Network Architecture in

Dell Active System 200 & 800

Jaiwant Virk: Networking Data Center Mktg

Mohan Mullapudi: Virtualization Solution Engineering

Table of Contents

Introduction

Dell™ Active Infrastructure is a family of converged infrastructure solutions that combine servers, storage, networking, and infrastructure management into an integrated and optimized system that provides general purpose virtualized resource pools. Active Infrastructure leverages Dell innovations including unified management (Active System Manager), converged LAN/SAN fabrics, and modular server architecture for the ultimate converged infrastructure solution. Active Infrastructure helps IT rapidly respond to dynamic business demands, maximize data center efficiency, and strengthen IT service quality.

The Active System 200 and 800 solutions, member of Dell Active Infrastructure family, are a converged infrastructure solution that has been designed and validated by Dell Engineering. The solutions are available to be racked, cabled, and delivered to your site to speed deployment. Dell Services will deploy and configure the solution tailored for business needs, so that the solution is ready to be integrated into your datacenter. Active System 200 and 800 are offered in configurations with either VMware® vSphere® (Active System 800v) or Microsoft® Windows Server® 2012 with Hyper-V® role enabled (Active System 800m) hypervisors. This paper defines the Reference Architecture for the VMware vSphere based Active System 800v solution and is also relevant to the Active System 200.

Both, Active System 200 and 800, offer converged LAN and SAN fabric design to enable a converged infrastructure solution. The end-to-end converged network architecture is based upon Data Center Bridging (DCB) technologies that enable convergence of all LAN and iSCSI SAN traffic into a single fabric. The converged fabric design of Active System reduces complexity and cost while bringing greater flexibility to the infrastructure solution.

Active System 200 and 800 differ in the server types they include. 200 includes Dell PowerEdge™ R720 rack servers while the Active System 800 has the PowerEdge™ M1000e blade chassis with Dell PowerEdge™ M I/O Aggregator, Dell PowerEdge™ M620 blades. Both solutions have Dell EqualLogic™ Storage, and Dell Networking 10GbE switches for LAN and SAN port aggregation. Both solutions include Dell PowerEdge™ R620 servers as management servers, Dell Active System Manager, VMware vCenter Server, EqualLogic Virtual Storage Manager for VMware, and Dell OpenManage™ Essentials .

Active System Network Layout

In this whitepaper we focus on how Active System rack integrates into the existing LAN infrastructure. We discuss the interoperability of Active System first in Dell Networking S4810 based infrastructures and later in Cisco networking environments. We discuss some common deployment scenarios that Active System is likely to meet in a data center. The Active System rack internally has a network of its own that binds compute, the blade chassis or rack servers depending on the model, with storage and LAN. How this gets connected into the external world, as seen in figure 1, and how it would scale beyond a single rack will be discussed in this paper.

From a pure networking perspective the value proposition of the Active System network is four-fold.

1. Ease of deployment and resulting TCO.
2. Highly interoperable with any vendor.
3. Ease of management and deployment.
4. Scalable and reliable design that grows with the business needs.
5. We start by emphasizing the ease of deployment and management for the Active System. Figure 1 shows the network connectivity inside the rack. The network is configured and ready to use for the customer. The VMs are initiated with an operational SAN.

An important component of Active Systems is the Dell Active System Manager which offers a unified management approach that is tailored to the Active Infrastructure family of systems and solutions. Active System Manager was purpose built to drive automation and accuracy in delivering workloads, while enabling end-to-end infrastructure management through a single robust and extensible console.

The Active System 200 can have up to 7 PowerEdge R720 rack servers, up to four EqualLogic storage arrays, two management servers, a management switch and two S4810 10GbE switches.

Workloads like [Microsoft SQL Server databases]() can run on this underlying platform and a pre-engineered solution is available from Dell.

