


Book of vSphere - vSphere Architecture

[PDF generated December 08 2021. For all recent updates please see the Nutanix Bible releases notes located at https://nutanixbible.com/release_notes.html. Disclaimer: Downloaded PDFs may not always contain the latest information.]

Node Architecture

In ESXi deployments, the Controller VM (CVM) runs as a VM and disks are presented using VMDirectPath I/O. This allows the full PCI controller (and attached devices) to be passed through directly to the CVM and bypass the hypervisor.



ESXi Node Architecture

Configuration Maximums and Scalability

The following configuration maximums and scalability limits are applicable:

- Maximum cluster size:64
- Maximum vCPUs per VM:128
- Maximum memory per VM:4TB
- Maximum virtual disk size:62TB
- Maximum VMs per host:1,024
- Maximum VMs per cluster:8,000 (2,048 per datastore if HA is enabled)

NOTE: As of vSphere 6.0

Pro tip

When doing benchmarking on ESXi hosts, always test with the ESXi host power policy set to 'High performance'. This will disable and P- and C- states and will make sure the test results aren't artificially limited.

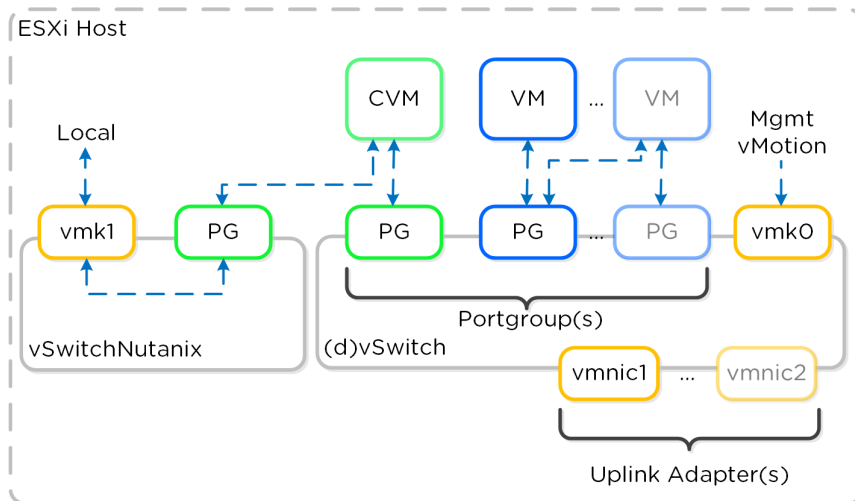
Networking

Each ESXi host has a local vSwitch which is used for intra-host communication between the Nutanix CVM and host. For external communication and VMs a standard vSwitch (default) or dvSwitch is leveraged.

The local vSwitch (vSwitchNutanix) is for local communication between the Nutanix CVM and ESXi host. The host has a vmkernel interface on this vSwitch (vmk1 - 192.168.5.1) and the CVM has an interface bound to a port group on this internal switch (svm-iscsi-pg - 192.168.5.2). This is the primary storage communication path.

The external vSwitch can be a standard vSwitch or a dvSwitch. This will host the external interfaces for the ESXi host and CVM as well as the port groups leveraged by VMs on the host. The external vmkernel interface is leveraged for host management, vMotion, etc. The external CVM interface is used for communication to other Nutanix CVMs. As many port groups can be created as required assuming the VLANs are enabled on the trunk.

The following figure shows a conceptual diagram of the vSwitch architecture:



ESXi vSwitch Network Overview

Uplink and Teaming policy

It is recommended to have dual ToR switches and uplinks across both switches for switch HA. By default the system will have uplink interfaces in active/passive mode. For upstream switch architectures that are capable of having active/active uplink interfaces (e.g. vPC, MLAG, etc.) that can be leveraged for additional network throughput.