Reference Architecture Guide

HP FlexFabric Reference Architecture Guide

Version 3.0

Applying HP Converged Infrastructure to Data Center Networks

Table of Contents

Introduction

The primary driver for evolving today's enterprise data center is efficiently deploying resources to support business applications. Apparent trends in consolidation and distribution of infrastructure are side effects of addressing this core goal. In fact, enterprise data center network architects and managers are now expected to build networks that can concurrently consolidate and geographically distribute resources. These include physical and application servers, data storage, management platforms, and network devices. This evolution did not happen overnight. It has been fueled by the accelerating needs of businesses to be more agile, to do more with less, and to increase their IT efficiency.

The current trend in network architecture is virtualization. Virtualization encompasses the ability to operate multiple servers concurrently on top of a single server hardware platform, sharing CPU, disk, interface, and network services. In this case, each virtual server operates as an independent entity on a single physical server. Virtualization does not dismiss the need for network convergence. Early deployments of virtualized servers operated statically on the platform they were deployed on. Today, virtualized servers are flexible in their deployment, with the ability to move to other physical server platforms. Moving with the virtualized server is their configured storage, data, and multi-modal communication functions. This type of infrastructure can now significantly reduce equipment, network facility and operational expenses.

In an effort to address this opportunity, traditional data center models have stretched to support autonomous data and storage networks with separate interface cards, switches, and cabling plants. This methodology has proven to be ineffective and costly when implementing and operating both centralized and geographically dispersed data centers.

To reduce complexity, data center architects adopted designs that utilize network convergence, where data and storage I/O are merged onto a single network. This converged approach can eliminate physical clutter and complexity, while making more efficient use of networking resources. However, the simplification, in many cases, is only true at the surface level. Data center networks today can be more complex below the surface than ever before. Virtualization has presented a substantial number of challenges which are driving a new evolution in data center network architecture that far exceeds the initial need for a converged network infrastructure.

HP has developed the FlexNetwork architecture, a component of the HP Converged Infrastructure, to reduce the complexity of data center networks, especially now that the primary requirement is networking virtual instances.

Enterprises can align their networks with their business needs -even as they change - by segmenting their networks into four interrelated modular building blocks that comprise the HP FlexNetwork Architecture: FlexFabric, FlexCampus, FlexBranch, and FlexManagement.

The HP FlexFabric architecture extends the HP Converged Infrastructure to the data center and provides the following benefits:

- Combines advanced, standards-based platforms and advanced networking technologies to optimize performance and latency in virtualized server environments
- Converges and secures data center network, compute, and storage in the physical and virtual worlds
- Reduces complexity
- Enables rapid businesses-aligned network provisioning
- Lowers total cost of ownership
- Consolidates multiple protocols into a single fabric that can easily flex with changing workloads

HP has developed a series of reference architecture content to define the technical aspects of the various FlexNetwork architectures. As part of that series, this document describes the HP FlexFabric reference architecture.

As part of the HP FlexNetwork architecture, this reference architecture guide provides design guidelines within an architectural framework for data center environments. Key business and technology drivers are discussed along with key networking requirements for data centers. Physical infrastructure and design models are discussed. Other aspects of networking at the data center is discussed, including data center interconnect, advantages and disadvantages to 1 to 3 tier data centers, data center security, multi-tenancy, network management, and much more.

HP believes simplification is the overriding key to success in networks supporting virtualization.

This approach allows data center architects and IT teams to develop new and more flexible data center models and methodologies. By doing so, IT can meet new demands head-on, rather than forcing businesses to adapt to technology

Overview of the Reference Architecture

This guide is intended for technology decision-makers, solution architects and other experts tasked with improving data center networking. It can serve as a baseline for network planning and design projects.

The Reference Architecture provides a high level best practice document from the perspective of Application and Virtualization network support within the data center. The Reference Architecture provides a holistic view in relationship to the data center, with flexible products for the enablement of the network. The accompanying "HP FlexFabric Deployment Guide" provides specific configuration examples, and best practices that should be followed in the data center.

Customers have seen the advantage of a dual vendor strategy within their networks. This reference architecture allows for, and incorporates, this dual vendor strategy, with representation not only with 3rd party devices such as Riverbed and F5, but with application of a cookie cutter approach for a Data Center design. This strategy allows for incorporation of HP and other vendors' switches and routers into the architecture as needed.

This document will frequently reference both technology trends which have, or are being driven through virtualization and standards, and those that have not. It will also introduce issues that confront data center architects in this fast-paced, results- driven and security-minded industry. It will provide guidance on network simplification for virtualized deployments that do not sacrifice performance or deployment flexibility. This guide also provides another level of detail to complement the HP Converged Infrastructure Reference Architecture Solution Block Design Guide.

This guide is divided into seven main sections:

- **Key business and technology drivers for a new data center network architecture:**

Brief introduction to technology drivers and data center deployment models

- **Key data center networking requirements:**

  Discusses virtualization technologies, including server deployment and virtualization considerations, as well as virtual machine (VM) mobility in Layer 2 networks. Also discusses drivers for Layer 2 networks, securing the virtual edge, managing and provisioning the virtual server edge, and converged network infrastructures

- **Data Center Network Design – HP FlexFabric Reference Architecture (FFRA):**

  Discusses "Top of Rack" and "End of Rack" considerations and provides descriptions to HP's one tier, two tier, and legacy data center models, including the logical layers of the design

- **HP Data Center Interconnection (DCI):**

  Discusses key considerations when deploying DCI, and HP's supported solutions based on the network resources available

- **HP Data Center Networking Portfolio:**

  Brief introduction to the HP FlexFabric product line

- **Partnering architecture:**

  Provides a brief overview of HP AllianceONE partner F5 and the solutions they provide

- **HP Networking Services:**

  Brief introduction to HP services

Please note, this document will be periodically updated, identified by date and version number.

Key business and technology drivers for a new data center network architecture

Large-scale data center consolidation

For many enterprise customers, the data center “is” the business. Mission-critical applications and services provide the foundation for day-to-day operations and delivery of end-customer services. The data center must deliver unquestioned availability and meet stringent service level agreements (SLAs). Exploiting server virtualization and low-cost computing power, customers are deploying more sophisticated applications on a larger scale. To reduce their complexity and improve operations in these deployments, customers are seeking to consolidate fragmented, dispersed facilities into fewer, centralized locations.

Today’s data center networks must be designed to deliver much higher levels of performance, scalability, and availability than ever before to meet service-level agreements and maintain continuity of operations. Beyond sheer performance, these data center networks must quickly recover from hardware- or software-related faults and protect against server, storage, network, and application vulnerabilities to maintain performance and minimize service disruptions.

Blade Servers and server virtualization technologies

The adoption of increasingly powerful multi-core-processor servers, higher-bandwidth interfaces, and blade servers is dramatically increasing the scale of data center deployments. Now, tens of

thousands of VMs are commonly deployed in a single data center to consolidate infrastructure and streamline operations. These large-scale solutions are dramatically increasing network performance requirements at the server edge and across the extended network. Likewise, virtualization and vMotion/Live Migration tools for moving virtual servers between machines in the same data center, or across geographically separated data centers are introducing high-volume machine-to-machine traffic flows. This impacts existing administrative practices, creating a new “virtual edge” which blurs the traditional boundaries between network and server administration.

New application deployment and delivery models

Traditional client/server software and infrastructure deployment models are being displaced by new application architectures and service delivery models that are reshaping the data center.

Web 2.0 mash-ups, service-oriented architecture (SOA) solutions, and other federated applications are being widely deployed to deliver integrated, content-correlated, context-specific information and services to end-users within the enterprise and beyond. These deployments drive new, bandwidth-intensive traffic flows within the data center and demand low-latency, high-performance server-to-server and intra-server, VM-to-VM connections.

At the same time, cloud computing and XaaS initiatives are introducing more stringent service level and security demands and driving requirements for a more agile and dynamic infrastructure. Now, employees, customers, and partners could be accessing your applications from almost anywhere—from headquarters, the campus, branch offices, or from any remote location—and the applications may be in a traditional or cloud data center.

The explosion of the cloud creates new opportunities and a new set of challenges. Networks must be faster and more flexible to support the needs of diverse mobile users, a fragmented security perimeter, and a constantly changing set of applications and devices.

To accelerate the move to the cloud, applications must be characterized prior to connecting them to the network and opening up access to users. Characterizing applications first enables definition of the necessary network resources, verification of resource availability, and aligns the resources with the application—allowing the network to deliver the expected service level.

Data center deployment models

The adoption of more virtualized, dynamic application environments is impacting traditional enterprise and hosted/multi- tenant data center designs. These methods enable new cloud-based delivery models that drive a whole new set of technology requirements across servers, storage, and networking domains. These increasingly popular models let enterprises provision applications more flexibly within traditional internal infrastructures, and enable hosted application and service providers to build entire businesses based on delivering services via a public cloud model. Given the range of use cases and options, customers often deploy a combination of architectures to address varied requirements and to optimize operations.

Table 1 summarizes some of the most important networking focus areas that emerge as customers pursue these diverse deployment models. While all these imperatives play some role across all the deployment models regardless of market or industry, certain initiatives figure more prominently in specific use cases.

Table 1 Data center deployment models and corresponding key networking imperatives

Deployment Model	Characteristics	Key Networking Focus Areas
Traditional Enterprise Data Center	DC services are a critical complement to the company's core business Complex application environment Security, cost and flexibility are key Evolving towards private cloud over time	Converged networking Virtualization scale-out Managing/provisioning the virtual server edge
Traditional Multi-tenant Data Center	DC services are the company's core business Complex application environment Security, SLAs and flexibility are key Evolving towards public cloud over time	Virtualization scale-out Securing the virtual server edge Managing/provisioning the virtual server edge
Multi-tenant XaaS/ Cloud Computing Data Center	DC services may be the company's core business Heavy use of blade servers Cost, latency and scalability are key	Virtualization scale-out Securing the virtual server edge Managing/provisioning the virtual server edge
High-performance Computing Data Center	DC services may be the company's core business Heavy use of blade servers Cost, latency, performance and scalability are key	Low latency

Key data center networking requirements

Virtualization scale-out: Layer 2 performance, scaling, high availability, and multi-site extension

Overview

Virtualization is now broadly accepted and deployed globally across data centers. With it, the operating systems, applications, and servers can work in a non-dedicated or loosely coupled manner, interacting as needed to meet enterprise needs. Virtualization provides several key benefits:

- **Higher efficiency and reduction in capital and operational expenditure:**

  As multi-core/multi-processor systems provide a substantial amount of computing performance, virtualization allows for the more efficient use of a given server and its resources. Capital expenditure on hardware and software, as well as operation and support costs, are reduced
- **Agility and flexibility:**

  Data center managers must keep operations running efficiently by rapidly deploying resources for business growth and for stress periods. Application developers look for flexibility in CPU and memory utilization. Virtualization provides the ability to do on-the-fly migration of virtual servers or applications across physical servers in the data center
- **Resiliency:**

Virtualization technology provides the ability to restack or shuffle applications in support of business continuity or even for routine maintenance

Virtualization creates the need for new data center network architecture methodologies. Traditional data centers with standalone servers were defined by rack space, switch port capacity, IP addressing, and sub-netting requirements. Bandwidth and packet forwarding performance have always been key elements, but now there are fewer physical servers and more processor cores within a given data center footprint. See Figure 1 ("VM footprint").

With a virtualized environment, a complete enterprise information system infrastructure can be deployed in less space than in the past. With virtualization, the number of servers requiring bandwidth in a given rack can easily quadruple the bandwidth requirements over more traditional data center topologies. See Table 2 ("Virtualization drives bandwidth and utilization").

Now, more than ever, data center architecture designs require thorough investigation and planning focused on bandwidth requirements, Ethernet switch performance, shared data storage resources, broadcast traffic, and geographic redundancy.

It is also important to note that as data centers become more virtualized, and networks converge, the need for designs providing performance and capacity are becoming critical. But the reality is that data centers are not "Greenfield" implementations. Data centers were previously designed with traditional three-tier network architectures which operate on a foundation of IP peering and single purpose rack mount application servers. This status quo is being turned upside down with the rapid deployment of virtual servers. The fact is, data no longer primarily moves between servers inside and outside of the data center, but now moves horizontally within and across data center virtual server and network boundaries. Traditional three-tier network architectures focused around Layer 3 IP are not well suited to support these virtualized deployments.

HP believes that in order to support low-latency, high-performance applications on virtual servers with converged I/O technologies, the solution should operate on collapsed Layer 2 networks with Layer 3 routing in the core or aggregation layers. With 80% of traffic flowing from server-to-server within the data center, support of East-West traffic has different design requirements than North-South Host-Client type traffic. As such, flat layer 2 networks are the preferred method of accomplishing the change in traffic direction. In addition, maintaining a larger Layer 2 domain provides the maximum flexibility to allow VM mobility with technologies like vMotion or Live Migration. HP is also driving technologies and standards focused around Data Center Bridging (DCB), and TCP in support of Internet SCSI (iSCSI), and Fibre Channel over IP (FCIP). It is only logical that with rack and processing utilization gains of 300 to 1000 percent with blade servers and VMs in a single rack, an enterprise can house their entire IT demands within that single rack. Therefore, server-to-server and VM-to VM communications are going to be best suited with flatter Layer 2 networks. HP is continuing to meet these demands with its Layer 2 network and Intelligent Resilient Framework (IRF) switching foundation strategy.

Figure 1 VM footprint

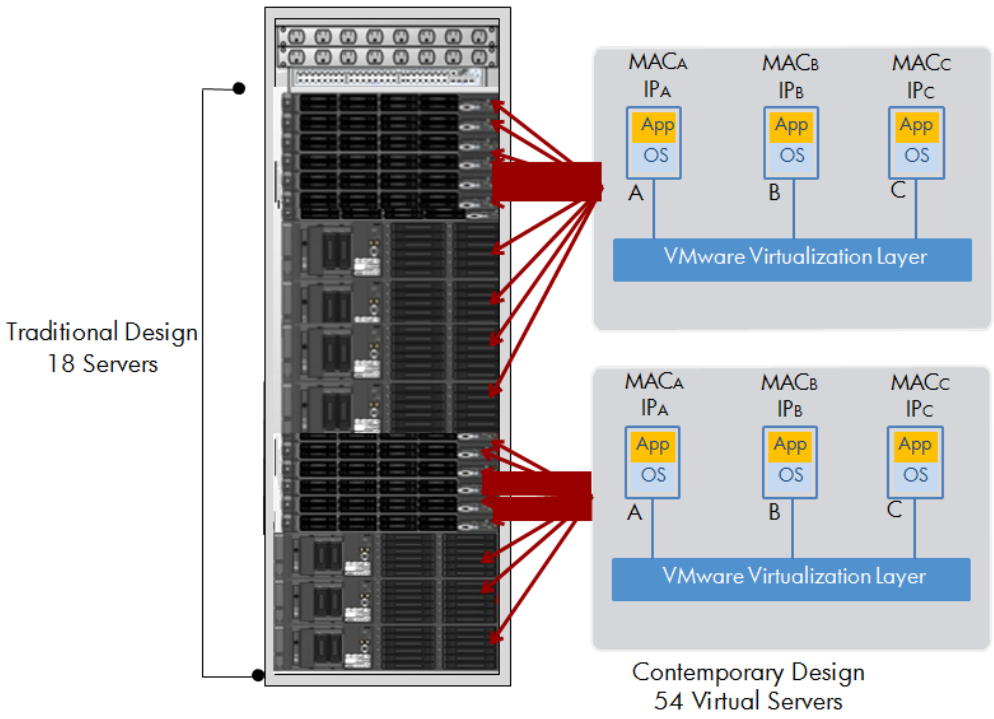

Table 2 Virtualization drives bandwidth and utilization

Server Type	# of Physical Servers	# of Virtual Servers	Server Bandwidth	Rack Bandwidth	Bandwidth & Footprint Delta
Rack Servers: • Single Quad Core Xeon • (30) VMs per Server					
No Virtualization	36	0	300Mbps	10.8Gbps	
Virtualized	36	1080	3Gbps	108Gbps	+900%
Blade Servers: • Single Quad Core Xeon • 16 Blades per Chassis • (30) VMs per Blade Server • 3 Chassis per Rack					
No Virtualization	48	0	300Mbps	14.4Gbps	
Virtualized	48	1440	3Gbps	144.0Gbps	+900%

Server deployment considerations

Virtualization of data, Web-based, and multimedia applications have driven the requirement for new servers, new platforms, and associated network storage. This requirement has resulted in unplanned growth in server complexity. IT departments are being charged with the improvement of computing

efficiencies in their data centers. The response has been server consolidation utilizing virtualization technologies implemented on blade servers.

Rack servers vs. blade enclosures

Rack servers

In a typical rack server environment, each rack is connected to two networks: The Ethernet general purpose local area network (LAN) and the storage area network (SAN). The Ethernet network has the following features:

- A single rack of servers may connect to top-of-rack (ToR) switches (typically deployed in pairs for redundancy) which uplink to a redundant aggregation network layer
- Several racks of servers may connect to end-of-row (EoR) switches which consolidate the network connections at the edge before going into the aggregation network layer

These types of networks have different interconnect implementations. Each server on the rack needs to have one set of interconnects to the LAN and one set to the SAN. In addition, each server has its own power, cooling, and management requirements. Each network has its own management process, as well as its own security and performance requirements. Some rack servers have now instituted splitting the physical frame into two separate server units. This type of design provides expanded capacity and processing power per rack unit, but lacks the network consolidation that blade servers possess.

Advantages

- **Simplicity:**

  One physical system sharing common CPU(s), memory, and storage controller
- **More expansion slots available for network and storage adapters:**

  Larger rack-mount servers (of more than 5U) have up to seven I/O expansion slots. Rack servers provide a large number of slots for network interface cards (NICs) or storage controllers in your virtual host for load balancing (LB) and fault tolerance
- **Traditional servers have a greater internal capacity for local disk storage:**

  If you're running several VMs on local disk storage, rack servers are a better choice, as they have more drive bays for internal disk storage
- **Processing power:**

  Rack servers can support eight or more CPU sockets. This works well when using a smaller number of powerful servers for virtual hosts
- **Serial, parallel, and USB I/O ports:**

  These ports allow you to connect external storage devices and optical drives. They can also accept hardware dongles used for licensing software

Disadvantages

- **Network consolidation:**

  Rack Servers require far more Ethernet ports for network connectivity which are significantly underutilized
- **Green data center:**

  These types of systems do not optimize the computing power per rack unit; they utilize more electrical power per rack unit and add to the overall requirement for cooling systems

Blade servers/enclosures

Blade servers can reduce the external switch count and significantly reduce the number of cables in a rack. The blade enclosure can include server-to-network interconnect modules, which can replace the access layer switch at either the ToR or EoR placements.

Advantages

- **Data center space:**

  Blade servers can increase rack density. Compared with traditional servers, this means up to 50 percent more servers in a standard 42U rack
- **Disk expansion blades:**

  Rack servers do not possess this feature, so they must connect to external small computer system interface (SCSI) array chassis or networked storage
- **Consume less power:**

  They are more energy efficient and require less cooling. The amount of power that blades consume is dependent on how full the blade chassis is. A fully-loaded chassis will consume far less power than an equivalent amount of traditional servers. If the chassis is not full, it will still consume less power than an equal amount of rack servers
- **Integrated Ethernet switches:**

  Ethernet cables plug into a chassis with a single connector, which makes cabling neater and eliminates the clutter of cables common with rack servers
- **Booting from SANs:**

  Blade servers require minimal resident disk storage. Virtual hosts can boot from the SAN, so no internal disk is required. The host performs a pre-boot execution environment (PXE) start up from the network and then connects to the SAN disk to access the files to continue the boot process

Disadvantages

- **Local expansion slots:**

  Relatively limited number of expansion slots for storage and/or internal blade NICs
- **Hard disk capacity:**

  Blade servers typically have a limited amount of localized hard disks (zero to four drives) and are dependent on SAN and/or network attached storage (NAS)
- **Expansion cost:**

  Blade server chassis have substantial capacity, but once full, the next incremental CPU blade requires another chassis and the additional integrated chassis components
- **With an Edge-Core SAN architecture, SAN designers selecting blade enclosures can use these approaches:**

  Replace an existing external Fibre Channel edge switch with an embedded blade network module. This reduces the number of physical connections at the server edge to both the network and storage media. Because the blades share the chassis power supply, power and cooling requirements are reduced

Server virtualization design considerations

Virtualization can enable an organization to create more agile IT services that can respond faster to business needs. Virtualization is not a single technology; rather, it is a way of designing technology

solutions. Many different aspects of IT infrastructure—such as servers, storage, networks, and clients—can be virtualized.

Server virtualization technologies—vSphere (VMware), Hyper-V (Microsoft®), XEN (Citrix)

Virtualization is the separation of a resource, application, or service from the underlying physical components of that service. For example, virtual memory allows applications to access more memory than is physically installed. Similarly, server virtualization can give applications the appearance that they have access to an entire server, including memory, CPU, and hard drives, while in reality, they may be sharing the physical hardware with other operating systems and applications. Each application and operating system combination is called a VM.

Data center administrators have the ability to utilize physical servers optimally by placing VMs across their physical system assets and infrastructure.

A key benefit of virtualization is that it provides the ability to run multiple operating system instances on the same physical hardware simultaneously and to share hardware resources such as disk, network, and memory. Running multiple instances that are independent of each other on a single physical hardware structure is known as partitioning.

VMware

This document is not intended to be a comprehensive discussion of VMware, please refer to the VMware website for specifics.

Because VMs are a true business driver for the data center architecture, we will provide a brief overview of VM's and concepts which need consideration when creating a network design, as VMware offers multiple components. The most important components are:

- **Hypervisor:**

  VMware ESX Server. Runs on physical servers and abstracts processor, memory, storage, and networking resources
- **VMware VM File System:**

  A file system for use by VMs
- **Virtual Center Management Server:**

  Allows for configuring, provisioning, and managing virtualized IT infrastructure
- **VMware vMotion:**

  Allows live migration of VMs running on one physical server to another with no downtime and complete transactional integrity
- **VMware High Availability:**

  In the event of server failure, automatically starts affected VMs on other servers with spare capacity
- **VMware Consolidated Backup:**

  Allows for a centralized, agent-free backup of VMs
- **VxLAN:**

  VXLAN is a method for "floating" virtual domains on top of a common networking and virtualization infrastructure. By leveraging industry-standard Ethernet technology, large numbers of virtual domains can be created above an existing network, with complete isolation

from each other and the underlying network

vSwitch

One of the most important components in the context of this document is called “virtual switch” or “vSwitch”. VMware provides a “vSwitch” as a software component running under the control of the hypervisor. It provides the necessary network connectivity between VMs running on the server, as well as between the VMs and the outside world. Each virtual switch provides modular support for functionalities such as Layer 2 forwarding, virtual local area network (VLAN) tagging and stripping, and Layer 2 security. The required modules for a switch are loaded at run time.

In many ways, a virtual switch behaves like a physical switch. For example, a virtual switch maintains a Layer 2 forwarding table, which contains associations between the port number and the media access control (MAC) address of any device connected to that port. As each Ethernet frame arrives at the virtual switch, it looks up the destination MAC and forwards the frame to the associated port. VLANs can also be set up on the virtual switch. Ports can be defined as either an access port (where the port is associated with a single VLAN) or as a trunk port, which allows multiple VLAN packets to traverse. A trunk port is typically used to connect two VLAN-aware devices, such as two switches, or a switch and a router. If the two switches were connected using access ports (same VLAN), then only traffic for that VLAN will pass between the two switches. If multiple VLANs are defined on both switches, a trunk port is needed to allow traffic for these VLANs to traverse both switches.

vSwitch ports also support promiscuous mode. Here, all traffic received on one port is copied onto another port regardless of VLAN membership or destination MAC. This is useful for network monitoring and sniffing applications, as well as intrusion detection systems (IDS) which require Layer 2 insertion. Unlike a physical switch, there is no need for frame flooding to map a MAC address with a switch port. Because each virtual switch port is connected to a virtual NIC in the VM, this association begins when the VM is started.

Unlike physical switches, VMware vSwitches cannot be interconnected on the same server. As a benefit, Ethernet loops cannot exist in the virtual topology. Therefore, the spanning tree protocol (STP) is not required and is not present in a vSwitch. Additionally, virtual switches cannot share the same physical Ethernet NIC and cannot have an entry in one table of a virtual switch port on another virtual switch. The Layer 2 hash tables are associated with each NIC and cannot be duplicated on other NICs.

An important concept in virtual switches is port groups, which are templates that hold sets of specifications for a port. A port group ensures that a particular VM has the same type of connectivity on any machine on which it is run. This is particularly important for vMotion, where a live migration of a VM from one physical server to another can take place without affecting application services.

The typical specifications carried by a port group include virtual switch name, VLAN IDs, Layer 2 security policies, and traffic shaping parameters. A VM can be associated with a particular port group, which then connects the virtual NIC to an appropriate port on the switch.

The connection between a virtual network and a physical network is via the physical Ethernet adapters on the host. These are called uplinks in VMware terminology and the virtual ports connected to them are called uplink ports. If a virtual switch has multiple VLANs configured, the link between the uplink port and the physical switch port must be in trunk mode.

In order to support VLANs, an element in either the virtual or the physical network has to be able to tag the Ethernet frames with IEEE 802.1Q tags.

There are three ways that this can be done:

- **Virtual switch tagging:**

  A port group is defined for each VLAN and the VM is associated with the relevant port group.

The port group tags all outbound frames and strips tags off all inbound frames

- **VM guest tagging:**

  An IEEE 802.1Q VLAN driver is installed in the VM, which tags all outgoing frames. The tags pass unchanged through the vSwitch

- **External switch tagging:**

  An external physical switch is used for tagging

There are some security policies which can be applied to the vSwitch to prevent VMs from seeing traffic destined for other nodes in the virtual network. By default, promiscuous mode is turned off. The administrator can also specify a MAC address lockdown, which prevents individual VMs from changing their MACs, thus preventing a node from seeing unicast traffic to other nodes. The administrator can also block forged transmissions to prevent VMs from spoofing other nodes.