Figure 1. Active System 200

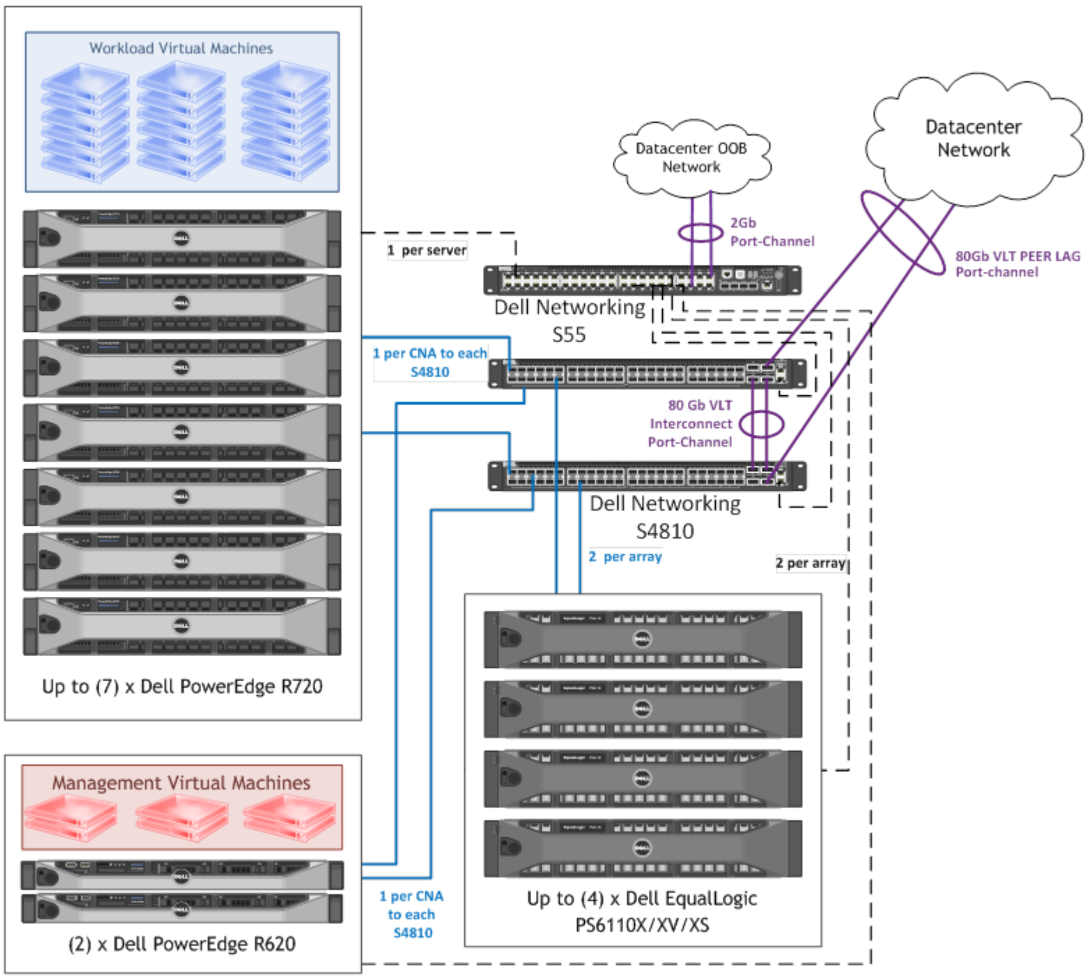

Inside the Active System 800 rack are up to two M1000e chassis with 16 M620 blade servers per chassis and two IO Aggregator modules, LAN & SAN converged Ethernet 10GbE Top-Of-Rack switches, up to eight EqualLogic storage arrays, a management switch, and two management servers.

A variety of pre-engineered workload solutions are available on the Active System 800. There is SQL server database, Microsoft SC VMM private cloud and VDI solution from Citrix and VMware with collaboration applications. A third party report on the performance of Active System 800 is available for download.

Figure 2. Active System 800

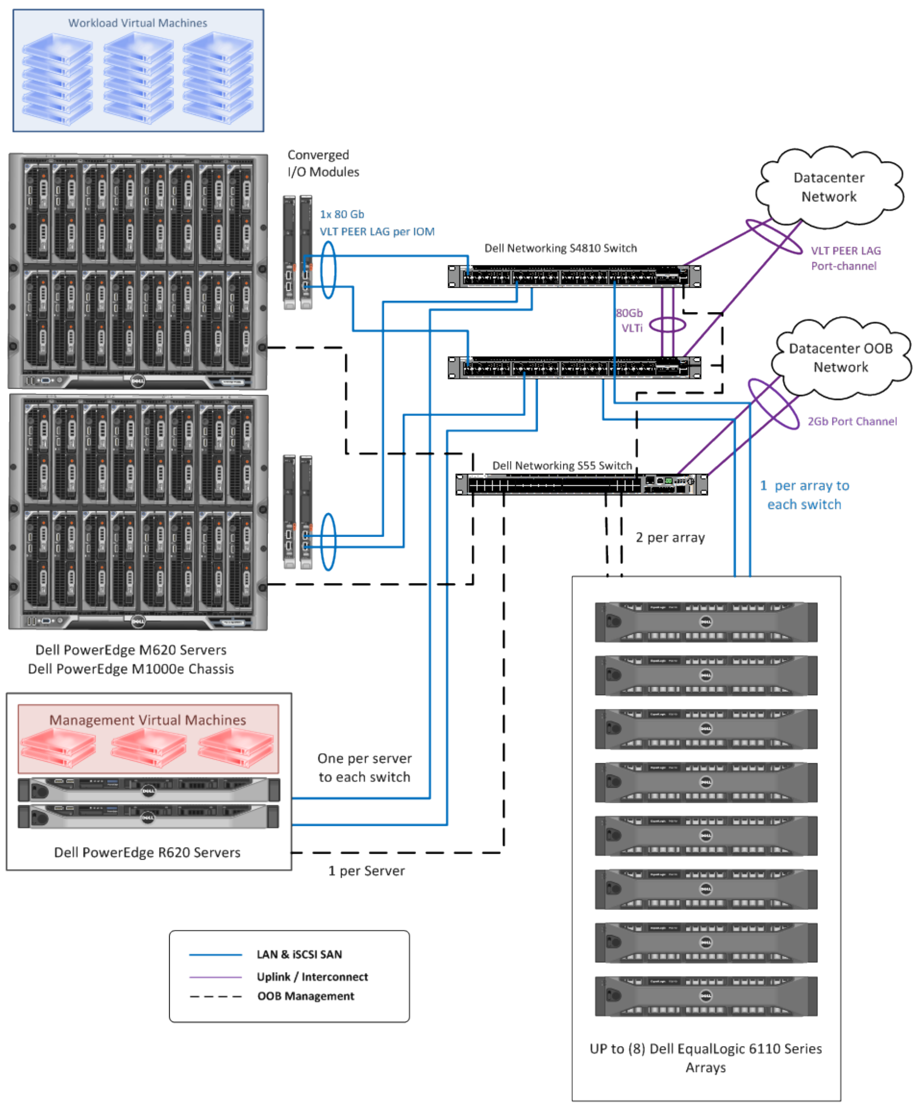

A brief description of the network inside the rack is given below.

The IO Aggregator supports 24 front-end ports and up to 32 internal ports facing the servers. In this Active System only 16 server facing ports are active and only the 2x 40G uplink ports are in use. The external facing 40G interfaces have been converted into 4x10GbE interfaces using the break-out cable that connects to 40G on one side and breaks out into 4 LC connectors. Check figure-3 for a picture of this break-out cable. For redundancy, these four LC strands connect into two different S4810 switches, two strands to each s4810 forming a LAG. The internal ports connect to the 16 half-height M620 servers. As the Active System 800 is scalable up to two M1000e chassis, all IO Aggregators, connect to the same pair of S4810s. The two S4810 switches connect to each other over an aggregated link of 2 x 40GbE interfaces, creating an inter-switch bandwidth of 80Gbits.

Active System 200 and 800 solutions leverage the new DCB standard, supporting the segmentation of LAN & SAN network traffic on the same infrastructure which otherwise would require completely separate hardware. The SAN traffic is aggregated on the S4810 along with the LAN traffic, where it's switched towards the EqualLogic iSCSI SAN consisting of up to 8 Arrays.

Active System and Dell Networking S4810 switch

The S4810 forms the backbone of this solution because it's quite a versatile switch. It uses 1/10/40GbE Ethernet media and is capable of being used as the top-of-rack inside the Active System as well as an aggregation or core switch for scaling out into multiple Active System solutions in the Data Center.

