Networking

Maxs

Per host: 1GbE VMNICs = 2-32 dependent on HW   10GbE VMNICs = 8 (or 6x10GbE & 4x1GbE)
PCI VMDirectPath devices = 8   Switches (vSS/vDS/VEM) = 248/16/1 ??????
vSS/vDS ports = 4096   Active ports (vSS/vDS) = 1016
Per vCenter: vDS switches = 32   vDS port groups = 5,000(256 ephemeral)   vDS ports = 30,000
Per switch: Hosts (per vDS) = 35   vSS port groups = 2   vSS switch ports = 4,088

Terminology: VMNICs - logical name for physical server NICs   vNICs - virtual NICs assigned to VMs
vSS - virtual Standard Switch   vDS - virtual Distributed Switch   dvPort(Group) - port (group) on a vDS
dvUplink - uplink VMNICs on a vDS   Network vMotion - tracking of VM's network state on a vDS

Shell Commands

`--help` for namespaces & commands relative to location

List VMNICs:	`esxcli network nic list`
List vSwitches:	`esxcli network vswitch standard list`
List vDS:	`esxcli network vswitch dvs vmware list`
List vSwitch Port Groups:	`esxcli network vswitch standard portgroup list`
List VMkernel ports:	`esxcli network ip interface list`
List VMkernel interfaces:	`esxcli network ip interface ipv4 get`
List VMkernel Default Gateway:	`esxcfg-route`
List hostname:	`esxcli system hostname get`
List DNS servers:	`esxcli network ip dns server list`
List DNS search domain:	`esxcli network ip dns search list`

`esxcli` does not support configuring vDS dvPorts and dvUplinks: use `esxcfg-vmknic` with `dvs-name`, `dvport-id` & `esxcfg-vswitch` with `dvp-uplink`, `dvp` options

Ethernet tagging: • EST (External Switch Tagging) - Default. No trunking required. 1-1 relationship from VMNICs to physical (access) switch ports. Each VMNIC only sees 1 subnet. VLAN ID of 0 or blank.
• VST (Virtual Switch Tagging) - Commonly used. VMNICs connected to a vSwitch can span several VLANs. Each Port Group has a VLAN ID of 1-4094. Set the VLAN ID to blank to use Native VLAN.
• VGT (Virtual Guest Tagging) - Rarely used. Install 802.1Q trunking driver software in VMs, vSwitch keeps tags given by VMs. VLAN ID of 4095 on vSS, VLAN policy on vDS. Avoid VLAN ID of 1 - native Cisco VLAN ID. Use VLAN 4095 with promiscuous mode to sniff other port groups (IDS/packet sniffer)

Jumbo frames: MTU > 1500 up to 9000 bytes. Enable per vSS/vDS. vNIC must be vmxnet2/3 or e1000

Link Discovery: vSS supports CDP (Cisco Discovery Protocol), vDS supports CDP or LLDP (Link Layer Discovery Protocol - 802.1AB). *Listen* (default), *Advertise* or *Both*.

PVLAN (Private VLAN): extension to VLAN standard to add further segmentation. Can reduce IP address wastage & solve VLAN ID limits. Not encapsulated. Primary PVLAN - Original VLAN divided into smaller groups. Secondary PVLAN - exists only within primary, has specific VLAN ID. Types: Primary is *Promiscuous* - connect with all VMs in primary. Secondary are *Community* - connect to themselves & VMs on promiscuous, or *Isolated* - connect with VMs on promiscuous.

NetFlow: Sends IP traffic records to collector for analysis. Traffic is intrahost, interhost or VM-physical

Port Mirror: Mirror ports intrahost or interhost. Cisco’s term is SPAN (Switch Port Analyzer).

NIOC (Network IO Control): prioritize egress traffic by type via dvUplink shares (low/normal/high-25/50/100) & host limits. Network Resource Pools: FT, iSCSI (not HW iSCSI), vMotion, Mgt, VR (SRM replication), NFS, VM, Custom (user defined). Supports 802.1p QoS priority tagging at MAC level.