Server-to-Network edge using VMs

VMs and virtual switches running on a physical server introduce new a complexity at the server edge and dramatically impact associated networks.

Challenges with VMs

- Managing VM sprawl and associated virtual networking
- Performance loss and management complexity of integrating software-based virtual switches with existing network management

These are significant challenges not fully addressed by any vendor today. HP is working with other industry leaders to develop standards to simplify and solve these challenges.

Virtual Ethernet Bridges

With the growing use of VMs and their associated virtual switches, a new level of complexity has been introduced at the server edge. Network management has to consider a new virtualized network with virtual switches. Typically, two different groups are now involved at the edge—the traditional network administrators, who are responsible for the physical network switches, and the server administrators, who are responsible for managing and configuring the virtual switches, which must be configured and managed manually.

Virtual switches do not have the traffic monitoring capabilities of physical access switches, so troubleshooting VM-to-VM traffic may be an issue. Additionally, they may lack some of the advanced security features provided by physical switches, and even when these are provided, they may not be fully compatible with their physical counterparts. This makes it difficult to provide consistent end-to-end security policies.

As the number of VMs in a server increases, the traffic through the virtual switches also increases. This growth drives greater CPU resource utilization on the physical servers used to handle the traffic.

The vNetwork distributed switch (VDS) was introduced with VMware vSphere 4.0 to respond to these complexities. A VDS treats the network as an aggregated resource, and provides a single abstracted view of all the individual vSwitches present in the virtual network. This abstracted view presents a single, large switch that spans multiple servers across the data center. Port groups that were earlier associated with a single vSwitch now become distributed virtual port groups that span multiple hosts and provide support for features such as vMotion.

The virtual switches in earlier versions of VMware had to be managed and configured separately.

With a VDS, only a single switch needs to be configured and managed.

Distributed virtual uplinks (dvUplinks) are a new concept with the VDS that provides an abstraction for the virtual NICs on each host. The NIC teaming, load balancing, and failover policies on the VDS and the distributed virtual port groups are applied to the dvUplinks rather than to individual virtual NICs. Each virtual NIC on a VM is then associated with a dvUplink, providing consistency in the way teaming and failover occur.

VM-to-VM Switching

With multiple Virtual Machines residing on a single physical server, there are requirements for efficient VM-to-VM traffic switching on the same server, and the ability to enforce physical switch policies on the virtual machines.

There are two models for enabling VM-to-VM switching:

- **Virtual Ethernet Bridge (VEB**) – A Virtual Ethernet Bridge is a virtual Ethernet switch. It can reside on a single physical server and provide a virtual network between virtual machines on that server, or it can be used to connect the virtual machines to an external network. Most VEBs are software-based. However, with the adoption of the PCI Single Root IO Virtualization (SR-IOV) standard, hardware- based virtual switches are built into NICs. Hardware-based VEBs provide better performance than software-based VEBs
- **Software-based VEBs** – The hypervisor creates Virtual NICs for each virtual machine, then creates one or more vSwitches that connect the virtual NICs to the physical NICs. Traffic received by a physical NIC is passed to the correct virtual NIC based on the configuration information held by the hypervisor. Traffic from a virtual NIC is treated in one of two ways, as shown below

  If the destination is external to the physical server or to a different vSwitch, the vSwitch forwards the traffic to the physical NIC

  If the destination is on the same vSwitch on the same physical server, the vSwitch forwards the traffic back to the target virtual NIC

Figure 2 vSwitch forwarding

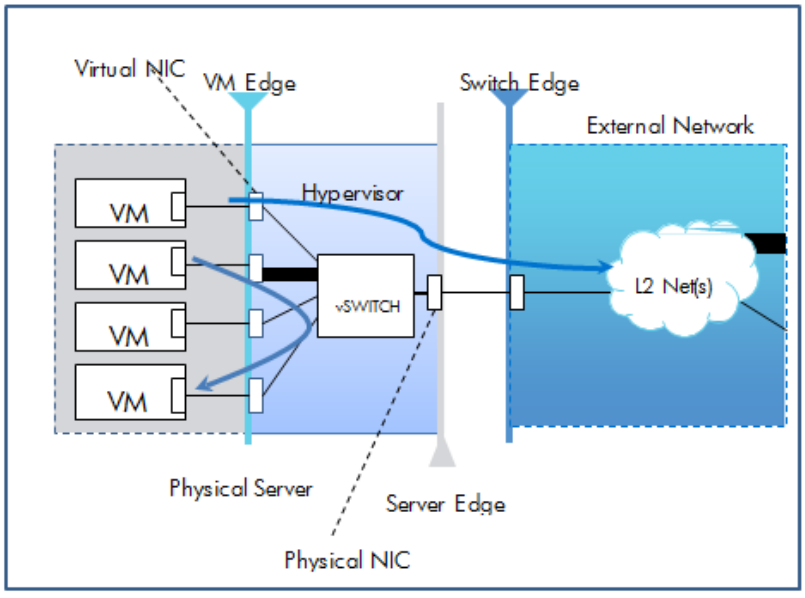

Advantages of Software-based vSwitches

- Good performance for horizontal, inter-VM traffic: Internal VM-to-VM traffic involves only memory copy operations, and the performance is limited by available CPU cycles and memory bandwidth
- Software-based vSwitches are standards-based, and can interoperate with a wide variety of external network infrastructure

Disadvantages of Software-based vSwitches

- Because the host server's CPU is used by the virtual switch, more CPU cycles are used for higher traffic. This limits the number of virtual machines that can be supported on a single server
- Lack of standard network monitoring capabilities such as flow analysis, statistics, and remote diagnostics: VM-to-VM traffic on the same vSwitch is not visible outside the physical server
- No way to enforce policies: External switches support port security, QoS, and access control lists. Even if vSwitches support some of these features, their management and configuration is not consistent with the way external switches are managed, thus making it difficult to create end-to-end policies in the data center

Issues with management scalability

As the number of virtual machines increases, the number of vSwitches also increases. Because each vSwitch must be managed separately, this adds to management overhead. Distributed vSwitches somewhat alleviate this problem, but do nothing for management visibility beyond the virtualized server. Single pane of glass management systems will have to address this issue with added functionality in their software.

Hardware-Based VEBs

The performance issues with software-based VEBs can be alleviated by moving the VEB functionality to the NIC. Single Root I/O Virtualization (SR-IOV) is the technology that makes this possible. SR-IOV is a standard developed by PCI-SIG that allows multiple PCI devices to be shared among virtual machines, and is implemented in hardware. SR-IOV creates many virtualized instances of a physical PCI device, and each virtualized instance is assigned to a VM. This allows the VM to access the PCI device directly for IO operations, without hypervisor mediation. This reduces the number of CPU cycles required to move traffic between VMs and from VMs to the external world.

In order for SR-IOV to work on a server, an SR-IOV NIC is required. This would have the virtual functions built into the hardware, making the virtualized instances possible. Additionally, the physical server's BIOS must be capable of recognizing the SR-IOV NIC. The hypervisor must be capable of loading the drivers to work with the SR-IOV NIC, and the guest operating systems must be able to perform IO directly with the SR-IOV NIC.

Figure 3 The SR-IOV process

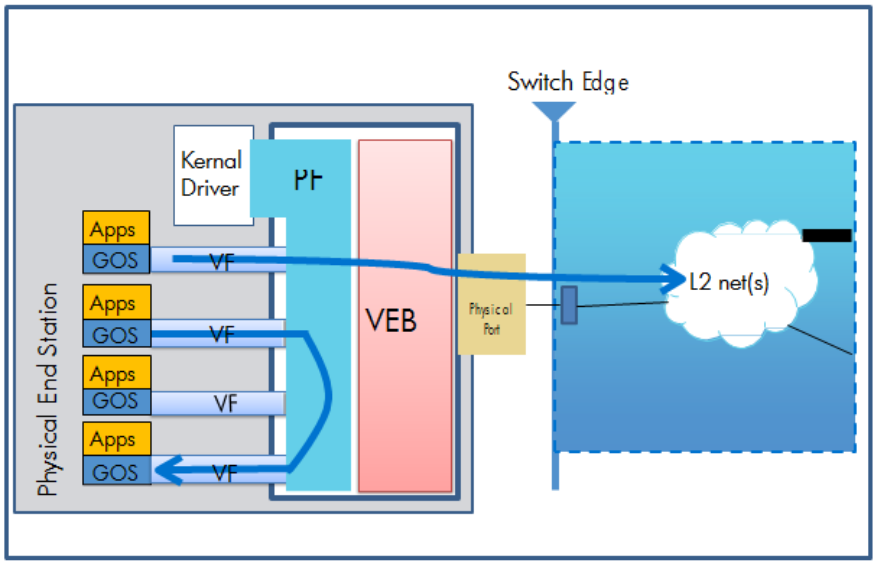

Advantages of using SR-IOV NICs

- Reduced CPU utilization compared with software based vSwitches
- Increased network performance due to direct I/O between a guest OS and the NIC

Disadvantages of using SR-IOV NICs

- As with software-based VEBs, SR-IOV NICs are not visible from the network. Because of limits on cost-effective NIC silicon, they often have fewer capabilities than software-based VEBs
- Similar to software VEBs, they do not offer advanced policy enforcement
- SR-IOV NICs usually have very small address tables, and they do not learn. This may increase the amount of flooding, especially if there is a large amount of VM-to-VM traffic
- Guest operating systems need to be able to support SR-IOV. This is currently available only in open source operating systems

Virtual Ethernet Port Aggregator (VEPA)

The IEEE 802.1 Working Group has agreed to base the IEEE 802.1Qbg EVB standard on VEPA technology because of its minimal impact and minimal changes to NICs, bridges, existing standards, and frame formats (which require no changes).

VEPA is designed to incorporate and modify existing IEEE standards so that most existing NIC and switch products could implement VEPA with only a software upgrade. VEPA does not require new tags and involves only slight modifications to VEB operation, primarily in frame relay support. VEPA continues to use MAC addresses and standard IEEE 802.1q VLAN tags as the basis for frame forwarding, but changes the forwarding rules slightly according to the base EVB requirements. In doing so, VEPA is able to achieve most of the goals envisioned for EVB without the excessive burden of a disruptive new architecture such as VN-tags.

Software-based VEPA solutions can be implemented as simple upgrades to existing software virtual switches in hypervisors. In addition to software-based VEPA solutions, SR-IOV NICs can easily be updated to support the VEPA mode of operation. Wherever VEBs can be implemented, VEPAs can be implemented as well.

VEPA enables a discovery protocol, allowing external switches to discover ports that are operating in VEPA mode and exchange information related to VEPA operation. This allows the full benefits of network visibility and management of the virtualized server environment.

There are many benefits to using VEPA:

- A completely open (industry-standard) architecture without proprietary attributes or formats
- Tag-less architecture which achieves better bandwidth than software-based virtual switches, with less overhead and lower latency (especially for small packet sizes)
- Easy to implement, often as a software upgrade
- Minimizes changes to NICs, software switches, and external switches, thereby promoting low cost solutions

MultiChannel Technology

During the EVB standards development process, scenarios were identified in which VEPA could be enhanced with some form of standard tagging mechanism. To address these scenarios, an optional "multichannel" technology, complementary to VEPA, was proposed by HP and accepted by the IEEE 802.1 Working Group for inclusion into the IEEE 802.1Qbg EVB standard. MultiChannel allows the traffic on a physical network connection or port (like an NIC device) to be logically separated into multiple channels as if they are independent, parallel connections to the external network. Each of the logical channels can be assigned to any type of virtual switch (VEB, VEPA, and so on) or directly mapped to any virtual machine within the server. Each logical channel operates as an independent connection to the external network.

MultiChannel uses existing Service VLAN tags ("S-Tags") that were standardized in IEEE 802.1ad, commonly referred to as the "Provider Bridge" or "Q-in-Q" standard. MultiChannel technology uses the extra S-Tag and incorporates VLAN IDs in these tags to represent the logical channels of the physical network connection.

MultiChannel supports:

- Multiple VEB and/or EVB (VEPA) virtual switches share the same physical network connection to external networks: Data center architects may need certain virtualized applications to use VEB switches for their performance and may need other virtualized applications to use EVB (VEPA) switches for their network manageability, all in the same physical server
- Direct mapping of a virtual machine to a physical network connection or port while allowing that connection to be shared by different types of virtual switches: MultiChannel technology allows external physical switches to identify which virtual switch, or direct mapped virtual machine, traffic is coming from.
- MultiChannel also allows direct mapping of a virtual machine that requires promiscuous mode operation (such as traffic monitors, firewalls, and virus detection software) to a logical channel on a network connection/port. Promiscuous mode lets an NIC forward all packets to the application, regardless of destination MAC addresses or tags, thus placing the port mirroring process burden on physical Ethernet switches that are better suited for this role.

Optional MultiChannel capability requires S-Tags and "Q-in-Q" operation to be supported in the NICs and external switches, and, in some cases, it may require hardware upgrades, unlike the basic VEPA technology, which can be implemented in almost all current virtual and external physical switches. MultiChannel does not have to be enabled to take advantage of simple VEPA operation.

MultiChannel merely enables more complex virtual network configurations in servers using virtual machines.

Figure 4 MultiChannel connectivity with virtual machines

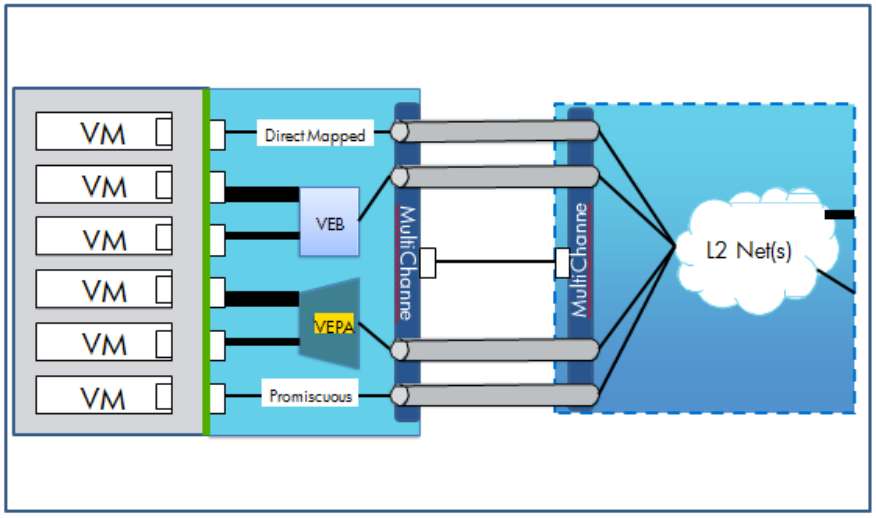

MultiChannel technology gives data center architects the ability to match the needs of their application requirements with the design of their specific network infrastructure:

- VEB for performance of VM-to-VM traffic
- VEPA/EVB for management visibility of the VM-to-VM traffic
- Sharing physical NICs with direct mapped virtual machines
- Optimized support for promiscuous mode applications

Layer 2 and VM mobility

Virtualization enables the ability to move VMs on the fly from one physical server to another, whether in the same rack or across data centers. Mainstream products which enable VM movement include vMotion from VMware or Live Migration from Microsoft. These products streamline application deployment, management of VMs, and physical servers. However, because VMs cannot be moved outside their Layer 2 network, there are implications that impact the choice of Layer 2 or Layer 3 solutions for virtualized server deployments.

Figure 5 Advantage of using vMotion on a Layer 2 network

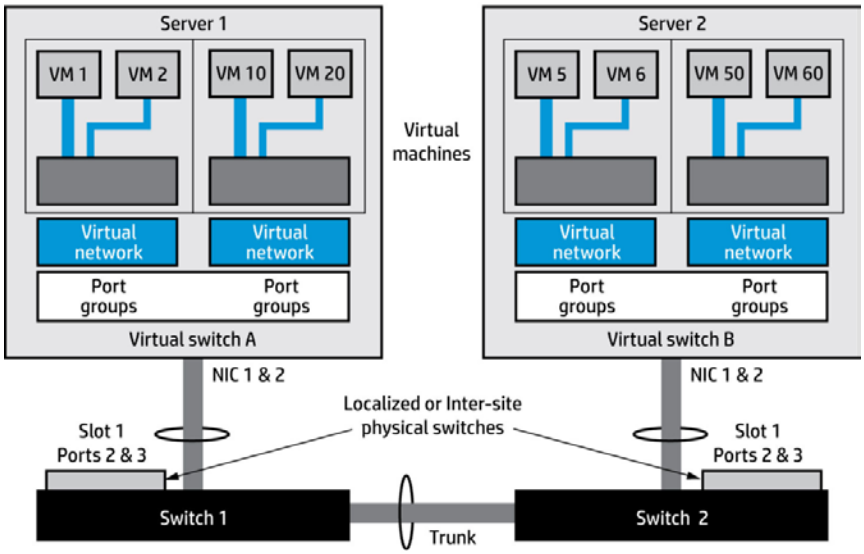

vMotion support—Moving virtual server from one physical server to another

Advantages

- VMs or servers can be moved without requiring the IP addresses to be changed
- Easily accommodate legacy applications/servers that require the IP address/default gateways to be embedded
- Eases the management of VMs/servers across a virtual data center or hosting during data center consolidation

Risks

- Topology changes can affect throughput
- STP by nature is limited and can operate at losses of more than 50 percent. This reduces the bandwidth and, consequently, increases completion times when moving VMs

Key Drivers for Layer 2 Networks

Application and server virtualization

The operational role of virtualized servers is ever-increasing in the data center. Virtualization of servers and applications is highly dependent on Layer 2 technology. The fact that it is not possible to move VMs outside of the Layer 2 network that contains it, is a key factor in this dependency. For this reason, the development of flexible Layer 2 architectures in the data center is part of the broad advantage of virtualization.

Network virtualization

Very much like the concepts behind application and server virtualization, the goal of network virtualization is to add flexibility and greater utilization of network resources. Today's data center networks are in need of virtual Layer 2 backbones that not only span one data center location, but can also operate across multiple geographically separated data centers. This need demands increased control of traffic flows, efficient use of bandwidth, and reduced network elements for virtual backbones. A major goal of network virtualization is to further enhance the gains made with

application and server virtualization. Providing a robust Layer 2 topology with technologies like IRF can further eliminate STP, while at the same time support both multipath forwarding and localized failure resolution.

Geographic redundancy

There are a substantial number of server and communication applications that require Layer 2 protocols to identify operational state through heartbeats. Layer 3 wide area networks limit the ability of servers to communicate in the same Layer 2 broadcast domains. Extending Layer 2 VLANs across physical network boundaries drives greater flexibility in data center and application deployment designs. Virtual Layer 2 backbones eliminate the need for Layer 3 tunneling to extend both Layer 2 heartbeat traffic and private IP address space between separate data center locations. These are key factors in designing for geographic redundancy.

Performance

Data centers have an ever-increasing need for more bandwidth. The deployment of Layer 2 virtualized backbones allows Ethernet switch clustering. These can reach far higher theoretical throughputs than their predecessors, which relied upon either STP or Layer 3 to manage traffic flow. In this new infrastructure, switches can be clustered in active/active configurations–virtually eliminating network links to servers that lie dormant in traditional designs. The introduction of technologies like shortest path bridging (SPB) and Transparent Interconnection of Lots of Links (TRILL) will further enhance the virtualized network backbone paradigm.

3-2-1-tier and flat Layer 2 network designs

In the past, traditional client/server-based data centers consisted of standalone servers supporting static applications. In this type of network structure, the traffic flowed upward and out of the network to communicate with computers and endpoints outside the data center. In comparison, the majority of network traffic in contemporary data centers is horizontal from server to server. This is due to inter-host, intra-VM, and synchronized inter-VM servers communicating with each other within the same network and also across geographically distributed networks.

Previous network design methodology supported Layer 3 solutions over Layer 2 solutions for the following reasons:

- Broadcast traffic could be isolated or limited so that the network traffic is reduced
- It is easier to scale the network with Layer 3 solutions
- The network is much more resilient with a Layer 3 solution compared to a Layer 2 solution
- The ability to address disparate needs in various customer environments

In the traditional three-tier data center network architecture, networks consisted of the access layer, aggregation layer, and core switches. In this legacy design, horizontal data flow between servers was restricted to travel at least two hops in the network.

With cloud computing and virtualization in the data center, it is estimated that by 2014 more than 80% of traffic in the data center will be server to server[1]. With the introduction of virtualization and the ability to move VMs from one physical server to another, using vMotion for example, retaining the original IP addresses and default gateway has become an important requirement. If a VM is moved from one data center to another, the IP address and default gateway has to be changed manually. In

[1] http://www.gartner.com/technology/reprints.do?id=1-17OKEMN&ct=111014&st=sb

a Layer 2 solution, this can be achieved without having to change any IP addresses, resulting in a simpler VM management environment. For a visual description, see the figure below ("Traditional three-tier vs. collapsed virtualized design").

There are challenges with Layer 2 based solutions that directly affect the move to a converged network infrastructure:

- A significant increase in broadcast traffic, larger broadcast domains, and loops
- Changes in the topology (STP/MSTP/RSTP). STP was designed when hubs were used to interconnect networks, and was intended to remove bridge loops in the network. STP helps ensure that there is a single active path to each network device. The STP process shuts down all other alternative paths. The consequence is that there is a loss of network bandwidth due to single path links
- Though multiple spanning tree protocol (MSTP) addresses some of the limitations, it needs careful design and implementation.
- Careful consideration needs to be taken with regards to the supported size of MAC tables on all switches in the Layer 2 domain, as well as the supported ARP table sizes on the switches where the routing will be performed (core)

Figure 6 Traditional three-tier vs. collapsed virtualized design

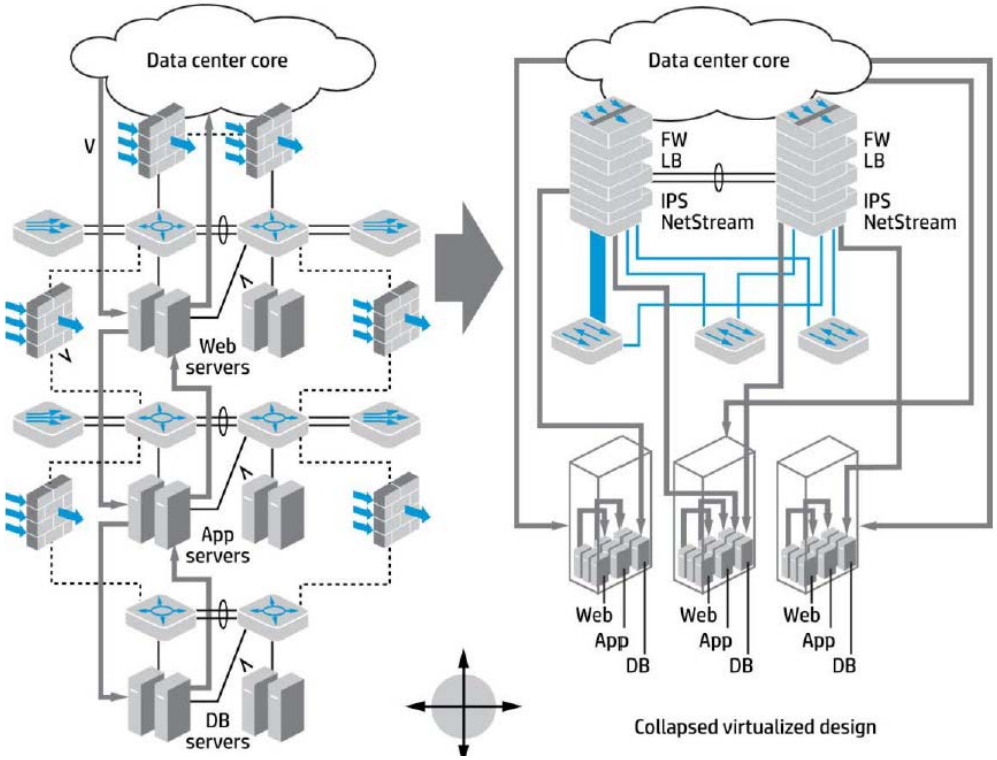

Layer 2 between Data Center Sites

Until recently, inter-site Layer 3-based solutions were preferred to Layer 2-based solutions for the following reasons:

- Easier to scale networks
- Isolation and limitation of broadcasts

- Layer 3 solutions are more resilient to failures compared to Layer 2 solutions
- Cost of Wide Area Network (WAN) connections

However, current data center deployments often require geographic redundancy. This means moving resources and services, such as VMs, from one data center to another quickly and seamlessly. To achieve this, their networks must reside in both data centers at Layer 2.

Key factors influencing a Layer 2-based solution are:

- The ability to move VMs from one data center to another (Long Range vMotion)
- The ability to move servers/machines without requiring IP addresses to change virtual data center hosting
- Data center consolidation often eliminates existing separate data centers
- Layer 2 application state through system heartbeats

Linking data centers to support a large Layer 2 network can be seen in the figure below ("How data centers can be linked to support a large Layer 2 network").