Figure 3.   S4810 switch

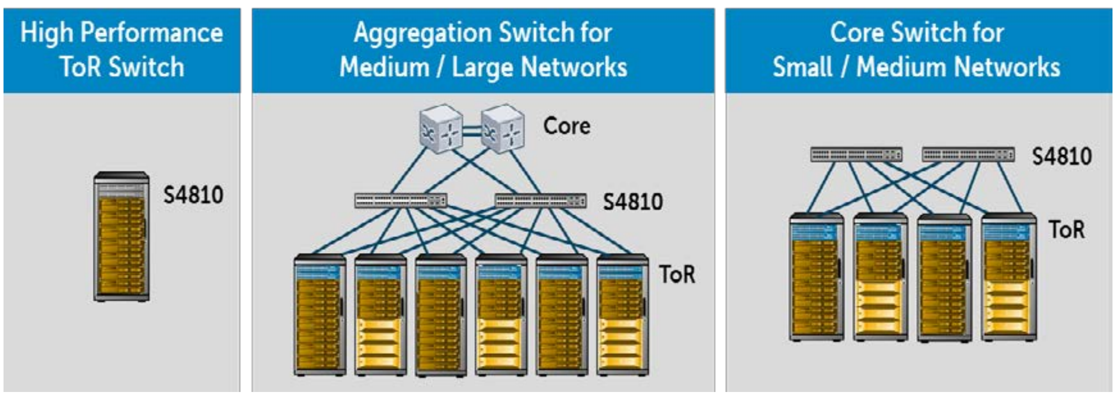

The software features combined with the hardware capabilities make this a versatile switch that can be used to consolidate 10GbE servers in the rack and build an aggregation of racks and the core. The S4810 can be the building block of a flat distributed core organized in a fat-tree non-blocking fabric that can scale up to tens of racks and hundreds of high-end servers. The 40GbE interfaces on the S4810 enable the next leap in capacity as servers become capable of 10GbE IO.

Connecting Active System into a Data Center

Let's discuss how the Active System 800 Infrastructure fits into a data center. We first discuss the physical media options available to connect an Active System into the Data Center Core. Next we discuss the network topology and technology options for a modern day data center.

Finally we discuss how multiple Active System 800s can be aggregated at a scale that suits a medium to large data center. Dell brings forth enormous value to its data center customers, by offering AS800 as a building block where multiple such racks put together can build a large data center in a fraction of time taken otherwise. This is invaluable for data centers that need to hit the ground running in a short time.

Physical Media Options to uplink the Active System

Connecting the Active System into the LAN is easy. Only the S4810 switches inside the rack enclosure need to be connected to the LAN core switches. The Active System rack enclosure has the front door showing the panels and a rear door for cabling. The S4810 switches can be seen mounted with the ports facing the rear door. The first step is to set up a console connection with these two switches as we may need some configurations done to get the uplink connections working.

The next step is to identify the ports that are going to be linked to the core switches. These ports on the S4810 that need to be uplinked can be 10GbE interfaces or the 40GbE interfaces located towards the right of the IO facing side. The 10G interfaces would need 10G SFP+ optical transceiver (if using fiber) or a Twin-Ax 10G cable that doesn't need an optic. For 40GbE, you would need a 40G-to-40G twin-ax or fiber optic cable if the core is 40G capable. The cabling options can be seen below.

Figure 4. Cabling options

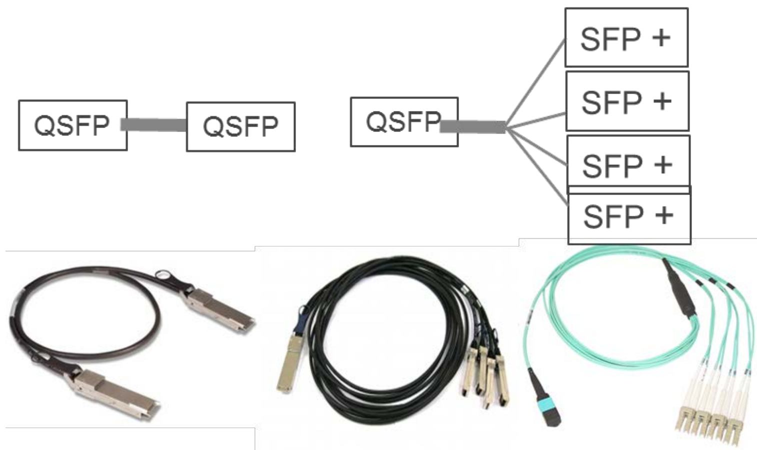

40G interfaces on the S4810 work in native 40GbE mode but can be converted into four 10GbE interfaces by a CLI command. If a breakout to 10GbE is needed, these are connected using a breakout cables as shown above. After the 40GbE interface has been converted to operate into four 10GbE mode, the respective fiber cables(shown in green) or twin-ax cables (shown in black) are used. When using fiber, a 40GbE optical transceiver is needed. The twin-ax has a transceiver built on the cable on both the 40GbE side and the four 10GbE sides.

Technology Options to uplink the Active System

Each network is different in some respects. Therefore not all configurations that we discuss here may be applicable but are based on a general usability pattern seen in most networks. We discuss two different topologies with which an Active System solution could be connected. The Layer 2 VLT based scale out and the equal-cost routed option. The figures in this section show a pair of s4810 switches in the core which fulfills the requirements at most small to medium size data centers. This pair of core switches would connect to the Active System racks as well as other network access infrastructure.

The extended Virtual Link Trunking (VLT) option

The large bridge network where certain VLAN domains need to be extended beyond a few racks is still a common requirement. Virtual Link Trunking, which eliminates the use of Spanning Tree, is a technology where a server or bridge uplinks a Trunk into two switches. The switches running this feature make themselves appear as a single switch to a connecting bridge or a server. We call this switch pair a VLT- pair. Both links from the bridge network actively forward and receive traffic which means full bandwidth utilization and redundancy is achieved. This technology is a perfect replacement for Spanning Tree based network designs.

We take this technology one step further by now creating multiple tiers of VLT. The VLT –Pair links into another VLT-Pair creating 2 tiers of VLT, running over 40GbE speeds. This helps us build ever larger layer-2 domains that can stretch across larger number of racks, aggregate bandwidth using the 40GbE and yet keep networks simpler. Simplicity is achieved by eliminating Spanning Tree as the primary enabler of redundancy while avoiding creation of loops.

