



Dell Networking Multi-tenant Data Center Reference Architecture

A Dell Reference Architecture

Dell PowerEdge Servers, Dell Compellent SC4020, Microsoft Server 2012 Hyper-V, and featuring Dell Networking 100GbE switches and Silver Peak WAN Solutions

Dell Network Solutions Engineering
October 2015

Revisions

Date	Description	Authors
October 2015	Initial Release	Jaiwant_Virk@Dell.com , Dell_Networking_Solutions@Dell.com , Dell@Silver-Peak.com

©2015 Dell Inc., All rights reserved.

Except as stated below, no part of this document may be reproduced, distributed, or transmitted in any form or by any means, without express permission of Dell.

You may distribute this document within your company or organization only, without alteration of its contents.

THIS DOCUMENT IS PROVIDED "AS-IS", AND WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED. IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE SPECIFICALLY DISCLAIMED.

PRODUCT WARRANTIES APPLICABLE TO THE DELL PRODUCTS DESCRIBED IN THIS DOCUMENT MAY BE FOUND

AT: <http://www.dell.com/learn/us/en/vn/terms-of-sale-commercial-and-public-sector-warranties> Performance of network reference architectures discussed in this document may vary with differing deployment conditions, network loads, and the like. Third party products may be included in reference architectures for the convenience of the reader. Inclusion of such third party products does not necessarily constitute Dell's recommendation of those products.

Please consult your Dell representative for additional information.

Trademarks used in this text:

Dell™, the Dell logo, Dell Precision™, PowerEdge™, PowerVault™, PowerConnect™, OpenManage™, Compellent™, FlexAddress™, and Force10™ are trademarks of Dell Inc. Other Dell trademarks may be used in this document. Silver Peak Systems™, the Silver Peak logo, Silver Peak Unity Boost™ and Silver Peak Unity EdgeConnect™ are trademarks of Silver Peak Systems, Inc. Intel®, Pentium®, Xeon®, Core® and Celeron® are registered trademarks of Intel Corporation in the U.S. and other countries. Microsoft®, Windows®, Windows Server®, Hyper-V®, Internet Explorer®, and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates. VMware®, Virtual SMP®, vMotion®, vCenter® and vSphere® are registered trademarks or trademarks of VMware, Inc. in the United States or other countries. IBM® is a registered trademark of International Business Machines Corporation. Broadcom® and NetXtreme® are registered trademarks of Broadcom Corporation. Qlogic is a registered trademark of Qlogic Corporation. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and/or names or their products and are the property of their respective owners. Dell disclaims proprietary interest in the marks and names of others.



Table of Contents

1	Introduction.....	5
1.1	The Network.....	6
1.2	The Management Network	7
1.3	The Virtualization Platform.....	7
1.4	The Storage.....	7
1.5	The Supporting Services (Security, Edge Routing, and SD-WAN)	7
2	The Underlay Network	8
2.1	Server and Networking Scaling	10
2.2	Dell S4048-ON as a Leaf Switches (ToR).....	12
2.3	Dell Z9100-ON as a Spine and WAN Edge Switches	12
2.4	Dell Storage SC4020 – iSCSI.....	16
3	Virtualizing the Network Fabric – Microsoft HNV, the Overlay	17
3.1	Deploying Hyper-V Network Virtualization Overlay.....	18
3.2	Logical and Virtualized Networks	21
3.3	Microsoft Windows Server Hyper-V Cluster Design.....	22
3.4	Microsoft Windows Server HNV Gateway	23
3.5	Dell Compellent SC4020 and Hyper-V Using iSCSI.....	24
3.6	Implementing the Private Cloud.....	25
3.7	Microsoft Server HNV Summary	25
3.8	Microsoft Server 2012 R2 and VMM Scaling Maximums	26
4	Services – Hybrid MPLS and Broadband Wide Area Networking with Silver Peak	28
4.1	Silver Peak SD - WAN Solutions	28
5	Validated Lab Environment	29
5.1	Application - Scenario One.....	31
5.2	Application - Scenario Two.....	32
6	Products used in this Solution.....	34
6.1	Dell Networking Products	34
6.2	Dell PowerEdge Server Products.....	35
6.3	Dell Storage Products.....	37
6.4	SD WAN Solutions.....	38



A Works Cited 40
B Additional Resources 41
C Support and Feedback 43



1 Introduction

Modern data centers typically contain multiple technologies that work together to enable the business applications. Broadly speaking, the most basic layer is comprised of the hardware infrastructure with multi-tiered software running on it. This hardware forms the basis of the three major processes: compute, network, and storage. Next generation infrastructures make use of a virtual network layer. It is on this virtualization layer that applications run. In addition to compute processes, applications need a storage service that they can access. At the base of the compute and storage processes is the network fabric that connects all the environments together.

This Data Center Reference Architecture (DCRA) provides a high-level view of one possible multi-tenant cloud solution and its major building blocks. The multi-tenant environment featured is built on the following components.

- Hardware infrastructure connected to a high-density Leaf-Spine Network architecture, which include
 - Multi-rate Ethernet switches that support 100G, 40G and 10G speeds
 - Top-of-Rack switch 10G/40G and spine switch 10G/25G/40G/100G
 - Dell PowerEdge rack and blade servers
 - Dell Compellent SC series storage
- Microsoft Server Virtualization Solution that supports
 - Microsoft Windows Server 2012 R2 with Hyper-V Network Virtualization (HNV)
 - Windows hypervisor, Hyper-V Virtual Machine Management Service
 - Dell Enterprise Storage Manager for Hyper-V and Compellent
 - Various components such as virtual machine bus (VMbus)
 - Managed with SCVMM (Software Center Virtual Machine Manager)
- WAN Optimization Solution - Silver Peak
 - VMs running in ROBO (Remote Office Branch Office)
 - Data Center to Campus or Remote Office environments – Silver Peak

This architecture is designed and documented to address the requirements for a multi-tenancy private, public or hybrid cloud environment using Microsoft's Hyper-V, VMM, and HNV applications. This design is based on the best-of-class network compute and storage infrastructure that scales horizontally according to current business needs.

Consult this document for a small data center that needs a few hundred virtual machines, or a very large organization that has thousands of servers with extensive virtualization. In addition, this document can help reduce the initial complexity during the planning and deployment of new technologies in the enterprise business. It highlights many of Dell's innovative enterprise components, solutions and their general relationships within a data center.

As information services, and infrastructure requirements grow and change and solutions continue to evolve, the need for reference architectures increases. The reference architecture (RA) serves as a tool for providing common information, guidance, and direction towards a solution development within an enterprise.



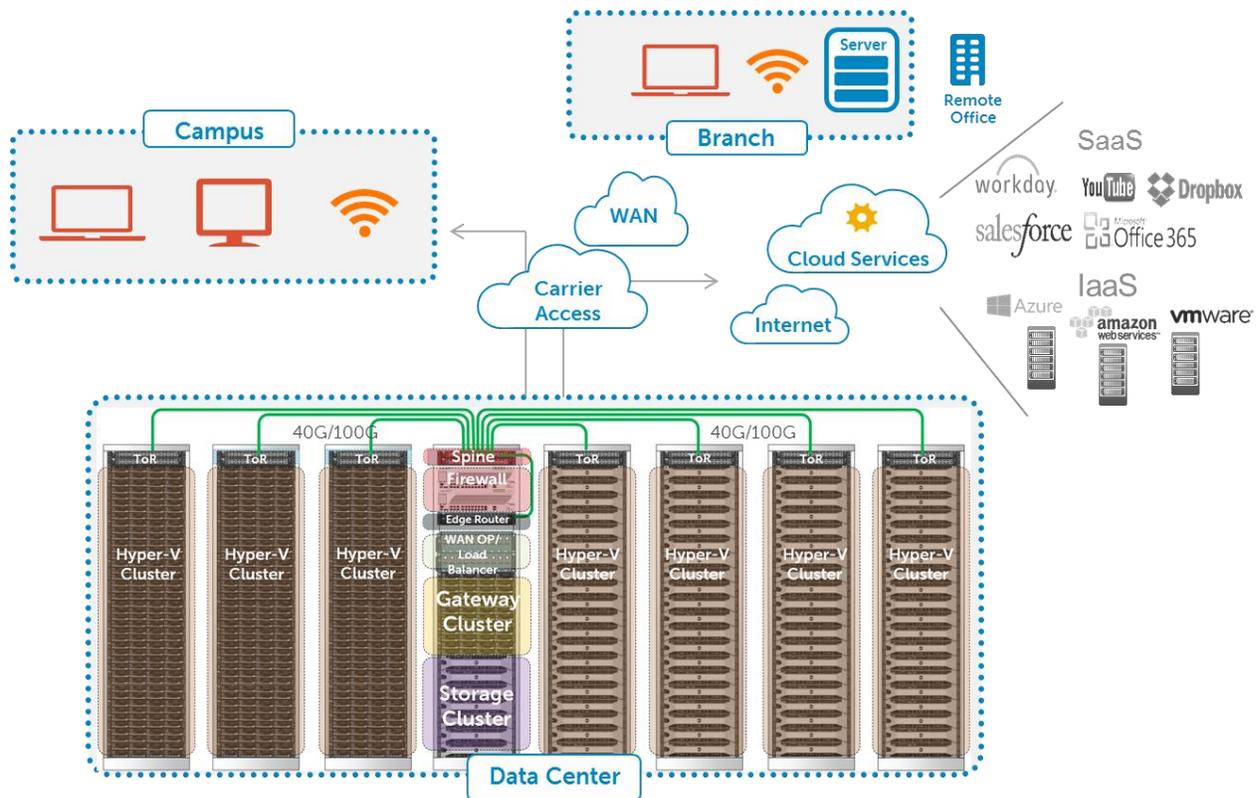


Figure 1

A typical data center reference architecture – Multi-tenant and multi-service environment

The following sections provide a look at the building blocks in this high-level architecture.

1.1 The Network

Within this reference architecture, the topic areas are arranged in a logical fashion beginning with the hardware implementation (underlying network). The spine and leaf architecture forms a baseline for all modern data center networks primarily due to its high-availability, and horizontal scaling features. It is characterized by fixed-form switching that packs high-speed interconnects with low latency and high throughput. The east-west (server-to-server/storage-to-storage) traffic gets full bi-sectional bandwidth over equal-cost links.

The network is normally a mix of layer 2 switching extending to the Top-of-Rack (ToR) and layer 3 routing extending beyond that. There are however, technologies that provide layer 2 extensions between servers in the same cluster and adjacent clusters. In the same cluster, Virtual Link Trunking (VLT) is used to provide multi-switch link trunking, which can be extended all the way to the spine layer for applications that need a larger layer 2 domain. VLT and its configuration/application design can be found in various documents including this knowledge base article on support.dell.com. [Knowledge Base: How to set up Virtual Link Trunking.\(Dell Inc., 2015\).](#)



1.2 The Management Network

A standard out-of-band management network is used to connect the server's iDRACs and switch management. In addition to being used to configure the solution, this network also provides management access to the solution components. Management infrastructures are designed in various ways depending on the data center and administration requirements. Based on the multiple methods available for this configuration and the understanding that most management infrastructures are already in place at the customer data center, the specific details on the management network are not included in this document.

1.3 The Virtualization Platform

The multi-tenant environment (Figure 1) provides for a private, public or hybrid cloud environment using the capabilities of Microsoft's Hyper-V with HNV, managed by System Center Virtual Machine Manager (SCVMM) and Dell Enterprise Storage Manager for Compellent.

Included in this Reference Architecture is Microsoft Windows Server 2012 R2 with the Hyper-V role enabled in a Multi-Tenant Cluster environment. The virtualization from HNV includes a tunneling capability based on the Network Virtualization Generic Routing Protocol (NVGRE) standard that creates an overlay network that VMs use in communication with each other. Overlay networks use encapsulation of payload frames into a regular Layer 3 packet that is routed inside the underlay network while maintaining its own private network that is invisible to the underlying network. This de-coupling of virtual network from the physical network is a key benefit of this encapsulation

1.4 The Storage

The storage service in this reference architecture is provided by the Dell Compellent Storage Center 4000 series. Dell Compellent storage is available in multi models that fulfill any storage needs. In this data center environment the focus is on a converged storage network configuration using iSCSI. With the addition of Dell Storage Enterprise Manager into the VMM architecture, the configuration (from a virtualized standpoint with Hyper-V) is more efficient and can be managed from a single location within VMM.

In many data center topologies, it is advantageous to have a completely separate storage network. This separation is due to protocol incompatibility with Fibre Channel SAN's and additional provides for clear boundaries for troubleshooting and management with iSCSI. Depending on how important it is for zero downtime with iSCSI the typical configuration will change. Adding Data Center Bridging (DCB) iSCSI or Fibre Channel over Ethernet (FCoE) to the mix and this becomes a different discussion around converging the networks together with very little separation.

1.5 The Supporting Services (Security, Edge Routing, and SD-WAN)

Network services, also known as layer 4 through 7 services, form the essential supporting features in a data center. In this document, we discuss security through Dell SonicWALL, WAN edge routing through Dell switches/routers and Software Defined WAN from Silver Peak. There are other major services that are required such as Application Delivery Controller (ADC) also known as Network-based load-balancers but this is out of scope for this RA.