TSO (TCP Segmentation Offload): enabled by default on VMkernel ports, allows very large frames (up to 64KB), even with smaller MTU. To enable VMs, use at least enhanced vmxnet vNIC.

NetQueue: enabled by default, allows certain VMNICs to spread processing across CPUs to improve ingress performance.

vSS & vDS options: Options can be overridden on vSS & dvPortGroups. Individual dvPorts can override options, but dvPortGroups can disallow overrides.

Options nomenclature: • vSS - *Properties* • vDS/dvUplinks - *Settings* • dvPortGroups - *Policies*.

General • *Number of uplinks* (vDS only) • *Number of ports* - vSS default - 120, dvPortGroup - 128 • *Port Binding* (dvPortGroups only): *Static* - when initially connected, *Dynamic* - when connected/powered-on, *Ephemeral* - no binding. Host can assign port if vCenter is down. • *MTU* - default 1500 (cannot override on Port Groups) see Jumbo Frames below • *Discovery Protocol* (vDS only) see Link Discovery below • *VLAN ID* (vSS PGs only)

Network Adapters (vDS only) • Host to dvUplinks mapping

Private VLAN (vDS only) • Primary to Secondary mapping

Netflow (vDS only) • Collector *IP Address* & Port • *vDS IP Address* - so collector interacts with vDS not hosts • *Active flow export timeout* • *Idle flow export timeout* • *Sampling rate* - 1 packet collected per sampling rate • *Process internal flows only* - just intrahost traffic.

Port Mirroring (vDS only) Add session to mirror • *Allow normal IO on destination ports* - port to receive normal IO as well as mirrored traffic • *Encapsulate VLAN* - create VLAN ID to encapsulate all frames if destination is an uplink port. If *Preserve original VLAN* unselected then if VLAN is present then it’s replaced not encapsulated • *Mirrored packet length* - limits size of mirrored frames • select *Ingress/Egress* • select *Port IDs* or *Uplink source & destination*.

Security • *Promiscuous mode* (default Reject) - only listens to traffic destined for its MAC address.
• *MAC Address Changes* (default Accept) - accepts inbound frames when VM changes MAC address.
• *Forged Transmits* (default Accept) - won't drop outbound frames if source MAC address is different

Traffic Shaping • *Status* (default Disabled) • *Average Bandwidth* (Kbps) • *Peak Bandwidth* (Kbps) • *Burst size* (KB). vSS can shape outbound traffic, vDS can shape traffic in & out (Ingress/Egress)

VLAN (dvPortGroup only) • *None* - access port • *VLAN* - set ID • *Trunk range* - restrict IDs on trunked links • *PVLAN*

Teaming & Failover • Load Balancing - spreads outbound traffic from vNICs across VMNICs/dvUplinks, incoming traffic is load balanced by physical switch. *Originating port ID* (default) - uses uplink based on where traffic entered. *ip hash* - based on source & destination IP address of each packet (use if physical switch ports are etherchannel). *Source MAC hash* - based on source MAC address. *Route based on physical NIC load* (vDS only) - based on current loads on dvUplinks. *Use explicit failover order* - uses first active uplink in list. • Network Failover Detection - *Link status only* (default) - detects cable pulls & switch power failures, not misconfigurations. *Beacon Probing* - can also detect some misconfiguration, but don't use with IP-hash load balancing & not supported with VGT. • Notify Switches - *No* or *Yes* (default) updates lookup tables. Disable for MS NLB in unicast mode. • Failback - *No* or *Yes* (default) uplink returns after recovering from failure. • Failover order - *Active*, *Standby* or *Unused* - Don't use standby uplinks with IP-hash load balancing.

Resource Allocation (dvPortGroup only) • select Network Resource Pool (see NIOC)

Monitoring (dvPortGroup & dvUplink only) - *Enable* or *Disable (default)* NetFlow (options on vDS)

Miscellaneous (dvPortGroup & dvUplink) • Port blocking - *No* (default) or *Yes* - shut down all ports.

Links: Troubleshooting Networking - http://communities.vmware.com/docs/DOC-9876