The S4810s inside each Active System come pre-configured for VLT, which allows multiple IO Aggregators to dual home into the S4180 VLT-Pair. The blade switches form a 40G LAG upwards into the VLT-Pair with 20Gb going to each switch in the VLT-Pair. Intra-chassis traffic, i.e server communication within the M1000e chassis, would never need to leave the blade IO aggregator, and inter-chassis traffic would be switched within the S4810 pair. At no point does the local traffic need to go out to core and then get switched back. Major benefits of this technology are:

1. Dual control plane on the access side that lends resiliency.
2. Full utilization of the active LAG interfaces.
3. Rack level maintenance is hitless and one switch can be kept active at all times.

Figure 5. Topology options

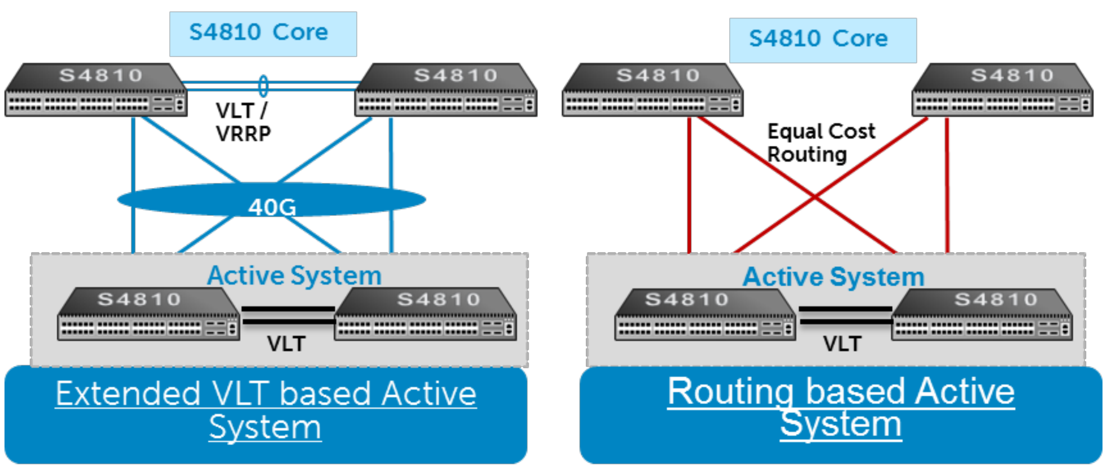

The routing option

The second most commonly used method of connecting racks into aggregation or core layers is when routing is enabled on the top-of-rack switches. Routing entails enabling routing protocols on both sides the Active System rack as well as the core side. Commonly used IGP routing protocols like OSPF and IS-IS are well suited for this purpose. The major benefits include:

1. This option is suited for a CLOS (Distributed Core) scale out. There is no need for the S4810s in Active System to connect with each other in this option. Optionally these could form a regular LAG for an alternate path.

2. Multiple paths to the destination in the routing table enable equal-cost multi-path capability. As each switch has at least two paths available, the traffic gets load balanced.
3. Limited domain of Layer-2 within the rack limits broadcast and other flooded traffic within the rack.
4. Network wide link states are visible to each Active System rack which helps in fast convergence in failures.
5. Rack level maintenance is hitless and each switch can be independently maintained.

Creating an Active System Scale Out

This section addresses the growth of infrastructure from a medium size scale to a larger core enabled for 40G.The Dell Z9000 switch would be an obvious choice to create a simple yet scalable network. Each Z9000 can support 32 line-rate 40GbE ports and could aggregate multiple racks with little or no over-subscription. When connecting multiple Active System racks, using the 40GbE uplinks from the rack, we can build a large fabric that supports multi-terabit clusters. Where each Active System forms the infrastructure of a public or private cloud, the scalability to accommodate multiple tenants is achieved at the core by building a routed infrastructure. The density of the Z9000 allows flattening the network tiers and creating an equal-cost fabric from any to any other point in the network.

Figure 6. Active System scale out

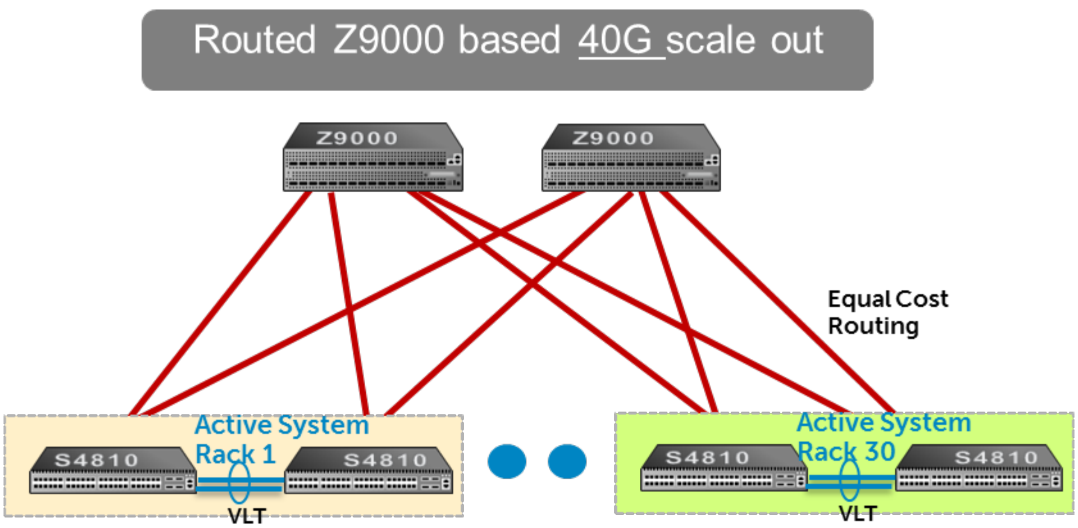

However, the option that can be used for large domain layer-2 requirements is the one that uses the extended VLT (eVLT) on the Z9000. The VLT-Pair thus formed can scale in terms of hundreds of servers inside multiple racks of Active System 800s. Using the base Active System 800 calculation, each of which has 32 servers, this Z9000 based aggregation can cover around 500 servers inside 15 or more Active System 800 racks. Each rack of Active System has 4 40GbE links to the Core network providing enough bandwidth for all the traffic destined from and to each rack.