Figure 7 How data centers can be linked to support a large Layer 2 network

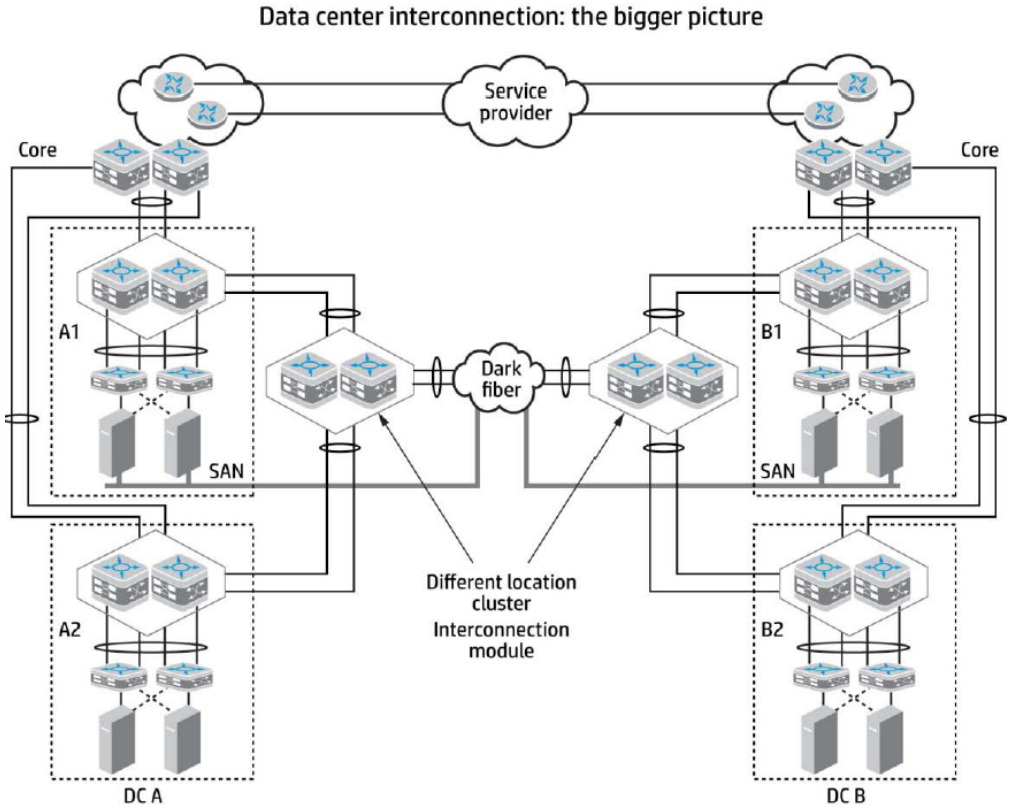

Layer 2 networking protocols

With Layer 2 networking, there are many important protocols existing, and ones being created, to accomplish a variety of evolving requirements within and across data centers.

- Spanning tree protocol (STP)
- Intelligent Resilient Framework (IRF)
- Transparent Interconnect of Lots of Links (Trill)
- Ethernet Virtual Interconnect (EVI)

- Virtual private LAN services (VPLS)
- Shortest Path Bridging (SPB)
- Pseudo Wire (PWE3)
- Virtual Lease Lines (VLL) - Multiprotocol Label Switching (MPLS)
- VPN Instance- Kompella & Martini
- VxLAN
- NVGRE
- VEPA
- And others under development…

Spanning Tree Protocol

The STP protocol is not a new networking protocol, STP is a true legacy protocol that has provided loop prevention within the network, as well as its share of broadcast storms and network crashes.

STP is generally the core of the three-tier network architecture. Many of the enhancements listed in this section were developed to move away from the limitations of STP as it tends to be inefficient and has a slow convergence period.

STP is a network protocol designed to provide a loop-free bridged Ethernet. It is a Layer 2 protocol, so it is not routable, and receives updates from other switches on their link conditions and path costs based on bridge protocol data units (BPDU).

The primary function of STP is to prevent bridge loops and broadcast proliferation. STP also allows networks to be designed with redundant links, providing automatic backup paths if an active link fails.

STP is standardized as IEEE 802.1d and creates a spanning tree within a mesh network of connected Layer 2 bridges and disables those links that are not part of the spanning tree, leaving a single active path between any two network nodes.

STP can take up to 50 seconds for switches to be notified of a topology change. This is unsatisfactory for use in a converged network infrastructure. An enhancement was made to this protocol which makes the old STP standard obsolete, called rapid spanning tree protocol (RSTP). Although definitely an enhancement, it can take up to two seconds to respond to topology changes in an Ethernet network. This is too long for latency-sensitive network data storage and VM applications.

Figure 8 STP

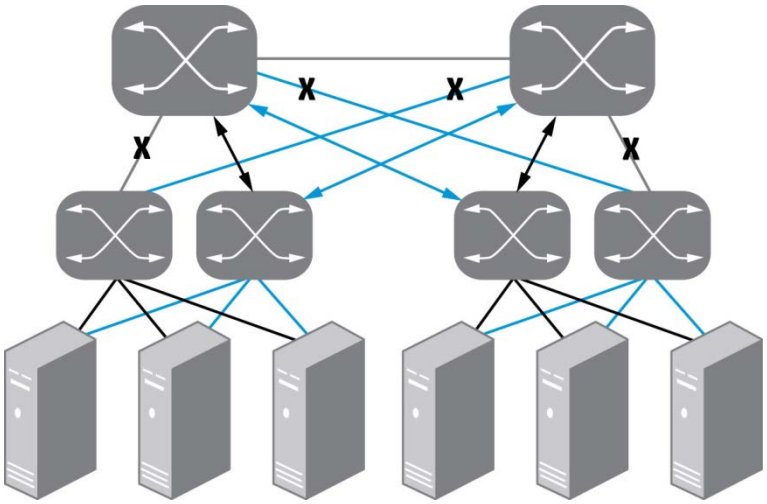

Intelligent Resilient Framework (IRF)

Traditional three-tier Ethernet networks were designed with hierarchical, modular Layer 2/3 Ethernet switches with STP and IP to accommodate for link path management, flow control, and destination routing. These types of designs are functional, but cannot meet the demands of today's converged network infrastructure.

Purpose-built HP networking solutions and technology are streamlining the design of next-generation data centers to help ensure superior resiliency, performance, and agility. One HP innovation is IRF, a technology that enhances ordinary Ethernet switching designs, allowing substantial improvements in the way Ethernet switches communicate. HP IRF provides the ability to flatten the data center and campus networks, eliminating the need for multiple tiers of aggregation switches and unutilized data paths. IRF provides a framework that enhances Ethernet and provides better link management, utilization, and redundancy.

Why is there a need for IRF?

In a network of interlinked switches, STP is used to detect and prevent loops—a highly undesirable, sometimes disastrous, situation that can occur when there are multiple active paths to the same switch. To eliminate loops, STP and its more modern variants, the RSTP and the MSTP, are designed to allow only one active path from one switch to another, regardless of how many actual connections might exist in the network. If the active path fails, the IRF protocol along with Multi Access Detection (MAD) automatically selects a backup connection and makes that the active path.

While STP is fairly effective in making a network resilient and simpler to configure, network convergence can still take several seconds, affecting applications that cannot handle that length of delay. In addition, the performance of STP is poor because it blocks all parallel paths except the one it has selected as active. Even when the network is operating normally, STP can reduce the effective bandwidth (possibly to a degree greater than 50 percent).

Figure 9 An example of the STP

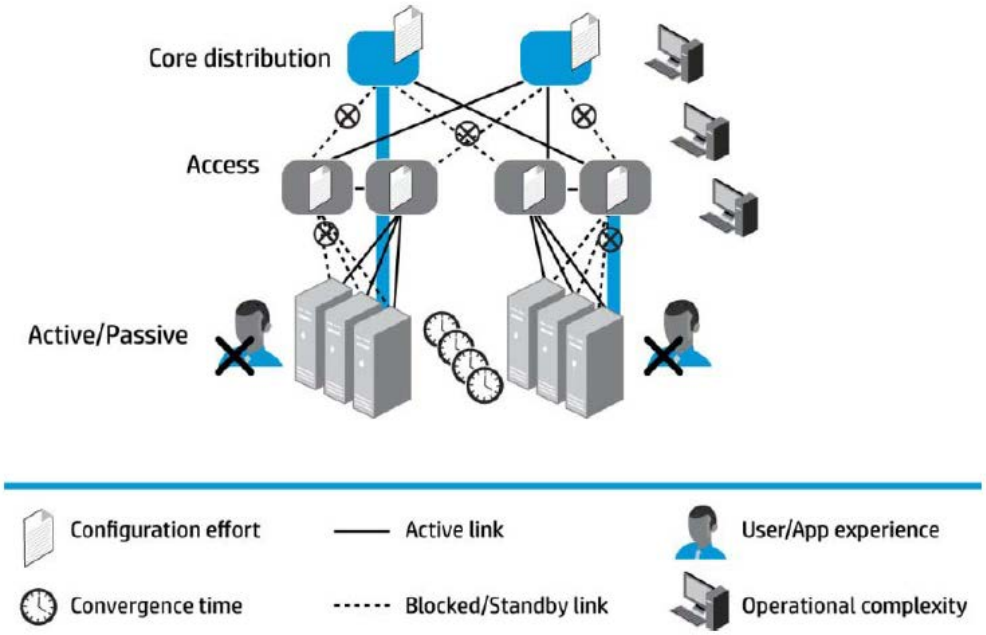

IRF can provide a network that is fully resilient, yet also simpler to setup and manage, faster to converge, and easier to scale. IRF simplifies network operations by consolidating management of multiple discrete devices into a single, easy-to- manage virtual switch, in every layer of the network.

IRF operational fundamentals

Think of IRF as a framework on which Ethernet operates, rather than a protocol that operates on top of Ethernet. This is not to say that you cannot still make use of the benefits of SPB or TRILL in a data center bridged Ethernet network, but IRF does so much more than Layer 2 Ethernet protocols. It was designed with much more in mind than just the replacement of STP.

IRF technology extends network control over multiple active switches. Management of a group of IRF-enabled switches is consolidated around a single management IP address, which vastly simplifies network configuration and operations. You can combine as many as nine HP Comware OS based switches to create an ultra-resilient virtual switching fabric comprising hundreds or even thousands of 1GbE or 10GbE switch ports.

This framework was designed to support the growing need for converged network infrastructures that demand:

- High capacity links like 10GbE, supporting Fibre Channel over Ethernet (FCoE)
- Rapid switch-to-switch topology updates and failovers (requires MAD)
- High performance frame forwarding fabrics
- Low latency communications
- Ethernet trunk aggregation with flow control for mission critical data types
- Ease of provisioning
- Reduced management complexity

One IRF member operates as the primary system switch, maintaining the control plane and updating forwarding and routing tables for the other devices. As mentioned before, IRF is a framework on

which Ethernet operates. It not only enhances the Layer 2 functionality, but also improves Layer 3 performance. If the primary switch fails, IRF instantly selects a new primary switch, preventing service interruption and helping to deliver network, application, and business continuity for critical applications. Within the IRF domain, network control protocols operate as a cohesive whole to streamline processing, improve performance, and simplify network operation.

Routing protocols calculate routes based on the single logical domain rather than the multiple switches it represents. Edge or aggregation switches that are dual-homed to IRF-enabled core or data center switches "see" associated switches as a single entity, eliminating the need for slow convergence technologies such as STP. And operators have fewer layers to worry about, as well as fewer devices to configure and manage.

Advantages of IRF

Design and operational simplification

With IRF, no longer do you connect to, configure, and manage switches individually. Configurations are performed on the primary switch and that configuration is distributed to all associated switches automatically, considerably simplifying network setup, operation, and maintenance.

Flexible topology

IRF provides a simplified, higher performing, more resilient, and flatter network design. IRF and HP Comware OS based switches allow enterprise networks to be designed with fewer devices and fewer networking layers—a big improvement over the low performance, high cost, and crippling latency of conventional multi-tier legacy solutions, which often rely on a variety of different operating systems and complex resiliency protocols. IRF provides a common resiliency mechanism in both Layers 2 and 3, eliminating the need for slower converging protocols like STP and VRRP. Also, unlike other Layer 2 Ethernet network frameworks that require the use of tunneling to extend VLANs, IRF allows network operations and design to get back to basics by providing a simple, yet flexible architecture based on standards-based IEEE 802.1Q tagging, Q-in-Q tagging, and basic but effective VLAN mapping. In addition, IRF supports advanced higher layer protocols like SPB, TRILL, and MPLS. This provides the most flexibility while maintaining industry-leading resiliency and performance. To flatten the network further, HP BladeSystem using HP Virtual Connect (VC) or the new HP 6125 Blade Switch can work with IRF to provide simplified server to server communications.

Higher efficiency

IRF's loop-free, non-blocking architecture keeps all links active, enabling highly efficient, high-bandwidth connectivity throughout the switching plane. No longer is it necessary to have switch–to-switch Ethernet trunks go unused in a standby mode in STP, or a partially utilized path with SPB or TRILL protocols. Existing networks that have legacy three-tier Ethernet architectures can be rapidly integrated into the IRF framework, immediately enhancing the performance and resiliency of the network.

Scalable performance

IRF and link aggregation control protocol (LACP) work together to boost performance and increase resiliency by bundling several parallel links between switches and/or servers, allowing scalable on-demand performance and capacity to support the demands of converged network infrastructures. In addition, IRF is fully supported with third-party devices, as it leverages standards-based link

aggregation for connectivity.

Faster failover

Should a network failure occur, IRF can deliver rapid recovery and network re-convergence in under 50 milliseconds (requires MAD) —much faster than the several seconds when STP or even RSTP are utilized.

Geographic resiliency

Geographic separation is no longer just the domain of the WAN. Data centers are ever-increasing in size with extremely large campuses. With the evolution toward converged network infrastructures and large, flattened L2 networks, the need to access server and I/O resources across a local campus involving substantial distances is becoming more common. IRF meets these demands, as well as provides connectivity into MANs and WANs. Within an IRF domain, the geographic location of switches does not matter. Switches can be extended horizontally and they continue to function as a single logical unit whether they are installed locally, distributed regionally, or even situated at distant sites. Moreover, employing IRF can enhance disaster recovery by linking installations up to 70 kilometers apart, providing the same fast failover as if they were sitting side by side within the data center.

In-service-software-upgrade

IRF delivers a network-based in-service-software-upgrade (ISSU) capability that allows an individual IRF-enabled switch to be taken offline for servicing or software upgrades without affecting traffic going to other switches in the IRF domain.

Transparent Interconnect of Lots of Links (Trill)

TRILL combines the simplicity and flexibility of Layer 2 switching with the stability, scalability, and rapid convergence capability of Layer 3 routing. All these advantages make TRILL very suitable for large flat Layer 2 networks in data centers.

The following are explanations of the basic TRILL concepts:

- **RBridge:**

  Routing bridge (RB for short) that runs TRILL. RBs are classified into ingress RBs, transit RBs, and egress RBs, depending on their positions in the TRILL network. A frame enters the TRILL network through an ingress RB, travels along transit RBs, and leaves the TRILL network through an egress RB
- **TRILL network:**

  A Layer 2 network comprised of RBs
- **Nickname:**

  Unique identifier of an RB in the TRILL network. TRILL automatically assigns nicknames to RBs
- **Link State Database:**

  The LSDB contains all link state information in the TRILL network
- **Link State Protocol Data Unit:**

  An LSP describes local link state information and is advertised between neighbor devices
- **Appointed VLAN-x Forwarder (AVF) and appointed port:**

TRILL supports VLANs. To avoid loops, TRILL requires all the traffic of a VLAN on a network segment to enter and leave the TRILL network through the same port of an RB. The RB is the AVF of the VLAN, and the port is the appointed port

- **Designated Routing Bridge:**

  The DRB corresponds to the DIS in IS-IS. It helps simplify network topology and appoints AVFs for VLANs on each RB

TRILL establishes and maintains adjacencies between RBs by periodically advertising Hello frames, distributes LSPs among RB neighbors, and generates an LSDB for all RBs in the network. Based on the LSDB, each RB uses the SPF algorithm to calculate forwarding entries destined to other RBs.

In a TRILL network, RBs compute a distribution tree for each VLAN according to the LSDB and use the distribution tree to guide the forwarding of multi-destination frames, which include multicast, broadcast, and unknown unicast frames in the VLAN.

Note that IRF and TRILL don't play in the same dimension. IRF must be seen as a "clustering" technology allowing multiple devices to be seen as one logical device, removing STP, VRRP from the network, with a single IP for the management. TRILL, on the other hand, provides a mechanism that allows every single node to have a tree rooted at itself, allowing the optimal (shortest path) distribution of traffic as well as multi-pathing for failure recovery.

IRF and TRILL are in fact not mutually exclusive. When used together TRILL and IRF can combine the best of both worlds, allowing for reduced routing table sizes and larger domains with faster recovery times.

Figure 10 Complimenting TRILL with HP IRF

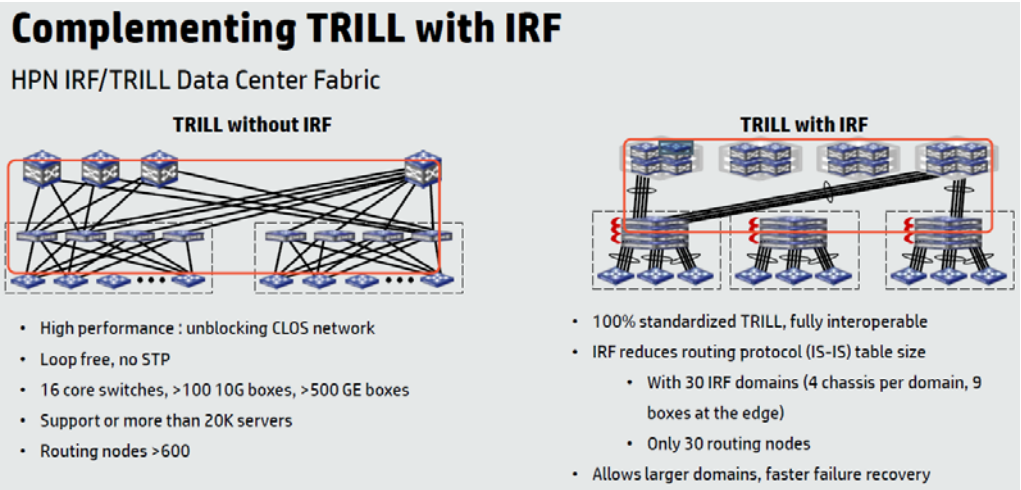

Ethernet Virtual Interconnect (EVI)

HP EVI, a new HP Virtual Application Network innovation and component of the HP Data Center Interconnect (DCI) solution, enables IT staff to simplify the interconnectivity of up to eight geographically disperse data centers. As a result, clients improve disaster recovery capabilities and can quickly respond to changing market conditions by easily moving virtual machines in minutes without network delay to any interconnected data center.

HP EVI runs over Internet Protocol (IP) transport and extends layer 2 domains across the networks of the connected data centers. By virtualizing and automating a layer 2 domain across data centers, HP EVI delivers the elements necessary to enable a software-defined networking (SDN) data center

infrastructure. It enables several data centers to work as one that is more responsive, with higher efficiency and solid high availability for business resiliency. With EVI, enterprises are able to accelerate the delivery workload mobility with remote vMotion, increase applications performance with load balancing, and achieve optimum degrees of high availability and disaster recovery for valuable data. When used along with HP's IRF switch virtualization technology, EVI delivers greatly enhanced reliability, resilience and faster remote vMotion capabilities. The combination of EVI and HP MDC brings multi-tenancy to Cloud-ready and remotely connected data centers

Virtual Private LAN services (VPLS)

VPLS delivers multipoint-to-multipoint connectivity between sites over a managed IP/MPLS network in a single bridged domain. Geographically dispersed sites can interconnect and communicate over MANs or WANs as if they are on the same LAN, promoting network resilience and response time. When combined with Ethernet, VPLS transforms the MAN into a virtualized Ethernet switch. This results in cost savings and operational efficiency while delivering all of the advantages of a true VPN and also provides the ability to migrate resources and services from one data center to another.

While MPLS is typically associated with IP VPNs, it is also the backbone transport mechanism for virtual private LAN services. Thus, it is possible to use the same infrastructure created for MPLS segmentation to create large Layer 2 networks.

Figure 11 VPLS

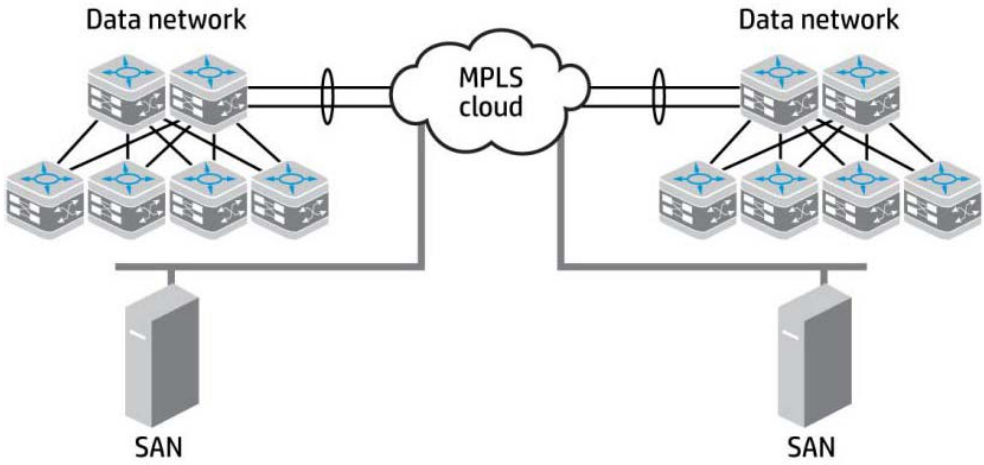

VPLS offers several benefits for connecting geographically separate data centers. Because VPLS provides a flat network, the deployment of services and management of the network are simplified. It operates on Layer 2 and makes VPLS present itself like an Ethernet switch. This allows network administrators to build higher-level networks, rather than building their networks around specific IP addresses.

VPLS uses edge routers that can learn, bridge, and replicate on a per-VPLS basis. These routers are connected by a full mesh of MPLS label switched path (LSP) tunnels, enabling multipoint-to-multipoint connectivity. Because the interconnection is created by direct access between the remote sites, there is less traffic overhead and improved performance for tasks such as disaster recovery and backup. Also, because data centers can be connected in a mesh, even if one data center loses a network link, all other sites will remain connected.

VPLS solutions also offer class of service (CoS) options, and the network administrator can configure VLAN tagging and define priorities.

Shortest Path Bridging (SPB)

SPB (**802.1aq Shortest Path Bridging)** is an emerging technology that greatly simplifies the creation and configuration of carrier, enterprise, and cloud networks, while enabling multipath routing.

SPB will be a replacement for STP protocols that wasted bandwidth by blocking traffic on all but one alternative path. It will allow all redundant paths in a network to be active and load sharing, and will provide much larger layer 2 topologies and faster convergence times.

SPB packets are encapsulated at the edge either in mac-in-mac *802.1ah* or tagged *802.1Q/802.1ad* frames and transported only to other members of the logical network. Unicast and multicast are supported and all routing is on symmetric shortest paths. Many equal cost shortest paths are supported.

Shortest Path Bridging-VID – SPBV

A primary feature of Shortest Path bridging is the ability to use Link State IS-IS to learn network topology. In SPBV the mechanism used to identify the tree is to use a different Shortest Path VLAN ID (VID) for each source bridge. The IS-IS topology is leveraged both to allocate unique SPVIDs and to enable shortest path forwarding for individual and group addresses.

Shortest Path Bridging-MAC - SPBM

SPBM reuses the PBB data plane which does not require that the Backbone Core Bridges (BCB) learn encapsulated client addresses. At the edge of the network the C-MAC (client) addresses are learned. SPBM is very similar to PLSB using the same data and control planes but the format and contents of the control messages in PLSB are not compatible.

Individual MAC frames (unicast traffic) from an Ethernet attached device that are received at the SPBM edge are encapsulated in a PBB (mac-in-mac) IEEE 802.1ah header and then traverse the IEEE 802.1aq network unchanged until they are stripped of the encapsulation as they egress back to the non participating attached network at the far side of the participating network.

SPB - key attributes and benefits for applications

Both SPBV and SPBM inherit key benefits of link state routing:

- The ability to use all available physical connectivity, because loop avoidance uses a Control Plane with a global view of network topology
- Fast restoration of connectivity after failure
- Under failure, the property that only directly affected traffic is impacted during restoration; all unaffected traffic just continues
- Rapid restoration of broadcast and multicast connectivity, because IS-IS floods all of the required information in the SPB extensions to IS-IS, thereby allowing unicast and multicast connectivity to be installed in parallel, with no need for a second phase signaling process to run over the converged unicast topology to compute and install multicast trees

Data Center Connected

Today's environments necessitate that users require global continuous access to data centers so they can access key applications. While this can be as simple as a single WAN connection, with the advent of Cloud Computing users are expecting to access their data, anytime, anywhere, demanding a wide range of connectivity options:

- High Performance Connectivity
- Multi-homed Internet and WAN Service Providers
- Data Center Interconnect

High performance connectivity options

Connectivity option becomes one of the key WAN requirements for any routers. Without the wide range of connectivity option you'll be limited to what services and connectivity options you can give to your users and customers. As such, HP's industry leading, multi-core, high performance routers are able to meet the most demanding Data Center network requirements. Whether you need High-Density 10GbE routing or a modular 1GbE routing platform, the HP routing portfolio has extensive support for various WAN connectivity options to meet the global demands for various connectivity options.

With the increasing demands for cloud services and the increased number of remote users accessing data, this places an enormous amount of pressure on the available bandwidth to serve mission critical applications. As such, QoS becomes an integral part of a good network design; a design that would allow the network traffic to be prioritized according to the SLA requirements of an organization. HP offers enhanced QoS options across the entire portfolio to enable your enterprise to make the most of its bandwidth.

Virtual Private Network (VPN)

VPNs are a primary means for organizations to allow remote users and sites to securely connect to the corporate resources by creating a private network through a public network infrastructure. When connecting remote sites or branch offices, site-to-site VPN technologies can be implemented on HP routers to permanently create a secure tunnel between the sites and the corporate headquarters allowing secure access of resources. HP routers also provide features for remote users such as sales people or remote workers, to create a remote access VPN where an individual user can establish a secure tunnel to access corporate resources as required.

