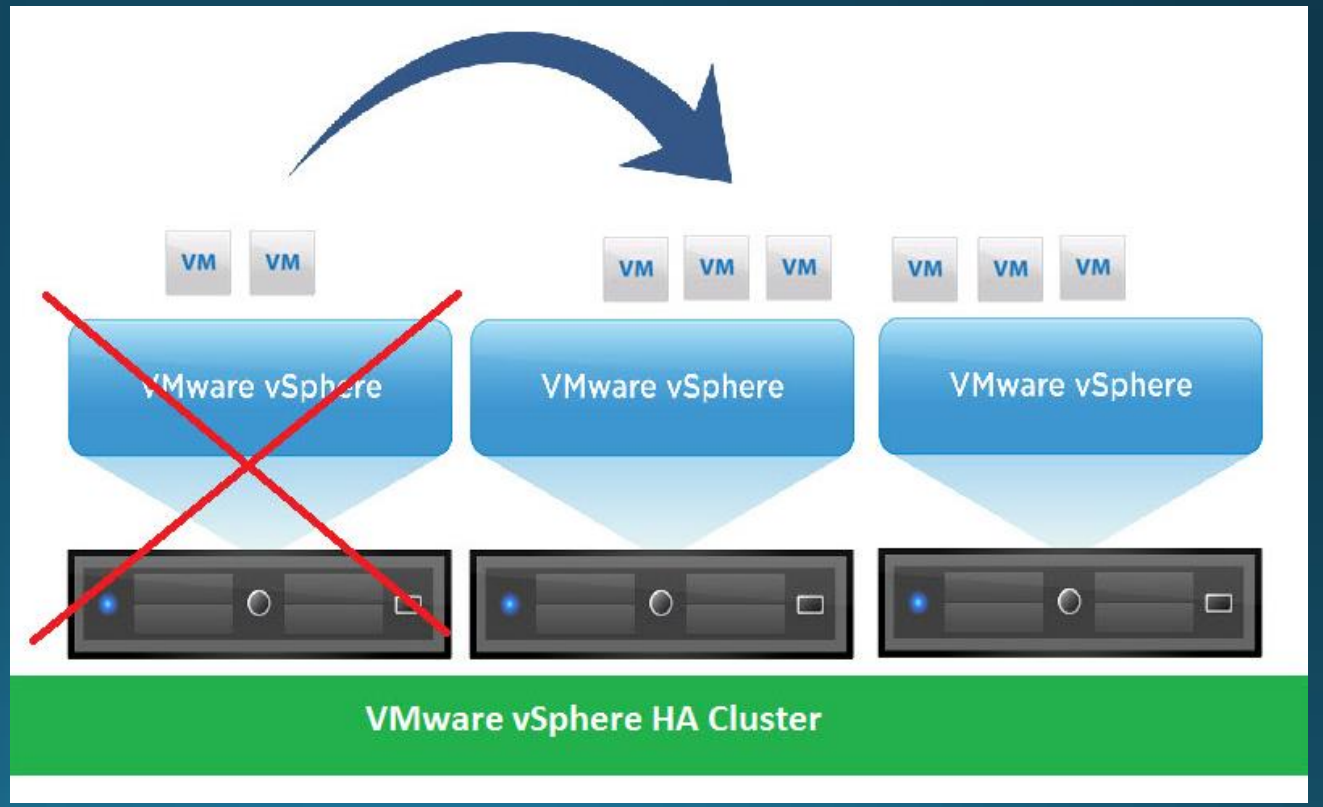

Source: https://masteringvmware.com/what-is-vmware-vsphere-ha/

DRS

- Workload balancing between different esxi hypervisors
- Initial placement and migrations
- VM / host centric
- Metrics include:
  - CPU
  - RAM
  - Network
- Automation options