Figure 7. Scale out on Layer-2 using VLT on Z9000 in core

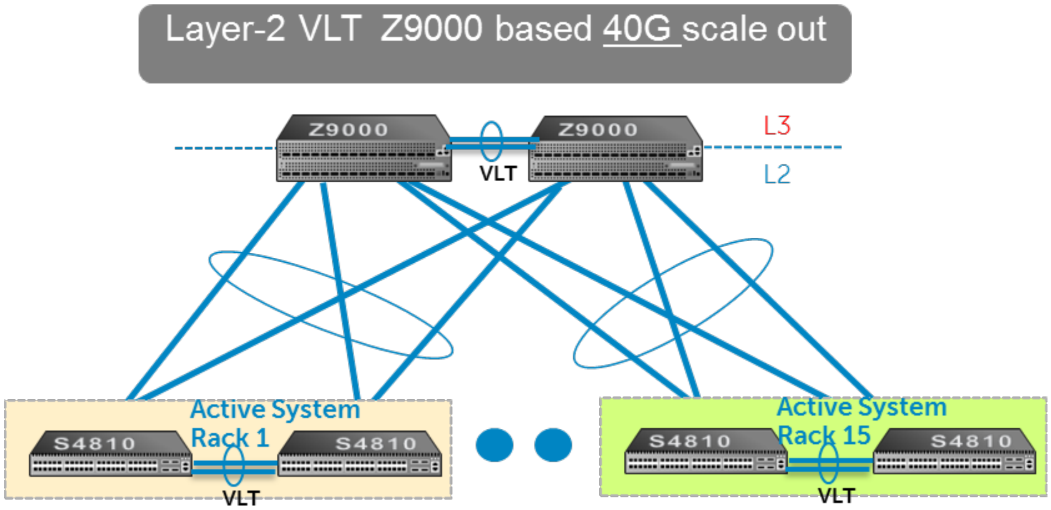

Large scale networks routinely deploy the Z9000 in HPC and cloud environments. Compared to a chassis (modular) switch, the density available on Z9000 gives the best ROI in terms of switching capacity, form-factor (2U), and power measured in fractions compared to modular chassis. Supported by FTOS operating system that has features suitable for high end switches and routers, the Z9000 becomes a great next-generation choice for a core switch.

Converged iSCSI and LAN

A major feature that enhances the Active System is the convergence of SAN and LAN over the same infrastructure.

The first step of convergence is the NPAR technology of the server's NIC. The NIC Partition technology can convert a 10GbE NIC/CNA into 4 virtual partitions. It can be used in both virtualized and non-virtualized environment. These NPARs are independent of the switches they connect to. Each 10GbE physical NIC can have up to 4 partitions, each with its own VLAN that maps into the physical switch.

Of these 4 partitions 1 belongs to management traffic, one for vMotion and one for iSCSI. Thus the iSCSI traffic is segregated on the server in a separate VLAN and a DCB priority is applied to it on the physical network.

DCB is an end-to-end concept. From the server the DCB is extended to the IO Aggregator, to the S4810 and to the EqualLogic storage arrays. Here's a snippet of what DCB looks like in the S4810 switch inside the Active System.

```
!interface TenGigabitEthernet 0/0
 no ip address
 mtu 12000
 dcb-policy input pfc
 dcb-policy output ets
!
```

```
 port-channel-protocol LACP

  port-channel 10 mode active
!

 protocol lldp

  advertise management-tlv system-name
 no shutdown
!
<SNIP>
!
qos-policy-output OTHER ets
 bandwidth-percentage 95
!
qos-policy-output iSCSI ets
 bandwidth-percentage 5
!
dcb-input no-pfc
 no pfc mode on
 pfc priority 4
!
dcb-input pfc
 pfc priority 4
!
priority-group OTHER
 priority-list 0-3,5-7
 set-pgid 2
!
priority-group iSCSI
 priority-list 4
 set-pgid 1
!
dcb-output ets
 priority-group OTHER qos-policy OTHER
 priority-group iSCSI qos-policy iSCSI
!
dcb-policy input stack-unit all stack-ports all pfc
!
```

Connecting Active System to Cisco Nexus 5548 core

It's likely that an Active System 800 is connected into a data center infrastructure that consists of Cisco or some other vendor. We discussed in the previous section the connectivity of Active System using Dell hardware; here we discuss the single most commonly observed use case of Active System connecting with Cisco.

In the figure 8 below, it can be observed that each S4810 in the Active System rack has a 2-port LAG/Port Channel linking into the two Cisco Nexus switches. The number of links is flexible and can be changed according to use cases. These two S4810s form a VLT link with each otherAlso in this example, 4 ports were used for the vPC "Peer Link". This number also can be changed to adapt to different configuration. The vPC Peer Link will be sharing VLAN traffic from the vPC uplinks, so they should be planned in an appropriate manner to avoid oversubscription. The “Peer Keep-Alive” link is suggested to be configured to utilize the management port (this is the default) and will send very little layer 3 traffic between the switches.

The two Cisco Nexus 5000 Series switches are configured with vPC and a high availability aggregation pair.

Figure 8. Active System with Nexus Core

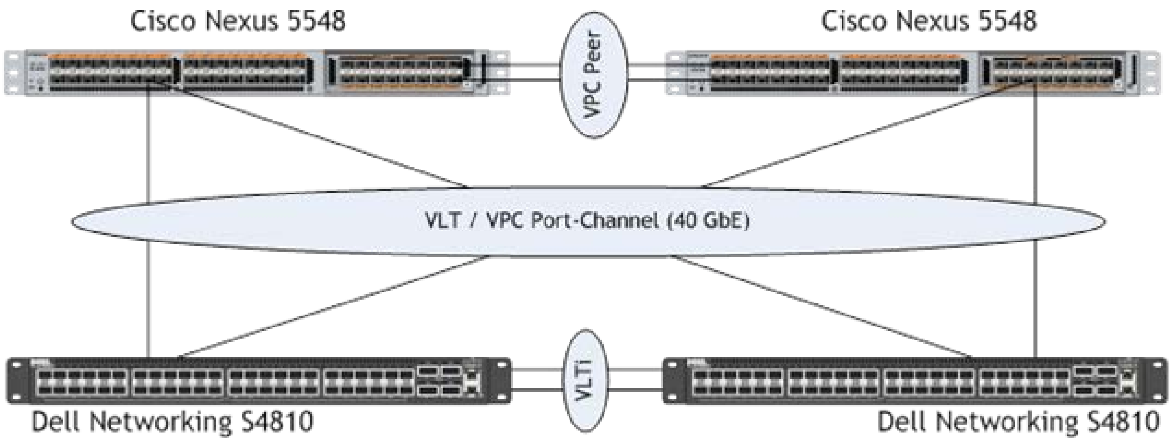

The features that need to be looked at in this setup are Layer-2 multi-path and spanning-tree.