2 The Underlay Network

This section covers the underlay or hardware used in the solution and assists in the high-level planning and configuring of the applicable hardware for this solution. This example of a pod-based solution consists of fourteen racks, each deployed with two data-network VLT leaf switches. Each rack can be treated and managed like an individual and modular unit consisting of a variable number of servers or storage arrays. The spine switches are connected to multiple WAN edge switches.

Note: A pod is defined as the maximum numbers of racks that can be connected to a single pair of VLT spine switches.



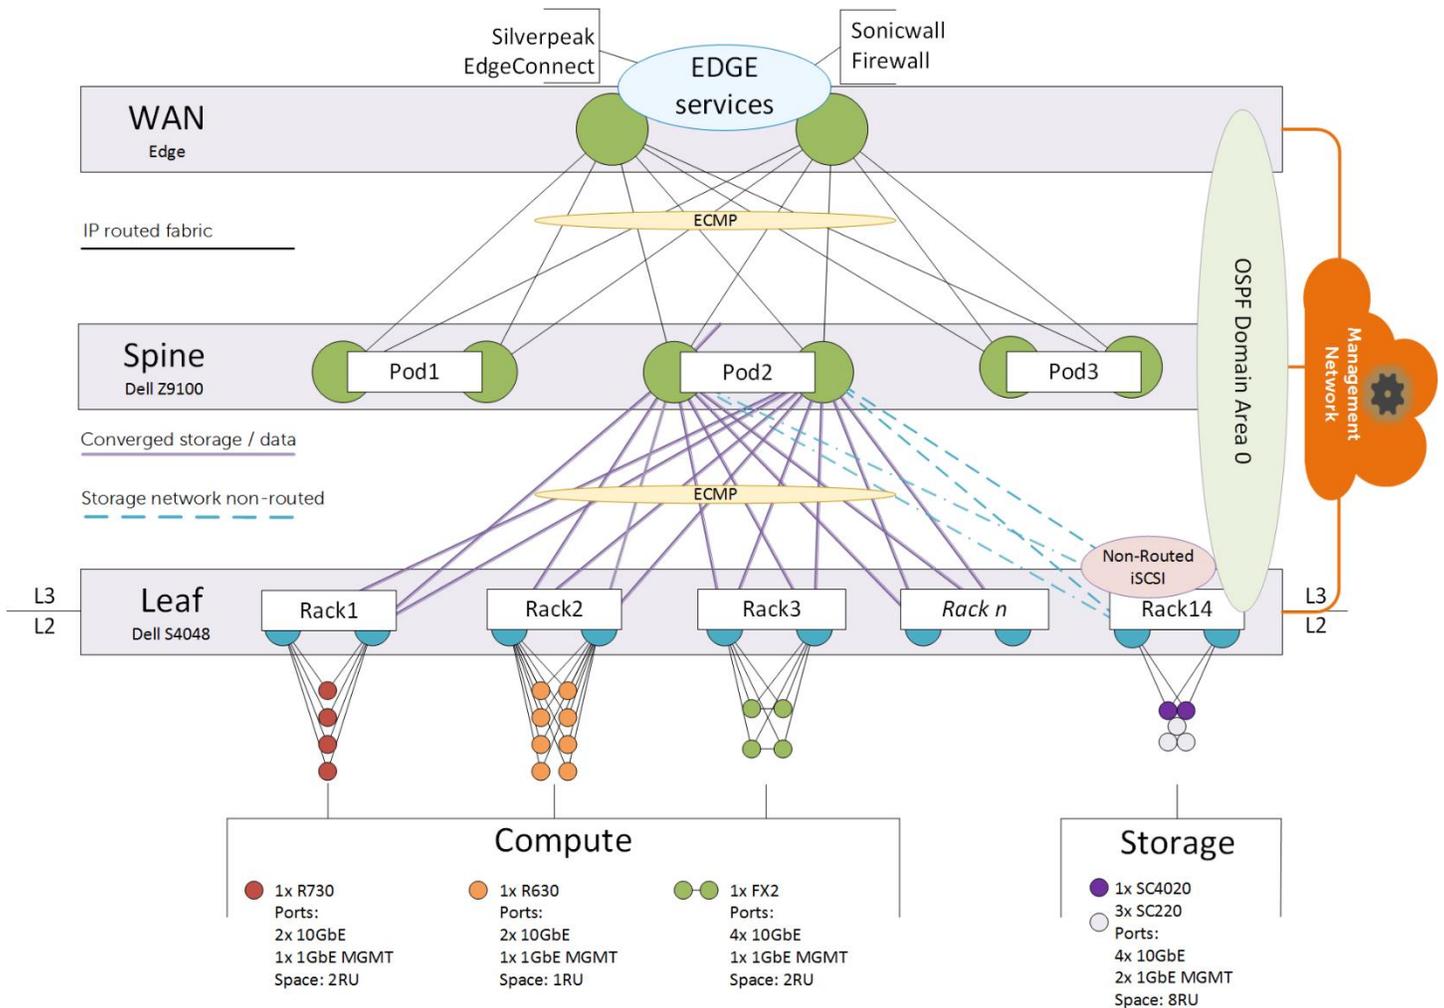
Dell Networking Z9100-ON multi-rate 100Gb L2/L3 switch

Figure 2

Dell's high performance (100G) switch, the Z9100, is used to implement the spine and WAN edge layer of this solution (Figure 2). Dell's ultra-low latency top-of-rack S4048 switches are deployed as leaf switches. VLT is implemented in the spine and leaf layers of this solution to offer high resiliency, availability and to effectively utilize multipath.

As shown in Figure 3, the Open Shortest Path First (OSPF) protocol with Equal-Cost Multipathing (ECMP) is configured between the leaf, spine, and WAN edge layer to provide for cross-pod connectivity. The sub-second convergence of ECMP on the Dell Networking switches, coupled with active-active VLT links, ensures that the failure of any link will be resolved by balancing the traffic over the alternate paths.

The bottom half of the illustration contains possible rack configurations for compute and storage. The deployment has assumed 42 rack units (RU) per rack with 36 RU being dedicated to compute or storage. The remaining space is reserved for additional infrastructure not covered in this guide. The red dots in Rack 1 represent an example of 2RU servers with dual 10GbE port adapters. In Rack 2, the orange dots represent 1RU servers with dual 10GbE adapters. The green dots (Rack 3) show the use of Dell PowerEdge FX2 each with dual FN410S IO Managers providing four SFP+ upstream connections per 2U. Rack 14 (far right) is an example of a storage pod. The storage pod consists of four Dell Compellent SC4020s, each with three SC220 storage arrays,



Whole Pod Topology

This Pod-based data center design incorporates Silver Peak Unity EdgeConnect as WAN solution. This WAN solution is the 'Industry's Most Complete, High-Performance Broadband and Hybrid WAN Solution'. For this solution's security, Dell SonicWALL Firewall has been implemented. Dell SonicWALL high-performance firewall appliances seamlessly integrate intrusion prevention, malware protection, application intelligence, control and visualization, IPsec VPN, SSL VPN, wireless controllers and other features to deliver comprehensive protection and maximum performance to address any sized deployment from enterprise to small and medium-sized businesses.



2.1 Server and Networking Scaling

This section provides details on the available server, storage, and networking port counts based on the products used for the pod design. Pod design can vary based on server, storage, and networking requirements by the customer.

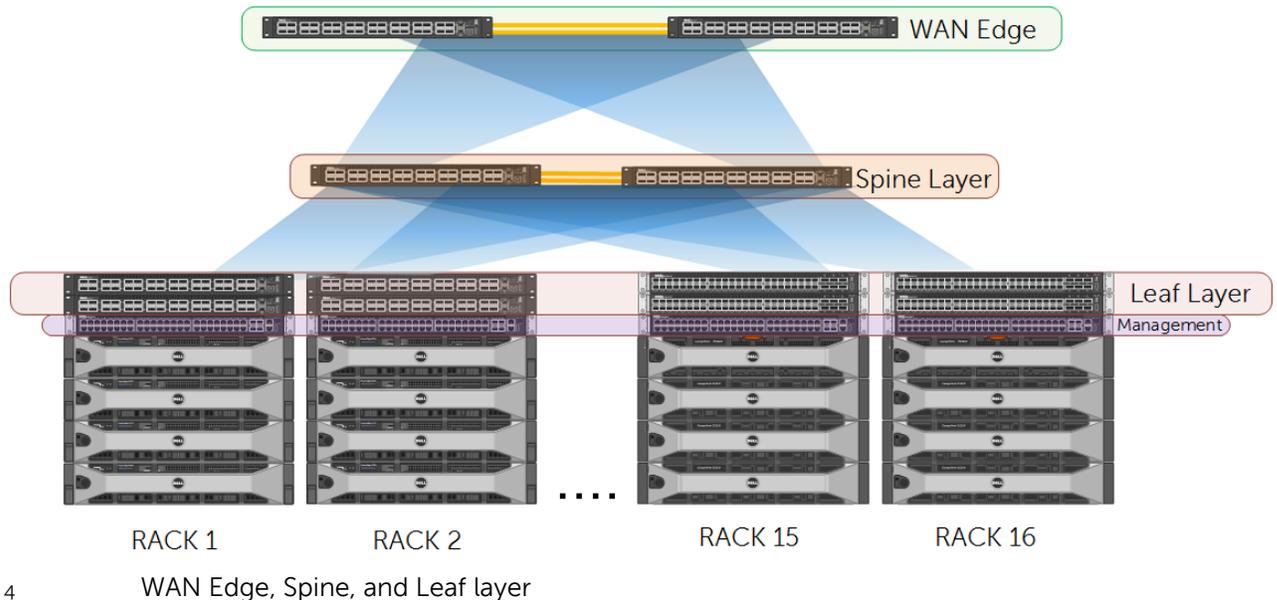


Figure 4

WAN Edge, Spine, and Leaf layer

As mentioned earlier, each rack is an individual unit in this pod-based data center. To provide an example of how these racks can be built, the following servers and storage have been selected.

- Dell PowerEdge Server 730 – The PowerEdge R730 server combines powerful processors, large memory, fast storage options, and GPU accelerator support to deliver outstanding functionality in a number of demanding environments with just 2U of rack space.
- Dell Power Edge Server 630 – This two-socket rack server delivers uncompromising density and productivity. The R630 is ideal for virtualization, driving large business applications or running transactional databases.
- Dell PowerEdge FX2 Enclosure – The PowerEdge FX2 is a 2U hybrid rack-based computing platform that combines the density and efficiencies of blades with simplicity and cost benefits of rack-based systems.
- Dell Storage SC4020 all-in-one array – The SC4020 combines dual controllers, 24 internal drive slots, and high performance of Fiber Channel or iSCSI network ports in a space saving 2U chassis. This storage chassis can support three SC220 extension modules, each of which is a 2U, 6GB SAS enclosure that supports up to 24x2.5" HDDs or an all-flash storage option in one chassis.

Table 1 presents additional details on each of the four rack types.

RACK VIEW							
	Component Size in Rack Units	Number of Servers per Rack	Rack Units available Networking Components	Bandwidth of Server link	Number of ports per Server	Total Bandwidth per fabric for Rack	Total Bandwidth for Rack
Rack Type 1 with PowerEdge R730 Servers	2	18	36	10	2	180	360
Rack Type 2 with PowerEdge R630 Servers	1	36	36	10	2	360	720
Rack Type 3 with PowerEdge FX2	2	16	36	10	4	360	720
Rack Type 4 with Dell Compellent SC4020 (+3 SC 220) Storage	8	4	32	10	4	80	160

Table 1 Rack View to Connections.

With the desire to limit this design to a single pair of spine layer switches, the architecture was backwards designed based on available ports for the different layers. This design ideology is just one way to implement an efficient data center.



2.2 Dell S4048-ON as a Leaf Switches (ToR)

The S4048-ON is a 10/40GbE top-of-rack 1U switch designed for high-performance data center and computing environments. Leveraging a non-blocking switching architecture, the S4048-ON delivers line-rate layer 2 and layer 3 forwarding capacity with ultra-low latency to maximize your network performance.

Table 2 ToR View for Port Counts

	Leaf Switch (ToR) View							
	10Gb Ports to Each ToR	Bandwidth for Rack Gb	40Gb Ports to VLT Spine Layer	Bandwidth to Spine Layer Gb	Over-subscription	40Gb iVLT ports for Leaf	ToR 10Gb unused ports	ToR 40Gb unused ports
Rack Type 1 with PowerEdge R730 Servers	18	180	1+1	40+40	4.5:1	2	30	2
Rack Type 2 with PowerEdge R630 Servers	36	360	1+1	40+40	9:1	2	12	2
Rack Type 3 with PowerEdge FX2	36	360	1+1	40+40	9:1	2	12	2
Rack Type 4 with Dell Compellent SC4020 (+3 SC 220) Storage	8	80	1+1	40+40	2:1	2	40	2

The leaf switch view of the topology shown in Table 2, illustrates the over-subscription and the number of ports for each type of rack associated with this design. Keep in mind that this is example data and should be adjusted for different data types and may or may not meet the needs of specific cases in the data center.

2.3 Dell Z9100-ON as a Spine and WAN Edge Switches

The Dell Z9100-ON is an access and aggregation switch/router product for 10/25/40/50/100GbE aggregation in Data Center networks. It can also be configured with multi-rate links for clustering (VLT, stacking) and uplinks to aggregation and core switches.