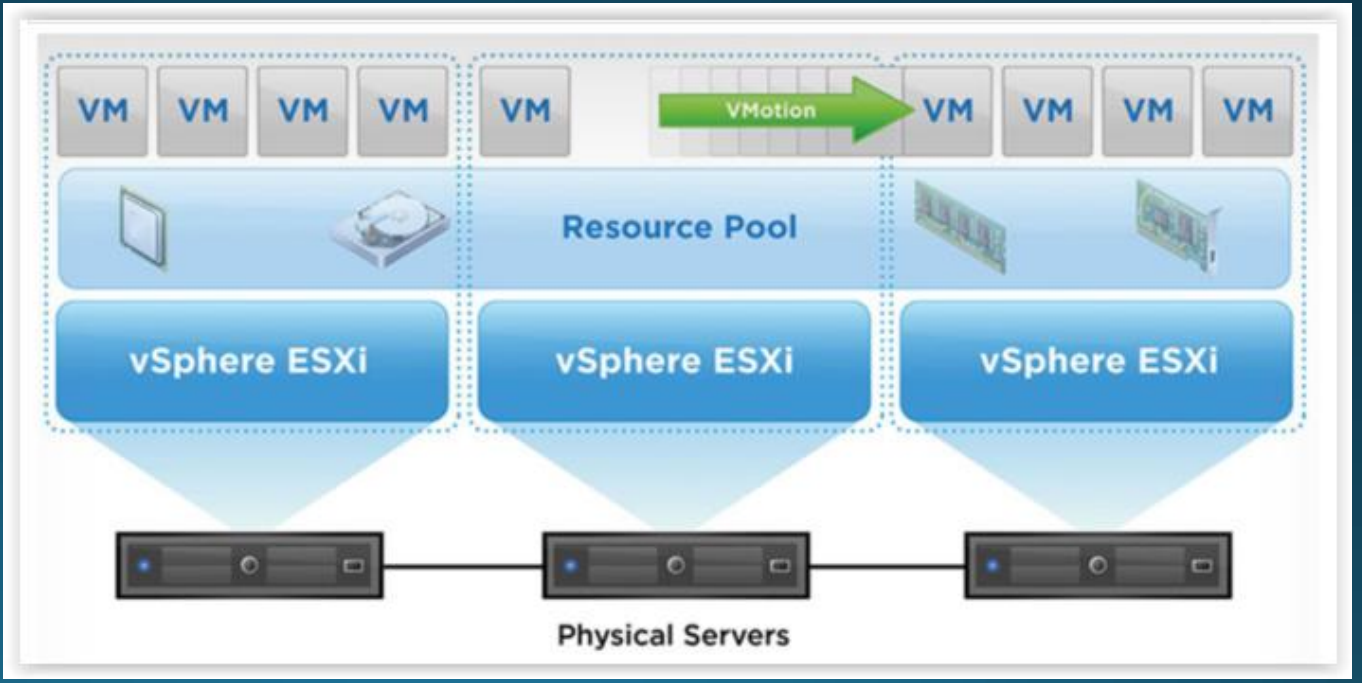

Source: https://cheatsheet.dennyzhang.com/cheatsheet-virtualization-a4

Affinity and anti-affinity rules

- Affinity / anti-affinity
- VM / Host DRS group
- VM-Host / VM-VM
- Preferential / Required