VPNs are a critical for ensuring the confidentiality, integrity, and availability of the information across distributed enterprise networks and HP routers supports a wide range of VPN technologies including the following VPN implementations below to enable secure communication to the Data Centers:

IP Security (IPSec)

IPsec is a security framework defined by the IETF for securing IP communications. It is a Layer 3 VPN technology that transmits data in a secure tunnel established between two endpoints. IPsec guarantees the confidentiality, integrity, and authenticity of data and provides anti-replay service at the IP layer in an insecure network environment.

In addition, IPsec can be combined with Generic Routing Encapsulation (GRE) to provide to encapsulate the IP packets and combining IPsec and GRE provides the necessary end-to-end security for site-to-site communication.

Dynamic Virtual Private Network (DVPN)

DVPN is a hub and spoke VPN architecture that allows multiple remote branch and regional offices (spoke) to establish site-to-site IPsec VPN tunnels to secure connectivity to the headquarters or data centers (hub). HP DVPN is an architecture rather than a protocol. HP DVPN helps enterprises to simplify the configuration and management of IPsec VPN tunnels (HP DVPN policies share the security access, management, and QoS policies) to easily connect thousands of remote branch and regional offices to the corporate headquarters or data centers.

DVPN provide superior OpEx as the deployment and maintenance is simplified. There is no longer a need to login to each VPN device to manually set up site-to-site VPN tunnels at each branch or regional office, corporate headquarter. Data center DVPN is an HP innovation to simplify secure WAN connectivity for the enterprise.

DVPN collects, maintains, and distributes dynamic public addresses through the VPN Address Management (VAM) protocol, making VPN establishment available between enterprise branches that use dynamic addresses to access the public network. DVPN allows enterprises to significantly reduce manual VPN configuration and complexity with a higher level of automation to securely connect branch, regional offices, campus/headquarters, and data center locations with secure site to site VPN connectivity. DVPN provides for greater IT agility and simplified configuration and deployment across distributed enterprise locations.

DVPN Advantages:

- **Automated VPN Tunnel Setup and Simplified Configuration:**

  With DVPN, you only need to configure your head end/DC router(s) once with very little configuration needed on your branch routers. DVPN automatically builds IPsec VPN tunnels between the branch, regional office, campus and data center locations. This approach eliminates the need for complex IPsec VPN configuration parameters at remote branch sites that may have little to no IT personnel onsite. If the address of the router is changed by the service provider, DVPN automatically rebuilds IPsec tunnels without a problem
- **Highly Scalable and Resilient:**

  Scales to 1000's of nodes per DVPN domain, with virtually infinite horizontal scalability. Provides active-active and fail-over capabilities for redundant, resilient VPN connectivity
- **Flexible Deployment Options Across Multiple Domains:**

  DVPN solutions can provide flexible topology options supporting combinations of large-scale hub and spoke and full-mesh/partial-mesh topologies on demand to satisfy the simplest or most complex enterprise WAN environments

Internet and WAN Service Providers

Despite 70-80% of network traffic traversing between servers within the data center, connectivity to the end-users, makes up about 20% of the traffic flow, built around a hierarchical model. This is the migration from a flat Layer 2 design at the server with high speed 10/40/100GbE links, to limiting WAN interconnects used by dispersed global users. When determining the level of high availability (HA) within the data center, one should note all variables; including diverse entrance facilities, diverse Internet providers, etc.

When providing external connectivity to the Data Center, the variations are via the Internet, an Intranet and Extranet. With the advent of cheap DSL and advances in security, DMVPN functionality

has continued to gain popularity. When larger dedicated WAN circuits are required, routers are a requirement for SONET, ATM, and Frame Relay.

Data Center Interconnection (DCI)

As enterprises race to adopt virtualized environments and Cloud computing, they need to deliver infrastructure as a service and to connect geographically dispersed data center to meet rising customer expectations.

HP routers and switches provide extensive support to enable the data center interconnect and below are some of the technologies which enable the DCI.

HP Ethernet Virtual Interconnect (EVI)

Many interconnect methods suffer from limitations, including transport dependency, complexity and lack of resiliency. HP EVI is designed to address these limitations by delivering responsive, efficient and resilient data center interconnect solution.

HP EVI, a new HP Virtual Application Network innovation and component of the HP DCI solution, enables IT staff to simplify the interconnectivity of up to eight geographically disperse data centers. As a result, clients improve disaster recovery capabilities and can quickly respond to changing market conditions by easily moving virtual machines in minutes without network delay to any interconnected data center.

HP EVI runs over Internet Protocol (IP) transport and extends layer 2 domains across the networks of the connected data centers. By virtualizing and automating a layer 2 domain across data centers, HP EVI delivers the elements necessary to enable a software-defined networking (SDN) data center infrastructure. It enables several data centers to work as one that is more responsive, with higher efficiency and solid high availability for business resiliency. With EVI, enterprises are able to accelerate the delivery workload mobility with remote vMotion, increase applications performance with load balancing, and achieve optimum degrees of high availability and disaster recovery for valuable data. When used along with HP's IRF switch virtualization technology, EVI delivers greatly enhanced reliability, resilience and faster remote vMotion capabilities. The combination of EVI and HP MDC brings multi-tenancy to Cloud-ready and remotely connected data centers.

Multi-Protocol Label Switching (MPLS)

MPLS encapsulates network layer packets with labels of short and fixed length. Here, multiprotocol means that MPLS supports multiple protocols, such as IP, IPv6, and IPX; label switching is to attach a label to a packet and to forward the packet according to the label without performing complex route searching and packet forwarding operations as IP does. MPLS was originally brought forth to combine the advantages of IP and ATM. IP is intended for connectionless control and data forwarding, while ATM is intended for connection-oriented control and data forwarding. MPLS adopts connectionless control and connection-oriented data forwarding. MPLS obtains link layer services from the link layer and provides connection-oriented services for the network layer. MPLS can obtain support from IP routing protocols and control protocols and, at the same time, support policy-based constrained routes. It possesses powerful and flexible routing functions and is capable of satisfying the networking requirements of various new applications.

MPLS is a protocol initially developed for increasing forwarding speed of routers. However, thanks to its inherent advantages, it has gained still wider applications in traffic engineering, VPN, and QoS and

is becoming an important standard for large-scale IP networks.

Multicast Virtual Private Network (Multicast VPN)

Multicast VPN is a technique that implements multicast delivery in MPLS L3VPN networks. An MPLS L3VPN is a VPN implemented based on the extension technologies of the Border Gateway Protocol (BGP) and MPLS. It comprises a set of customer sites that are interconnected only by means of an MPLS provider backbone network. The VPN can be regarded as a set of policies that control the interconnections between these sites.

MPLS Layer 2 Virtual Private Network (MPLS L2VPN)

For traditional VPN, the transmission of the data flow between private networks on the public network is usually realized via such tunnelling protocols as GRE, L2TP, and PPP. MPLS-based VPN can connect geographically scattered branches of private networks by using LSP. MPLS-based VPN also supports the interconnection between different VPNs. MPLS L2VPN provides L2 VPN services based on MPLS networks. From the perspective of users, such an MPLS network is an L2 switching network, through which the L2 connections can be established between different sites.

Virtual Private LAN Service (VPLS)

VPLS, also called “Transparent LAN Service (TLS)” or “virtual private switched network service”, can deliver a point-to-multipoint L2VPN service over public networks. With VPLS, geographically-dispersed sites can interconnect and communicate over a MAN or WAN as if they were on the same LAN.

VPLS provides Layer 2 VPN services. However, it supports multipoint services, rather than the point-to-point services that traditional VPN supports. With VPLS, service providers can create on the PEs a series of virtual switches for customers, allowing customers to build their LANs across the Metropolitan Area Network (MAN) or WAN.

Securing the virtual server edge

Setting technology aside for the moment, a data center is a location where individuals consider their data safe from theft and environmental disaster. A data center’s ability to provide a reliable and safe infrastructure for applications and data is critical. If a system is unavailable for some reason, business operations will be impacted. Because data centers house critical data, it would be logical that, in addition to taking other precautions to protect servers physically, information security should possess the same importance.

Data security is designed or implemented to prevent or minimize security breaches. Traditional enterprise security controls do not take into account the complexity of the network and the systems that they are designed to protect. As in any control, security can have a negative impact on performance and data center design flexibility.

Key security implications - compliance, convergence, and consolidation

Within a network, internal and external threats continue to grow while our networks become more complex. Today, there are three key security implications that increase the complexity of data center security—compliance, convergence, and consolidation.

Compliance

Compliance has gained traction as organizations deal with regulations to mitigate the risk in data loss and application downtime. Companies need to comply with regulations such as the Payment Card Industry-Data Security Standard (PCI-DSS) for the retail industry and the Health Insurance Portability and Accountability Act (HIPAA) for the healthcare industry. These regulations deal with data security and privacy, and data centers need to ensure that adequate security measures are in place.

Regulatory compliance and industry controls

- **The Health Insurance Portability and Accountability Act 1996 (HIPAA):**

  Applies to health plans, health care clearinghouses, and those health care providers who electronically conduct certain financial and administrative transactions that are subject to the transactions standards adopted by the Department of Health and Human Services
- **The U.S. Sarbanes-Oxley Act of 2002 (SOX):**

  Requires CEOs and CFOs to attest to the accuracy of corporate financial documents, as well as provide IT and networking controls and their audit as per Section 404 of the Act
- **The Gramm-Leach-Bliley Act of 1999 (GLBA):**

  Applies to banking and financial industries control. Institutions must provide clients a privacy act that explains information collection, sharing, use, and protection
- **The Payment Card Industry Security Standards Council (PCI-DSS):**

  PCI-DSS is a security standard which describes requirements for security management, network architecture, software design and other critical protective measures. The standard is intended to help organizations proactively protect customer account data
- **The Federal Information Security Management Act 2002 (FISMA):**

  Federal information security guidelines
- **The North American Electric Reliability Corporation (NERC):**

  Standards and guidelines established for the national electric grid and its operation. IT and network controls are a major component of the guidelines for security
- **DOD Directive 8500.1:**

  Security directive established by the Department of Defense for information assurance in 2002
- **The international accord on banking operations (BASEL II):**

  Of which information security and assurance are major components

All of the outlined regulatory and industry standards affect the operations and design of security systems for data centers. As converged network infrastructure and virtualized systems become more prevalent in the data center, it is important to identify the risks, the controls, and the weaknesses. It is important to then close the gaps and continue to reassess the network's security readiness.

Data centers need to adhere to compliance standards by providing that adequate measures are in

place to ensure data security and privacy. The steps involved in meeting this requirement are:

- Identifying the vulnerabilities and threats, and setting up the appropriate security controls
- Quantifying the risk exposure. This includes identification of the vulnerabilities and threats, identification and definition of their compensating controls, and the calculation and documentation of risk
- Creating an IT risk management plan
- Ensuring that all the IT compliance requirements are met
- Ensuring that all the personnel are trained in IT security awareness

Convergence

The convergence of multiple security functions into products that support a single function has evolved to include policy management integration. The three security product functions—firewalls (legacy as well as next generation firewalls), intrusion prevention systems (IPS), and content security gateways—are converging into an integrated perimeter and data center security function. In addition to these, other security product functions such as security information and event management (SIEM) tools, and network behavior anomaly detection (NBAD) tools are also being integrated into policy management functions. Integration of these different types of products simplifies the management of network security and reduces the cost of overall security. However, integrated security products cannot compromise security for the sake of integration and convergence.

Figure 12 Integrated policy management

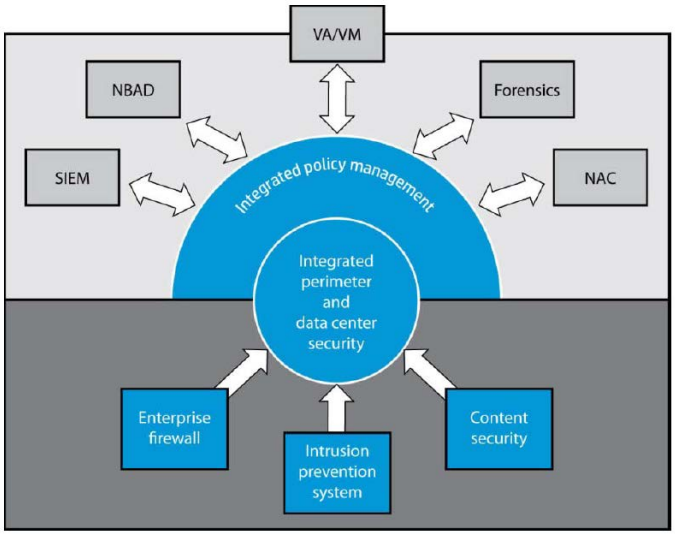

In the context of data center security, convergence refers to the integration of various tools such as firewalls, intrusion prevention systems, and monitoring systems into an integrated tool. While this may simplify management and monitoring, care should be taken to ensure that security is not compromised.

There are two elements of convergence that need to be considered:

- **Data centers can no longer rely on a single security perimeter around the enterprise network:**

  Individual assets, such as a VM or a host in the virtualized environment, need to be protected. In order to achieve this, purpose-built intrusion prevention systems need to be considered that

would inspect both ingress and egress traffic and block threats automatically. Additionally, the products should support automated updates so that newly discovered threats are addressed and the network vaccines are updated automatically

- **Division and convergence of security responsibilities:**

  Network teams are usually responsible for network security and the operations/server teams for OS security. With the introduction of virtualization, some network elements are no longer physical devices and may form part of the virtualized environment, such as the vSwitch. As this is a software component, it would usually fall within the responsibilities of the operations team. The operations team may not necessarily be aware of how to implement network security and, as a consequence, security at the vSwitch may not be enforced correctly. To address issues such as these, the network team, the server team, and the security team need to work together closely to look at all aspects of security

Consolidation

The third key security trend is the consolidation of data center infrastructure. In the past, data centers were built for the purpose of bringing teams in the organization closer to the applications and data that they required, resulting in a dispersed data center infrastructure. New applications such as Web applications, new protocols such as IPv6, and new traffic types such as voice and video traffic increase the complexity of the security function. The consolidation of data types and security functions onto similar network equipment in data centers requires IT teams to not only understand the technology and the management functions, but understand the impact they have across technical teams.

The complexity and criticality of data center security has only increased. Hackers not only seek to cause disruption in the data center, but in many cases cause governmental, commercial, and individual mayhem and harm. In the past, these threats were typically handled by IT security teams with a single security perimeter comprised of both firewalls and antivirus products.

However, threats are now more sophisticated and many times are caused by professionals who seek to steal data, modify it, and illegally distribute it for financial, political, or other gains. In addition, the rise of new applications and services such as file sharing, database access, and video conferencing has contributed to a new set of security weaknesses. Data centers can no longer rely on a single security perimeter around the network; separate security perimeters around individual assets in the network and in the data center need to be built.

Figure 13 Security perimeters around individual assets in the network in the data center

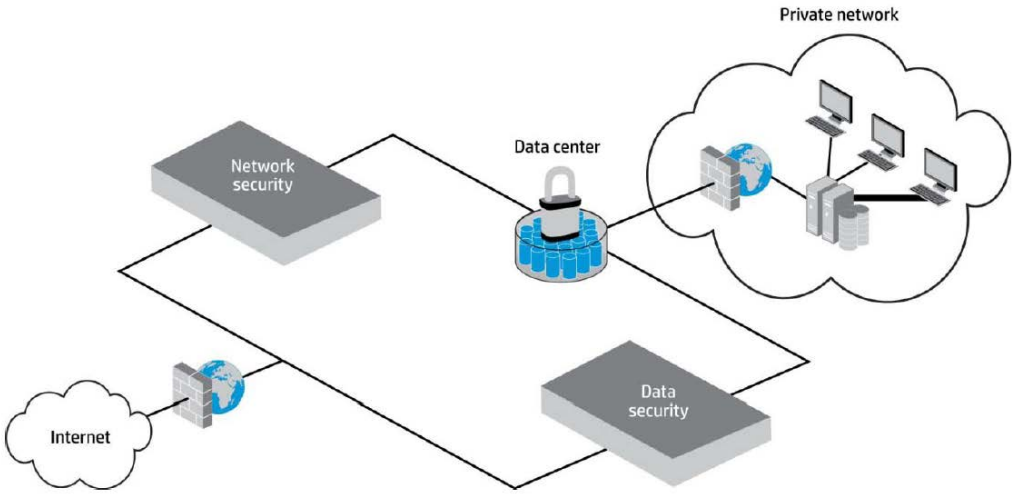

Without reviewing the advantages of converged network infrastructure and server virtualization, it is important to note an associated negative trend that goes along with the increase in virtualized data center resources. This trend is not necessarily caused by negligence, but more through a lack of understanding regarding the impact of virtualization on underlying security. Gartner states that "60 percent of virtualized servers will be less secure than the physical servers they replace through 2012[2]."

In an attempt to cut costs, companies are consolidating their data centers by reducing the number of geographically separate locations. In addition, the number of physical servers in each data center is being reduced due to virtualization technologies. Here, we take a look at some of the challenges introduced by virtualization.

[2] http://www.gartner.com/it/page.jsp?id=1322414

Figure 14 Virtualized servers

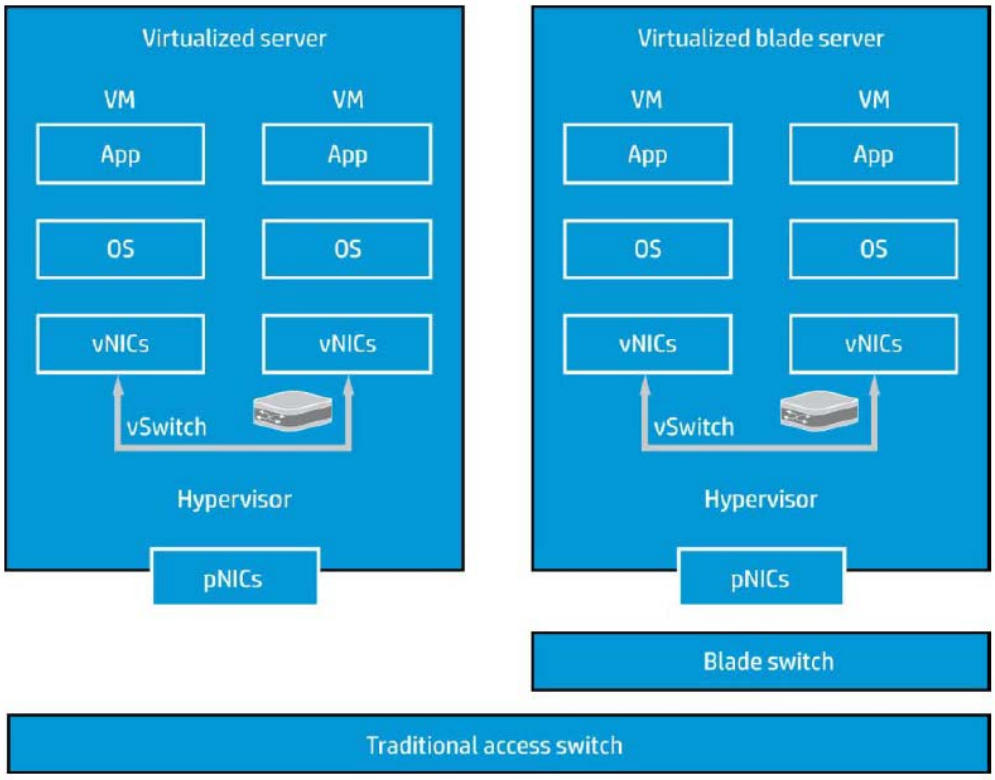

A virtualized environment comprises a physical server (host) that hosts multiple operating systems (guests) and allows them to run concurrently. As discussed earlier, management of VMs is carried out by a hypervisor, which presents itself to the guest operating system as a virtual operating platform and monitors the execution of the guest operating system. Multiple guest operating systems share the same virtualized resources, such as the CPU and the Ethernet port.

Challenges introduced by hypervisor and virtual switch

Hypervisor security

Securing the hypervisor is critical, as a compromised hypervisor could lead to a situation where access to the hosted operating systems may be made available to attackers. Another possibility is a denial of service (DoS) attack on the hypervisor, which could affect all the hosted virtual systems.

Virtual switch security

The virtual switch is a software component that allows communication between VMs and other systems. The packets are examined and sent to other VMs or to the physical switch, depending on the destination. In a virtual environment, it is important to consider that the network team may not have visibility to all of the network traffic. For example, the communication between two VMs on a single physical server would not reach the physical switch and, as a result, the traffic may not even be detected.

Mobile workloads also require consideration. VMs may move from one physical server to another due to situations such as an increased load that dynamically shifts the VM onto another physical server

supporting the hypervisor, or a physical server failing, as depicted in the following figure.

Figure 15 Increased load on a physical server or a physical server failing

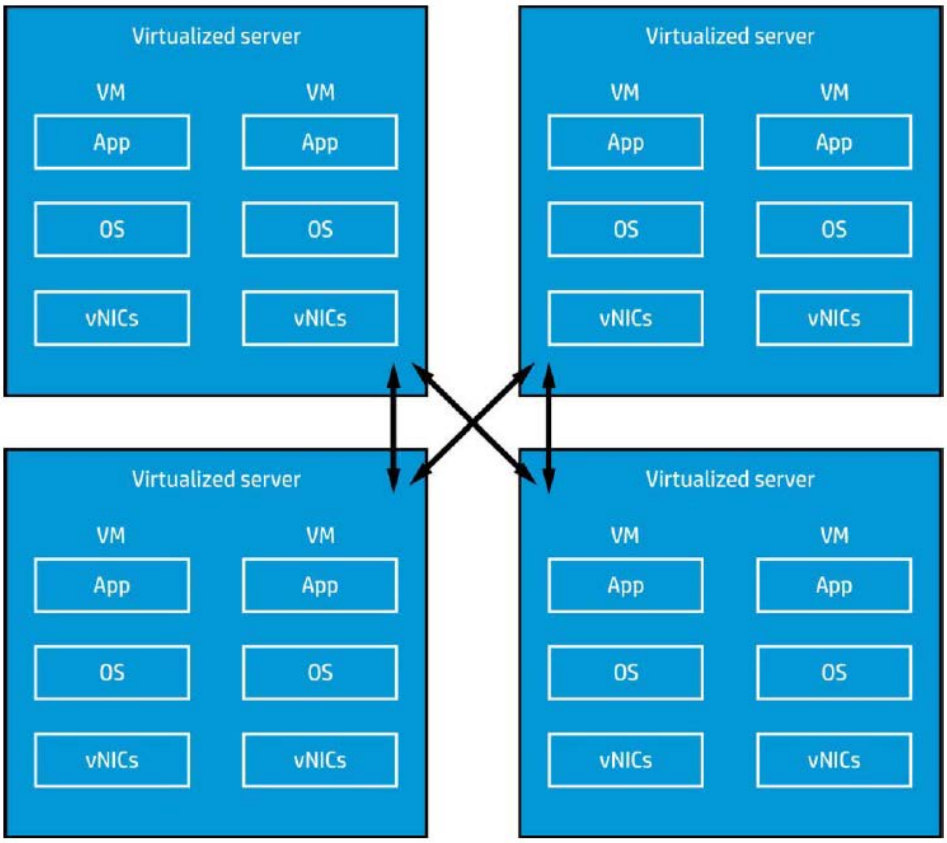

Due to this kind of mobility, unless adequate controls and monitoring methods are established, it may be difficult to implement policies. Policy and zone definitions must follow the mobile workloads. In particular, security policies must be maintained for the different workloads at rest, operating, and in motion.

Figure 16 IPS and Secure Sockets Layer (SSL) in the converged network infrastructure

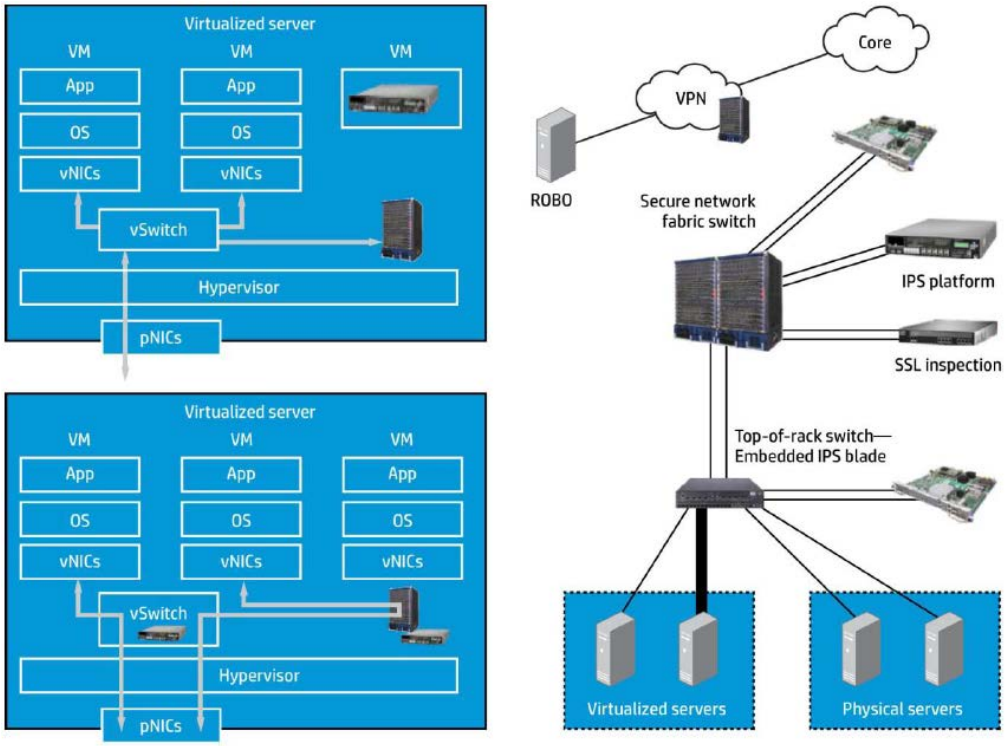

The figure above depicts:

- An IPS device, external to the physical server that hosts multiple VMs at the network level
- A virtual IPS (vIPS) that resides on the host

Virtual and physical IPSs can be used separately, but combining the two provides the advantage of monitoring traffic at the network as well as the virtual layers. One such product that provides a superior platform is the HP TippingPoint IPS solution.