Port flexibility allows connections using QSFP28, QSFP+, SFP28, and SFP+ optics (using breakout cables). High-density 10/25/40/50/100GbE ports allow network administrators to deploy to wide variety of environments as well as to plan for future connectivity.



2.3.1 Spine Layer Switches

Typically, end-of-row switches have been chassis-based products such as C-Series switches. With the density and economics of a Z9100, it is likely that the switch will see adoption in middle-of-row or end-of-row applications. In a middle-of-row application, two or three racks can have rack servers terminate on a single Z9100 switch. This reduces the number of top-of-rack devices in the access layer. In an end-of-row role, a Z9100 can aggregate multiple blade server chassis to an end-of-row switch and then hand off to a core network over 10/40G links. A Z9100 can be positioned with MXL/IOA blades in a Dell blade server chassis as the aggregation point for blade switch uplinks. Usually the uplinks from blade chassis are dual-homed across two switches.

A data center network can be designed as a leaf-and-spine architecture based on a Clos fabric using Z9100 switches to provide a non-blocking network for high-performance data centers. A core fabric can be based on Z9100 switches in the spine with a separate layer of leaf switches. Options for these leaf switches include Dell S4048, S6000, S6100, or Z9100 switches. Alternately, the spine layer could function as a leaf layer with servers directly connected to it. A primary benefit to this architecture is its ability to scale. Leaf and spine nodes can be added as the demands of the organization grow. Leaf-and-Spine architecture has greater resiliency than a conventional core architecture. Because of the distributed nature of the fabric, the impact of the loss of a single node is significantly lower than it would be on a conventional core design.

The logical topology between leaf and spine can be layer 2 or layer 3. Load distribution of traffic between the leaf and spine trunks can be achieved with layer3 protocols like OSPF using ECMP or using simple layer 2 LAG. Recent hashing enhancements in Dell Network Operating System (DNOS) prevent polarization of flows within the fabric and provide fair distribution of flows between leaf and spine layers. OSPF, Border Gateway Protocol (BGP) or Intermediate System to Intermediate System (IS-IS) can be used as the layer 3 control plane protocol and VLT for layer 2 deployments. Using the Z9100, physical links between the leaf and spine can be 100GbE. With the Z9100, these interconnect links can be 100GbE as migration to 40GbE servers and 100GbE optics starts to occur and cable offerings mature. Manageability and resiliency in a distributed core model are improved compared to centralized core architecture, since the failure of a single node within a Clos network cannot bring down the entire switching fabric.

Illustrating the view at the rack and leaf layer switches, Table 3 captures the over-subscription from Leaf to Spine and the port count utilized to have the active fabric VLT links configured.

Table 3 Top or Rack to Leaf/Spine Port Usage

	40G Ports from Leaf to Spine	Total Bandwidth GB	100G ports to Spine Edge	100G ports to iVLT Spines	Spine unused ports	Over Subscription Leaf to Spine
Leafs to Spine Bandwidth (14 Racks)	28	1120	2x100	2x100	0	5.600



2.3.2 WAN Edge Switches

Distributed Data Center core (Active Fabric) represents a key application for the Z9100. The WAN edge layer acts as an interface between the internal protocols like OSPF and the external protocols like BGP. In this example of a pod-based solution, WAN edge layer switches are connected to a SonicWALL Firewall and to a Silver Peak WAN device. There are multiple WAN edge layer switches, which are connected to multiple spine VLT switches. This WAN edge layer is responsible for the inter-pod communications. ECMP and LAG enable networks to be efficient and resilient to failures.

2.3.3 Leaf and Spine - Edge using OSPF and VLT

In simple terms, VLT allows two physical switches to appear as a single logical switch. With physical links as a port-channel, connecting two individual switches configured with VLT would logically group it as single entity only for the access switches which connect to the VLT domain. Both the VLT peers have their own configuration, control and data planes. VLT allows creating port-channels across two switches, eliminating the need for spanning tree blocking states across its connectivity and effectively utilizing its entire VLT links connected to the access/ToR switches. The access device could be a switch, switch stack, dual NIC server or any other device supporting LACP port-channels. High-availability and redundancy in the network is achieved by its duality in physical connectivity.

In this scenario OSPF is configured on VLAN interfaces as the broadcast network (default OSPF network). No other OSPF network (point-to-point) type is supported between the leaf and spine Layers. Since all the other VLT peers form adjacency with the Core VLT peers, as a best practice configure the primary node of the core VLT domain as Designated Router (DR) and secondary node as Backup Designated Router (BDR) for optimal routing adjacencies. By configuring all the VLAN interfaces in the primary VLT peer of the aggregation layer, the priority could be incremented to act as OSPF DR. The number of OSPF adjacencies from each peer depends upon the VLAN subnets.

With OSPF converged, all the nodes have the routes and next-hop details needed to reach other nodes. Albeit the VLAN interfaces are configured with OSPF, as a same broadcast domain, and underlying layer 2 dependency is achieved. For instance, VM migration within the rack and across VLT domains could be orchestrated.

It should also be noted that ECMP is used for path choice across the separate pods in the active fabric configuration. The ECMP feature enables flow-based load-balancing of traffic across the multiple routes where the cost of reaching the next hop is equal. Larger scale networks routinely deploy this feature, which is usually available only in high-end switches and routers, to balance traffic over multiple links ensuring high-availability and full bandwidth availability.

The following diagram gives an overview of the OSPF, ECMP and Routed VLT topology. Please see these references for applicable deployments of the Routed VLT, and VLT topologies.

http://en.community.dell.com/techcenter/extras/m/white_papers/20441720

http://en.community.dell.com/techcenter/networking/m/networking_files/20440236

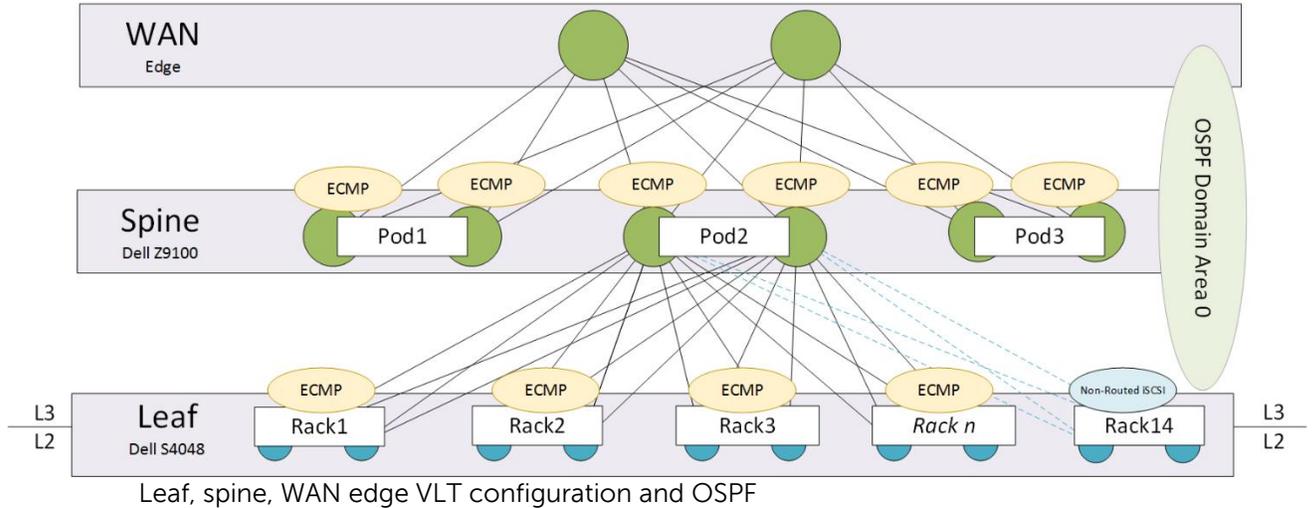


Figure 5

2.3.4 Hybrid WAN, MPLS and BGP

As more applications move to the cloud, companies are finding that Multiprotocol Label Switching (MPLS) is no longer sufficient to address today's WAN requirements. Silver Peak's Unity EdgeConnect enables enterprises to dramatically reduce the cost and complexity of building a WAN by leveraging broadband in order to connect users to applications. Silver Peak's Unity EdgeConnect is the preferred WAN device in this pod-based solution.

BGP's emphasis on security and scalability make it the routing protocol of choice between the ISPs. Other protocols are concerned solely with finding the optimal path towards all the known destinations. Because the agreements between the ISPs result in complex routing policies, BGP cannot take this simplistic approach. In this solution, BGP has been enabled between the WAN edge routers and Silver Peak WAN devices. The routers within the autonomous network maintain two routing tables: one for the interior gateway protocol (OSPF in this case) and one for BGP.



2.4 Dell Storage SC4020 – iSCSI

Dell's approach to modern storage architectures starts with flexible, modular designs that easily integrate across our broad ecosystem of internal and external capabilities. The result is storage solutions that specifically address workload and data center challenges, and deliver real advances in capabilities and value. Dell Products support various storage protocols such as iSCSI, SAS, FC, FCoE etc.

FC and FCoE require dedicated networking hardware and a specialized skill set is needed to configure and maintain it. In addition, supporting storage interconnectivity over long distances is expensive. In order to simplify and use existing networking switches (vs. dedicated hardware), the Dell Storage SC4020 utilizing iSCSI is the storage component of this pod based solution. The iSCSI network is comprised of Dell PowerEdge R630 and R730 servers as iSCSI initiators, which relay SCSI commands to the iSCSI targets on SC4020 Dell Storage arrays. Dedicated storage racks can serve all the storage requirements for the servers within a pod or WAN edge routers can be used to access different pods. Dedicated layer 2 VLANs segregate iSCSI networks from normal Ethernet traffic.

This Dell Enterprise Storage Manager is used to enable this configuration with VMM in the overlay. This topic will be covered in detail in [Section 3](#).

Note: For additional information on Dell Enterprise Manager for Compellent, please see: [Dell Enterprise Manager Configuration](#) (Dell Inc., 2013).

+ Additionally see this document for specific steps to deploying SMI-S and Dell Enterprise Manager with SCVMM: [Dell Enterprise Manager SCVMM configuration](#) (Dell Inc., 2015)



3 Virtualizing the Network Fabric – Microsoft HNV, the Overlay

The next part of the solution discusses the overlay applied to the physical hardware (the underlay).

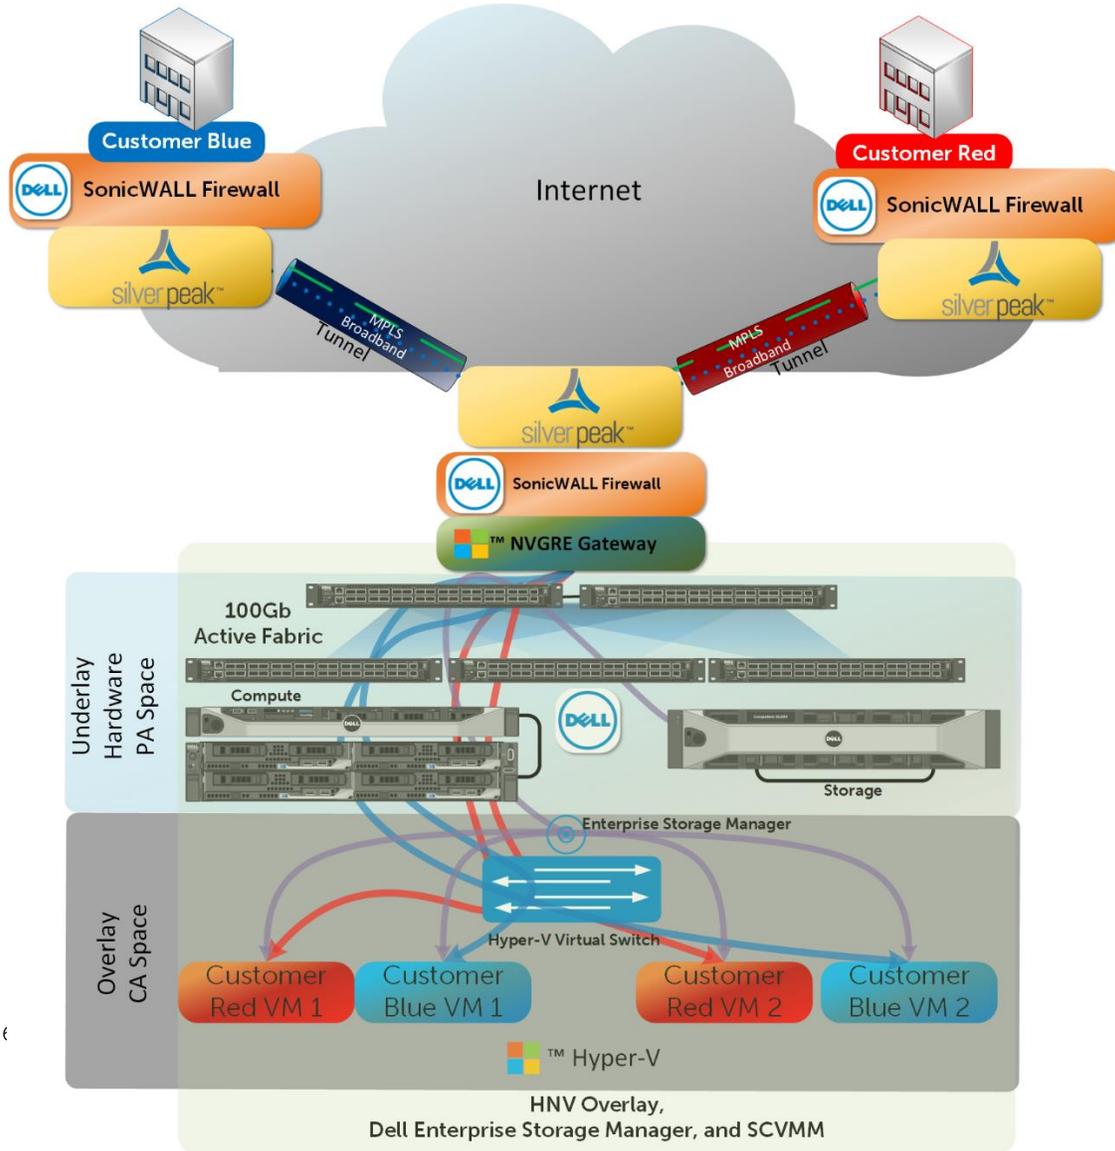


Figure 6

Overview

The current trend in business is to move computing toward the cloud. As this trend continues, a key component for successful implementation is adopting a multi-tenant software-defined networking (SDN) solution. Microsoft Hyper-V Network Virtualization (HNV) enables a multi-tenant solution and enables an Infrastructure as a Service (IaaS) solution for customers. This allows tenants to bring their own IPs, and if



they desire, their entire network topology. This is possible because the virtualized network is abstracted from the underlying physical environment.