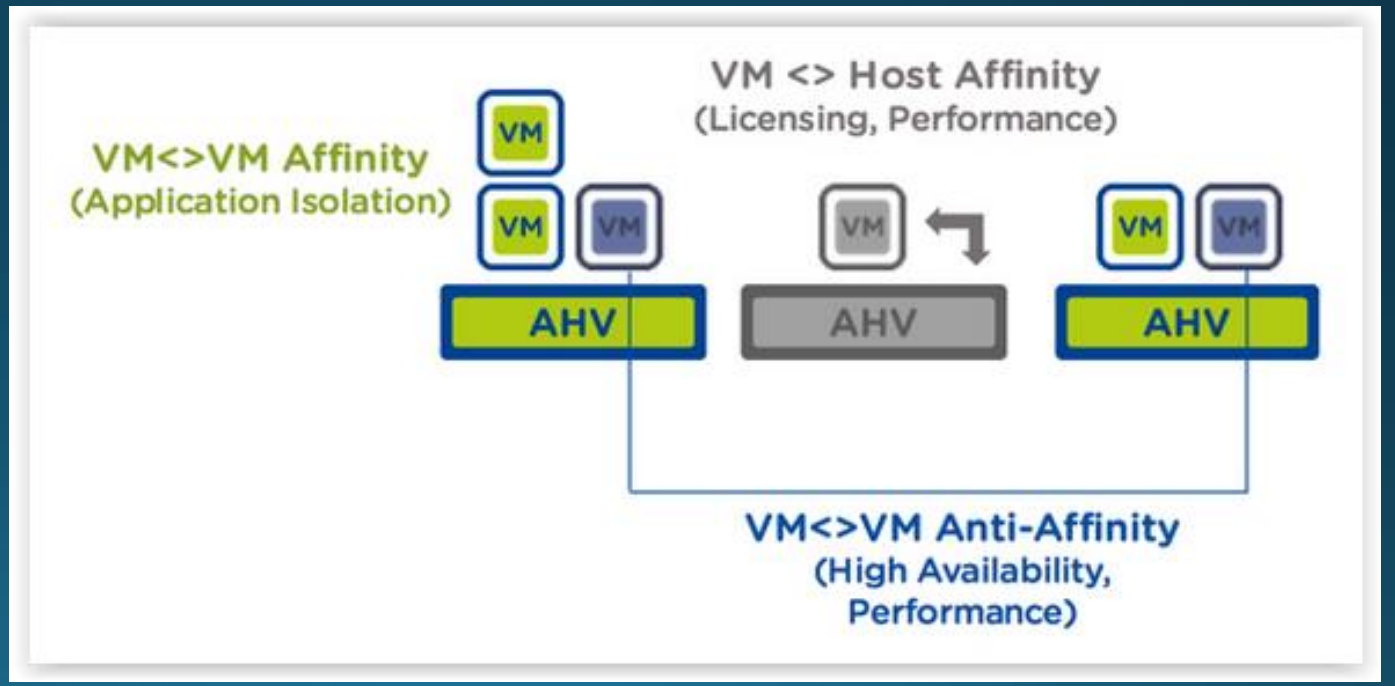

Source: http://www.joshodgers.com/tag/anti-affinity/

Fault Tolerance

- Lossless recovery from outages
  - Data
  - Transactions
  - Connections

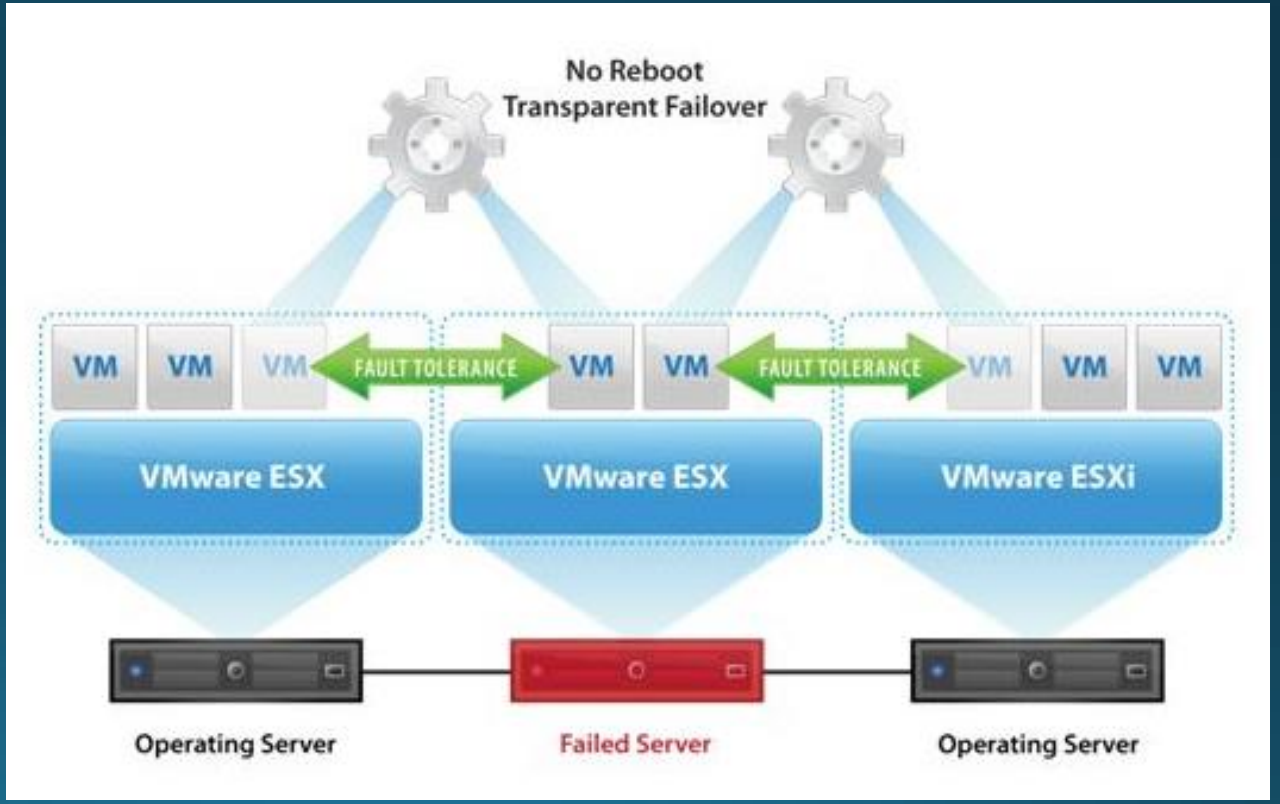

Source: https://www.climb.co.jp/blog_vmware/vmware-4419

Memory management

- Memory reclamation
- Baloon driver
- Guest OS swapping
- Hypervisor swapping

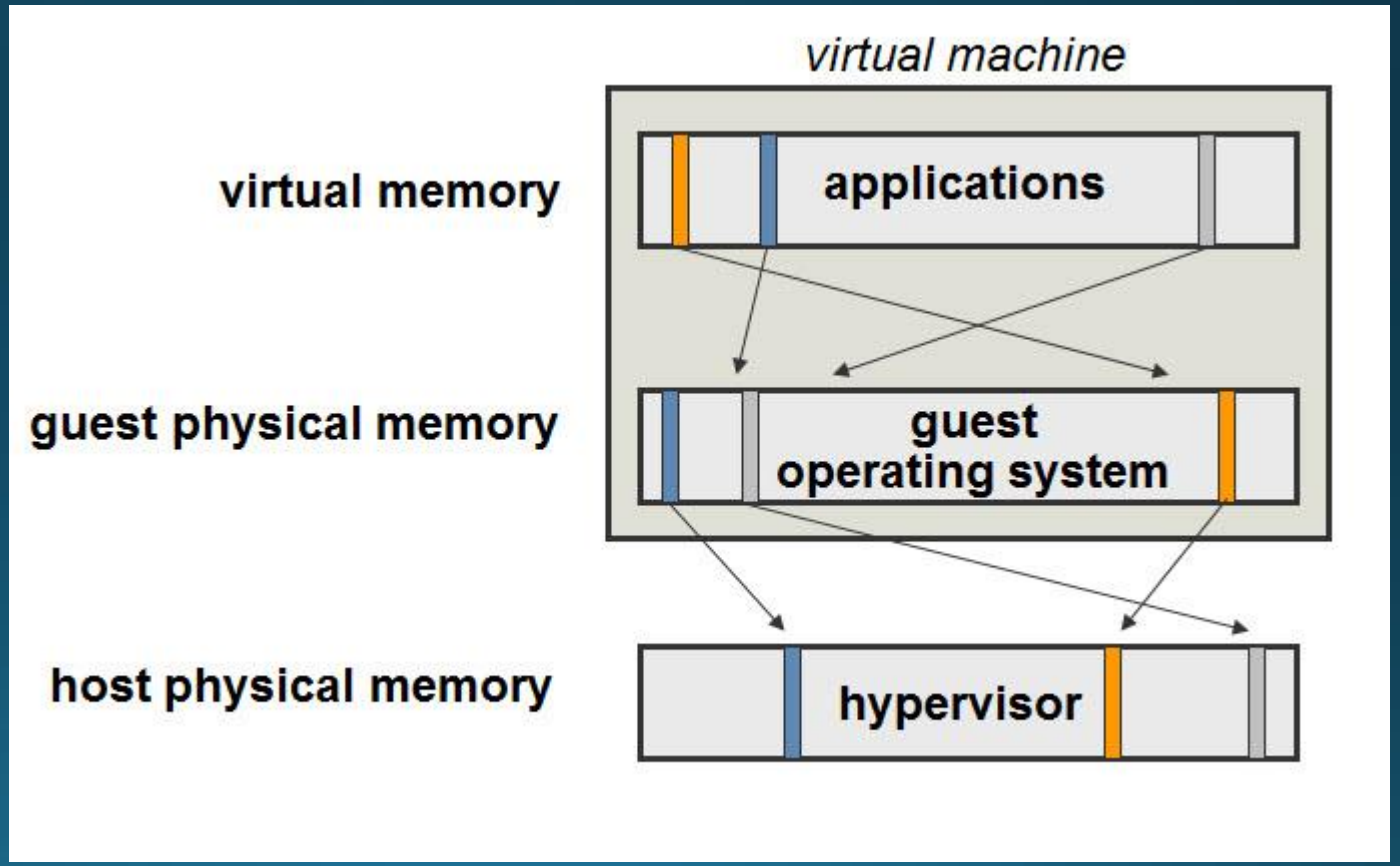

Source: http://www.vfrank.org/2013/09/18/understanding-vmware-ballooning/

Migrating Virtual Machines

Migration Type	Virtual Machine Power State	Change Host or Datastore?	Across vCenter Server Instances?	Shared Storage Required?	CPU Compatibility
Cold	Off	Host or datastore or both	Yes	No	Different CPU families allowed
Suspended	Suspended	Host or datastore or both	Yes	No	Must meet CPU compatibility requirements
vSphere vMotion	On	Host	Yes	Yes	Must meet CPU compatibility requirements
vSphere Storage vMotion	On	Datastore	No	No	N/A
Shared-nothing vSphere vMotion	On	Both	Yes	No	Must meet CPU compatibility requirements