As discussed above with respect to mobile workloads, the policies should be ideally applicable to workloads at rest, while operating, and in motion. A product that can integrate disparate management tools and provide the management of resources, services, and users simplifies the task of implementing policies. One such product is HP Intelligent Management Center (IMC).

Managing and provisioning the virtual server edge

Network infrastructure convergence and virtualization have allowed multiple networks and systems to be collapsed into unified entities. This trend, from the outside looking in, may appear to have simplified the tasks and processes associated with data center systems and network management, but in many cases it has added complexity.

With converged infrastructure, responsibility and ownership are key issues. As systems and network technologies merge, the responsibility to install, provision, maintain, and monitor still exists. This is not a new topic; it has been seen in the merger of voice and data into the enterprise network and has been addressed by multiple management experts and forums. This section focuses on the identification of the technical hurdles and the tools needed to manage systems and networks in the converged infrastructure data center.

The converged infrastructure data center will not be comprised 100 percent of VMs and DCB networks. Currently, there appear to be fewer physical entities to manage in data centers, and hardware consolidation is going to continue. But, keep in mind the term “convergence.” As the systems and network services converge, so must the management platforms. They will also become true converged platforms.

Single-pane-of-glass management is that converged platform. This is what data centers have been trying to achieve with simple network management protocol (SNMP) and agent-driven management software suites. At the same time, the roles and responsibilities of the data center network and systems managers have blurred. The lines of responsibility are eroding and management systems are being required to evolve alongside that merging of responsibilities. Since the role of data center technical teams is changing, management systems too must change. When evaluating management platforms, there are important questions to ask:

- Do these management systems really provide a simpler management platform to work with, or do they just provide the ability to view and receive alerts from every device in the network?
- Do these platforms deliver an interface that allows for native system management routines aggregated together as though they are one integrated software system?
- Are these just portals that aggregate the systems and network management tools that multiple vendors provide onto a manager of managers?

A fully functional management platform must provide the following:

- A single view of all systems with access to alarms and configuration tools
- Secure communications to the managed endpoints
- Support for secure access to individual VMs and the underlying virtualization layer
- Multiple types of system and network reporting schemas
- Flexible workflow configurations that management teams can customize for their organizations
- Flexible user, group, and role/responsibility security configurations
- A native application look, feel, and performance
- Availability and resilience

These capabilities will be addressed in greater detail later in this document with a focus on what is currently being used by data center managers and the path to single-pane-of-glass management.

Virtualization management

One of the objectives of convergence and virtualization is to make it more manageable. The fewer distinct elements that need to be managed, the less complexity is required. Unfortunately, the management of converged and virtualized network elements did not automatically converge at the same time. Again, HP’s major goal is to simplify the network without impact to performance and flexibility. When we look at the virtualized network, storage, and server platforms through respective tools, they should provide a combined view of the independent pieces, thereby reducing the complexity of management.

HP IMC can provide the data center operations team a combined view of the entire infrastructure, enabling them to provide efficient end-to-end business management to address the stringent demands of today’s mission-critical enterprise IT operations.

The HP vision for cloud networks

The explosion of the cloud creates new opportunities and a new set of challenges. Networks must be faster and more flexible to support the needs of diverse mobile users, a fragmented security perimeter, and a constantly changing set of applications and devices.

An integrated module to IMC, the HP Virtual Application Network Manager (VAN) is designed to characterize applications, virtualize the network to align the resources needed for delivering the application, and automate the orchestration. With VAN, IT teams can provision network services faster, more consistently, more securely, and reduce downtime due to manual errors.

VMware management

HP IMC integrates with the VMware vCenter and allows for centralized management of hundreds of VMware ESXi/ESX hosts and thousands of VMs, delivering operational automation, resource optimization, and HA to IT environments. Using a single management client for all tasks, administrators can provision, configure, start, stop, delete, relocate, and remotely access VMs.

HP, along with VMware, provides solutions for creating standardized and repeatable processes. This includes managing capacity and performance, controlling configuration changes, protecting business-critical applications against outages, and self-service provisioning for development and testing environments. Users of other virtualization platform vendors must rely on third-party products to receive the same breadth of virtualization management functionality.

Server management

There are multiple tools that aid in the support of server management. While there are multiple tools available, the overriding goal is to fully integrate management tool capabilities through APIs and/or Web-enabled applications into an integrated management system like HP IMC. VMware virtual servers and their virtual network components are integrated into IMC, providing a single interface to view, provision, move, and manage the virtualized network.

Traffic flow analysis and capacity planning

Analyzing network traffic to determine which applications, servers, and clients are consuming network resources is a best- practice method for understanding application delivery problem root cause issues. Traditional methods of analyzing network consumption, such as link or device utilization, do not provide IT staff with adequate insight into why network resources are being consumed. HP IMC's centralized network platform provides the intelligent collection, analysis, and reporting services that increase the insight that can be gleaned from flow data to circumvent issues, help ensure QoS, and expedite remediation.

NetFlow, sFlow, and Netstream

The flow data produced by networks generally comes in one of three common formats—NetFlow, sFlow, or NetStream. Standardization around these formats makes it possible for routers and switches to send their flow data to a wide variety of collectors and analysis tools, and to be combined

with flows in multi-vendor networks for wider analysis. Flow data has now become an important part of network scalability and performance, particularly in busy router cores and edge devices that handle a large number of concurrent, short-duration flows.

- **NetFlow** is the oldest of the flow formats. It originally served as a caching algorithm in network devices, which helped optimize network efficiency. As this data was already being collected, it made sense to eventually export it for analysis and reporting purposes, which could be done without much additional overhead to the network device. NetFlow has spawned various iterations—it is now up to version 9—as well as similar formats optimized for different purposes and processing
- **sFlow** was created as a standard in 2001 for high-speed networks based on sampled data rates rather than 100 percent packet capture. sFlow was developed exclusively as a monitoring technology. It is scalable and can provide more detailed statistics on all Layers 2-7 throughout the network. As a result, it has gained wide acceptance from network vendors. sFlow is fully implemented in all HP Networking devices today
- **NetStream**, a flow format created by 3Com (now HP) for its enterprise networking products, includes additional flow details. NetStream provides detailed flow information, is compatible with NetFlow and is implemented in HP Networking routers and available on high-end switching platforms via an add-on module

HP IMC, as the collection and analysis engine, can handle flow data in all of these formats as well as other formats from a wide range of devices from many manufacturers in order to provide network-wide visibility.

SNMP

With the rapid development of the Internet, two problems were introduced:

- An increase in the number of networks and network devices makes it difficult for administrators to monitor the status of all the devices in time and identify and correct network faults
- Network devices may be from different vendors, providing an independent set of management interfaces (e.g., command lines), making network management more complicated

SNMP is an application-layer protocol between an network management system (NMS) and agents introduced to solve the above-mentioned issues. It defines the standardized management framework, common languages in communication, security, and access control mechanisms used in monitoring and managing devices in a network. The administrators can query device information, modify device parameters, monitor device status, and enable automatic detection of network faults and generation of reports by using SNMP.

SNMP provides the following advantages:

- A TCP/IP-based standard protocol, with UDP as the transportation layer protocol
- Automatic network management. Administrators can search and modify information, find and diagnose network problems, plan for network growth, and generate reports on network nodes
- A combination of simple request-reply mode and active notification mode, timeout, and re-transmission mechanism
- Few packet types and simple packet format, which facilitates resolution and implementation

HP IMC has a built-in SNMP collector that allows it to receive SNMP messages and statistics from any SNMP device in the network, including switches, routers, servers, NAS, etc.

Single-pane-of-glass management

HP IMC is a next-generation management software which provides the data center operations team with a comprehensive platform that integrates network technologies and provides full fault, configuration, accounting, performance, and security (FCAPS) management functionality.

Built from the ground up to support the Information Technology Infrastructure Library (ITIL) operational center of excellence IT practices model, IMC's single-pane-of-glass management paradigm enables efficient end-to-end business management to address the stringent demands of today's mission-critical enterprise IT operations.

Configuration management—backup

Configuration management can be defined as the ability to control changes to the operating status of a managed device, as well as the ability to detect and/or prevent unauthorized changes in the operational status of a device. Maintaining an accurate inventory of last known hardware, software, and configuration information enhances this function.

To manage a data center, the operations team must have an up-to-date configuration inventory across all types of devices, irrespective of vendors. HP IMC has earned bragging rights; it supports 6,000 devices from more than 220 manufacturers, of which 1,400+ are from Cisco!

Figure 17 Configuration of over 5,786 devices from more than 150 manufacturers

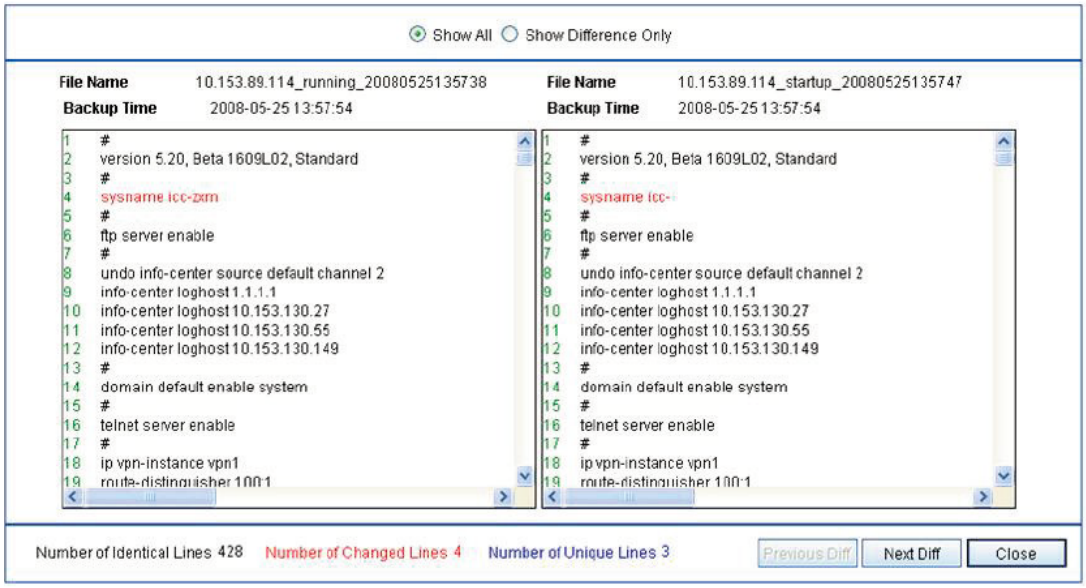

The time required to roll out network changes and the likelihood of configuration errors are both greatly reduced with IMC's powerful bulk configuration functionality. Configuration baselining helps ensure that changes to the stable network configuration are flagged promptly.

The powerful configuration comparison feature provides the rapid identification of configuration differences, enabling the system administrator to either accept the new configuration or roll back to the original stable configuration. Additional functionality includes bulk backup and restore, an extremely flexible agent management function that enables the system administrator to have total

control of the upgrade process.

The ability for change detection and the system inventory also act as informational resources for the fault management process. Understanding how a device is supposed to be configured enables the database operations staff to recognize and correct faults. Change recommendations can only be made after the current configuration status is completely known, understood, and stable.

Configuration management is both reactive (changes can be identified over time) and proactive (control can be established to prevent configuration changes). Small configuration changes can go unnoticed on a LAN or WAN and they can completely collapse an entire nationwide network. Arguably, configuration management is the most important component of the FCAPS model.

Traffic analysis and capacity planning

As the enterprise network infrastructure expands to support different types of traffic and users, traffic management becomes critical, and complete visibility into a network's behavior becomes more important and more challenging. What is or is not happening throughout the network grid—including application performance, bandwidth utilization, network congestion, and appropriate prioritization of user and application traffic—are questions that often go unanswered.

In today's connected business environment, straightforward and effective traffic management from the network core to the network edge is essential. Enterprises need a network infrastructure that scales to meet new business needs and manages added complexity in a cost-effective manner. In addition, the data center operations team is expected to control the network in such a way that it is transparent to users. Essential information assets need to be instantly available around the clock. However, this is impossible to achieve without the right tools to make smart, informed decisions.

Most network administrators do not have simple, affordable tools that can quickly answer the following questions, regardless of the size of the network:

- Is network performance slowing down or becoming congested?
- Is a NIC chattering, effectively clogging the network?
- What is the current network usage, and what has it been in the past hour?
- Which network routers are most active or over-utilized?
- Why is a server slow or inaccessible?
- Which users and applications are driving network traffic?
- Which users and applications are starving for bandwidth?
- How much bandwidth do I need for new applications?

HP IMC Network Traffic Analyzer (NTA) is a graphical network monitoring tool that utilizes industry-supported flow standards to provide real-time information about the top users and applications consuming network bandwidth.

Figure 18 Applications consuming network bandwidth 1

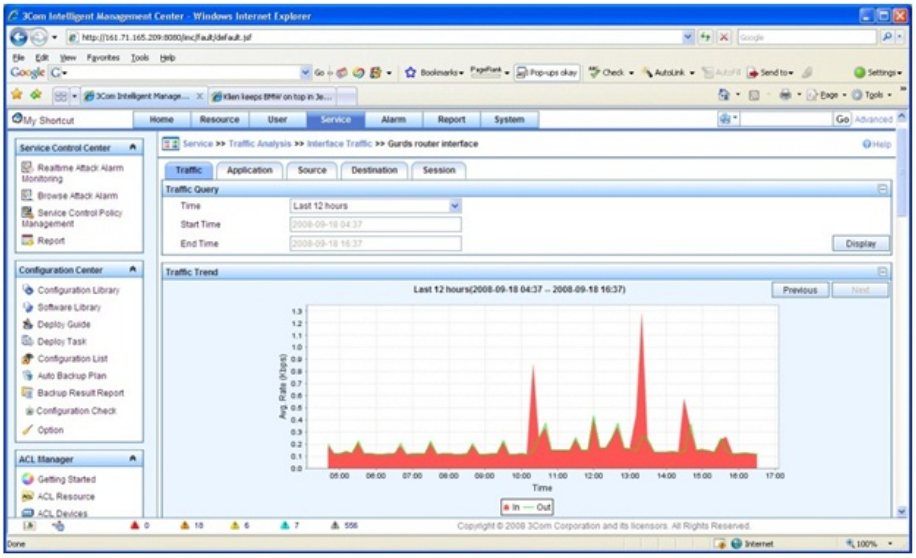

Figure 19 Applications consuming network bandwidth 2

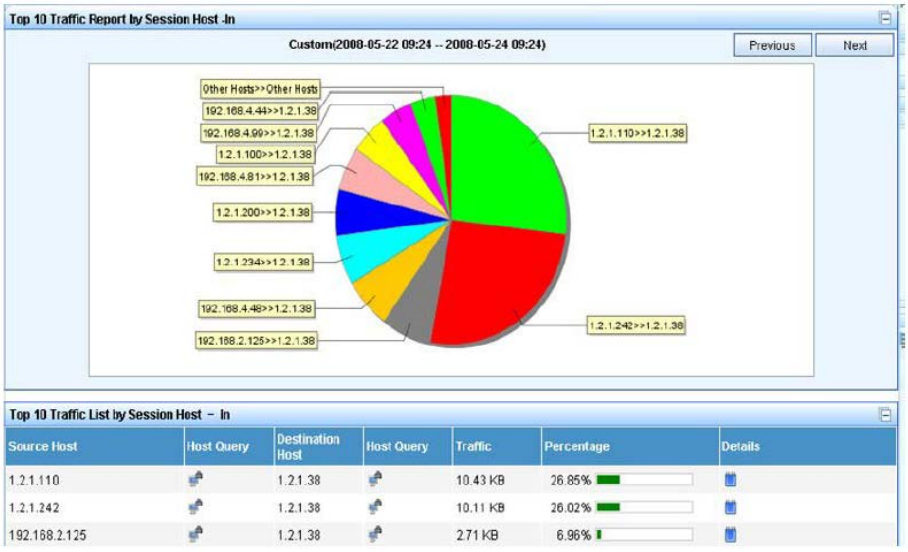

A reliable solution for enterprise and campus network traffic analysis, HP IMC NTA statistics help network administrators better understand how network bandwidth and resources are being used, as well as which source hosts carry the heaviest traffic. This information is invaluable in network planning, monitoring, optimizing, and troubleshooting—IMC NTA identifies network bottlenecks and applies corrective measures to help ensure efficient throughput.

Out-of-band management

Out-of-band management (OOBM) operates on a management plane, which is separate from the data plane used by data traffic on the switch and by in-band management traffic. That separation means

that out-of-band management can continue to function even during periods of traffic congestion, equipment malfunction, or attacks on the network. In addition, it can provide improved switch security. A properly configured switch can limit management access to the management port only, which prevents malicious attempts to gain access via the data ports.

Network OOBM typically occurs on a management network that connects multiple switches. It has the added advantage that it can be carried out from a central location and does not require an individual physical cable from the management station to each switch's console port.

In a typical data center installation, ToR switches connect servers to the data network, while the management ports of those switches connect to a physically and logically separate management network. This allows network administrators to manage the switches even if operation on the data network is disrupted.

In the illustration below, the switches face the hot aisle of the data center, allowing easy connection to the network ports on the backs of the servers.

Figure 20 Network out-of-band management in a data center

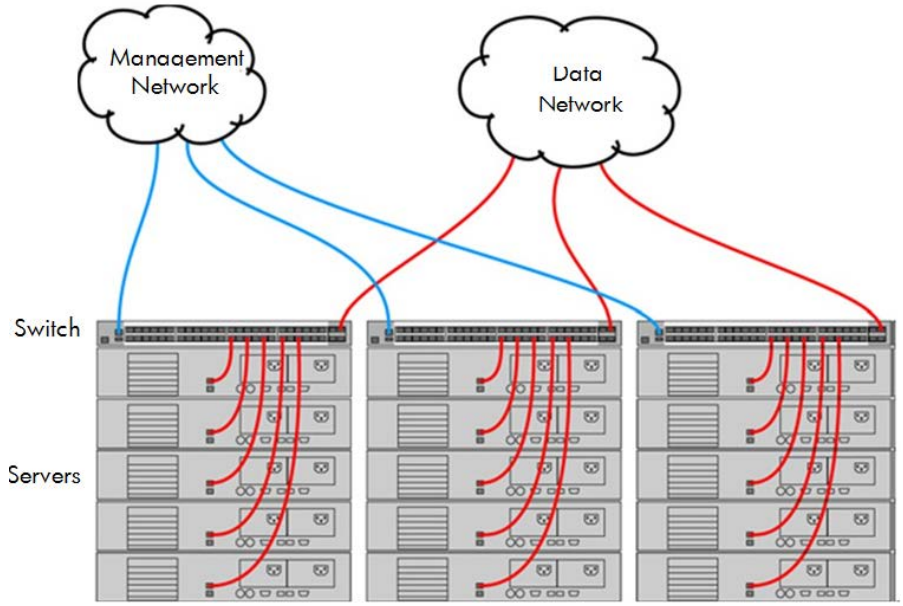

For even more control, the serial console ports of the switches can be connected to the management network through a serial console server (essentially, a networked serial switch), allowing the network administrators to view the console activity of each switch at boot time and to control the switches through the console ports, as well as through the management ports.

Port mirroring

Port mirroring, also known as switched port analyzer (SPAN), is a method of monitoring network traffic. In port mirroring, the packets passing through a mirroring port/VLAN are copied to another port (called the monitor port) connected to a monitoring device for packet analysis. The administrator can select to port mirror inbound, outbound, or bi-directional traffic on a port/VLAN as needed. A network administrator uses port mirroring as a diagnostic tool or debugging feature, especially when fending off an attack. Port mirroring can be managed locally or remotely.

In local mirroring, the mirroring ports/VLANs and the monitor port are located on the same device, while in remote mirroring traffic on the mirroring ports, or ports in the mirroring VLANs, is mirrored to the monitor port located on another device.

Network traffic monitoring is needed for packet analysis or IPS deployment. However, monitoring all the traffic in a large switching network is difficult, so port mirroring can be configured to copy selected traffic of a port or ports to a specific port for monitoring.

Converged network infrastructure: unifying data and storage networks

Convergence is a technical term historically used to express the combining of voice and data onto the same network fabric. Now expressed as a converged network infrastructure, it encompasses the sharing of network resources between data and storage networks. This trend constitutes a move towards a unification of data and storage networks.

Networked attached storage

The requirement for users to access shared data on the network has grown immensely over the past decade. Fibre Channel (FC) was developed to meet the high speed and reliability needs of enterprise customers and was very expensive. A more cost-effective solution was to use the existing IP infrastructure and attach storage devices to it (NAS).

NAS devices can scale to multi-terabyte capacity, but typically don't meet the performance and reliability requirements of today's large scale data centers. Protocols like network file system (NFS), common internet file system (CIFS) (Samba), etc. are the foundation of NAS.

NAS is a dedicated storage device that operates in a client/server mode. Clients connect to NAS via the LAN. The protocols supported for NAS are NFS, and CIFS and use cases are shown below.

- NFS and SMB may have lower performance in comparison with SANs. NAS is connected over the LAN using TCP/IP
- Network file system (NFS)—UNIX®/Linux®
- SMB – (Server message block) Windows remote file system (drives) mounted on the local system (drives). This evolved from Microsoft NetBIOS, NetBIOS over TCP/IP (NBT), and CIFS
- Samba—SMB on Linux (making Linux a file server for Windows clients)

Advantages

- No distance limitations
- Bridges the SCSI and serial advanced technology attachment (SATA) logical unit capacity limitations
- Provide multi-client access to files
- Snapshot, cloning and replication capabilities
- Can provide simultaneous access to files for both Windows and NFS clients
- Less complex to deploy and manage than SAN
- Wide range of available devices from consumer oriented to high-end enterprise solutions

Disadvantages

- Performance (although recent advances have closed or eliminated this gap)
- Latency
- Older protocols are less secure than the newer versions (NFSv4 and SMB2)

Weakness

- NFS and CIFS are inherently insecure protocols. Encapsulation of NFS and CIFS traffic into tunnels only adds to the speed and latency issues

Figure 21 Typical NAS configuration on a LAN

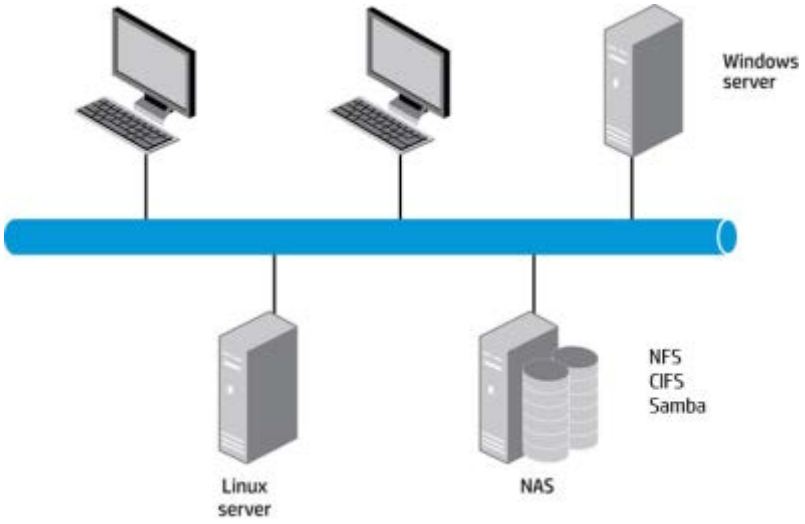

In storage systems, the read/write transactions span from the magnetic domains on the hard disk to the operating system kernel. The OS treats the storage as locally connected, though the data might actually be traversing multiple hops. A storage network mimics the physical path normally traversed by data, from the systems PCI bus, and then to the processing engine and operating system OS. The speed of the storage I/O operations is critical and delays can result in operating systems hanging or even crashing. Also, it is critical to ensure that networks utilized to transport storage traffic are reliable and resilient to common errors.

Figure 22 Generic network transport architecture

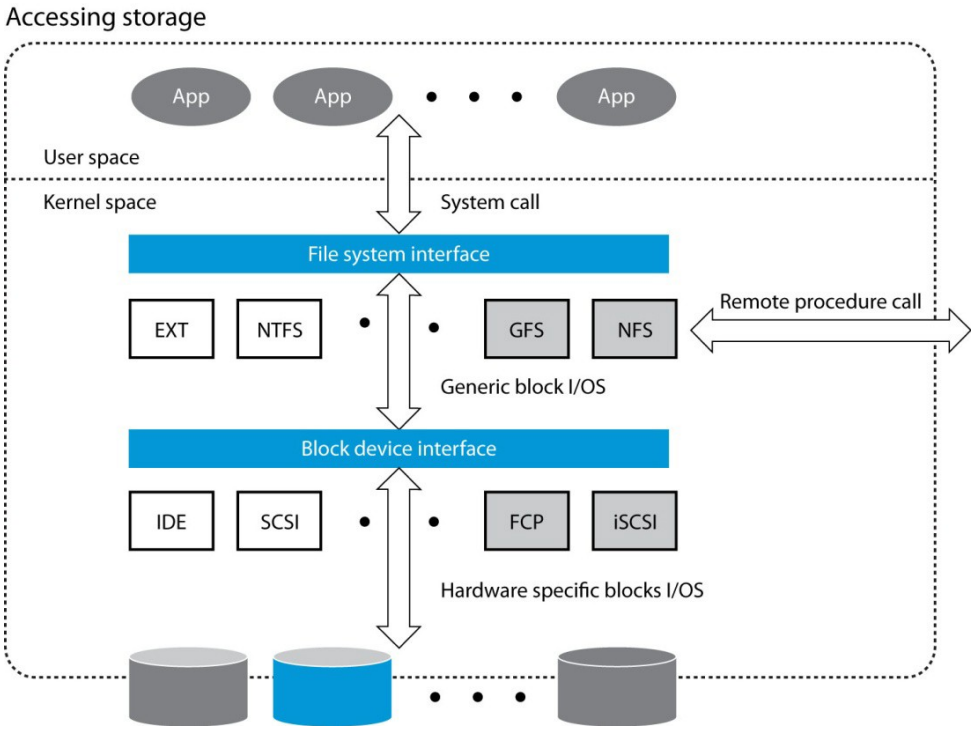

Storage area networks (Non-IP)

SANs are specialized, dedicated high-speed networks joining servers and storage. It can include disks, disk arrays, tapes, etc. Storage (data store) is separated from the servers and the I/O is offloaded from the server's processors. SAN technology provides high capacity, HA, high scalability, ease of configuration, and ease of reconfiguration. Fibre Channel is the de-facto standard SAN architecture, although other network standards can be used. SANs differ from NAS in that the Fibre Channel network does not carry any other traffic except data storage traffic.