The vPC feature, comparable to the Dell networking VLT feature, virtualizes the path from the access switch to the core. The Cisco pair could additionally be running HSRP for gateway redundancy. There is no interoperability needed on HRSP with Active System, as it’s operated between the Nexus pair and is independent of what connects to this pair.

Dell S4810 and Cisco Interoperability

Cisco Nexus 5K’s default spanning tree protocol is the R-PVST enabled upon bootup. On Dell S4810 switch, any one of PVST/MSTP/RSTP has to be enabled on the switch manually. In the above design if there’s no LAG between the Active System S4810s then there’s no need to participate in the Nexus R-PVST. Otherwise we enable PVST on the S4810s with the root bridge always on the Cisco pair. You can achieve that by manually setting a low bridge priority number (ie. 4096) on one of the Cisco switches.

The following table shows the interoperability between Cisco switches and Dell S4810s on spanning-tree and other active protocols.

Table 1. Interoperability checklist

Nexus 5K/ CAT IOS	Dell S4810	Status	Note
R-PVST	PVST+	Standard based Interoperability	Dell PVST+ is Rapid in nature*
MST	MST	Standard based Interoperability	
LLDP	LLDP	Standard based Interoperability	
OSPF	OSPF	Standard based Interoperability	
LACP LAG	LACP LAG	Standard based Interoperability	
vPC	VLT	No Interoperability	Both are proprietary features for HA purpose.

* sub-second convergence timing

As shown the spanning tree seamlessly interoperates between Dell switches and Cisco. The standards based features like LACL and OSPF are bound to interop as they are implemented according to a common standard. Of particular note is the HA feature from both Cisco and Dell for the purpose of using full bandwidth across Link-Aggregation and avoiding spanning tree altogether. Cisco has the vPC feature and Dell has an equivalent Virtual Link Trunking (VLT). This feature is exclusively run between the pair of switches of same vendor, to which the Layer-2 switch or server connect, which are oblivious to the fact the LAG terminates into two different switches. This is what allows the switches inside Active System 800 to connect to the Core infrastructure running on Cisco. The vPC and VLT can both work together and switch pairs get connected to each other, forming a single logical LAG, with no loops. Therefore the need to run spanning-tree from ToR to Agg/Core is optional.

Summary

This paper touched upon the details of the networking part of Active System 200 and 800 solutions from Dell. We discussed the value proposition of the Dell solution and how the networking elements from Active System plug into the LAN infrastructure. We also discussed how the LAN infrastructure could be different for customers and the options available for different scenarios. There are design options available on how to connect the Active System into the Aggregation or a flat core network infrastructure. The options discuss Link Aggregation with VLT (Virtual Link Trunking), Layer-3 routed Equal Cost Routing and large scale Z9000 based Core where multiple Active System connect into the Data Center Core. The S4810 is a versatile switch that can work at the access, aggregation and the core levels. It has the density and features to optimally work in any environment. Lastly the switching inside the blade server is the most compelling with Dell's IO Aggregator module that provides local switching to the blade servers. The management of Active System 200 and 800 is controlled by the Active System Manager that provides a single pane of glass for the configuration and lifecycle management needs.

Finally we discussed how the Active System would work effortlessly in an existing Cisco environment and interoperate on all protocols. The details of the configurations and outputs are shown in the appendix section.

Appendix

Dell Networking S4810 Core with VLT configuration

Following CLI and output snippets are from Dell Networking S4810 as the core, into which one or more Active Systems would connect. Refer to the diagram in Figure 5, this would be the pair on the top. A similar configuration can be followed for Z9000 switches in Figure 7.

Following these steps we will configure VLT on the pair of S4810s that connect the Active System racks together. To configure virtual link trunking, you must create a VLT domain, configure a backup link and interconnect trunk, and connect the peer switches in a VLT domain to an attached access device (switch or server).

Figure 9. S4810 VLT interconnect

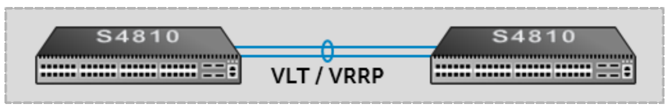

As a first step however, RSTP should be configured as a best practice on the S4810 as well as the S55 the management switch. Here's snippet from the S4810. Note the bridge priority on the switch; this is configurable depending on the fact where the root needs to lie.

Caution: This snippet is not an example from inside the Active System Top-of-Rack switch. Follow these instructions for core or aggregation switches outside the rack.

```
S4810_VLTpeer1(conf)#protocol spanning-tree rstp
S4810_VLTpeer1(conf-rstp)#no disable
S4810_VLTpeer1(conf-rstp)#bridge-priority 4096

#Repeat the same on VLTPeer2 with a different bridge priority to make it the
root.

S4810_VLTpeer2(conf-rstp)#bridge-priority 0
```

The next figures show a sample VLT configuration. VLT works over an ICL or primary link and a backup link. In absence of a direct path to the destination, the ICL link would carry the traffic to the peer. The backup link is only for heartbeat status from the peer and no data traffic flows over it.

Configure a VLT domain

```
vlt domain <domain id >
```

Fix the VLT system parameters to avoid negotiations (for faster convergence)

```
primary-priority <value> (suggestion: 1 for the primary, 8192 for sec)
system-mac mac-address <value> (same MAC address on all VLT peers)
unit-id <value> (suggestion: 0 for the primary, 1 for secondary)
```

Avoid picking a random MAC addresses that could be reserved or multicast.

Prepare your port-channel for VLTi (interconnect) configuration. To become a VLTi the port-channel should be in default mode (no switchport).

Note: The system will automatically include needed VLANs to be tagged into the VLTi. You do not need to manually tag VLANs on the VLTi.