VMware NSX

VMware NSX

- Network virtualization
- Switching
- Routing
- Load balancing
- Firewall

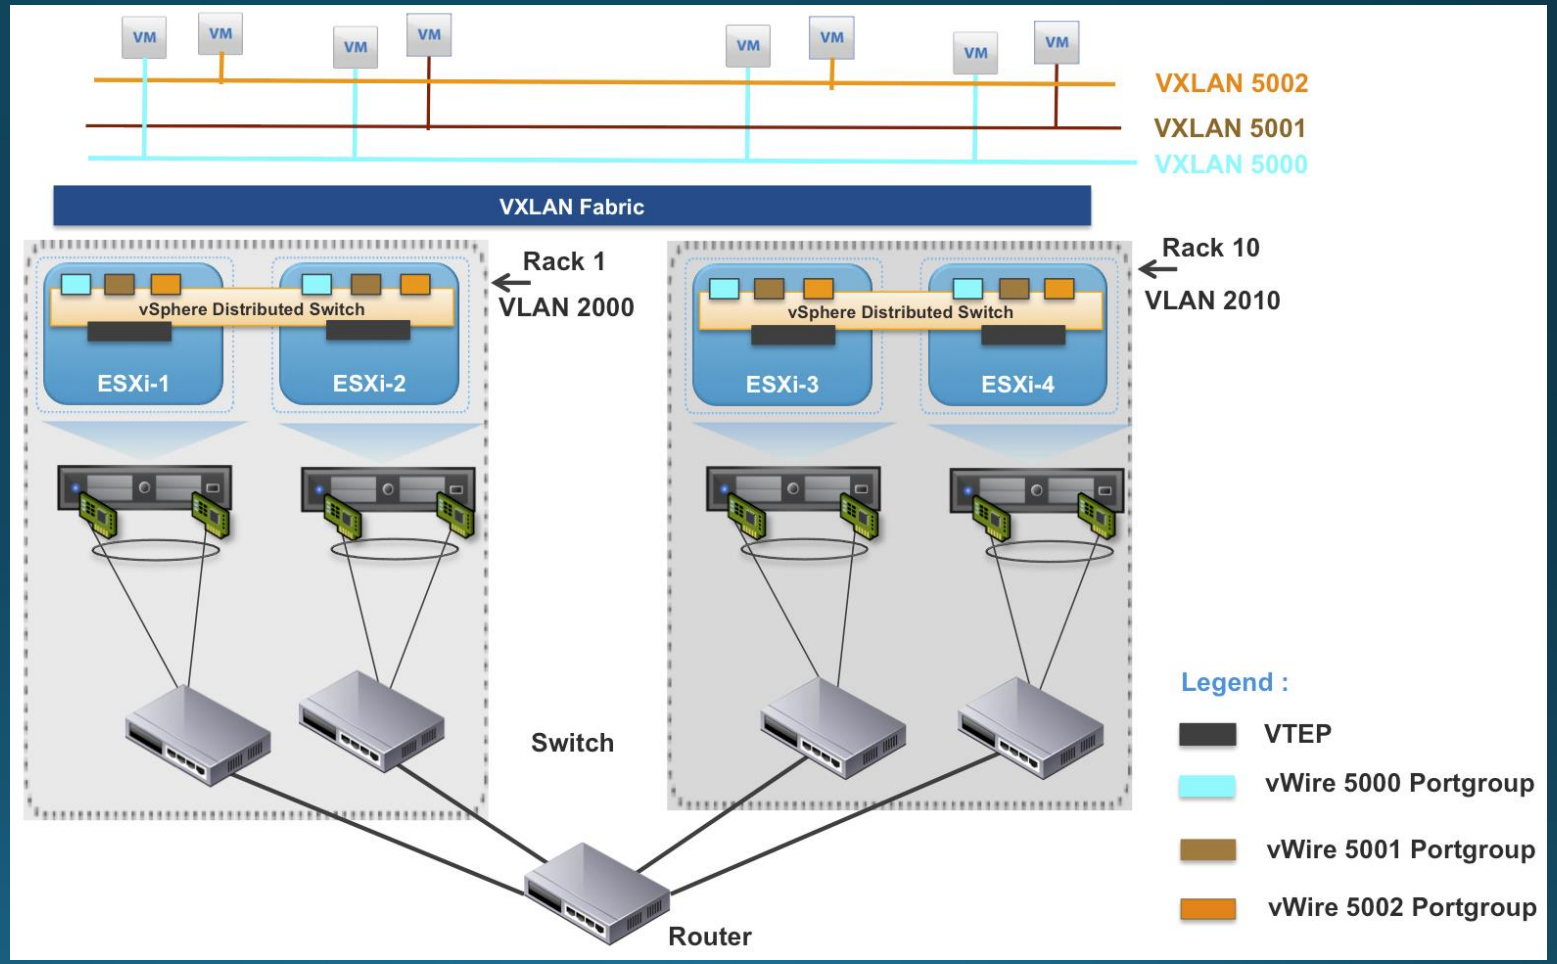

Source: https://blogs.vmware.com/vsphere/2013/01/useful-vxlan-commands-in-esxcli-5-1.html