The concept of network virtualization has existed for quite some time, with the Virtual Local Area Network (VLAN) being the original network virtualization protocol. The limitations of VLANs include the following.

- **Scalability** There is a maximum of 4,094 VLANs available. It is a common practice to use one subnet per VLAN. Most customers require many VLANs.
- **Size** A VLAN is limited to a single L2 network and cannot be routed. Stretching a L2 network beyond one or two racks in a datacenter has numerous complications, primarily due to spanning tree limitations.
- **Configuration** VLANs must be defined on every device in order for VLAN segmentation to work properly. This limits the portability of a VM, as moving the VM would require that all network settings are identical on the target physical server.

To address these limitations, network virtualization, including Microsoft HNV, provides a solution to route layer 2 networks over routable layer 3 IP networks. Similar to server virtualization that allows multiple virtual machines to operate on a single physical server with the illusion of isolation, network virtualization achieves the same abstraction from the physical network. With server virtualization, the benefits include:

- Isolation
- Efficient use of physical resources
- Easier migration of workloads across those physical resources

Network virtualization strives to deliver these same benefits to the network. Multiple tenants can operate their own, isolated, virtual networks on the same physical underlay. This includes the tenants IP space, routers, security appliances, load balancers, etc. Tenants can use the same physical underlay network, thereby improving the efficiency of those components. These virtual networks can be moved to another physical location without requiring any changes to the tenant's virtual network. A complete decoupling is achieved between the physical and virtual environments.

The benefits of network virtualization enable a scalable and flexible solution that can be used to provide a framework for hosting numerous tenants. By allowing the tenants to bring their own topologies, devices, and services, the provider can focus on the performance of the underlying physical network and not the constant change control required to provide services in a traditional, VLAN defined, multi-tenancy environment.

3.1 Deploying Hyper-V Network Virtualization Overlay

Windows Server 2012 R2 running Hyper-V along with the System Center 2012 R2 Virtual Machine Manager (SCVMM) is used by many customers and partners to offer a multi-tenant IaaS solution to their customers. Providing network isolation in such a way that the virtual machines from one customer cannot talk with the virtual machine of another customer requires careful planning in the compute, storage, network, and the network service layers. Building on the Active Fabric defined in the previous section, compute and storage building blocks are now introduced to construct the framework for network virtualization (overlay



network). In Figure 7 the workflow of deploying a Microsoft enabled virtualized network is outlined in general steps. Ensuring that each of the steps is completed will ensure a smooth deployment.

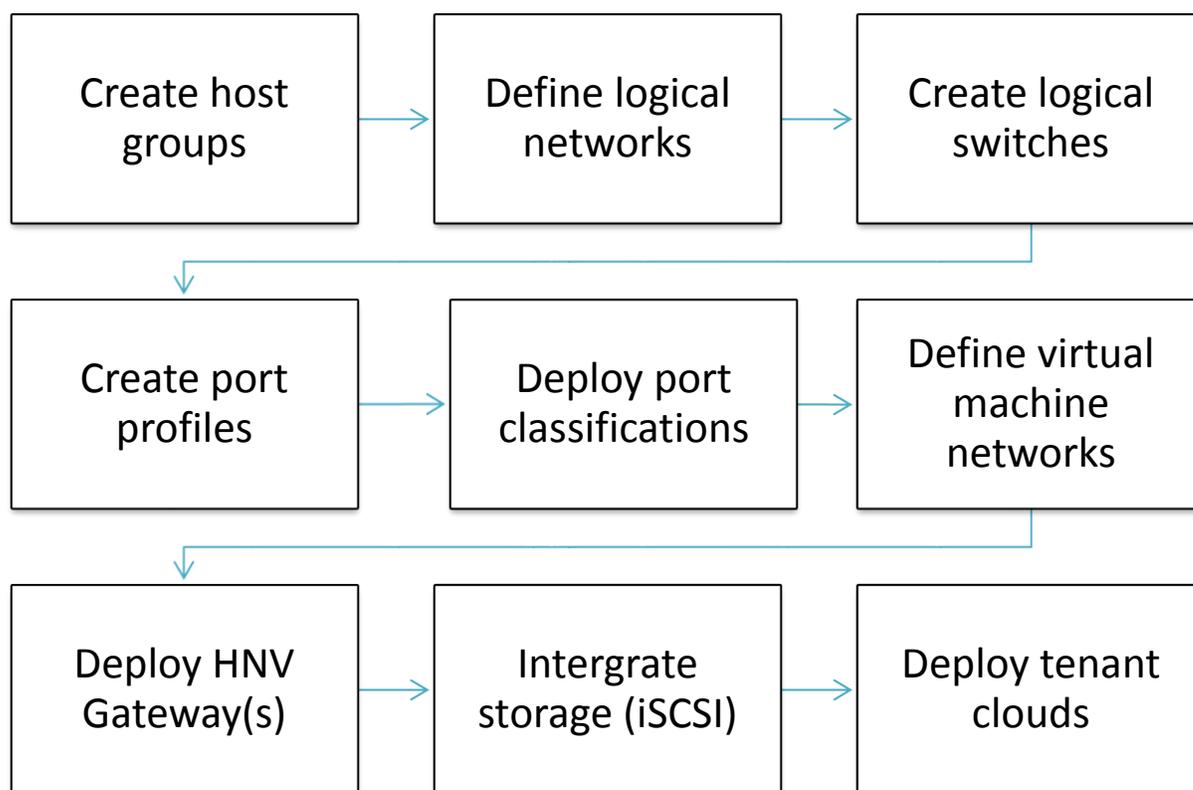


Figure 7

SCVMM Deployment Workflow

Before moving onto the overlay configuration of this RA, it is important to take a moment to plan for the required VLANs in the environment. Even with tenants being able to tunnel through the data center environment VLANs still serve the purpose of providing broadcast segmentation in the environment. This is particularly important with any iSCSI and management networks. With Microsoft Server Hyper-V clusters, each separate component (heartbeat, migration, etc.) can also benefit from separate VLANs. In this deployment, the following networks are used:

- Two VLANs for isolating iSCSI traffic
- One VLAN for the Provider Address (PA) Space
- One VLAN for SCVMM host management
- One VLAN for SCVMM virtual machine migration
- One VLAN for Hyper-V cluster monitoring

Depending on the architectural decision of where the L2/L3 boundary is set, the number of VLAN IDs may increase to reflect separate VLANs for each rack in the datacenter or the VLANs can have SVIs associated with the VLANs at the Spine layer to allow routing.

As shown in the workflow (Figure 7), a virtualized network solution built on Windows Server 2012R2 and SCVMM depends on a number of different components. It is important to understand the role these components play in the solution and how they are interconnected, especially if you need to troubleshoot issues with connectivity or the topology must be changed to meet updated business requirements.

Figure 8 shows a simplified representation of the different layers and components that create a virtualized network solution based on Windows Servers and SCVMM. In the diagram, the established physical network fabric is shown to the left; Hyper-V hosts with connections to the physical network are shown to the right. Along the top of the illustration, the component is listed for each column while the bottom text is how the solution connects to the previous column.

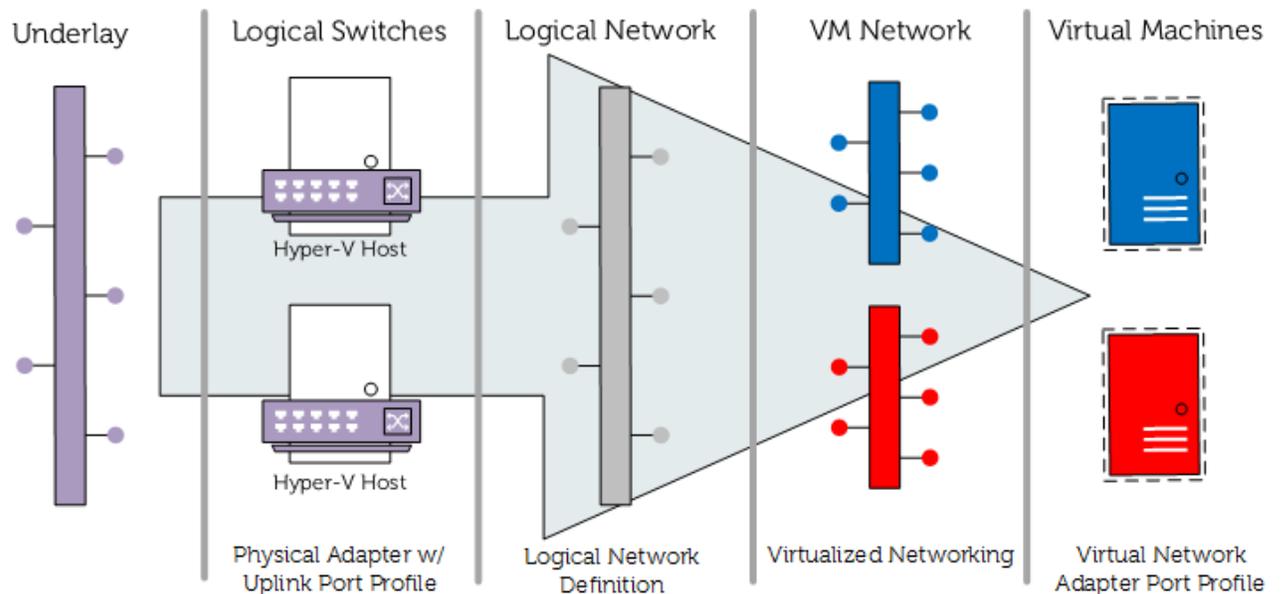


Figure 8

HNV Overlay Components and Connections

The following list defines the components illustrated above as well as additional components not shown.

Uplink port profiles - These are applied to physical network adapters as part of a logical switch deployment. They define the set of logical networks that should be associated with those network adapters, including Windows Server NIC teaming.

Network adapter port profiles - These are templates that define offload settings, bandwidth policies, and security settings for virtual network adapters.

Logical switches - These bring together uplink port profiles, network adapter port profiles, and port classifications together for any given physical or logical network.

Logical networks - These represent an abstraction of the underlying physical network infrastructure and enable the modeling of the network-based connectivity requirements.

IP address pools - These allow VMM to automatically allocate static IP addresses to Windows-based virtual machines. (Not shown in diagram)

Virtual machine networks -These provide the network interface through which a virtual machine will connect to a particular virtual network using NVGRE.

3.2 Logical and Virtualized Networks

To provide an example, logical networks and VM networks are shown in Figure 9 and Figure 10. As shown, most logical network definitions have virtual counterparts, except for the Provider Network, which contains the Provider Address (PA) space, and is used by the VM networks Tenant Blue and Tenant Red via NVGRE. Typically, a network is defined on both sides of the Hyper-V virtual switch if a virtual machine needs direct access to the equivalent physical network. In this case, the Edge network needs to allow external network access and the Storage network needs to provide Dell Compellent storage management through SCVMM.