Supported SAN topologies:

Point-to-point

- The simplest topology for very limited connectivity needs
- Guarantees "in order" delivery and full bandwidth access
- Can handle any multi-path connectivity to a set of disks because there are no other elements in the communication path

Arbitrated loop (not recommended for new designs)

- Designed to scale to a limited number of nodes (up to 127)
- Low cost (no interconnecting switch devices required)
- Arbitration protocol is designed to manage media sharing across nodes. This may be disruptive when a node gets added/removed from the loop and the loop initialization protocol starts
- An arbitrating hub can be used instead of a distributed protocol

Switched fabric

- A switching element is added to the nodes to allow interconnections via point-to-point links
- Extends number of devices (potentially thousands) and greater distances can be achieved
- A scalable, robust, and reliable architecture, but the cost of the interconnection devices can add up

Figure 23 Typical SAN networks

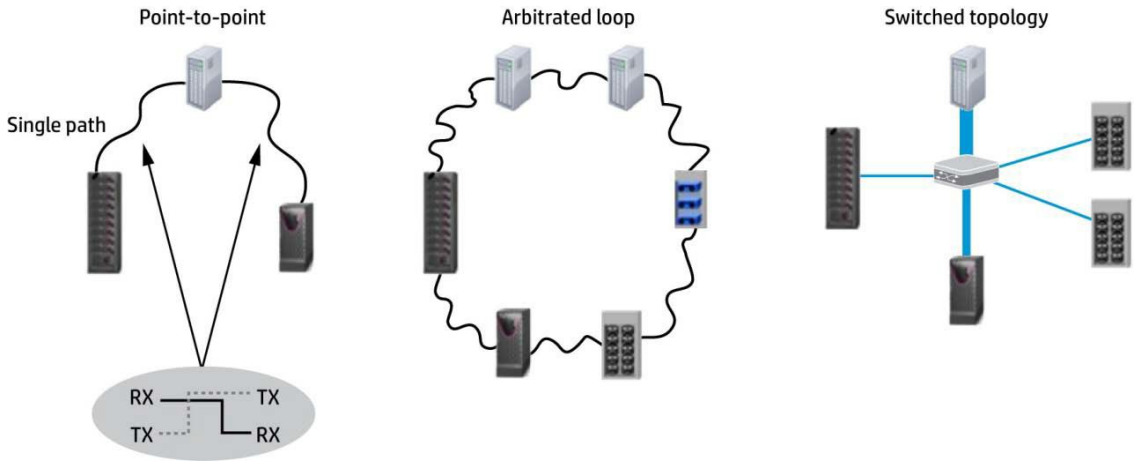

Storage area networks (IP & Ethernet-enabled)

SANs deliver a substantial number of advantages to the data center. However, the biggest disadvantages have been the need for additional Fibre Channel interface cards, cables, additional devices to manage, and distance limitations. IP SAN technologies are now ported onto the IP layer. However, the SAN must have a reliable transport to enable the communication of SCSI commands to and from SCSI controllers and peripherals. There are two IP-enabled SAN standards in use today.

Both iSCSI and FCIP provide the following benefits in common with the legacy Fibre Channel standard:

- Enable storage consolidation
- Data sharing
- Non-disruptive scalability for growth
- Provide SAN extension over MAN/WAN networks
- Improved backup and recovery
- Good performance (relative to the network structure)
- Data integrity
- Disaster tolerance
- Ease of data migration
- Lower cost of ownership

A substantial number of IT managers deployed iSCSI and FCIP with early acceptance of the performance without changing any of their network structure. As the adoption of the technology into their data centers became more prevalent, additional IP SAN systems were deployed and network performance dropped. Network managers have had to adapt their network designs to compensate for the convergence of storage traffic onto their networks, just as they had to do with voice traffic.

The issue now becomes the network that storage traffic rides on and the transport protocol. A data center or enterprise IT architect has to identify whether they have enough bandwidth, switch packet

forwarding capacity, and QoS configured to support acceptable storage performance.

Network architects have to identify and plan for the following in support of IP SANs:

- The need for true non-blocking Ethernet interconnects modules
- Higher performance routers
- Oversubscription ratios on routers and switches (WAN & LAN)
- Physical or logical Ethernet separation (Will this affect the ROI?)
- QoS requirements (Now a true need for classes of traffic)
- VLAN separation
- Flow control
- Redundant interconnects and paths (Impact of STP?)
- Jumbo frames need to be supported end-to-end

Storage traffic on the LAN has provided the following:

Benefits

- Ubiquitous technology
- Operates on known platforms (IP and Ethernet)
- Standards-based solutions
- Commodity pricing driven by large vendor community and customer demand
- Operational efficiency
- Ethernet is universally deployed
- Technical staff that know how to manage IP and Ethernet
- Reliable transport
- Wide Area Networking
- Removal of distance limitations
- Flexible deployment designs
- Aids in disaster recovery initiatives and data replication
- Strong long-term viability as Ethernet speeds increase to 10GbE, 40GbE, and 100GbE

Negatives

- TCP slow start impacts I/O latency and throughput
- The sliding window algorithm can cause substantial re-transmissions during heavy congestion periods
- Data re-transmissions can affect applications dependent on timely execution of data requests
- Difficult, expensive, and inefficient integration with other types of IP SANs and legacy Fibre Channel systems

iSCSI

One of the reasons that the Fibre Channel protocol became prevalent is that it is a very lightweight, high-performance protocol that maps locally to the SCSI initiators and targets. But Fibre Channel protocols are not routable and can be used only within a relatively limited area, such as within a

single data center.

The iSCSI protocol sought to improve that limitation by moving the SCSI packets along a typical Ethernet network utilizing TCP/IP. The iSCSI protocol serves the same purpose as the Fibre Channel in building SANs, but iSCSI runs over the existing Ethernet infrastructure, and avoids the cost, additional complexity, and compatibility issues associated with Fibre Channel SANs.

An iSCSI SAN is typically comprised of:

- Software or hardware initiators on the host server connected to an Ethernet network
- Storage resources (targets)
- The iSCSI stack at both ends of the path is used to encapsulate SCSI block commands into Ethernet packets for transmission over IP networks. Initiators include both software- and hardware- based initiators incorporated on host bus adapters (HBAs) and NICs

Figure 24 iSCSI is SCSI over TCP/IP

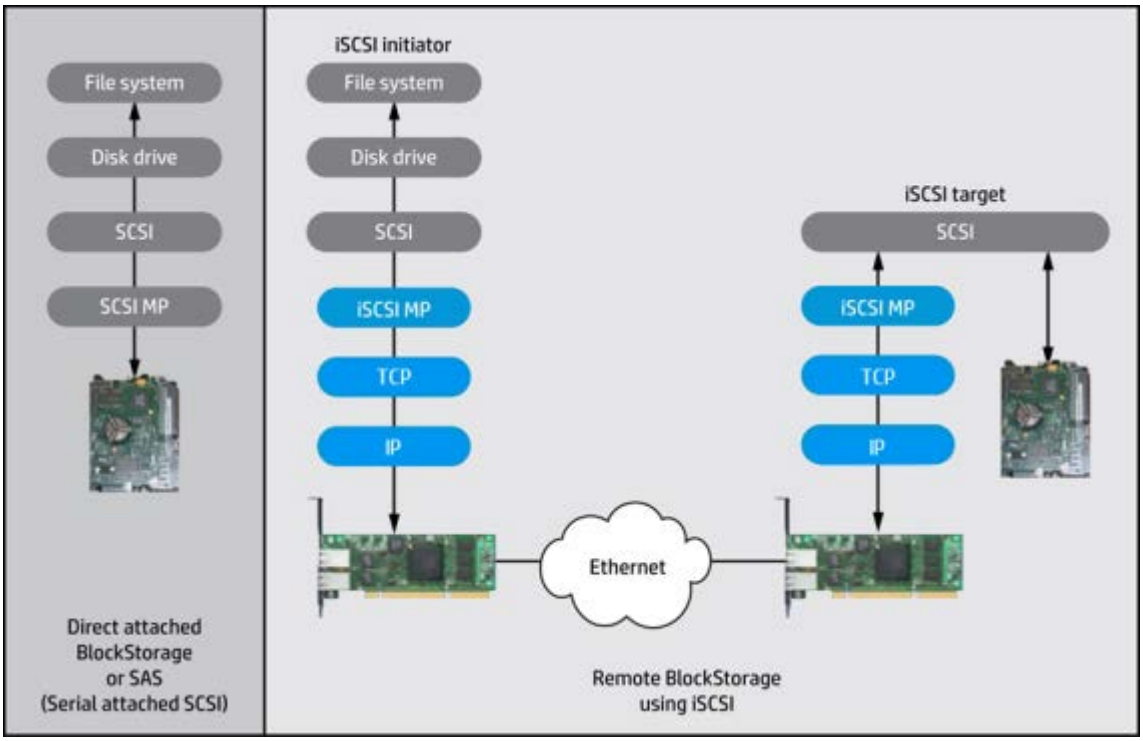

It is possible to create a lossless Ethernet network using older IEEE 802.3x Ethernet mechanisms. This is very useful in focused configurations like HP Left Hand storage clusters. However, if the network is carrying multiple traffic classes, which many data centers do, the existing mechanisms can cause QoS issues, which limit the ability to scale a network and impact performance.

The newer generations of iSCSI technology solve these issues:

- High-performance adapters are now available which fully offload the protocol management to a hardware-based iSCSI HBA or NIC. This is called full iSCSI offload
- Using 10GbE networks and 10GbE NICs with iSCSI SANs generates performance comparable to a Fibre Channel SAN operating at 8 GbE
- New centralized iSCSI boot management tools provide mechanisms for greater scalability when deploying large numbers of servers

With 10GbEbased iSCSI products, iSCSI becomes a more viable solution for converged networks for small-medium businesses as well as enterprise data centers. For customers with iSCSI-based storage targets and initiators, iSCSI can be incorporated today with their existing Ethernet infrastructure. iSCSI does not require the new Data Center Bridging (DCB) infrastructure. But if present, iSCSI will benefit from the QoS and bandwidth management offered by DCB. Because iSCSI administration requires the same skill set as a TCP/IP network, using an iSCSI network minimizes the SAN administration skills required, making iSCSI a good choice for Greenfield deployments or when there are limited IT teams and budgets. iSCSI can be enhanced even more with HP IRF technology, which is covered elsewhere in this document. iSCSI will continue to be an integral part in many converged network infrastructure designs with the enhancements that HP brings to Ethernet switching technology.

The future of SANs

New technologies and standardization have made it possible to converge Fibre Channel traffic and network traffic at the server edge. It will take time and a multi-step process to converge further into the core of data center networks. This highlights the need for planning for the future, to enhance the networks supporting the new demands. iSCSI is greatly improving, especially with the increase in network connectivity speeds at 10GbE. However, most storage designs in data centers are based on Fibre Channel technology, and as a result, FCoE and DCB are critical for converging large scale data centers.

The complexity of guaranteeing low-latency/high-performance delivery for each type of data traffic, and specifically storage traffic, continues to be the critical component of data center network design. For example, applications using NAS can require network performance that is measured in real-time.

The server to network edge

By implementing convergence first at the server edge, data center architects will be able to take advantage of the standards that are in place, while avoiding the risk of implementing technology that will have to be replaced or modified as standards solidify.

The server-to-network edge is also becoming increasingly complex due to the sprawl of VMs, which is covered in greater detail in another section of this document. However, the management of VMs and their demand for I/O resources drive the need for DCB.

The HP approach to technologies at the server-to-network edge is based on using industry standards. This helps ensure that new technologies will work within existing customer environments and organizational roles, yet will also preserve customer choice. The HP goal for customers is to enable a simple migration to advanced technologies at the server-to-network edge without requiring an entire overhaul strategy for the data center.

The server-to-network edge refers to the connection points between servers and the first layer of both LAN and SAN switches. The most common networks used in enterprise data centers are Ethernet for LANs and Fibre Channel for SANs.

Different topologies are one of the reasons that HP is focusing on the server-to-network edge. Administrators should be allowed to maintain similar processes, procedures, data center structures, and organizational governance roles while reaping the benefits of reduced infrastructure at the server-to-network edge.

Traditionally, the server edge is the most physically complex part of data center networks; the server

edge has the most connections, switches, and ports (especially in environments using only rack-based servers). For most enterprise data centers using Fibre Channel SANs and Ethernet LANs, each network type has its own requirements:

- Unique network adapters for servers
- Unique switches
- NMSs designed for each network
- Different organizations to manage the networks
- Unique security requirements
- Different topologies
- Different performance requirements

Each of these differences add complexity and cost to the data center, especially for enterprise data centers that have large installed Fibre Channel and Ethernet networks.

HP VC is redefining SAN Networking infrastructure with Direct-Attach Fibre Channel for 3PAR Storage Systems – using Flat SAN technology. This reduces the traditionally complex multi-tier SAN infrastructure and flattening it. HP VC FlexFabric Modules can be directly connected using a Fibre cable to a3PAR SAN Storage System which reduces latency and the cost of the SAN Fabric. One additional benefit is faster provisioning with the elimination of the requirement to configure zoning.

Some data center architects use non-Fibre Channel storage. In this case, iSCSI solves these challenges without requiring a new Ethernet infrastructure. However, iSCSI technology is still evolving and has experienced a slow rate of adoption, especially in enterprise data centers. The question is whether to enhance the converged network infrastructure with iSCSI or utilize DCB. Either way, the demand for bandwidth in data centers and their encompassing racks continues to grow with the widespread adoption of server virtualization.

Other topics to consider

Switches with a non-blocking architecture should be used. In most commonly used switches, the backplane capacity may be less than the sum of the capacity of all the ports.

- A dedicated high-speed network with either a physical or logical separation is recommended. Flow control in the Ethernet protocol is very minimal, and several standards are emerging to tackle the complexities of storage traffic
- Redundancy is also a critical factor in designing a network for storage traffic. There are several pieces to keep in mind when designing a network with redundant switches:
    - Multiple paths add fault-tolerance and performance since more than one physical path exists between the computer system and storage devices through the buses, controllers, switches, and bridge devices connecting them
    - An IRF-based network extends the control plane across multiple active switches, enabling interconnected switches to be managed as a single common fabric with one IP address; this increases network resilience, performance, and availability while simultaneously reducing operational complexity
- The support for jumbo frames is the biggest imposition of iSCSI on Ethernet switches due to the characteristic block size of storage volumes, which at 2KB is larger than the nominal Ethernet frame size of 1.5KB

Section Summary

The fundamental nature of data center computing is rapidly changing. Today's data center networks must evolve to support on-demand, virtualized IT environments. HP delivers the foundation for the data center of the future, today, by providing a unified, virtualization-optimized infrastructure. HP Networking solutions enable the following:

- Breakthrough cost reductions by converging and consolidating server, storage, and network connectivity onto a common fabric with a flatter topology and fewer interconnect modules
- Predictable performance and low latency for bandwidth-intensive server-to-server communications
- Improved business agility, faster time to service, and higher resource utilization by dynamically scaling capacity and provisioning connections to meet virtualized application demands
- Removal of costly, time-consuming, and error-prone change management processes
- Modular, scalable, industry standards-based platforms and multi-site, multi-vendor management tools to connect and manage thousands of physical and virtual resources from a single-pane-of-glass

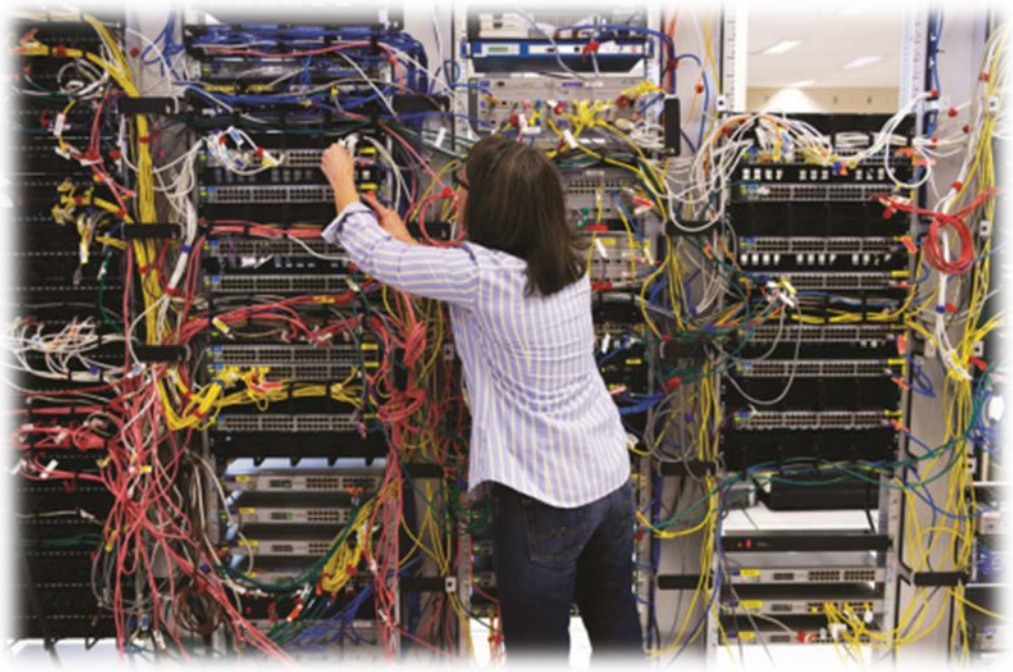

Data Center Network Design - HP FlexFabric Reference Architecture (FFRA)

This section introduces HP's reference architectures for the virtualized and converged network infrastructure. The FlexFabric Reference Architecture is an overall discussion of the use of current HP technologies to fulfill the vision of virtualization.

The HP FFRA specifically refers to HP Networking products; however, the overall HP data center solutions can fulfill the following visions:

- **Best-of-breed servers:**

  Meeting the demands of virtualized and non-virtualized data center deployments. For example, HP BladeSystem meets the demands of the converged network infrastructure through enhancements in processing density, VC LAN and SAN connectivity with its reduced cabling, and enhanced ability to function in flatter Layer 2 network designs to meet the demands of server virtualization
- **Best-of-breed technology partnerships:**

  HP has a strong product and interoperability alliance with VMware and Microsoft. As your virtualization deployments continue to grow, the alliance with these partners will continue to add performance, flexibility, and simplified management of VMs and their attached networks
- **Best-of-breed Layer 2/3 networking components:**

  To meet the demands of virtualization and converged I/O, all of which are enhanced by HP's standards-driven architecture and IRF
- **Utilization of best-of-breed network security products:**

  This not only secures the network, but is also designed with the security challenges of the virtualized network in mind. HP TippingPoint IPS product suite can be deployed at the core, or at other various locations within the data center. There is also the option to operate them on virtual appliances using vController and vIPS which extend network security to the virtualized

network. Security is also controlled with discovery and management of VMs using the Security Management System (SMS). Network security can be complex and involve a great deal of management overhead. HP security products specifically address this problem and simplify network security. For example, HP TippingPoint products work in unison with HP IMC, to allow a single-pane-of-glass management system to manage multiple security controllers throughout the data center

- **Best-of-breed management tools:**

  Orchestration is highlighted in HP's network and server hardware in the data center through the use of HP IMC. The solution cohesively integrates fault management, element configuration, and network monitoring from a central vantage point. With support for third-party devices, IMC enables network administrators to centrally manage all network elements with a variety of automated tasks, including discovery, categorization, baseline configurations, and software images. IMC also provides configuration comparison tools, version tracking, change alerts, and more

  HP Insight Management is a complete suite of server lifecycle management capabilities that can flexibly operate from embedded on-system utilities, your preferred CMS, and now even from the cloud. Managing ProLiant servers with Insight Management delivers increased efficiency and precise control of your server infrastructure resources
- **Utilization of best-of-breed storage and storage networking technologies:**

  From HP StorageWorks division, with enhancements to FCoE, iSCSI, and the future with DCB. All of these are enabled through the use of VC, DCB-enabled ToR, EoR, and core switches for aggregation
- **Utilization of best-of-breed power and cooling:**

  HP Data Center Smart Grid leads the industry in lower power consumption and device monitoring

Evolving network designs

Access layer switching

In current data center deployments, there are two access layer switching deployment models: ToR and EoR.

Top-of-rack

As the name implies, in a ToR placement, servers within a rack connect to an access layer switch generally placed at the top of the rack, as in shown below ("ToR placement").

Figure 25 ToR placement

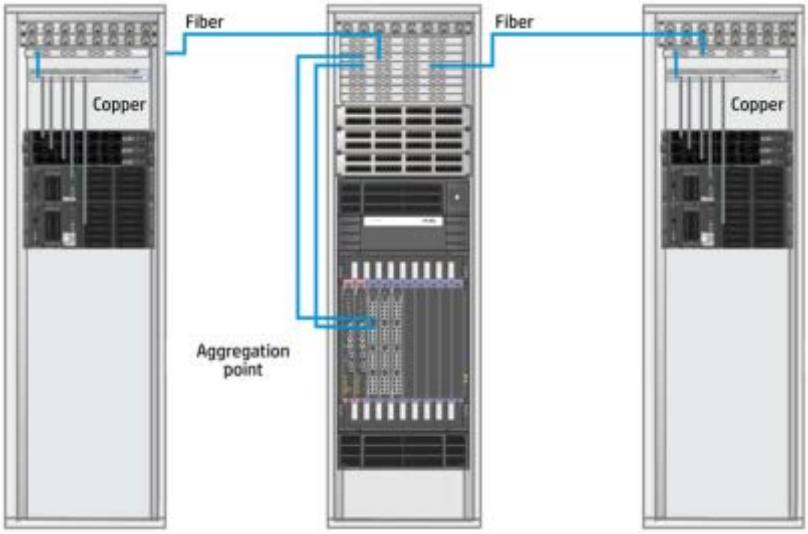

In a ToR design, servers connect to an access switch via copper or fiber within the rack, while the access switch connects to other consolidation or backbone switches within the data center via multi-mode fiber (MMF). Often, the data center backbone consists of high-capacity aggregation or distribution switches with Layer 2/3 capabilities. Ethernet copper is generally restricted to relatively short runs inside the rack, while connections outside the rack can be made with smaller form factor multi-mode fiber. This reduces the overhead weight, as well as having the ability to be connected to different capacity interfaces to support the bandwidth requirements of the rack.

Shorter copper Ethernet runs in the rack allow for multiple choices on cable types which can support various required speeds dictated by the systems in the rack. In many cases, server to switch connections can be up to 10GbE connections with support for integrated I/O. Longer runs from the rack to core usually utilize 10GbE MMF, however more bandwidth could be provided by using 40GbE or 100GbE ports. HP currently offers 40GbE ports on some HP switches, and we will be expanding 40GbE coverage to more switches in the future. HP will also be offering 100GbE ports for even higher bandwidth links in the future.

Advantages

- **Issue isolation:**

  Each rack, or group of racks using IRF, with a ToR configuration is treated as an independent module. Any issues or outages that occur with an access layer switch typically affect only the servers within that rack which are connected to that access switch
- **Traffic isolation:**

  Because each access switch is a home run back to a backbone aggregation/core switch, traffic can be monitored and isolated to an identified switch port within a specific rack
- **Physical disasters:**

  A potential physical disaster affecting cabling at the ToR will have adverse effects on the ToR rather than an entire row

Disadvantages

- **Number of switches:**

  Each rack adds to the number of switches that need to be managed as independent entities

- **Capital outlay:**

  This design also takes into consideration that every rack may not be populated with access switches when the data center is built. A general rule of thumb is to only install equipment in a rack when there is a demand for it. The ToR placement allows the data center manager to populate the rack as required with little planning except for available port capacity on the aggregation switch
- **Additional rack to rack hops:**

  ToR switches can introduce

End-of-row

In an EoR placement, a rack containing the switching equipment is typically placed at either end of a row of cabinets or racks. Bundles of cabling provide the connectivity from each server rack to the switching equipment rack. Servers are usually connected to a patch panel inside each server rack. The copper or fiber bundles are home run to another patch panel in the rack containing the access switches. The switches are then connected to the patch panel via patch cables. EoR switches are typically connected back to the core with a series of fiber patch cables.

EoR does not imply that the network racks have to be placed at the end of the row. There are designs where network switch racks are placed together in the middle of cabinet/rack rows. Placing the switch rack in the middle of a row limits the length of cables required to connect the furthest server racks to the nearest network rack.

Unlike the ToR model, where each rack is treated as an independent module, in the EoR placement model, each row is treated as an independent module.