VXLAN encapsulation

- Encapsulation protocol; extends L2 over L3 by using MAC-in-UDP encapsulation

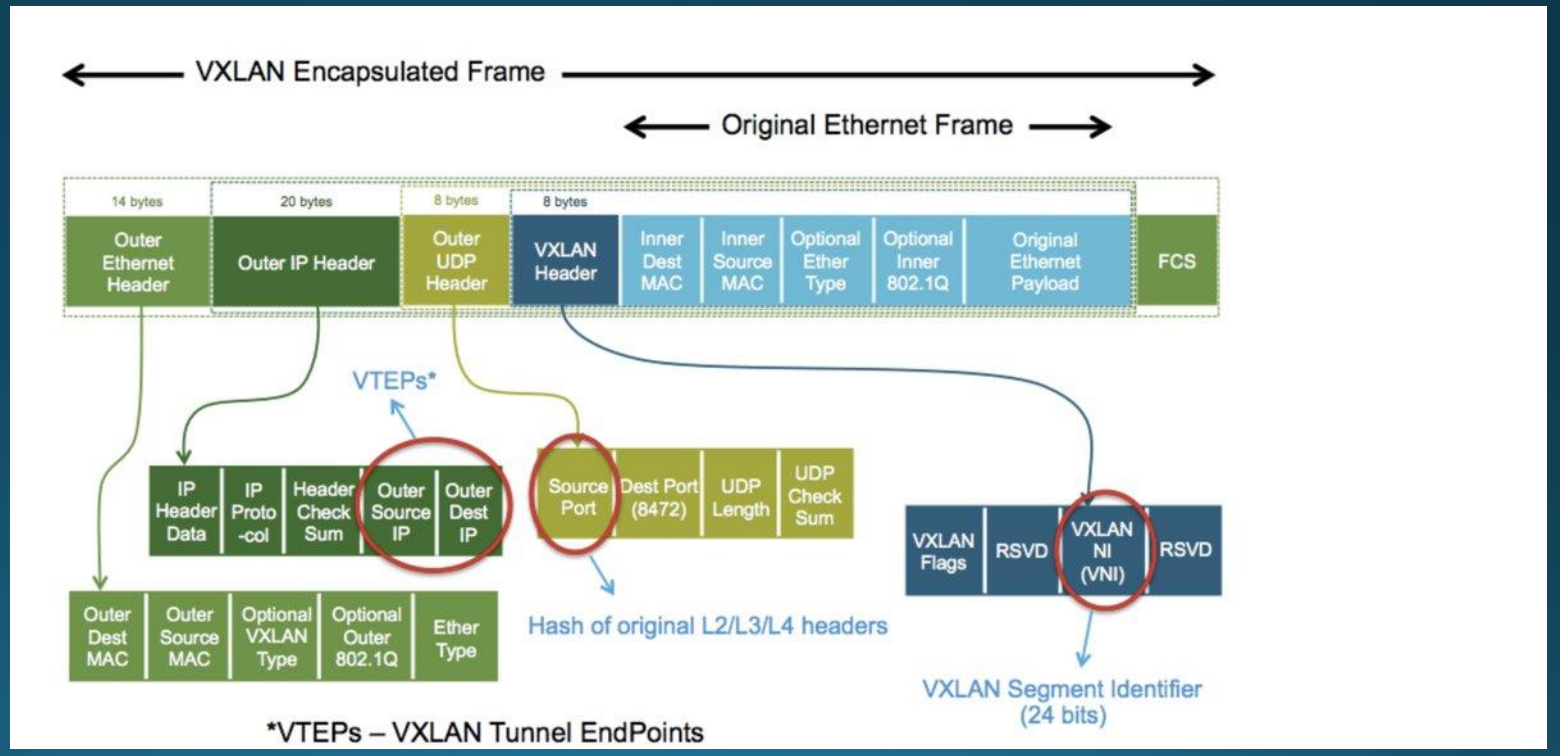

Source: https://virtuallyahead.wordpress.com/2018/05/27/configure-vxlan-for-nsx/

VXLAN packet flow

- Initial state

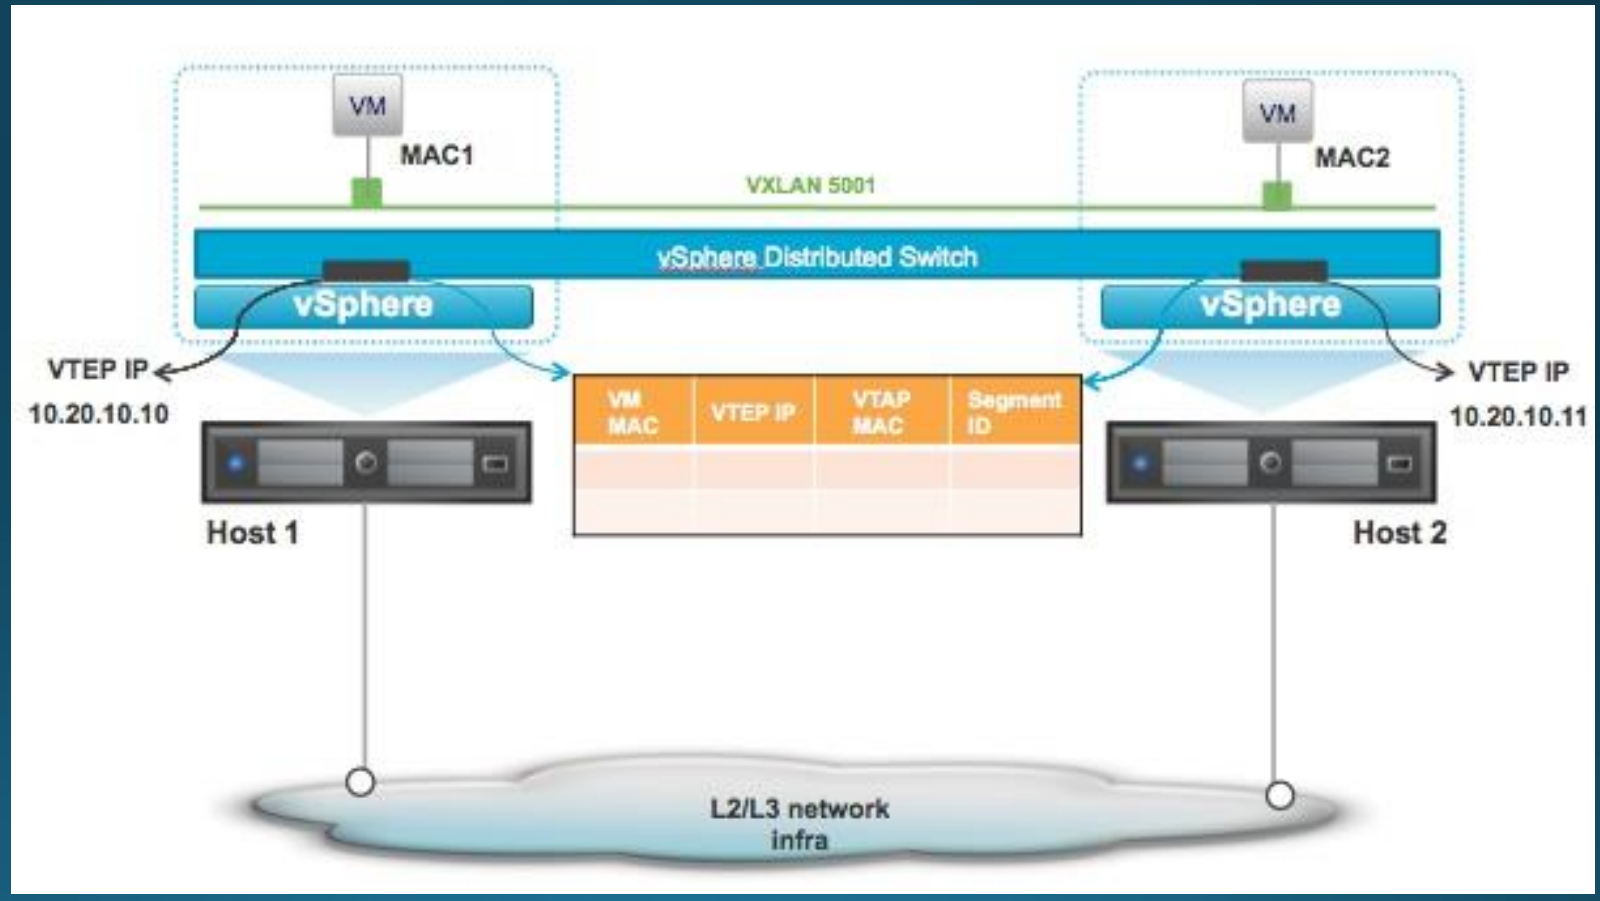

Source: https://blogs.vmware.com/vsphere/2013/05/vxlan-series-how-vtep-learns-and-creates-forwarding-table-part-5.html