Logical Networks and IP Pools (6)

Name	Network Compliance	Type	Description
Edge	Fully compliant	Logical Network	Used by HNV Gateways to enable external network access
Host Cluster	Fully compliant	Logical Network	Used by Window Cluster Services
Host Management	Fully compliant	Logical Network	Used by SCVMM to provide a management network
Host Migration	Fully compliant	Logical Network	Used by HNV to enable virtual machine migration
Provider Network	Fully compliant	Logical Network	PA space to allow NVGRE between tenant subnets
Storage Network	Fully compliant	Logical Network	Used for iSCSI traffic

Figure 9

SCVMM Logical Networks Example

VM Networks and IP Pools (8)

Name	Gateway Connection	Description
Edge	No	Edge of HNV network; Used with HNV Gateway clusters for access to external networks.
Host Cluster	No	Cluster communication network; carriers heathbeat
Host Management	No	Network connectivity to management network
Host Migration	No	Used for virtual machine and storage migration
Storage Network	No	iSCSI transport fabric
Tenant Blue	Yes	Virtual network enabled by NVGRE across PA network with a VSID
Tenant Red	Yes	Virtual network enabled by NVGRE across PA network with a VSID

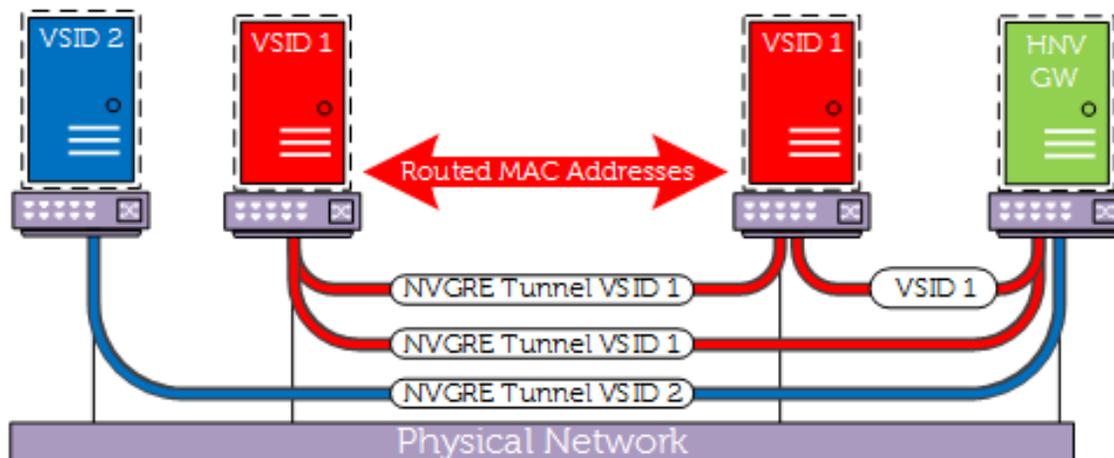
Figure 10

SCVMM Virtual Machine Networks

Virtualized networks use NVGRE, which involves creating virtual layer 2 topologies on top of a physical layer 3 network. Connectivity for the virtual networks is provided by tunneling Ethernet frames in Generic



Routing Encapsulation (GRE) over IP over the physical network fabric. Each virtual network is assigned a Virtual Subnet Identifier (VSID), which is carried in the GRE header using the key field defined in RFC 2890.



NVGRE Tunneling in a HNV Environment

Figure 11

As shown in Figure 11 each tenant has a VSID assigned automatically by SCVMM when deployed. This ID is inserted in the GRE header ensuring tenant isolation across the physical network. The VSE in the diagram is a HNV Gateway cluster running Windows Server 2012 R2. When the Blue tenant VMs require network access outside of the virtual subnet, the request is sent to the VSE, where the GRE header is removed and the packet is either forwarded across a S2S VPN tunnel to the tenant's remote location or alternatively NAT is used to allow Internet access. All of this is handled automatically by the HNV module, which is inserted into each of the virtual switches in the environment and requires no administrative effort.

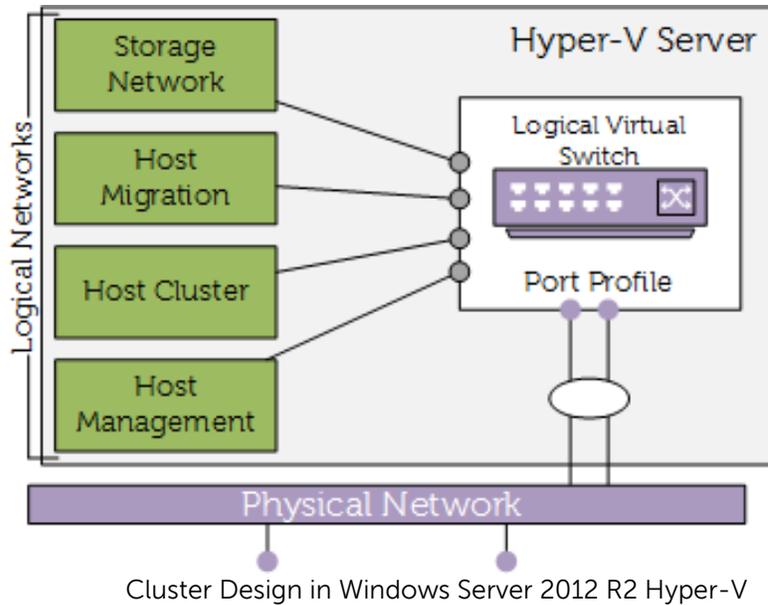
Note: For additional information on the NVGRE protocol please visit: <https://tools.ietf.org/html/draft-sridharan-virtualization-nvgre-08>

3.3 Microsoft Windows Server Hyper-V Cluster Design

When implementing virtualization, numerous services are consolidated onto fewer pieces of hardware. That hardware needs to be clustered to ensure resiliency for each individual Windows Server Hyper-V host. This redundancy is created using Microsoft Clustering Service. While SCVMM is used to create and deploy the virtual networks, System Center Configuration Manager (SCCM) is used to create the Hyper-V clusters.

By leveraging the failover cluster feature, Hyper-V servers can share storage resources enabling virtual machine migration. In the event of a node failure in a cluster, all virtual machines can be evacuated to other node members in the cluster. To enable this feature, four logical networks are leveraged: Storage Network, Host Migration, Host Cluster, and Host Management (Figure below). Each of these logical networks is segmented in individual VLANs to provide isolation throughout the data center. Once the target number of clusters has been created, an HNV gateway service is deployed to allow tenants access to their resources from the public Internet or their respective remote offices/campuses.





3.4 ^{Figure 12} Microsoft Windows Server HNV Gateway

Windows Server 2012 R2 together with SCVMM gives hosting providers a multi-tenant gateway solution that supports multiple host-to-host VPN tenant connections, Internet access for tenant virtual machines by using a gateway NAT feature, and forwarding gateway capabilities for private cloud implementations.

The HNV gateway is a virtual machine running on a Windows Server 2012 R2 Hyper-V cluster. There is a caveat, the Hyper-V cluster hosting the HNV gateway can only perform this service (HNV gateway). This cluster is dedicated to being only an ingress/egress point to the virtual networks. In the illustration in Figure 13, the HNV gateway encapsulates/decapsulates Ethernet data frames inbound/outbound to external networks.

In this RA, the HNV Gateway is deployed in route mode allowing tenants to terminate S2S tunnels directly to the gateway. These tunnels can be encrypted using standard tunnel encryption or in the scenario where a Silver Peak EdgeConnect is utilized to allow the Silver Peak device to provide tunnel encryption. In the event that a customer has multiple locations, iBGP can be enabled and used across these S2S tunnels to provide routing in the event that one of the tunnels fails.



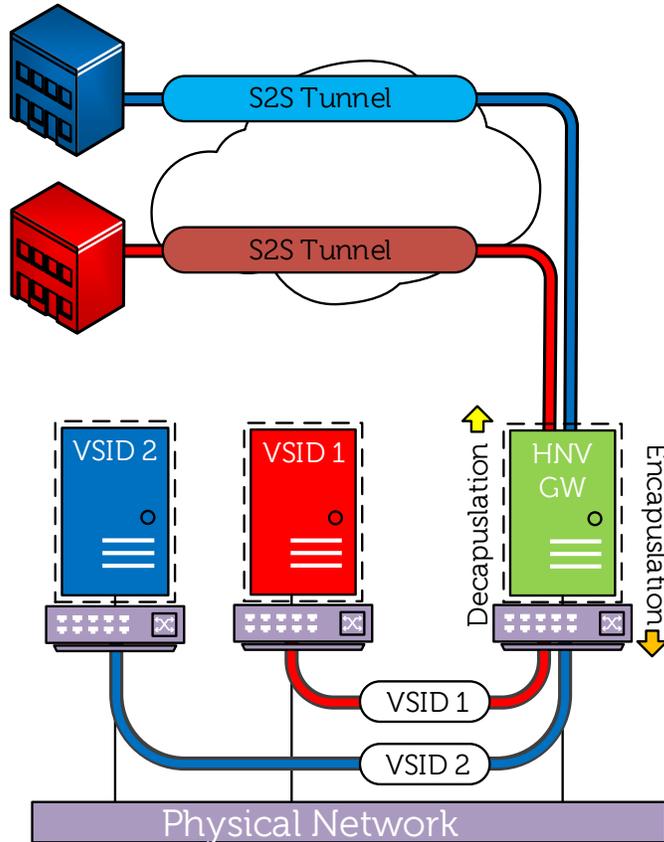


Figure 13 HNV Gateway Example

Note: For capacity planning purposes, one gateway cluster is recommended per 100 tenants.

3.5 Dell Compellent SC4020 and Hyper-V Using iSCSI

SCVMM is capable of integrating with nearly any type of storage solution using Storage Management Initiative Specification (SMI-S). SMI-S provides the ability to classify and assign storage to individual hosts or tenant clouds. Using the Dell Compellent SC4020, iSCSI is available in the underlay fabric enabling block level storage throughout the fabric.

To allow SCVMM to manage the Dell Compellent SC4020 an additional VM, Dell Enterprise Manager, needs to be deployed. This acts as the intermediary between SCVMM and the Compellent Storage Center. Enterprise Manager allows you to monitor, manage, and analyze multiple Storage Centers. It also behaves as a storage provider allowing SCVMM to view and manage LUNs through the SCVMM console. This is done using SMI-S. A storage provider, which is linked to the Enterprise Manager, can be loaded into SCVMM, which then displays the volumes and contained LUNs on the Compellent Storage Center instances. Each volume can be assigned a customer classification like Gold, Silver, and Bronze. Figure 14 shows a Compellent SC4020 with the classifications of Gold, Silver, and Bronze created. In this case, the Gold classification is expanded showing SSDs with multiple LUNs. Using classifications allows SCVMM to automatically create and allocate LUNs to hosts based on tenant requirements.

Name	Type	Size	Available Capacity	Assigned	Description	Provisioning Type
Bronze	Classification	0 GB	0 GB			
Gold	Classification	4,918.44 GB	4,788.80 GB			
Compellent SC4020 SSD	Storage pool	4,918.44 GB	4,788.80 GB			
DCRA FX2 Cluster Volume 1	Logical unit	1,024.00 GB		No	DCRA FX2 Cluster Volume 1	Thin
DCRA FX2 Cluster Witness	Logical unit	1.00 GB		No	DCRA FX2 Cluster Witness	Thin
TenantBlue	Logical unit	500.00 GB		No		Thin
Silver	Classification	0 GB	0 GB			

Storage managed by SCVMM 2012 R2

Figure

Note: For additional information on Dell Enterprise Manager for Compellent, please see: [Dell Enterprise Manager Configuration](#) (Dell Inc., 2013).

+ Additionally see this document for specific steps to deploying SMI-S and Dell Enterprise Manager with SCVMM: [Dell Enterprise Manager SCVMM configuration](#) (Dell Inc., 2015)

3.6 Implementing the Private Cloud

Once all the resources (compute, network, and storage) are implemented, the last decision to plan is the size and scope of tenant clouds. A cloud is a further SCVMM abstraction of the total number of resources available in the deployment. A tenant cloud consists of a number of key resources and configurations, including the following:

- Capacity (amount of memory, processors, and storage resources)
- Classifications of storage
- Available networks

Once the clouds are created and assigned to specific tenants, they will have their own quotas within the capacity of that cloud. However, to allow a tenant to use these resources various roles need to be created within SCVMM to allow those administrators the appropriate security access. To provide self-service capability commonly associated with the private cloud, the System Center App Controller (SCAC) application needs to be deployed. SCAC provides a web-based front end to expose the clouds defined in SCVMM to the user. App Controller allows tenants basic management function such as the creation and management of virtual machines and snapshots. Additionally it provides integration into the Microsoft Windows Azure environment, allowing a customer to migrate virtual machines to and from the Azure cloud.

3.7 Microsoft Server HNV Summary

When all of the components outlined in this section are brought together, the resulting converged network looks similar to the diagram in Figure 15. In this topology, there are two tenants, Red and Blue. Both tenants use the network ID 10.10.10.0/24 and through the deployment of Windows Server HNV both

tenants can exist in the datacenter on the same physical underlying network fabric. This enables each tenant to extend their campus or branch office networks into the datacenter.

Beyond tenant's virtual subnets, SCVMM is part of the management network and has connections to all layers of the HNV solution including storage via Dell Enterprise Manager and the NVGRE Gateway. Our storage solution, the Dell Compellent SC4020 along with Dell Enterprise Manager, can be managed through the SCVMM interface and provision volumes/LUNs for each tenant.