Advantages

- **Number of switches:**

  There are fewer switches to manage as separate entities
- **Network pre-planning:**

  A data center can be pre-planned and cabled without deploying a substantial number of switches into individual racks
- **Capital outlay:**

  Fewer switches means less capital outlay
- **Fewer rack to rack hops**

Disadvantages

- **Issue Isolation:**

  A potential physical disaster affecting cabling or the EoR aggregation switch can have adverse effects on not just one rack of servers, but an entire row
- **Traffic isolation:**

  Because each EoR switch is a home run back to a backbone aggregation/core switch, traffic will be for an entire row rather than just within a specific rack
- **Cabling:**

  The magnitude of cabling home runs back to the EoR switch makes for substantial cabling bundles

Figure 26 EoR

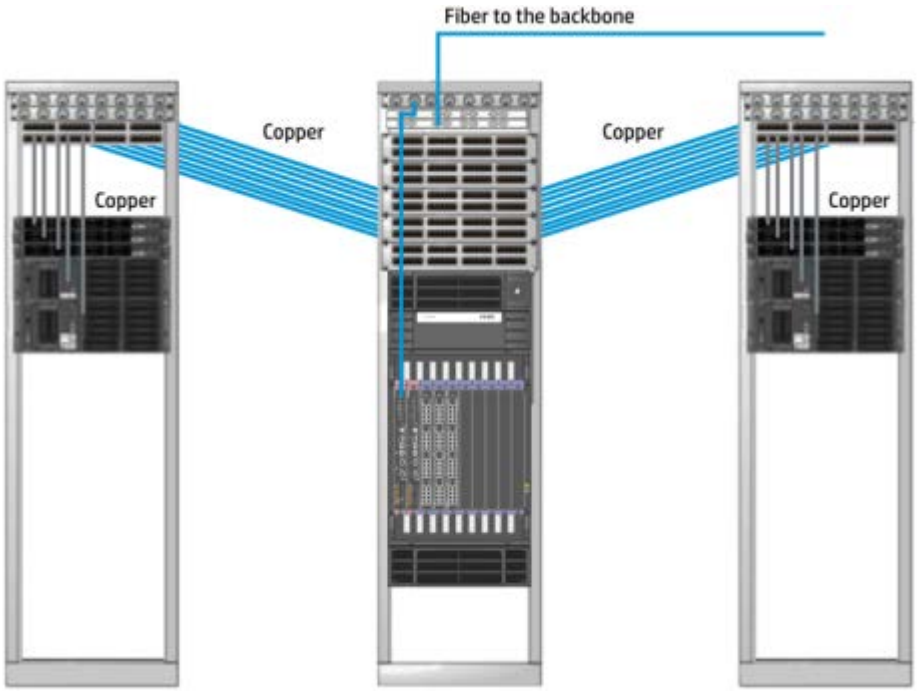

Cabling trends—fiber vs. copper

Early client/server-based data center networks primarily addressed static applications and email. Data centers today support new applications and traffic patterns that are dynamic and demand new approaches for real-time application access, as well as guaranteed performance, low latency, and any-to-any communication patterns. The ever-increasing demand for bandwidth and throughput, as well as the promise of converged network infrastructure, justifies the need for 10GbE or higher speeds.

Converging the network infrastructure today

Data center administrators can address complexity at the server-to-network edge by consolidating server I/O in various ways:

- Combining multiple lower-bandwidth connections into a single, higher-bandwidth connection
- Converging different network types (LAN and SAN) to reduce the amount of physical infrastructure required at the server edge

Administrators can already use HP VC Flex-10 technology to allocate one 10GbE server port into 4 individual Ethernet connections or functions. Each connection can be adjusted dynamically to meet the demands of the workload. The VC FlexFabric Module takes this a step further by allowing one connection per port to be used as either a Fibre Channel or iSCSI function.

Advantages

- Reduces management requirements
- Reduces the amount of physical infrastructure (number of NICs, HBAs and number of managed switch modules)

- Simplifies the amount of cross-connect in the network
- Reduces power and operational costs

The goals of a converged network are:

- To simplify and flatten the typical network and storage topology
- To improve QoS
- To reduce host/server adapters
- To reduce cabling requirements
- To reduce switch ports and costs
- To enable simpler, more centralized management

Blade server one-tier design

This type of network deployment signifies the current pinnacle of network virtualization. Server blades allow for substantial compute density per rack, row, and data center. HP has optimized the BladeSystem server portfolio to support the vision and reality of virtualization. The network design optimizes the reality of high performance networking with simplicity. It allows flexibility in VM networking and converged I/O options. This approach is at the forefront of network design for virtualization, since it utilizes both the IRF framework and VC. It operates well while allowing for seamless management and troubleshooting of VMs.

Objectives

When determining if this deployment type is right for you, keep in mind the overall objectives:

- **Layer 2 Flexibility:**

  Because the overall goal is to support VMs, a flattened Layer 2 network allows VMs to be moved without the need for IP address changes. Keep in mind that a flattened Layer 2 design supports long range vMotion using methods discussed later in this document (see the HP FFI section)
- **Reduced management complexity:**

  Flattening the network with large core switches reduces the number of devices that have to be provisioned and monitored when setting up VLANs for VMs
- **Zero need for STP/RSTP:**

  Coupled with IRF, this network design is loop-free, thus the need for STP/RSTP is eliminated
- **Frame forwarding and packet forwarding:**

  Performance is optimized in this flat IRF framework. VMs communicate between each other, with their converged I/O resources served by an optimal Layer 2 network, VLANs, and LACP
- **VLAN management:**

  A flatter Layer 2 network is ideal for VMs. It provides support for network extension across the data center and between sister data centers, the framework for VM mobility
- **Centralized security:**

  Layer 2 networks allow the IP-based IPS devices to be placed into the core layer, where VLANs

are aggregated. Spanning of VLANs across a flattened Layer 2 data center network provides greater flexibility in rack mount server connectivity in disparate rack locations without the need for the creation of complex IP-based access control lists (ACLs)

Where is a blade server one-tier network best suited?

Anywhere cost, performance, flexibility, and real estate are at a premium!

The flatter the network, the more efficient it is. Flatter networks provide greater performance than two-tier/three-tier models.

Architecture

Typical blade server 1-tier deployments will usually consist of two to four core switches utilizing IRF, which then connect to blade-servers housed in C-Class enclosures. These types of deployments are able to extend VLANs across the entire data center and are optimized for virtualized environments where VMs may be moving from rack to rack or data center to data center.

Rack and logical diagrams

Figure 27 Blade server one-tier architecture rack view

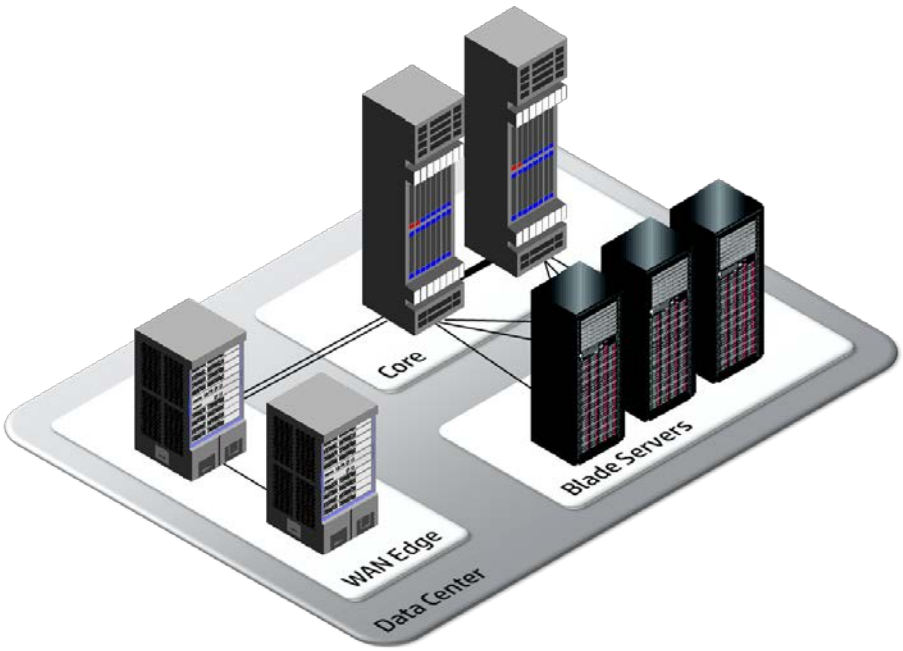

Figure 28 Blade server 1-tier logical

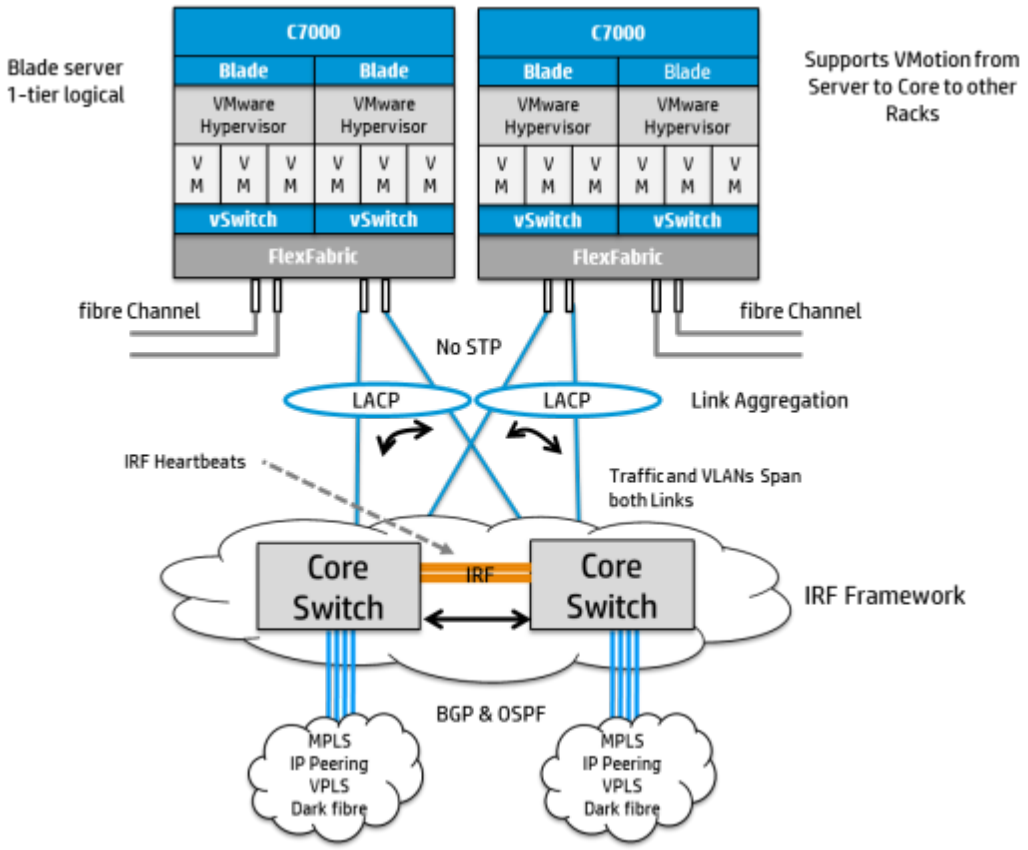

Logical layers of this design

Virtualization

The network edge physically starts at the Virtual Connect FlexFabric modules, but really it starts internally in the VMs and configured virtual switches. HP 12500 Switches are able to communicate natively with VMware virtual switches, allowing updates of ARP tables between physical and virtual switches. The VMs and virtual switches can provision VLANs, which in turn interoperate with the IRF fabric, allowing seamless VM movement with vMotion and high performance frame forwarding. All of the switches in an IRF framework support 4,096 VLANs, which provides substantial network flexibility to configure a substantial number of VM groups onto their own virtual network segments.

Combining LACP and IRF in this design provides high-speed link aggregation with re-convergence times at sub-50 ms in the event of a link failure. It also allows links to be aggregated and utilized for higher bandwidth from the converged network adapters across all switches to forward traffic.

This design also supports the dynamic storage of worldwide names. This feature allows a FlexConnect module to be configured with the required IP and VLAN information once. In the event of a device failure, the replaced FlexConnect device will gracefully provision itself with the original configuration. This feature provides advantages when a specific component can have an impact on a substantial number of VMs.

This network design supports long-range vMotion connectivity, enabling clustering or synchronizing VMs between disparate locations. The switches will allow VPLS connectivity between switches in geographically separated data center locations. Be aware that long-range fiber segments using MPLS can add latency that limits the ability for VMs to use converged I/O resources between data center

locations (see the HP FFI section). Long distance WAN networks generally address normal server communications and disaster recovery efforts.

Security

Although not depicted in the diagrams, HP TippingPoint IPS devices can be attached to the core switches. Because VLANs can be configured to span all the way to the core switches, all network security can be centralized in a fully redundant N+1 configuration by placing an IPS blade in each of the core switches.

Orchestration

This network configuration supports the full feature complement of HP IMC from the core switches all the way down the network to the physical servers themselves. This provides a single-pane-of-glass management platform to support the servers, VMs, virtual switches, IRF frameworks, IP routing, and security devices.

The HP IMC VAN module enables a more agile model by eliminating unnecessary steps in virtualizing business environments. Time to VM deployment is accelerated through the upfront definition of policies for VM connectivity and provisioning automation. These policies allow for rapid deployment and follow the workload if it is moved, paused, and/or resumed.

Support for the eco-centric data center

HP servers and switches lead the industry in limiting power consumption without compromise on performance. Every component in this design from the rack up supports a "sea of sensors" which can be monitored using HP System Insight Manager (SIM).

Simplified two-tier design (ToR)

"If you have switches with adequate capacity and you've got the right ratio of input ports to trunks, you don't need the aggregation layer, which provides marginal value at best. You can avoid a lot of complexity, cost, and extra heat while simplifying design and troubleshooting."

Joe Skorupa, Research VP, Gartner, Inc.

Objectives

When determining if this deployment type is appropriate, keep in mind the overall objectives:

- **Layer 2 flexibility:**

  Flattening the network with Layer 2 better suits VMs residing on servers which require the

ability to communicate with other VMs that may be in the same rack or row, within the same data center facility, or to geographically separate data centers. Since the overall goal is to support VMs, a flattened Layer 2 network allows VMs to be moved without the need for IP address changes. Keep in mind that a flattened Layer 2 design supports long-range vMotion using VPLS

- **Reduced management complexity:**

  Flattening the network with large aggregation switches reduces the number of devices that have to be provisioned and monitored when setting up VLANs for VMs

- **Less dependence on STP/RSTP:**

  HP recommends switches that support IRF, but many data centers already have switches that may not support IRF. In this case, using IRF in the core will certainly help optimize and simplify the network. In either case, fewer Ethernet switches means less link state management between servers, ToR switches, and aggregation switches within the data center network

- **VLAN management:**

  A flatter Layer 2 network is ideal for VMs. It provides the framework for VM mobility extension across the data center and between geographically separate data centers

- **Centralized security:**

  Layer 2 networks allow the IP-based IPS devices to be connected to the core switches, where VLANs are aggregated. This means fewer network security devices to provision and monitor. It allows for less complexity in security policies because they require the provisioning of fewer devices. Spanning of VLANs across a flattened Layer 2 data center network provides greater flexibility in server connectivity in disparate rack locations because it removes the need for complex IP-based ACLs

Where is 2-tier best suited?

When utilizing a mix of rack and blade servers.

The flatter the network, the more efficient it is. Flatter networks provide greater performance than three-tier/legacy models.

Architecture

Typical 2-tier ToR deployments will utilize two to four core switches utilizing IRF, which then in turn connect to various ToR switches. These ToR switches, which may or may not be utilizing IRF, will then connect to servers within the same rack.

Layer 2 2-tier ToR designs extend VLANs across the entire data center and are optimized for virtualized environments where VMs may be moving from rack to rack or data center to data center.

Layer 3 2-tier ToR designs can isolate VLANs within each rack, and still provide for good scalability and expansion capabilities.

Rack and logical diagrams

Figure 29 Simplified two-tier architecture rack view

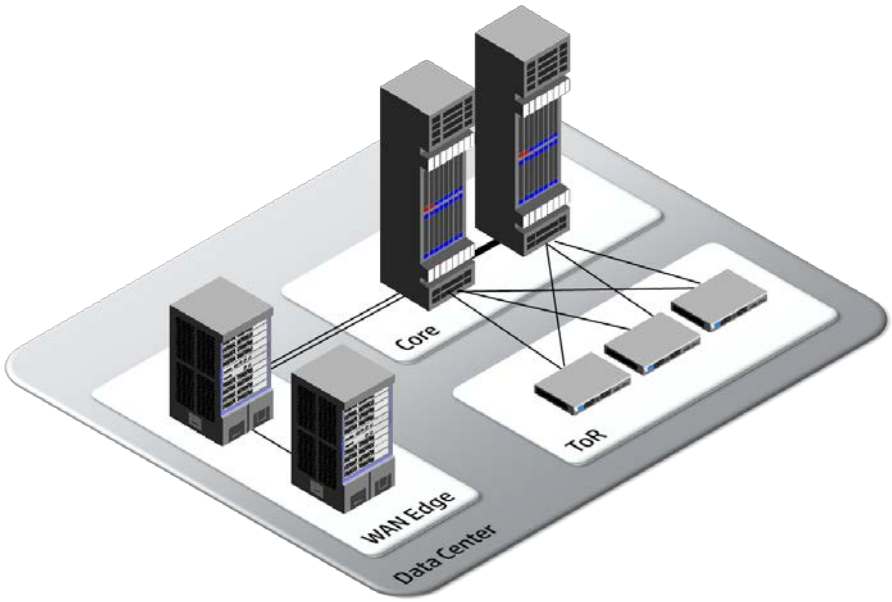

Figure 30 Simplified two-tier logical

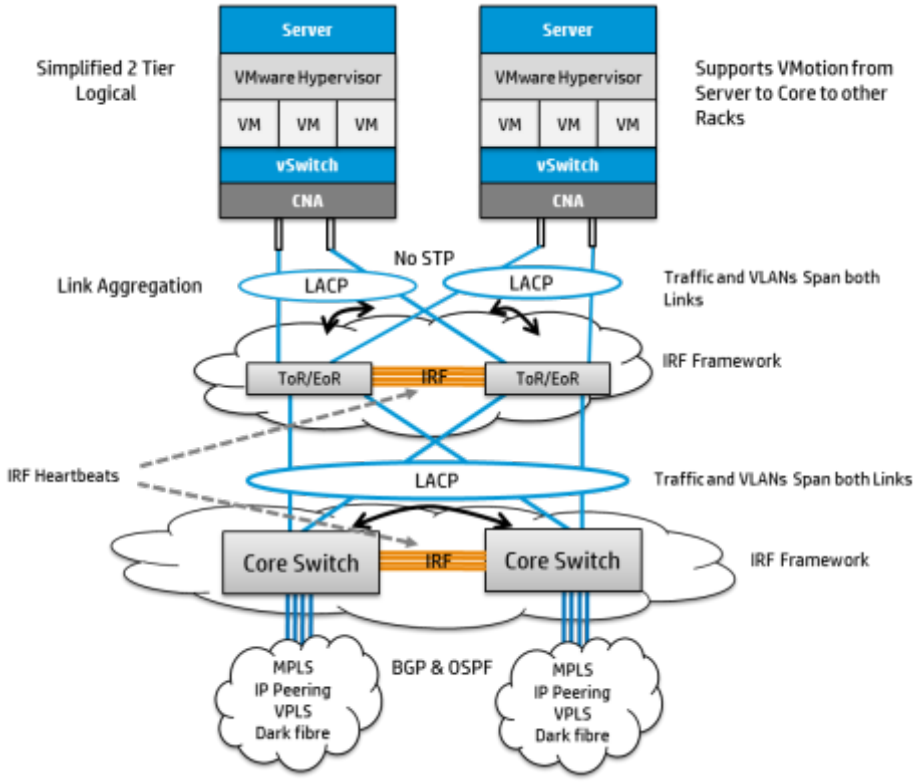

Logical layers of this design

Virtualization

The network edge physically starts at the converged network adapters, but really it starts internally in the VMs and configured virtual switches. The VMs and virtual switches can provision VLANs, which in turn interoperate with the IRF fabric, allowing seamless VM movement with vMotion and high performance frame forwarding. All the HP switches in an IRF framework support up to 4,096 VLANs, which provides flexibility to configure a substantial number of VM groups into their own virtual network segments.

Combining LACP and IRF in this design provides high-speed link aggregation with re-convergence times at sub-50 ms in the event of a link failure. It also allows links to be aggregated and utilized for higher bandwidth from the converged network adapters across all switches to forward traffic.

The network design does not require the STP/RSTP protocol. But in the event that existing third-party switches are already deployed in the racks, STP protocols are supported through standards interoperability.

This network design supports long range vMotion connectivity, enabling clustering or synchronizing VMs between disparate locations. The switches will also allow VPLS connectivity between switches in geographically separated data center locations. Be aware that long-range fiber segments using MPLS can add latency that limits the ability for VMs to use converged I/O resources between data center locations (see the HP FFI section). Long distance WAN networks generally address normal server communications and disaster recovery efforts.

Security

Although not depicted in the diagrams, HP TippingPoint IPS devices can be attached into various layers in the network. Since VLANs can be configured to span all the way to the core switches, all network security could even be centralized by placing an IPS device into the core layer.

Orchestration

This network configuration supports the full features of HP IMC from the core switches to the physical servers. This provides a single-pane-of-glass management platform to support the servers, VMs, vSwitches, IRF frameworks, IP routing, and security devices.

The HP IMC VAN module enables a more agile model by eliminating unnecessary steps in virtualizing business environments. Time to VM deployment is accelerated through the upfront definition of policies for VM connectivity and provisioning automation. These policies allow for rapid deployment and follow the workload if it is moved, paused, and/or resumed.

Support for the eco-centric data center

HP servers and switches lead the industry in limiting power consumption without compromise on performance. Every component in this design from the rack up supports a “sea of sensors” which can be monitored using HP SIM.

Three-tier design

This type of network deployment can be deployed as a Greenfield or it fits well into data center networks where added bandwidth, 10GbE port capacity, and simplified management are paramount. It also helps ensure the interoperability of legacy deployed EoR and ToR. Although the depicted design focuses on HP switches, IRF permits third-party switches to be inserted at any level and will interoperate using standards-based networking.

HP has optimized its network technologies to support this type of deployment design.

Objectives

When determining if this deployment type is appropriate, keep in mind the overall objectives:

- **Layer 2 flexibility:**

  Because the overall goal is to support VMs, a flattened Layer 2 network allows VMs to be moved without the need for IP address changes. This design also provides the flexibility to leverage Layer 3 separately where needed to mitigate broadcast domains and control traffic between segments. Also, a flattened Layer 2 design supports long range vMotion using VPLS
- **Optimized management complexity:**

  Leveraging IRF in the various layers of the network allows what was typically a complex management scheme to be implemented in a 3-tier framework. IRF allows each layer to be managed as a single entity. This provides a clear depiction of the devices and traffic in the data center
- **Less dependence on STP/RSTP:**

  Ideally, all switches should support IRF. However, existing switches may not support IRF. In this case, using IRF in the core will certainly help optimize and simplify the network. In either case, fewer Ethernet switches means less link state management between servers, ToR switches, and aggregation switches within the data center network
- **VLAN Management:**

  A flatter Layer 2 network is ideal for VMs. It provides support for network extension across the data center and between geographically separate centers, the framework for VM mobility
- **De-centralized security:**

  The 3-tier model allows security to be distributed throughout the network allowing for a granular view of the individual segments within the data center. With de-centralized management, one can deploy many smaller HP TippingPoint IPS systems in the distribution layers and configure different security profiles depending upon the application requirements

Why 3 Tier?

When utilizing a mix of rack and blade servers and when needing to incorporate into an existing infrastructure.

Architecture

Typical 3-tier legacy deployments are similar to the 2-tier deployments however these solutions will have an added aggregation layer positioned in between the core and ToR layer. This aggregation layer will generally consist of chassis based switches utilizing IRF.

In these types of deployments, Layer 3 routing typically occurs from the aggregation layer to the core and WAN. This way, Layer 2 domains can be isolated to each aggregation layer device and the ToR switches that connect to it.