VXLAN encapsulation

- ARP request

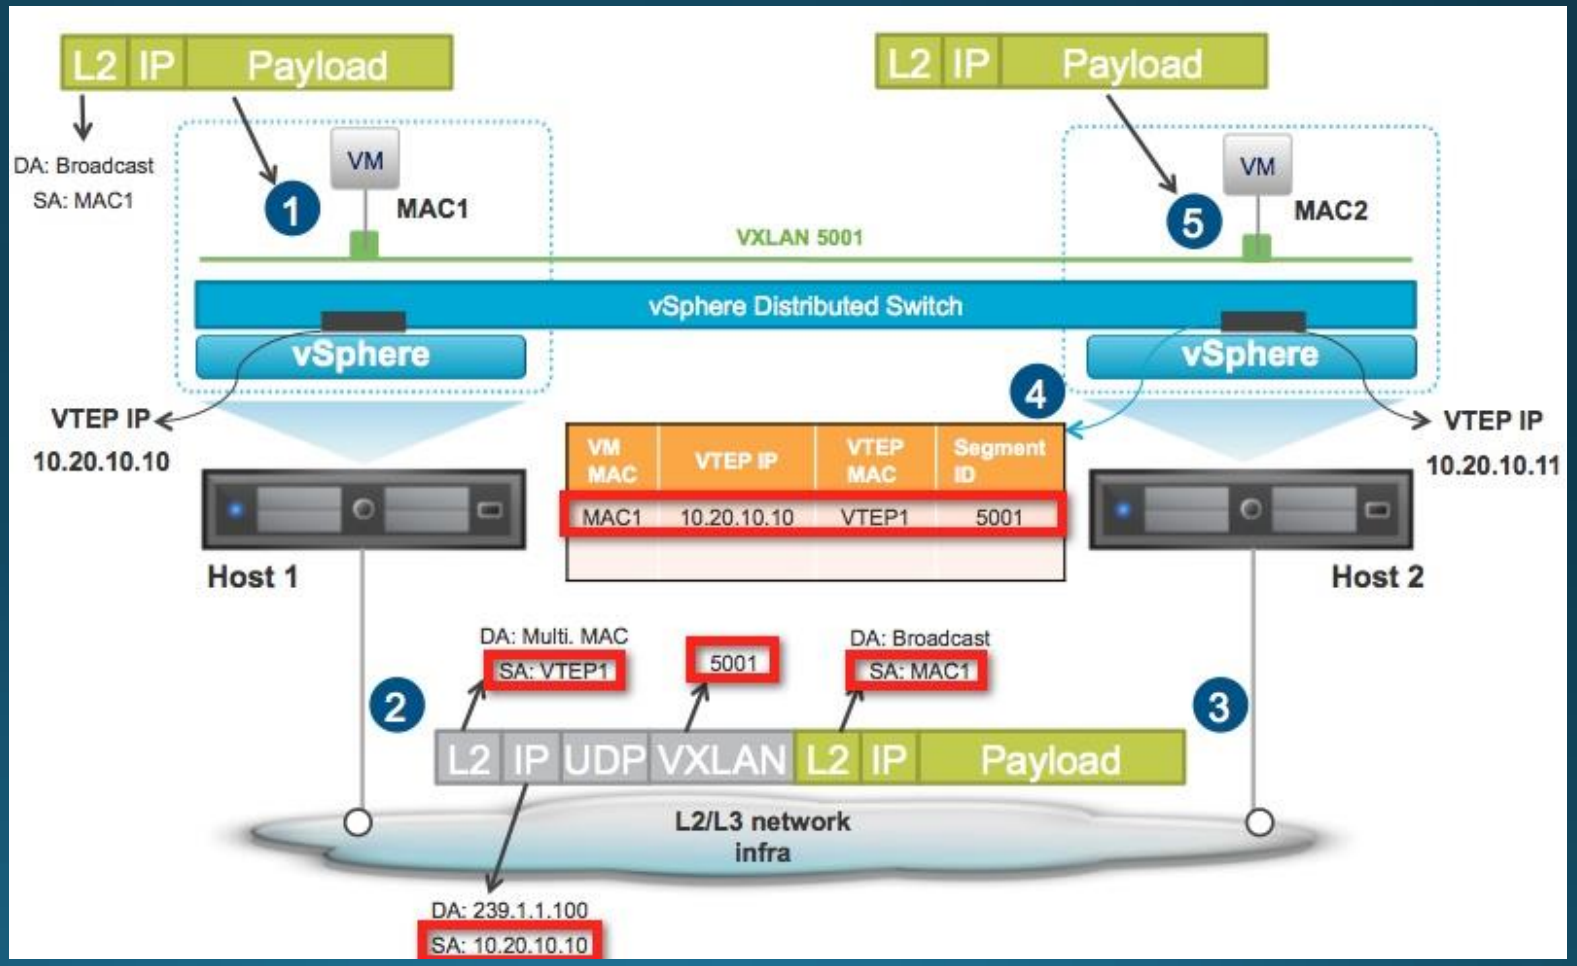

Source: https://blogs.vmware.com/vsphere/2013/05/vxlan-series-how-vtep-learns-and-creates-forwarding-table-part-5.html

VXLAN encapsulation

- ARP reply

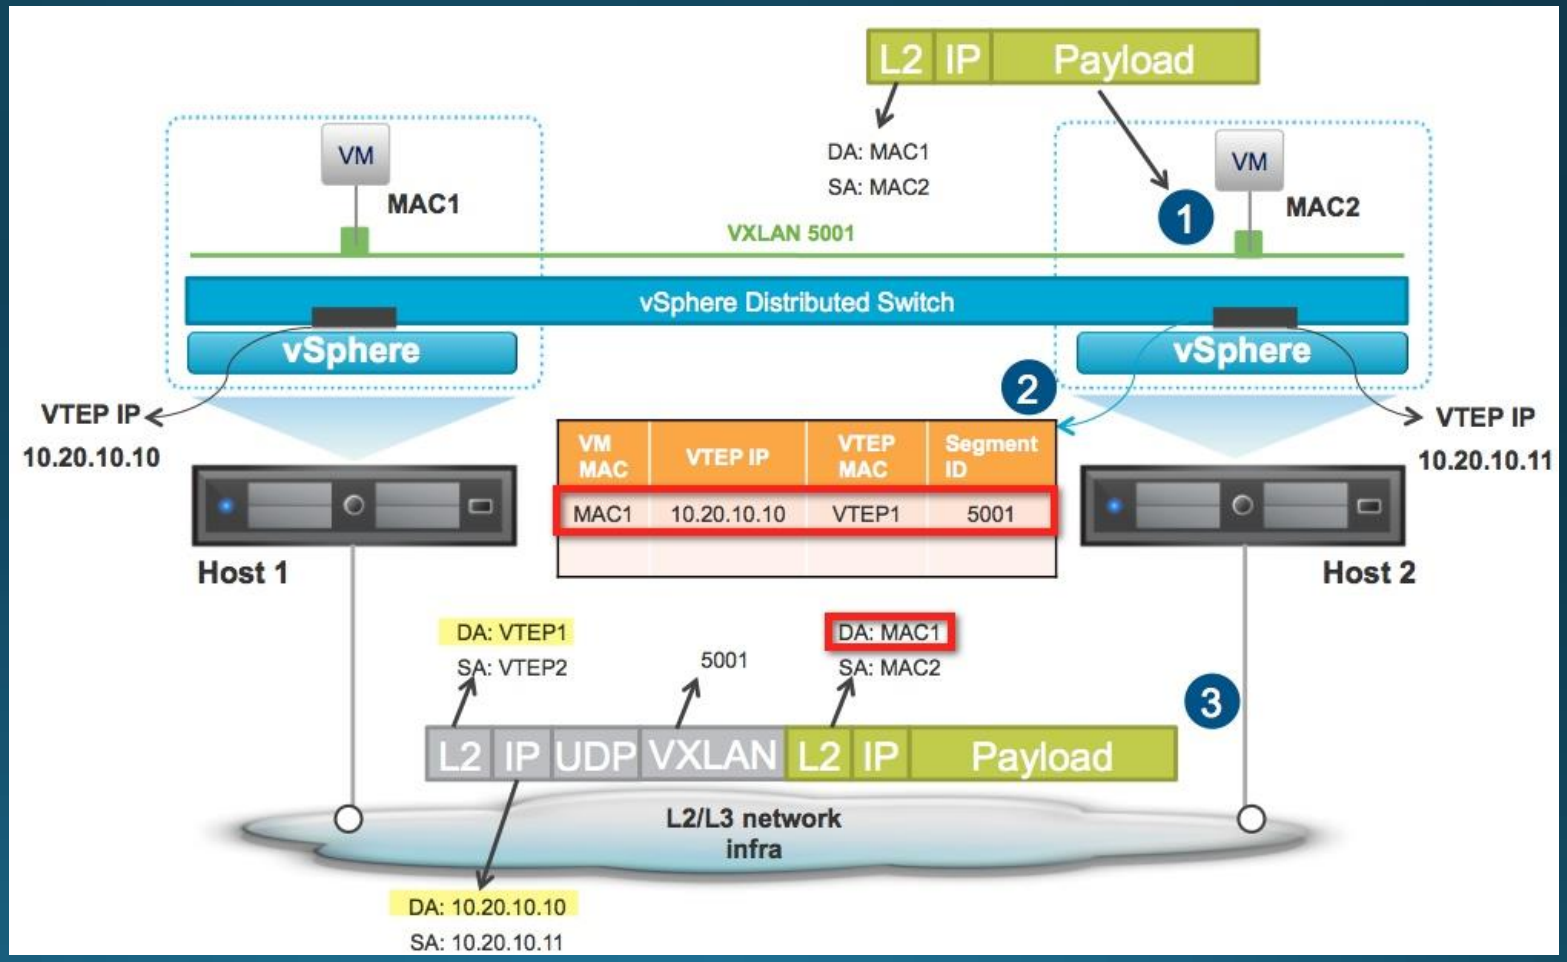

Source: https://blogs.vmware.com/vsphere/2013/05/vxlan-series-how-vtep-learns-and-creates-forwarding-table-part-5.html