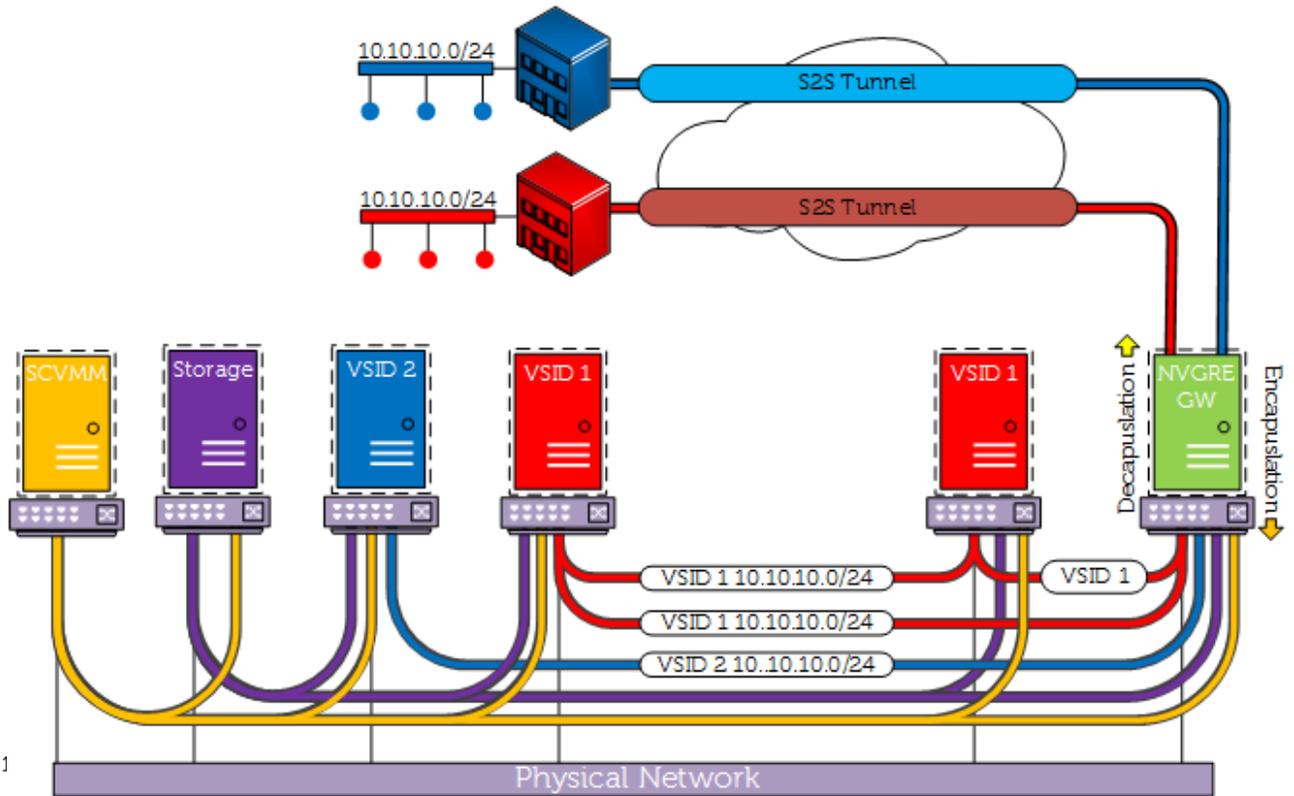


Figure 1

HNV Topology

3.8 Microsoft Server 2012 R2 and VMM Scaling Maximums

The following tables (Table 4-Table 6) contain the Microsoft approved VM maximums capable with Windows Server 2012 R2 Hyper-V and Microsoft Cluster Services. This information may provide assistance in planning your virtual compute and network environments. For additional information on Windows Server Hyper-V scaling visit: <https://technet.microsoft.com/en-us/library/JJ680093.aspx>.



Table 4 Hyper-V Maximums

Hyper-V Component	Maximum
Logical processors	320
Virtual processors per logical processors	320
Running virtual machines per server	1024
Virtual processors per server	2048
Memory	4 TB
Virtual switches	No limit
Virtual network switch ports per server	No limit

Table 5 Hyper-V Failover Cluster - Maximums

Hyper-V Failover Cluster Component	Maximum
Nodes per cluster	64
Running virtual machines per cluster	8,000 per cluster

Table 6 Virtual Machine - Maximums

Virtual Machine Component	Maximum
Virtual processors	64
Memory	1TB
Virtual hard disk (.VHDX)	64TB
Virtual IDE disks	4
Virtual SCSI controllers	4
Virtual SCSI disks	256
Virtual Fibre Channel adapters	4
Snapshots	50
Virtual network adapters	12



4 Services – Hybrid MPLS and Broadband Wide Area Networking with Silver Peak

4.1 Silver Peak SD - WAN Solutions

With many traditional data center based services residing in the cloud, many customers would like to leverage broadband connectivity rather than expensive Multi-Protocol Label Switching (MPLS) connectivity, which is expensive and can take months to deploy. However, a gap remains in that a single broadband connection may not provide the same level of availability, reliability, or security as MPLS. Silver Peak Unity EdgeConnect addresses these issues, allowing customers to leverage the economics of an all-broadband or hybrid MPLS + broadband WAN connection, but with similar or better reliability, availability, and security as MPLS.

Delivered in both physical and virtual form factors, Silver Peak Unity EdgeConnect creates a virtual network overlay, so that enterprises do not have to replace existing routers and firewalls at branch offices. With Unity EdgeConnect, customers can move to a broadband WAN at their own pace, whether site-by-site, or via a hybrid WAN approach that leverages both MPLS and broadband Internet connectivity (cable, DSL, LTE, etc.). This is often referred to as software-defined WAN (SD-WAN).

Unity EdgeConnect is an overlay solution, which means connectivity decisions are made independent of carriers, avoiding lengthy procurement and deployment delays for faster time to service. Connectivity can be rapidly extended, moved, or changed where and when the business demands. With a zero-touch, plug-and-play deployment model, Unity EdgeConnect can be deployed at a branch office in seconds, automatically connecting with other Silver Peak instances in the data center, other branches, or in cloud Infrastructure as a Service (IaaS) with the likes of Amazon, Microsoft Azure, and VMware's vCloud Air.

4.1.1 Unity EdgeConnect Deployment

With Unity EdgeConnect, customers can move to a broadband WAN at their own pace, whether site-by-site, or via a hybrid WAN approach that leverages both MPLS and broadband Internet connectivity (cable, DSL, LTE, etc.). Silver Peak has made installing Unity EdgeConnect simple and automated with zero-touch provisioning. When deploying the physical appliance, someone at the branch needs to simply plug in the appropriate WAN and LAN ports and the power cord and turn it on. The device immediately registers with the Unity Orchestrator located at a regional hub or a data center and indicates the branch is ready to be added to the WAN. The WAN administrator accepts the request to join the network and pre-determined business intent policies are automatically assigned to the device. For the virtual appliance, the software needs to be downloaded on any Dell PowerEdge (or other X86) server running any of the major hypervisors (ESXi, Xen, Hyper-V or KVM).



5 Validated Lab Environment

The lab environment that was used to validate portions of this architecture is listed in Table 7.

Table 7 Validated lab environment

Dell Product Solution	Description
PowerEdge R630, R730, and FX2s with FC630 Blade Servers	Compute for underlay
Compellent SC4020	Storage for VMs within the Hyper-V Cluster
Dell Networking S4048	Networking switches stacked for storage implementation
Dell Networking Z9100	Networking switches in Active Fabric configuration for Multi-Tenant Hyper-V Hybrid Cloud
PowerEdge R630 and Silver Peak EdgeConnect VM appliance	DUT WAN physical/virtual appliances for connectivity and optimization of MPLS/broadband links
Ixia IxLoad	Workload generation

Much of the validation environment focused on the availability of a strict multi-tenant configuration in which tenants never know of or see available resources of other tenants. In order to validate this, multiple VMs were set up for two separate tenants (Customer Blue, and Customer Red). Figure 16 diagrams this environment. East-west traffic within the overlay environment for each individual tenant was then tested. It was verified that this traffic between a tenant's VMs was always accessible with very quick access and very low latency. This could only be observed in our testing with multiple VM implementation.

The multi-tenant environment was then checked by verifying that Customer Red could not reach Customer Blue's environment unless special configurations were put into place by the administrator to bypass typical security and configuration. In other words if Customer Red was hosting a website available to the public this could be configured to be accessible by anyone including Customer Blue but that traffic would exit the overlay and come back in as any other public traffic on the internet.

While the whole solution is based on customer multi-tenancy, it also provides great benefits to the branch office and campus environments by gaining access back into the data center via WAN optimization and acceleration. Silver Peak's Unity EdgeConnect is used in this solution to implement these capabilities in a manner that can be easily benchmarked for performance.



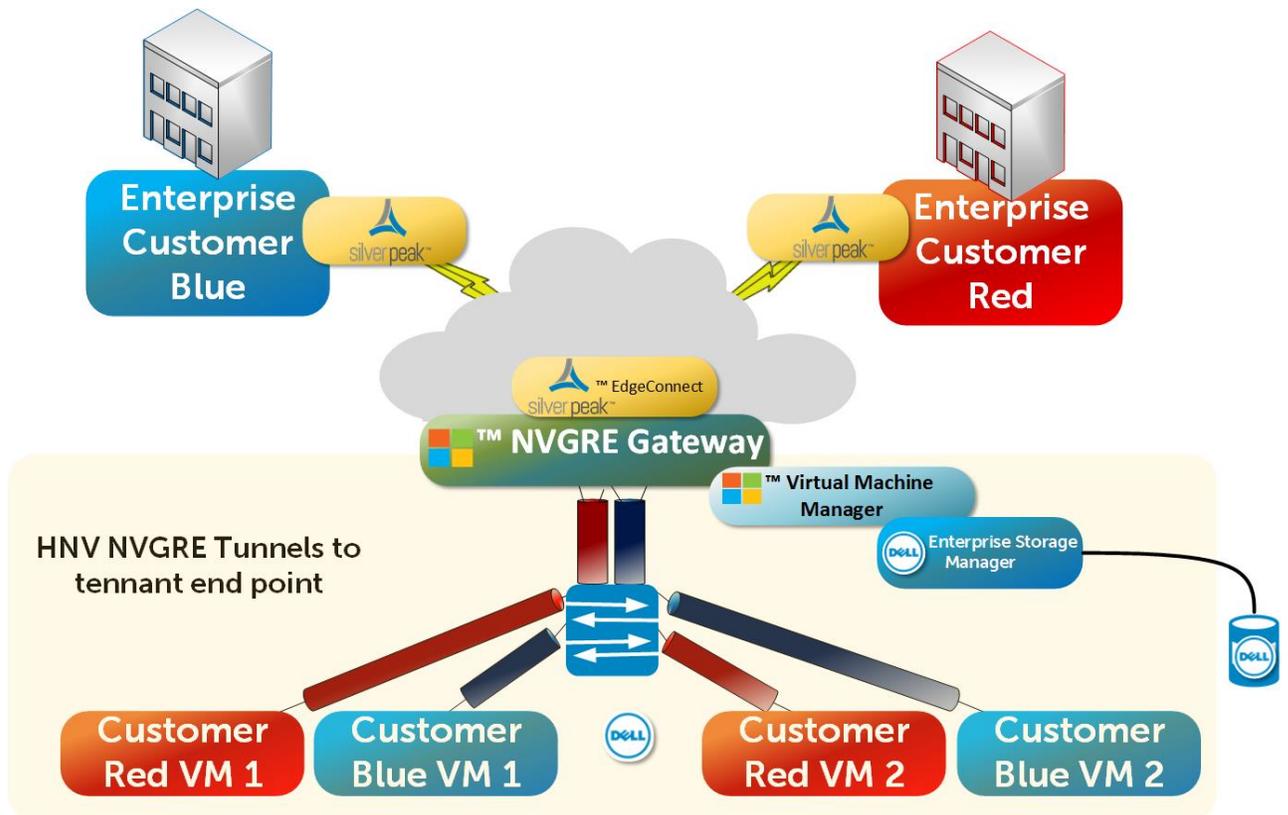


Figure 16 Validation scenario

For this reference architecture, Silver Peak partnered with Dell to set up a three-node Hybrid WAN environment consisting of:

- **Data Center** – with Unity EdgeConnect EC-L physical appliance and optional Unity Boost Software (for Acceleration and Bandwidth Optimization). Connectivity into this Data Center location via dual MPLS and Internet 1Gbps WAN links
- **Campus** - with Unity EdgeConnect EC-V and optional Unity Boost Software (for Acceleration and Bandwidth Optimization) virtual appliances running on a Dell PowerEdge host running Microsoft Hyper-V (for Acceleration and Bandwidth Optimization). Connectivity to this Campus location was via a 100Mbps Internet Broadband link and a 100 Mbps MPLS circuit
- **Branch** - with Unity EdgeConnect EC-V and optional Unity Boost Software (for acceleration and bandwidth optimization) virtual appliances running on a Dell PowerEdge host running Microsoft Hyper-V. Connectivity to this branch location was via a 100Mbps Internet broadband link and a 100Mbps MPLS link.

The following two scenarios were run to demonstrate the capabilities of Silver Peak Unity EdgeConnect and Dell Networking end-to-end solutions.



5.1 Application - Scenario One

Scenario one consisted of video streaming with dynamic path control and failover from MPLS to broadband (Figure 17).