Configure the core VLT peering relationship across the port-channel that will become the VLT interconnect (VLTi)

Note: it is recommended to build the VLTi port-channel statically to minimize negotiations in the VLT domain core.

```
(conf-vlt-domain)# peer-link port-channel <LAG-ID>
```

Configure the VLT backup link (used for health checks)

```
(conf-vlt-domain)# back-up destination <ip-address>
```

The backup link should be a different link than the VLTi and if possible following a diverse path. This could be the management interface IP address.

Note: It is recommended that VLTs that are facing hosts/switches should be preferably built by LACP, to benefit from the protocol negotiations. However static port-channels are also supported

```
interface port-channel <id-number>
vlt-peer-lag port-channel <id-number>
```

Sample VLT configuration on peer1

```
S4810_VLTpeer1(conf)#vlt domain 999
S4810_VLTpeer1(conf-vlt-domain)#peer-link port-channel 100
S4810_VLTpeer1(conf-vlt-domain)#back-up destination 10.11.206.35 << Peer
management IP address

S4810_VLTpeer1(conf-vlt-domain)#exit

S4810_VLTpeer1(conf)#interface ManagementEthernet 0/0

S4810_VLTpeer1(conf-if-ma-0/0)#ip address 10.11.206.23/16

S4810_VLTpeer1(conf-if-ma-0/0)#no shutdown
S4810_VLTpeer1(conf-if-ma-0/0)#exit
```

```
S4810_VLTpeer1(conf)#interface port-channel 100 << The VLTi link

S4810_VLTpeer1(conf-if-po-100)#no ip address
S4810_VLTpeer1(conf-if-po-100)#channel-member fortyGigE 0/56,60
S4810_VLTpeer1(conf-if-po-100)#no shutdown
S4810_VLTpeer1(conf-if-po-100)#exit

S4810_VLTpeer1(conf)#interface port-channel 110 << The LAG to the L2 switch
S4810_VLTpeer1(conf-if-po-110)#no ip address
S4810_VLTpeer1(conf-if-po-110)#switchport
S4810_VLTpeer1(conf-if-po-110)#channel-member fortyGigE 0/52
S4810_VLTpeer1(conf-if-po-110)#no shutdown
S4810_VLTpeer1(conf-if-po-110)#vlt-peer-lag port-channel 110
S4810_VLTpeer1(conf-if-po-110)#end

S4810_VLTpeer1# show vlan id 10
Codes: * - Default VLAN, G - GVRP VLANs, P - Primary, C - Community, I -
Isolated
   Q: U - Untagged, T - Tagged
x - Dot1x untagged, X - Dot1x tagged
G - GVRP tagged, M - Vlan-stack, H - Hyperpull tagged

NUM      Status      Description                  Q Ports
10       Active                                   U Po110(Fo 0/52)
                                                  T Po100(Fo 0/56,60)
```

S4810 inside Active System

The S4810 switches inside the Active System rack would already have the relevant configurations if you have purchased the Active System solution and it has been installed for you. This is for your reference if you are following the Active Systems reference architecture.

The S4810 switch configuration details on those inside the Active System check out the next section. Spanning-Tree Root Bridge would lie on the VLT master S4810 of the core no inside the rack. The two VLT domains, inside and outside the rack, need to be separate IDs.

Cisco Nexus 5548 Core with vPC configuration

Following CLI and output snippets are from a Cisco Nexus 5548 operating as the core, running vPC and the Active System connected to this core.

S4810 inside Active System

###########Configuration-S4810-01#####	###########Configuration-S4810-02####
protocol spanning-tree pvst no disable ! interface Port-channel 12 no ip address	protocol spanning-tree pvst no disable ! interface Port-channel 11 no ip address

mtu 2500 portmode hybrid switchport no shutdown ! interface TenGigabitEthernet 0/49 no ip address mtu 2500 ! port-channel-protocol LACP port-channel 12 mode active no shutdown ! interface TenGigabitEthernet 0/50 no ip address mtu 2500 ! port-channel-protocol LACP port-channel 12 mode active no shutdown ! interface TenGigabitEthernet 0/51 no ip address mtu 2500 ! port-channel-protocol LACP port-channel 12 mode active no shutdown ! interface TenGigabitEthernet 0/52 no ip address mtu 2500 ! port-channel-protocol LACP port-channel 12 mode active no shutdown ! interface Vlan 11 no ip address mtu 2500 tagged Port-channel 12 no shutdown ! interface Vlan 12 no ip address mtu 2500 tagged Port-channel 12 no shutdown !	mtu 2500 portmode hybrid switchport no shutdown ! interface TenGigabitEthernet 0/49 no ip address mtu 2500 ! port-channel-protocol LACP port-channel 11 mode active no shutdown ! interface TenGigabitEthernet 0/50 no ip address mtu 2500 ! port-channel-protocol LACP port-channel 11 mode active no shutdown ! interface TenGigabitEthernet 0/51 no ip address mtu 2500 ! port-channel-protocol LACP port-channel 11 mode active no shutdown ! interface TenGigabitEthernet 0/52 no ip address mtu 2500 ! port-channel-protocol LACP port-channel 11 mode active no shutdown ! interface Vlan 11 no ip address mtu 2500 tagged Port-channel 11 no shutdown ! interface Vlan 12 ip address 10.11.1.2/24 mtu 2500 tagged Port-channel 11 no shutdown !