VXLAN encapsulation

- Packet delivery

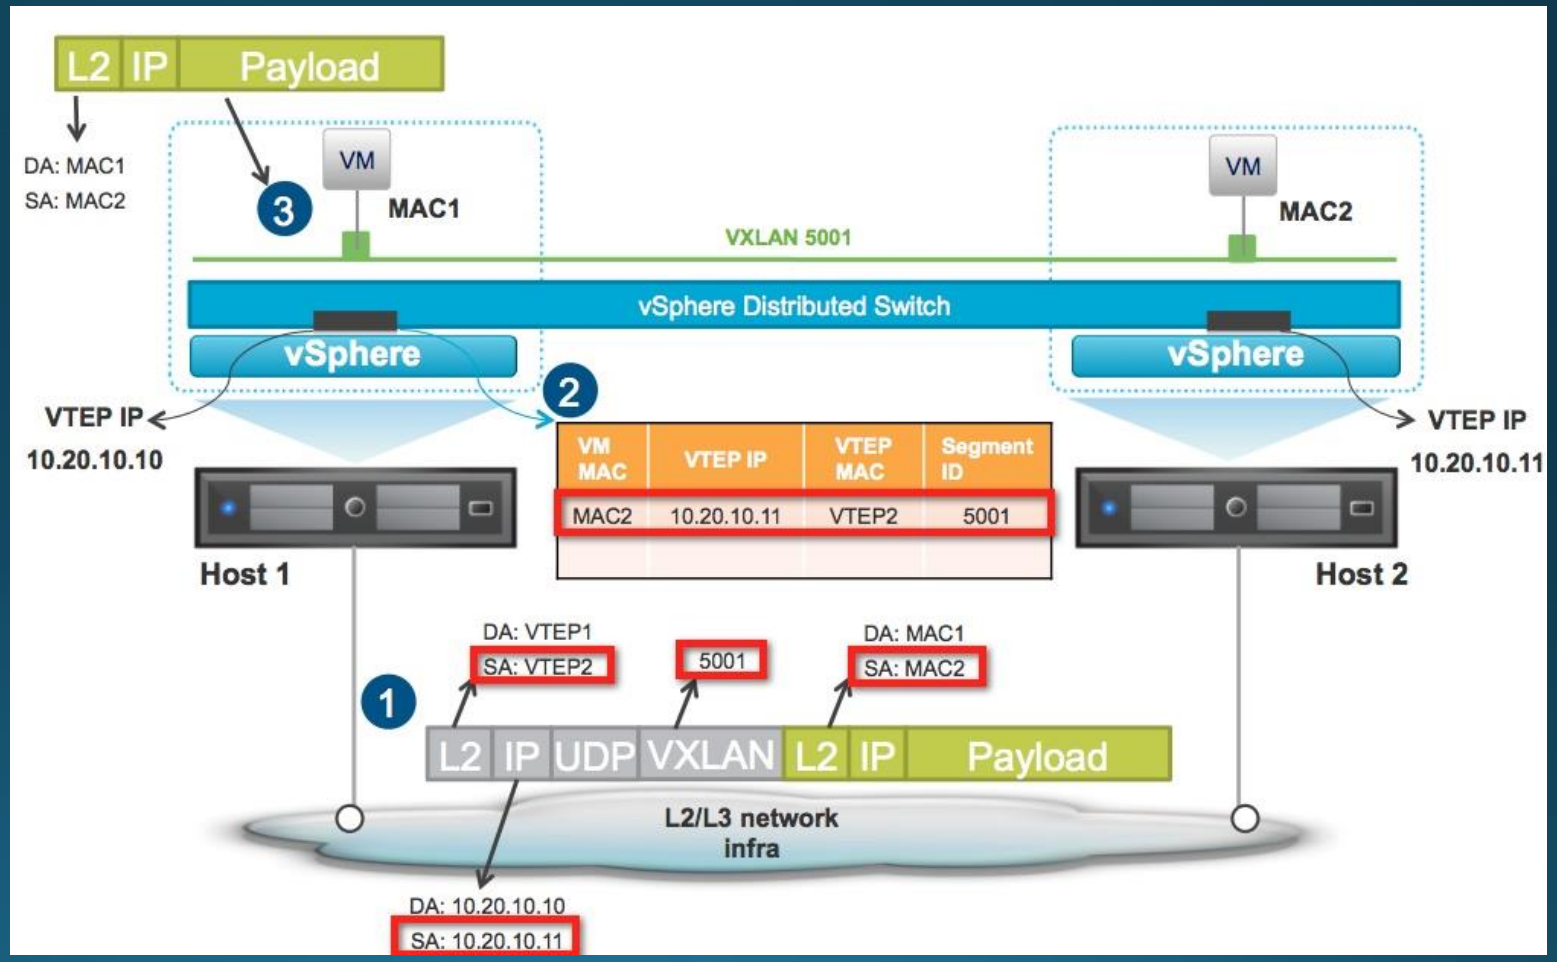

Source: https://blogs.vmware.com/vsphere/2013/05/vxlan-series-how-vtep-learns-and-creates-forwarding-table-part-5.html

DR Replication

- Intercept source VM write I/O-s
- Replicate I/O-s to DR site
- Store I/O-s to journalised log