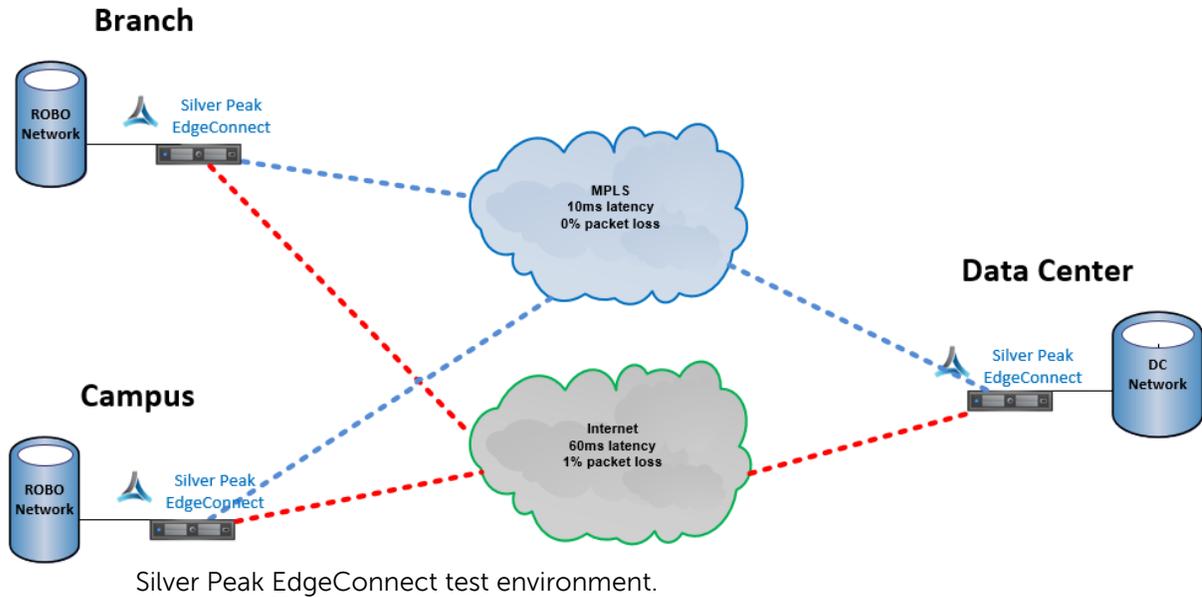


Figure 17

In this test, 300 simultaneous video streams were transmitted via HTTP from the Data Center to the Campus location – initially over the MPLS connection (Figure 18). The MPLS connection was then disconnected. Silver Peak Unity EdgeConnect advanced dynamic path control (DPC) was able to fail over the streaming video connection to the Internet broadband circuit without a degraded experience for the end user (Figure 19). By using Silver Peak’s unique WAN path conditioning on the Internet broadband link, packet loss was virtually eliminated, retaining the quality of the video over this link.

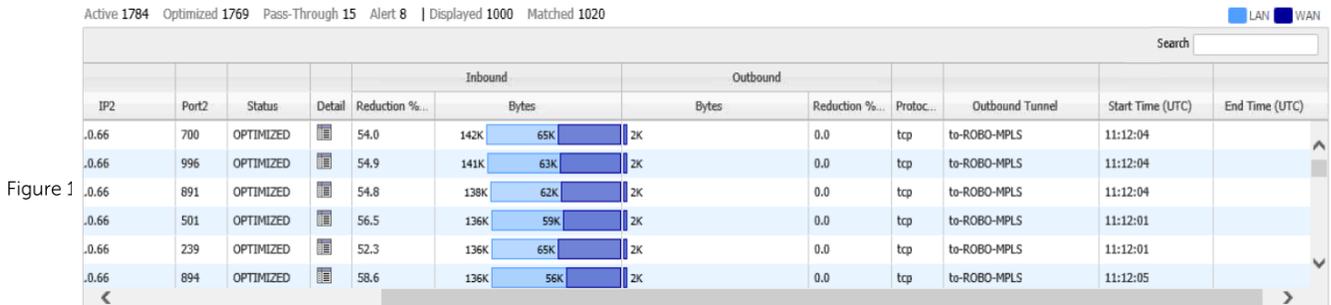


Figure 18

HTTP video streams traversing the MPLS link to the branch office



Active 10448 Optimized 10445 Pass-Through 3 Alert 2 | Displayed 1000 Matched 10448

Search

	IP2	Port2	Status	Detail	Inbound		Outbound		Protoc...	Outbound Tunnel	Start Time (UTC)	End Time
					Reduction %...	Bytes	Bytes	Reduction %...				
0	10.22.0.66	258	OPTIMIZED		93.9	24M	1.5M	401K	27.5	tcp	to-ROBO-INTERNET	10:39:24
	10.22.0.63	80	OPTIMIZED		48.0		421K	2.1M	89.8	tcp	to-ROBO-INTERNET	10:43:29
	10.22.0.63	80	OPTIMIZED		47.9		427K	1.9M	90.5	tcp	to-ROBO-INTERNET	10:45:40
	10.22.0.63	80	OPTIMIZED		49.8		426K	2.1M	89.9	tcp	to-ROBO-INTERNET	10:44:43
7	10.22.0.63	80	OPTIMIZED		49.6		426K	1.9M	90.8	tcp	to-ROBO-INTERNET	10:46:13
	10.22.0.63	80	OPTIMIZED		49.2		416K	2.0M	90.0	tcp	to-ROBO-INTERNET	10:45:32

HTTP video streams after failing over from the MPLS link to the Internet link

5.2 Application - Scenario Two

Figure 19

Scenario two involved remote office collaboration. Specifically file transfer from data center to remote office with Silver Peak Unity EdgeConnect Boost. In this scenario, five files were transferred from the data center to the remote office via multiple FTP streams without Silver Peak optimization (Unity Boost) enabled. The non-optimized transfer over the 100Mbps maxed out at about 15Mbps throughput (Figure 20). Next, Silver Peak Unity EdgeConnect's Boost capability was enabled. As can be seen from Figure 21, with Unity Boost enabled, the optimized transfer peaked out at more than 500Mbps on the LAN, while only using 15Mbps of the WAN bandwidth - resulting in a 30X overall increase in available bandwidth.

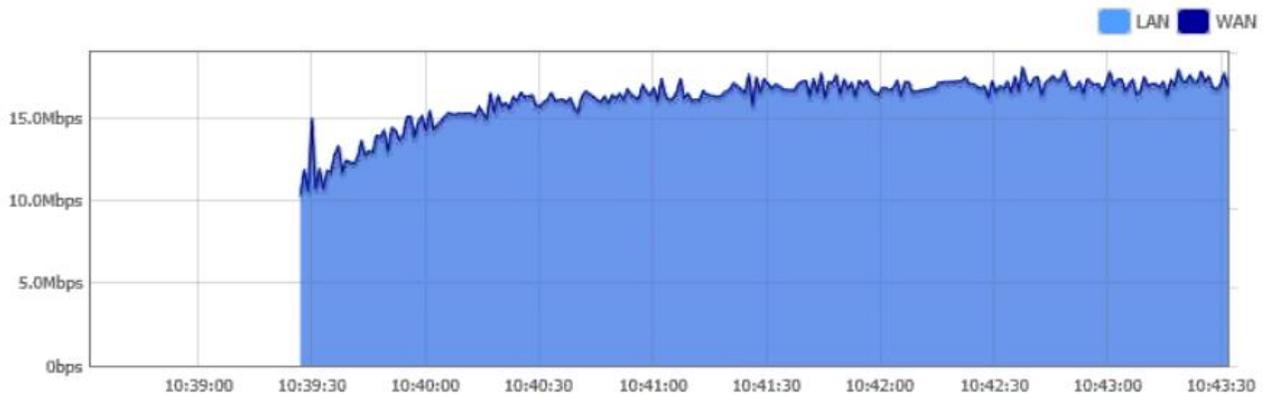
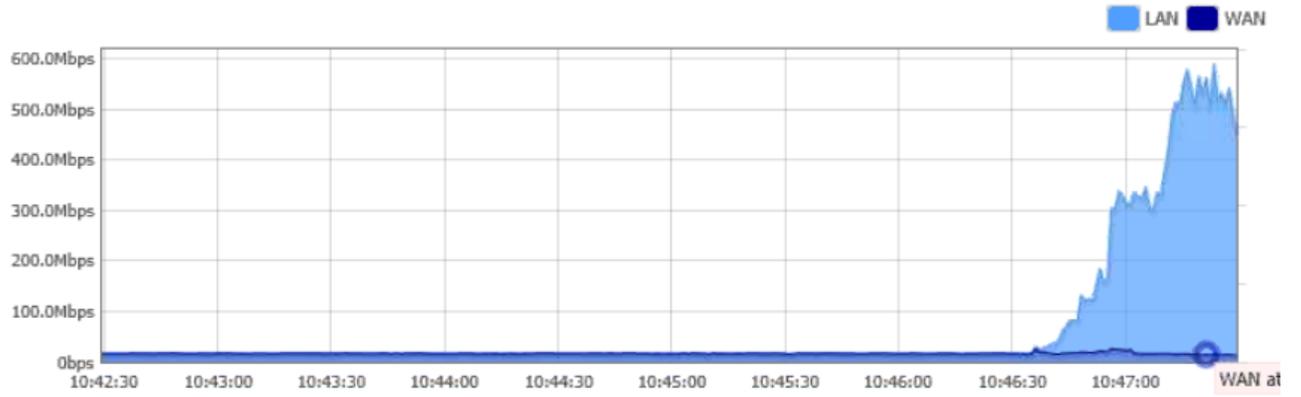


Figure 20

Throughput graph of non-optimized FTP traffic streams





Throughput graph of Silver Peak optimized FTP streams

Figure 21



6 Products used in this Solution

Dell Networking, Dell Compellent Storage, and Dell 13th Generation PowerEdge Servers are used in this architecture. The products are shown with some key features in the following subsection.

6.1 Dell Networking Products

6.1.1 Dell Networking Z9100-ON



Dell Networking Z9100-ON - 10/25/40/50/100 GbE Multi-rate top-of-rack ON switch

Figure 22 The Dell Networking Z9100-ON is a 10/25/40/50/100GbE top-of-rack (ToR) fixed port switch purpose-built for applications in high-performance data center and computing environments. Leveraging a non-blocking, cut-through switching architecture, the Z9100-ON delivers line-rate L2 and L3 forwarding capacity to maximize network performance. The compact Z9100-ON design provides industry-leading port density of either 32 ports of 100GbE, 64 ports of 50GbE, 32 ports of 40GbE, 128 ports of 25GbE or 128 ports 10GbE and two SFP+ ports of 10GbE/1GbE/100MbE to conserve rack space while enabling a smaller footprint and simplified migration to 100Gbps in the data center core. Priority-based flow control (PFC), Data Center Bridge exchange (DCBx) and enhanced transmission selection (ETS) make the Z9100-ON ideally suited for DCB environments. In addition, the Z9100-ON incorporates multiple architectural features that optimize data center network flexibility, efficiency, and availability, including redundant, hot-swappable power supplies and fans. These new offerings provide the flexibility to transform data centers and offer high-capacity network fabrics that are easy to deploy, cost-effective and provide a clear path to a software-defined data center. The Dell Z9100-ON supports the industry standard Open Network Install Environment (ONIE) for zero-touch installation of alternate network operating systems. This document refers to this ON switch preloaded with the Dell Networking OS. Characteristic of any ONIE device, other ONIE load images may be loaded by the operator.

6.1.2 Dell Networking S4048-ON



Dell Networking S4048 – 10/40 Gb data center switch

The S4048-ON is a 10/40GbE top-of-rack 1U switch designed for high-performance data center and computing environments. Leveraging a non-blocking switching architecture, the S4048-ON delivers line-rate layer 2 and layer 3 forwarding capacity with ultra-low-latency to maximize network performance.

The compact design of the S4048-ON allows for industry-leading density of 48 dual-speed 1/10GbE (SFP+) ports as well as six 40GbE QSFP+ uplinks to conserve rack space and simplify the migration to 40Gbps in the data center core. Each 40GbE QSFP+ uplink can support four 10GbE ports with a breakout cable delivering 72 potential 10-gigabit ports.

The S4048-ON also supports the open source Open Network Install Environment (ONIE) for easy installation of alternate network operating systems such as Dell's own industry-hardened OS 9, Cumulus Networks Linux OS, and Big Switch Networks Switch Light OS. This provides the option to build an application-agnostic network infrastructure and simplify data center network management, orchestration, and automation.

6.2 Dell PowerEdge Server Products

6.2.1 Dell PowerEdge R630 Server



Figure 24 Dell PowerEdge R630 server

With computing capability previously only seen in 2U servers, the PowerEdge R630 two-socket 1U rack server delivers an impressive solution for virtualized environments, large business applications, or transactional databases. The PowerEdge R630 server is versatile and highly configurable for a variety of solutions, delivering the latest Intel Xeon processor E5-2600 v3 product family, 24 DIMMs of high-performance DDR4 memory and a broad range of local storage options.

Build a blazingly fast transactional database with the R630 24 all-flash drive configuration. Combine the density, scalability and performance of 24 1.8" SSDs with dual-PERC capability to deliver a superior solution that can boost performance by storing an entire database. When combined with the SanDisk DAS Cache solution, the R630 can help increase OLTP transactions and concurrent usage for latency-sensitive applications, such as e-Commerce or financial trading.