interface Vlan 20 no ip address mtu 2500 !untagged TenGigabitEthernet 0/1-32 !untagged Port-channel 1,10,12 no shutdown !	interface Vlan 20 no ip address mtu 2500 !untagged TenGigabitEthernet 0/1-16 !untagged Port-channel 11,20 no shutdown !

```
############## LLDP Connections####################
S4810_01#sho lldp neighbors
 Loc PortID   Rem Host Name     Rem Port Id                Rem Chassis Id
 -------------------------------------------------------------------------
 Te 0/49      Nexus5548-Botto...Eth1/2                     54:7f:ee:53:3e:89
 Te 0/50      Nexus5548-Botto...Eth1/1                     54:7f:ee:53:3e:88
 Te 0/51      Nexus5548-Top     Eth1/2                     54:7f:ee:56:55:49
 Te 0/52      Nexus5548-Top     Eth1/1                     54:7f:ee:56:55:48

##############
S4810-02#sho lldp neighbors
 Loc PortID   Rem Host Name     Rem Port Id                Rem Chassis Id
 -------------------------------------------------------------------------


 Te 0/49      Nexus5548-Botto...Eth1/26                    54:7f:ee:53:3e:a1
 Te 0/50      Nexus5548-Botto...Eth1/25                    54:7f:ee:53:3e:a0
 Te 0/51      Nexus5548-Top     Eth1/26                    54:7f:ee:56:55:61
 Te 0/52      Nexus5548-Top     Eth1/25                    54:7f:ee:56:55:60


##############Spanning tree Output#################
S4810_01#sho spanning-tree pvst vlan 11 brief
VLAN 11
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 4107, Address 547f.ee53.3ec1  <<< Cisco is the Root
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID    Priority 32768, Address 001e.c9f1.053e
Configured hello time 2, max age 20, forward delay 15
Bpdu filter disabled globally

Interface                                             Designated
 Name       PortID   Prio Cost    Sts       Cost       Bridge ID        PortID
---------- -------- ---- ------  --------- ------- ------------------ --------
Po 12      128.13   128  1400    FWD       1401    32779 0023.04ee.bff4  144.10

Interface
 Name       Role    PortID   Prio Cost    Sts    Cost    Link-type Edge BpduFilter
---------- ------ -------- ---- ------- ------ ------- --------- ---- ----------
Po 12      Root    128.13   128  1400    FWD    1401    P2P       No   No

S4810_01#sho spanning-tree pvst vlan 12 brief

VLAN 12

Executing IEEE compatible Spanning Tree Protocol

Root ID    Priority 4108, Address 547f.ee53.3ec1  <<< Cisco is the Root
```

```
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID    Priority 32768, Address 001e.c9f1.053e
Configured hello time 2, max age 20, forward delay 15
Bpdu filter disabled globally

Interface                                             Designated
 Name       PortID   Prio Cost    Sts       Cost        Bridge ID        PortID
---------- -------- ---- ------  -------- ------- -------------------- --------
Po 12      128.13   128  1400    FWD       1401    32780 0023.04ee.bff4  144.10

Interface
 Name       Role   PortID   Prio Cost    Sts     Cost    Link-type Edge BpduFilter
---------- ------ -------- ---- ------- ------ ------- --------- ---- ----------
Po 12      Root   128.13   128  1400    FWD    1401    P2P       No   No
```

Primary Nexus switch

```
###########Interface Config######################
Nexus5548-Bottom# sho run interface po 12

!Command: show running-config interface port-channel12
!Time: Tue Jul  3 23:03:54 2012

version 5.1(3)N2(1a)

interface port-channel11

  switchport mode trunk

  switchport trunk allowed vlan 11-12

  vpc 11

interface port-channel12
  switchport mode trunk
  switchport trunk allowed vlan 11-12
  vpc 12

###########VPC Status#########################
Nexus5548-Bottom# sho vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 500
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status: success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : primary, operational secondary
Number of vPCs configured         : 2
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled

vPC Peer-link status
```

```
--------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po500  up     1,11-12,20

vPC status
---------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
11     Po11        up     success     success                    11-12
12     Po12        up     success     success                    11-12
###########Spanning tree########################
Nexus5548-Bottom# sho spanning-tree summary
Switch is in rapid-pvst mode
Root bridge for: VLAN0001, VLAN0011-VLAN0012, VLAN0020
Port Type Default                        is disable
Edge Port [PortFast] BPDU Guard Default  is disabled
Edge Port [PortFast] BPDU Filter Default is disabled
Bridge Assurance                         is enabled
Loopguard Default                        is disabled
Pathcost method used                     is short

Name                   Blocking Listening Learning Forwarding STP Active
---------------------- -------- --------- -------- ---------- ----------
VLAN0001                     0         0        0          1          1
VLAN0011                     0         0        0          3          3
VLAN0012                     0         0        0          3          3
VLAN0020                     0         0        0          1          1
---------------------- -------- --------- -------- ---------- ----------
4 vlans                      0         0        0          8          8


##############VPC Config########################
Nexus5548-Bottom# sho run vpc

!Command: show running-config vpc
!Time: Tue Jul  3 23:07:14 2012

version 5.1(3)N2(1a)
feature vpc

vpc domain 500
  peer-keepalive destination 172.25.188.60 source 172.25.188.61
  ip arp synchronize

interface port-channel11

  vpc 11

interface port-channel12

  vpc 12

interface port-channel500
  vpc peer-link
```

Secondary Nexus switch

```
###########Interface Config######################
Nexus5548-Top# sho run interface port-channel 11
!

interface port-channel11
  switchport mode trunk
  switchport trunk allowed vlan 11-12
  vpc 11
!

interface port-channel12
  switchport mode trunk
  switchport trunk allowed vlan 11-12
  vpc 12

###########VPC Status##########################
Nexus5548-Top# sho vpc brief
Legend:
                (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id                     : 500
Peer status                       : peer adjacency formed ok
vPC keep-alive status             : peer is alive
Configuration consistency status: success
Per-vlan consistency status       : success
Type-2 consistency status         : success
vPC role                          : secondary, operational primary
Number of vPCs configured         : 2
Peer Gateway                      : Disabled
Dual-active excluded VLANs        : -
Graceful Consistency Check        : Enabled

vPC Peer-link status
---------------------------------------------------------------------
id   Port   Status Active vlans
--   ----   ------ --------------------------------------------------
1    Po500  up     1,11-12,20

vPC status
----------------------------------------------------------------------------
id     Port        Status Consistency Reason                     Active vlans
------ ----------- ------ ----------- -------------------------- -----------
11     Po11        up     success     success                    11-12
12     Po12        up     success     success                    11-12



###########Spanning tree########################

Nexus5548-Top# sho spanning-tree summary
Switch is in rapid-pvst mode
Root bridge for: none
Port Type Default                        is disable
```

```
Edge Port [PortFast] BPDU Guard Default  is disabled
Edge Port [PortFast] BPDU Filter Default is disabled
Bridge Assurance                         is enabled
Loopguard Default                        is disabled
Pathcost method used                     is short

Name                   Blocking Listening Learning Forwarding STP Active
---------------------- -------- --------- -------- ---------- ----------
VLAN0001                      0         0        0          1          1
VLAN0011                      0         0        0          3          3
VLAN0012                      0         0        0          3          3
VLAN0020                      0         0        0          2          2
```

Further reading

Dell Active Systems Wiki

Active Fabric in a Small and Medium Data Center: For an example of configurations on multi-tiered VLT.

Dell Active System Manager