6.2.2 Dell PowerEdge R730 Server



Dell PowerEdge R730 server

The Dell PowerEdge R730, an extremely versatile two-socket 2U rack server for mid-sized and large enterprises, such as data warehouses, e-commerce, virtual desktop infrastructure (VDI), databases and high-performance computing (HPC) as a data node, is a key to strong datacenter compute performance. With the Intel Xeon processor E5-2600 v3 product family and state-of-the-art DDR4 memory workloads of all types can be delivered with maximum performance.

PowerEdge servers allow the construction and management of highly efficient infrastructures for data centers and small businesses. Accelerate time-to-production with automated deployment processes that use fewer manual steps and reduce the potential for error. Improve IT productivity with innovative server management tools like iDRAC Direct and iDRAC Quick Sync to deliver in-depth system health status and speed deployment, and optimize data center energy usage with improved performance-per-watt computation.

6.2.3 Dell PowerEdge FX2 Server Chassis



Dell PowerEdge FX2 server chassis and FC630 blade servers

Tailor a data center according to needed workloads and redefine data center agility with the revolutionary FX architecture. The FX2 converged 2U rack chassis allows for incrementally adding or swapping compact, modular IT building blocks, including servers, storage and IO, to fit any computing environment

The FX2 provides more choice, less complexity. Simplify building and managing data centers, public, private, and hybrid clouds by bringing advanced management technologies to converged infrastructure.

Right size your IT with the FX2. Easily and rapidly scale the resources your workloads and business require.

Fit the infrastructure to the job. Tailor infrastructure precisely, with the right power, storage, and connectivity to meet specific workload needs.

6.3 Dell Storage Products

Dell storage products play a vital role in the data center. These products provide resilient and redundant storage for any type of data center needs.

6.3.1 Dell Compellent SC4020 Storage



Compellent SC4020 without bezel and with bezel

Figure 27

The SC4020 is ideal for standalone SANs or departmental/branch office deployment. It was designed to keep future scalability and large-scale growth capabilities in place through deployments of various sizes. The dual controller, 24 internal drive slots and high-performance Fibre Channel or iSCSI network ports come standard in the space-saving 2U chassis.

6.3.2 Dell Compellent SC220 Storage



Compellent SC220 without Bezel and with Bezel

Figure 28

Dell Compellent offers customers an agile, high-performance storage platform that delivers ongoing economic value. With its open, scalable hardware platform, Compellent allows for the intermix of the latest industry-standard drive technologies in the same array with the energy-efficient SC200 and SC220 expansion enclosures. The introduction of new and emerging technologies can be introduced over time and on the fly to support your SAN and NAS environments on demand without disruption or downtime.

6.4 SD WAN Solutions

6.4.1 Silver Peak Unity EdgeConnect

Key features of Unity EdgeConnect include:

- **Dynamic Path Control (DPC)** – real-time traffic steering over any broadband or MPLS link based on company-defined business intent policies. In the event of an outage or brownout, DPC automatically fails over to the secondary connection in about one second.
- **WAN Hardening** – data is secured edge-to-edge via 256-bit AES encrypted tunnels. No unauthorized outside traffic is allowed to enter the branch. WAN hardening secures branch offices without the appliance sprawl and operating costs of deploying and managing dedicated firewalls.
- **Path Conditioning** – overcomes the adverse effects of dropped and out-of-order packets that are common with broadband Internet and MPLS connections. Path Conditioning provides private-line-like performance over the public Internet.
- **Cloud Intelligence** – delivers real-time updates on the best performing path to reach hundreds of software as a service (SaaS) applications ensuring users connect to their applications in the fastest, most intelligent way available.

6.4.1.1 Boost Application Performance on the Fly

When branch offices are deployed as part of a broadband or hybrid WAN, customers may require higher performance for specific applications, for example, accelerating replication data over distance for disaster recovery. Silver Peak Unity Boost is an optional performance pack that accelerates application performance as needed and includes:

- **Latency Mitigation** – TCP and other protocol acceleration techniques minimize the effects of latency on application performance and significantly improve application response time across the WAN.
- **Data Reduction** – WAN compression and deduplication eliminate the repetitive transmission of duplicate data. Silver Peak software inspects WAN traffic at the byte level and stores content in local data stores. Advanced fingerprinting techniques recognize repetitive patterns for local delivery, granular visibility, and control with Unity Orchestrator.

6.4.1.2 Unity Orchestrator: Centralized Management and Zero-Touch Provisioning

Unity Orchestrator is included with Unity EdgeConnect, and provides customers with unprecedented levels of visibility into both legacy and cloud applications, and the unique ability to centrally assign business intent policies to secure and control all WAN traffic. Key features of Unity Orchestrator include:

- **Single Screen Administration** – enables quick and easy implementation of network-wide business intent policies, which eliminates complex and error-prone policy changes at every branch.
- **Granular Real-Time Monitoring and Historical Reporting** – provides specific details into application, location, and network statistics, including continuous performance monitoring of loss, latency, and packet ordering for all network paths. All HTTP and native application traffic are



identified by name and location, and alarms and alerts allow for faster resolution of service provider issues.

- **Bandwidth Cost Savings Reports** - document the cost savings for moving to broadband connectivity.

6.4.1.3 Unity EdgeConnect Packaging

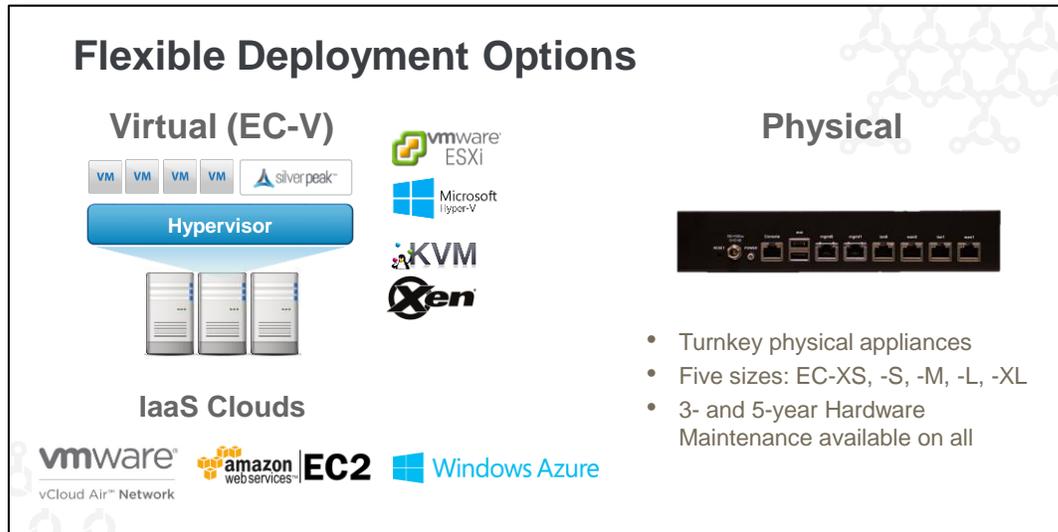


Figure 29 Unity EdgeConnect can be deployed as a physical or virtual appliance

	EC-XS	EC-S	EC-M	EC-L	EC-XL
Typical Deployment	Small Branch	Large Branch	Small Hub	DC/Large Hub	DC/Large Hub
WAN BW	2-200Mbps	10-1000Mbps	50-2000Mbps	1-6Gbps	2-10Gbps
Interfaces	4xRJ-45 10/100/1000	6XRJ45 10/100/1000	4 X RJ45 2x1/10G Fiber	4 X RJ45 2 X 1/10G Fiber	4 x 1/10G Fiber

Figure 30 Unity EdgeConnect hardware appliance options

Silver Peak solutions have been deployed in the real world in thousands of locations globally and can scale from the smallest locations to the largest data center applications.

For more information on Silver Peak solutions, please see the following link: <http://www.silver-peak.com>



A Works Cited

TechCenter Extras: Hyper-V 2012 R2 Best Practices for Dell Compellent Storage Center. Dell Inc. (2013, October 28) Retrieved from Dell TechCenter:

http://en.community.dell.com/techcenter/extras/m/white_papers/20437923

Dell Enterprise Manager 2015 R1 Administrator's Guide. Dell Inc. (2015, June 1) Retrieved from Dell.com:

http://www.dell.com/support/manuals/us/en/19/Topic/storage-sc2020/EM2015R1_AG-v4/en-us/GUID-B569C2CD-B82C-4F91-B4FE-A9096291B253

Dell Storage SC4020 all-in-one array. Dell Inc. (2015) Retrieved from Dell.com:

<http://www.dell.com/us/business/p/dell-compellent-sc4020/pd>

Knowledge Base: How to set up Virtual Link Trunking (VLT) on Dell Networking Force10 Switches. Dell Inc. (2015, October 1) Retrieved from support.dell.com:

<http://www.dell.com/support/article/us/en/19/HOW10660/EN>

PowerEdge R730. Dell Inc. (2015). Retrieved from Dell.com: [http://www.dell.com/us/business/p/d/shared-](http://www.dell.com/us/business/p/d/shared-content~data-sheets~en/Documents~Dell-PowerEdge-R730-Spec-Sheet.pdf.aspx?layoutvariation=modal&modaltyp=box&position=center&modalwidth=600&modalheight=600&modalscroll=yes&ovrcolor=%23000000&ovropac=50&modaltarge)

[content~data-sheets~en/Documents~Dell-PowerEdge-R730-Spec-Sheet.pdf.aspx?layoutvariation=modal&modaltyp=box&position=center&modalwidth=600&modalheight=600&modalscroll=yes&ovrcolor=%23000000&ovropac=50&modaltarge](http://www.dell.com/us/business/p/d/shared-content~data-sheets~en/Documents~Dell-PowerEdge-R730-Spec-Sheet.pdf.aspx?layoutvariation=modal&modaltyp=box&position=center&modalwidth=600&modalheight=600&modalscroll=yes&ovrcolor=%23000000&ovropac=50&modaltarge)



B Additional Resources

www.Dell.com is the site for all of your Enterprise Product needs.

Support.Dell.com is focused on meeting your needs with proven services and support.

DellTechCenter.com - [Networking Guides](#) is an IT Community where you can connect with Dell Customers and Dell employees to share knowledge, best practices, and information about Dell products and installations.

Referenced or recommended publications:

- Dell EqualLogic Configuration Guide:
<http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/19852516>
- Microsoft SQL Server 2008: Disk Partition Alignment Best Practices for SQL Server:
<http://msdn.microsoft.com/en-us/library/dd758814.aspx>
- Coming of age: The new economics of 100GbE
http://www.bitpipe.com/data/document.do?res_id=1439913427_144&nobrand=true&src=5421293&asrc=EM_SRU_46981625&uid=9758775&utm_medium=EM&utm_source=SRU&utm_campaign=20150831_SRU+Transmission+for+08%2F31%2F2015+%28UserUniverse%3A+1703023%29_myka-reports%40techtarg.com
- Learning Hyper-V
By: Vinícius R. Apolinário
Publisher: Packt Publishing
Published Date: May 30, 2015
Web ISBN-13: 978-1-78439-961-0
Print ISBN-13: 978-1-78439-986-3
Pages in Print Edition: 384
- Mastering Hyper-V 2012 R2 with System Center and Azure
By: John Savill
Publisher: Sybex
Published Date: April 14, 2014
ISBN-13: 978-1-118-82833-5
Web ISBN-10: 1-118-82818-6
Print ISBN-13: 978-1-118-82818-2
Pages in Print Edition: 576



- Microsoft System Center Virtual Machine Manager 2012 Cookbook
By: Edvaldo Alessandro Cardoso
Publisher: Packt Publishing
Published Date: March 26, 2013
Print ISBN-13: 978-1-84968-632-7
Web ISBN-13: 978-1-84968-633-4
Pages in Print Edition: 342
- Microsoft System Center: Building a Virtualized Network Solution
By: Nigel Cain; Alvin Morales; Michel Luescher; Damian Flynn; Mitch Tulloch
Publisher: Microsoft Press
Published Date: February 18, 2014
Print ISBN-13: 978-0-7356-8310-5
Pages in Print Edition: 136
- Optimizing and Troubleshooting: Hyper-V Networking
By: Mitch Tulloch
Publisher: Microsoft Press
Published Date: July 15, 2013
Print ISBN-10: 0-7356-7900-2
Web ISBN-10: 0-7356-7913-4
Web ISBN-13: 978-0-7356-7913-9
Print ISBN-13: 978-0-7356-7900-9
Pages in Print Edition: 132
- Optimizing and Troubleshooting: Hyper-V Storage
By: Mitch Tulloch
Publisher: Microsoft Press
Published Date: July 15, 2013
Print ISBN-10: 0-7356-7898-7
Print ISBN-13: 978-0-7356-7898-9
Web ISBN-10: 0-7356-7909-6
Web ISBN-13: 978-0-7356-7909-2
Pages in Print Edition: 152
- Silver Peak Unity EdgeConnect information to Optimize or Accelerate your Data Center to remote site connections
<http://www.silver-peak.com>



C Support and Feedback

Contacting Technical Support

Support Contact Information

Web: <http://Support.Dell.com/>

Telephone: USA: 1-800-945-3355

Feedback for this document

We encourage readers of this publication to provide feedback on the quality and usefulness of this deployment guide by sending an email to Dell_Networking_Solutions@Dell.com