Rack and logical diagrams

Figure 31 Legacy three-tier architecture rack view

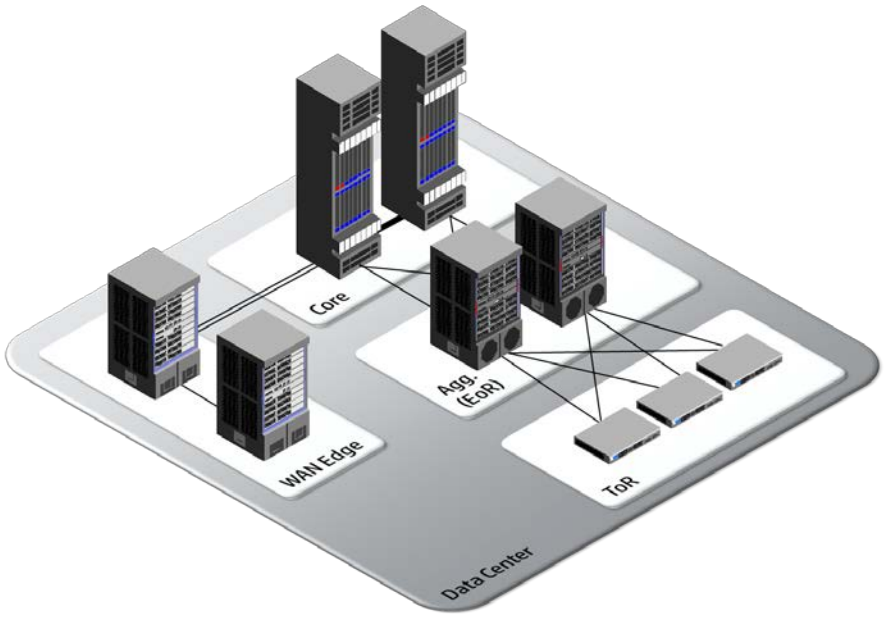

Figure 32 Three-tier logical

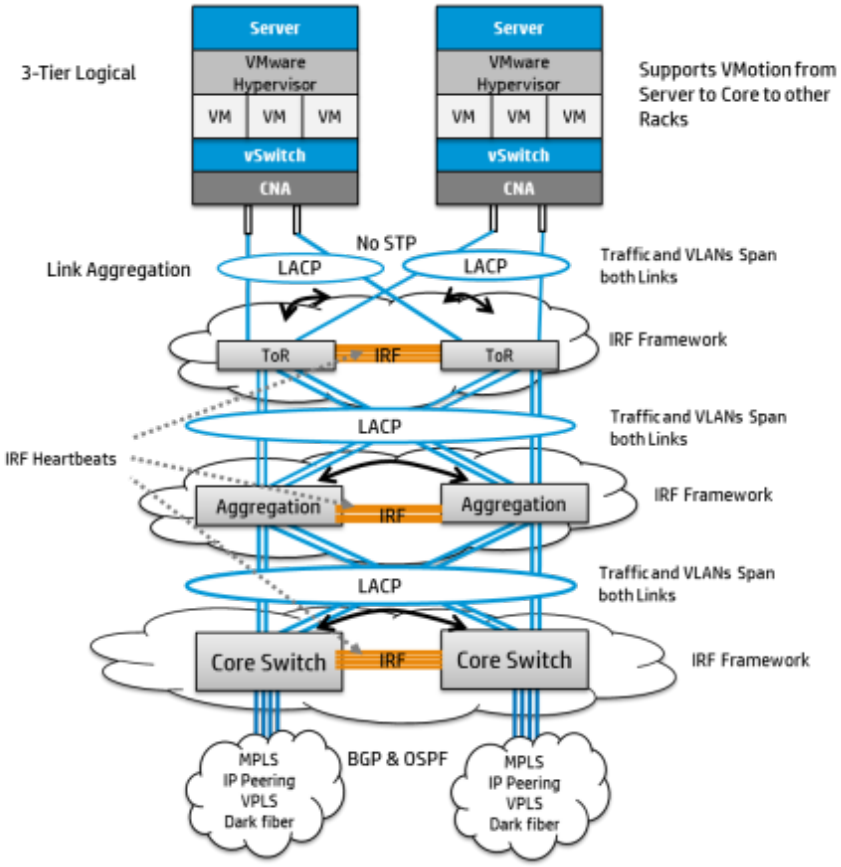

Logical layers of this design

Virtualization

The network edge physically starts with the converged network adapters, but really it starts internally in the VMs and configured virtual switches. The VMs and virtual switches can provision VLANs, which in turn interoperate with the IRF fabric, allowing seamless VM movement with vMotion and high performance frame forwarding. All the switches in an IRF framework support up to 4096 VLANs, which provides substantial network flexibility to configure a substantial number of VM groups onto their own virtual network segments.

In comparison with a two-tier network deployment, the aggregation switches add additional port aggregation and cabling route flexibility. The HP switches are also IRF-capable and can be connected into their own IRF framework. The aggregation switches function well when the rows within a data center are extremely long and the data center architects want to transition away from ToR switch deployments. This approach aggregates port connectivity into high capacity switches and removes multiple points of management.

The network design does not require the STP/RSTP protocol. However, in the event that existing third-party switches are already deployed in the racks, STP-family protocols are supported.

Combining LACP and IRF in this design provides high-speed link aggregation with re-convergence times at sub-50 ms in the event of a link failure. It also allows links to be aggregated and utilized for higher bandwidth from the converged network adapters across all switches to forward traffic.

This network design supports long-range vMotion connectivity, enabling clustering or synchronizing of VMs between disparate locations. The switches also allow VPLS connectivity between switches in geographically separated data center locations. Be aware that long range segments using MPLS can add latency that limits the ability for VMs to use converged I/O resources between data center locations (see the HP FFI section). Long distance WAN networks generally address normal server communications and disaster recovery efforts.

Security

Although not depicted in the diagrams, HP TippingPoint IPS devices can be added into various layers in the network. Since VLANs can be configured to span all the way to the core switches, all network security could even be centralized in a fully redundant N+1 configuration by connecting the IPS devices to each of the core switches.

Orchestration

This network configuration supports the full feature complement of HP IMC, from the core switches to the physical servers. This provides a single-pane-of-glass management platform to support the servers, VMs, vSwitches, IRF frameworks, IP routing, and security devices.

Additionally, the HP IMC VAN module enables a more agile model by eliminating unnecessary steps in virtualizing business environments. Time to VM deployment is accelerated through the upfront definition of policies for VM connectivity and provisioning automation. These policies allow for rapid deployment and follow the workload if it is moved, paused, and/or resumed.

Support for the eco-centric data center

HP servers and switches lead the industry in limiting power consumption without compromise on performance. Every component in this design from the rack up supports the HP “sea of sensors” which can be monitored using HP IMC.

HP Converged Infrastructure provides a complete end-to-end virtual computing solution, encompassing racks, servers, storage, networking technologies, security, and management.

TRILL based designs

TRILL HP supported technology which combines the simplicity and flexibility of Layer 2 switching with the stability, scalability, and rapid convergence capability of Layer 3 routing. All these advantages make TRILL very suitable for large flat Layer 2 networks in data centers.

TRILL and IRF can be deployed in the same environment. IRF can be seen as a “clustering” technology allowing multiple devices to be seen as one logical device (node), removing STP, VRRP from the network, with a single IP for the management. TRILL, on the other hand, provides a mechanism that allows every single node to have a tree rooted at itself, allowing the optimal (shortest path) distribution of traffic as well as multi-pathing for failure recovery.

Objectives

When determining if this deployment type is appropriate, keep in mind the overall objectives:

- **Layer 2 flexibility:**

  Because the overall goal is to support VMs, a flattened Layer 2 network allows VMs to be moved without the need for IP address changes.
- **Reduced management complexity with IRF:**

  IRF and TRILL are in fact not mutually exclusive. When used together TRILL and IRF can combine the best of both worlds, allowing for simplified management, reduced routing table sizes and larger domains with faster recovery times.
- **Zero need for STP/RSTP:**

  TRILL switches (RBridges) eliminate the need for STP in a large bridging domain by running a link state protocol in which connectivity is broadcasted to all the Rbridges. In this way each RBridge knows about all the other RBridges, and the connectivity between them. This gives RBridges enough information to compute optimal paths for unicast traffic, as well as calculate distribution trees.
- **VLAN Management:**

  A flatter Layer 2 network is ideal for VMs. It provides support for network extension across the data center and between geographically separate centers, the framework for VM mobility
- **Centralized security:**

  Layer 2 networks allow the IP-based IPS devices to be connected to the core switches, where VLANs are aggregated. This means fewer network security devices to provision and monitor. It allows for less complexity in security policies because they require the provisioning of fewer devices. Spanning of VLANs across a flattened Layer 2 data center network provides greater flexibility in server connectivity in disparate rack locations because it removes the need for complex IP-based ACLs

Why TRILL?

TRILL is an evolutionary step in Ethernet technology designed to address some of the shortcomings within Ethernet, specifically spanning tree and loop prevention. TRILL is able to use Layer 3 multipathing and shortest path routing techniques to create large flat Layer 2 domains, so that clients and VMs can move and migrate without having to change their IP addresses.

Architecture and scalability

TRILL architectures can be as simple as a few devices meshed at a single layer or they can scale to many devices meshed across different layers. As shown below a TRILL network may consist of many access layer devices which are meshed to 2, 4, or more distribution/core layer devices. In this example, the core/distribution layer devices would be configured as the designated routing bridges (DRB). The DRBs will be responsible for creating a distribution tree which will guide the forwarding of multi-destination frames, which include multicast, broadcast, and unknown unicast frames in the VLAN. Layer 3 routing in this example would be handled by the routers located at the WAN edge.

IRF is not specifically shown in the below image, but IRF and TRILL are not mutually exclusive. In the real world, each individual device shown could actually be a stack of devices virtualized into a single device using IRF.

Rack and logical diagrams

Figure 33 TRILL architecture rack view

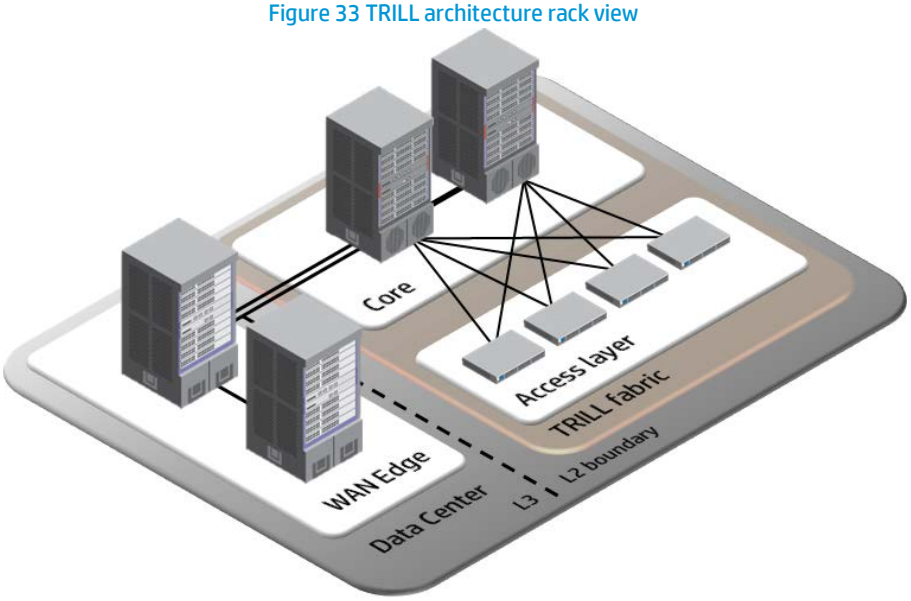

Figure 34 TRILL logical

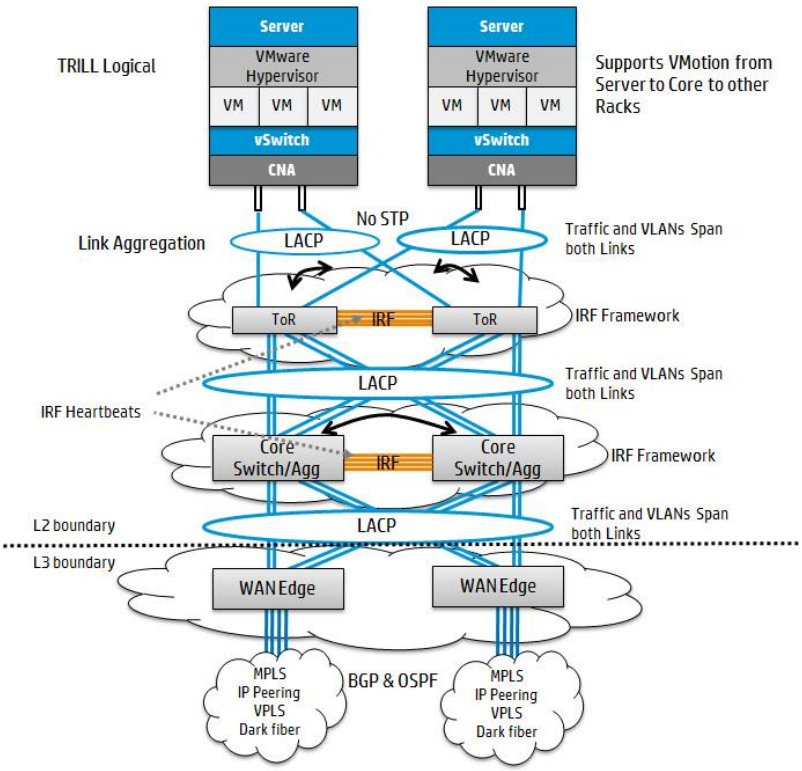

Logical layers of this design

Virtualization

The network edge physically starts at the converged network adapters, but really it starts internally in the VMs and configured virtual switches. The VMs and virtual switches can provision VLANs, which in turn interoperate with the IRF fabric, allowing seamless VM movement with vMotion and high performance frame forwarding. All the HP switches in an IRF framework support up to 4,096 VLANs, which provides flexibility to configure a substantial number of VM groups into their own virtual network segments.

Combining LACP and IRF provides high-speed link aggregation with re-convergence times at sub-50 ms in the event of a link failure. It also allows links to be aggregated and utilized for higher bandwidth from the converged network adapters across all switches to forward traffic.

This network design supports long range vMotion connectivity, enabling clustering or synchronizing VMs between disparate locations. In this solution, routers at the WAN edge will allow VPLS connectivity between routers in geographically separated data center locations. Be aware that long-range fiber segments using MPLS can add latency that limits the ability for VMs to use converged I/O resources between data center locations (see the HP FFI section). Long distance WAN networks generally address normal server communications and disaster recovery efforts.

Security

Although not depicted in the diagrams, HP TippingPoint IPS devices can be attached into various layers in the network. Since VLANs can be configured to span all the way to the core switches, all

network security could even be centralized by placing an IPS device into the core layer.

Orchestration

This network configuration supports the full features of HP IMC from the core switches to the physical servers. This provides a single-pane-of-glass management platform to support the servers, VMs, vSwitches, IRF frameworks, IP routing, and security devices.

The HP IMC VAN module enables a more agile model by eliminating unnecessary steps in virtualizing business environments. Time to VM deployment is accelerated through the upfront definition of policies for VM connectivity and provisioning automation. These policies allow for rapid deployment and follow the workload if it is moved, paused, and/or resumed.

Support for the eco-centric data center

HP servers and switches lead the industry in limiting power consumption without compromise on performance. Every component in this design from the rack up supports a “sea of sensors” which can be monitored using HP SIM.

Multi-Tenancy

An emerging and related datacenter requirement is the ability to logically partition a single physical device, such as the HP 12500, into many virtual devices.

This capability—referred to as multi-tenancy, gives an administrator the flexibility to set-up multiple customers or lines of business on the same physical hardware while ensuring dedicated and discrete management, security and network services separation and isolation from other “tenants”.

HP Multi-Tenant Device Context (MDC)

HP has developed MDC which is a true multi-tenant solution enabling organizations to build a single network that can serve multiple clients/tenants. MDC provides true isolation in a multi-tenant environment by allowing the memory of a single device to be separated into protected partitions running different instances of the switch. This allows for full separation of forwarding databases and traffic in each MDC for those different tenants. As an example, in the Enterprise context, those tenants could be the following departments: Finance, Marketing, R & D and Legal.

Figure 35 HP Multi-tenant Device Context

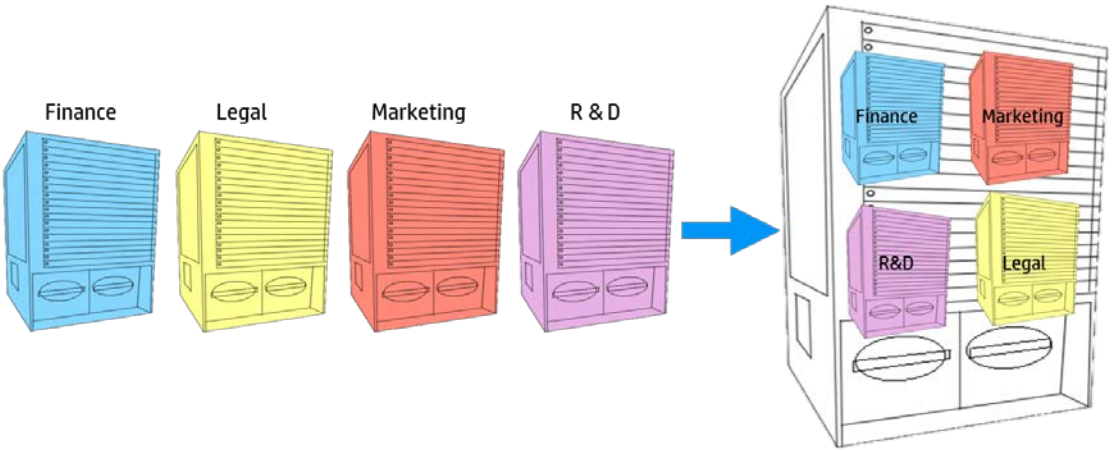

MDC provides complete separation and increased resiliency, especially when used with IRF across a group of switches in the core of a network. MDC would exist in all the switches in that group. In this way the organization could have complete transparent failover in the event of a single device failure, or when a single device needs to be taken offline for service.

Serving four tenants in a single device, with full separation instead of building four separate networks, means organizations can use up to 75% less equipment, spend up to 75% less money, and consume up to 75% less power in cooling and space, and administration.

Provider Bridging

IEEE 802.1ad, sometimes referred to as 802.1QinQ, is the enabling technology for Provider Bridging. **Provider Bridging** enables an Ethernet frame to carry two VLAN tags which allows the Service Provider to offer an Ethernet LAN service over their existing network where the tenant can maintain their own VLANs.

Figure 36 802.1ad QinQ frames

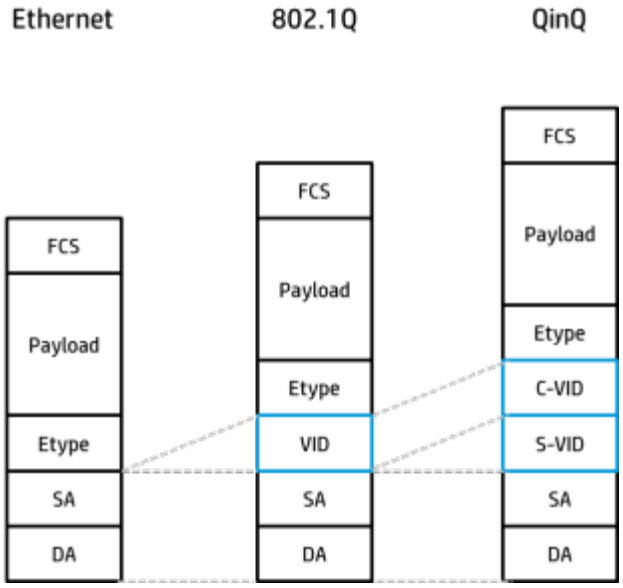

In the 802.1ad frame, the Outer Tag is the VLAN allocated to the Tenant in the provider network and as such is called the S-TAG (Service Tag) which contains the Service VLAN ID (S-VID). The Inner Tag

represents the VLANs belonging to the tenant and is called the C-TAG (Customer Tag) and contains the Customer VLAN-ID (C-VID).

Forwarding in Provider Bridging networks requires switches to performed two new operations

- Tag pop – which removes the outer VLAN tag
- Tag push – which adds an outer VLAN tag

While QinQ allows for each customer to be contained to a unique VLAN it does not abstract the customer and provider network entirely. Switch forwarding decisions are still based on the S-Tag and the Destination MAC address, therefore, the Provider Network needs to learn the MAC addresses of all customer devices. Furthermore QinQ does not restrict the flow of Layer 2 control frames for protocols such as Spanning Tree. To address these shortcomings Provider Backbone Bridging (PBB) was developed.

Provider Backbone Bridging

IEEE 802.1ah-2008, PBB or MAC-in-MAC , extends the work achieved in 802.1ah by providing a hierarchical network infrastructure that completely abstracts the Service Provider Backbone from the Customer Network. PBB networks are deployed to aggregate existing PBB networks

When a frame arrives at a Backbone Edge Bridge (BEB) it is encapsulated and the BEB adds the Backbone Source, Destination MAC (B-SA and B-DA) and Backbone VLAN. The I-SID or Service Identifier Field is a 24-bit Customer Specific Identifier which supports 2^{24} or ~16 Million tenants.

Figure 37 IEEE 802.1ah-2008, PBB

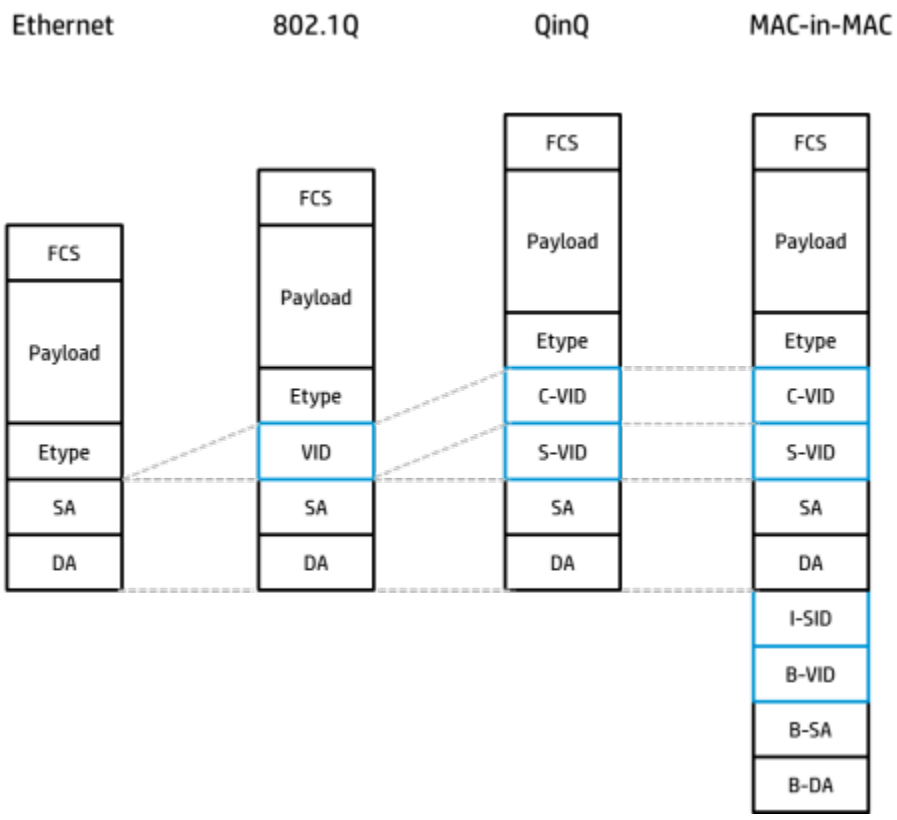

Virtual Routing and Forwarding

Virtual Routing and Forwarding allows one router to maintain several separate instances of a routing table. Conceptually, it is not the control plane that is virtualized – it is the data plane that is virtualized allowing forwarding decisions to be segregated by customer or any other administrative grouping required by the network engineer. VRFs generally go hand in hand with MPLS VPN although they can be used separately without any problems. VRF lite, for example is a description of the Virtual Routing and Forwarding technology being used without MPLS VPNs.

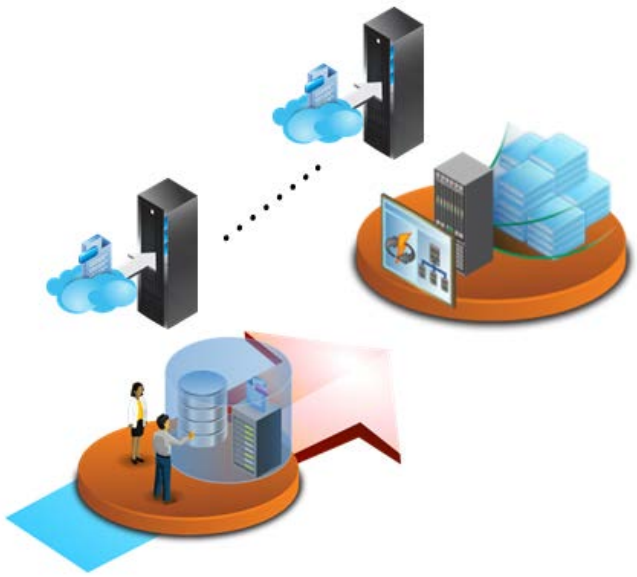

HP Data Center Interconnect (DCI) - connecting geographically dispersed data centers

As IT technology has been maturing, many organizations have been reducing costs and increasing efficiency by following the trend to consolidate their data centers. This consolidation has been driving organizations to heavily leverage virtualization in the data center as well as private clouds and multi-tenant environments which support different business units. This allows organizations to increase their availability and help them to move and/or deploy new workloads onto resources that can best serve them.

These advancements mean that connecting geographically dispersed data centers is now more important than ever before. Geographic data center interconnections allow the IT designer to put in place disaster avoidance and disaster-recovery mechanisms that increase the availability of the applications. Geographic dispersion also enables shared resource utilization, optimization of application response, and allows the flexible mobility of workloads and services.

Unfortunately, most interconnect methods suffer from limitations which include transport dependency, complexity and lack of resiliency. The HP DCI solutions are designed to address these limitations by delivering responsive, efficient and resilient data center interconnect solution.

DCI features and benefits include:

- **Mobile Workloads and long distance vMotion:**

  DCI establishes reliable connections between data centers that serve as a platform for fast and reliable vMotion between distant data centers. The added speeds of up to 80% improvement of vMotion enhance the value of workload mobility and make applications more available to users
- **Multi-tenant enabled Layer 2 extension solution:**

  The HP Ethernet Virtual Interconnect (EVI) solution can leverage HP Multi-tenant Device Context (MDC) technology to partition a single HP 12500 into multiple logical devices, which provides up to 75% reduction of the number of physical platforms leading to CapEx and OpEx reductions
- **Follow-the-sun support:**

Organizations can deliver a better experience with up to date information by adopting a follow the sun model which would be leveraged by call centers, engineering, and development teams

- **Bursty and seasonal traffic patterns:**

  Organizations could load balance heavy or bursty workloads among geographically dispersed data centers making more effective use of data center resources

- **Investment Protection:**

  DCI solutions can be customized to fit the customer's exact environment, from IP to MPLS to DWDM. DCI solutions can work with customers' existing IP networking infrastructure without requiring changes. This protects networking investments and allows easy and seamless deployment

- **High Availability:**

  The DCI solution establishes active/active links between data centers. Along with HP IRF, DCI establishes reliable links with link aggregation and failover capabilities. This gives users uninterrupted access to applications in the event of disruption of service at one data center

- **Disaster Recovery and Data Replication:**

  Data can be replicated across data centers using DCI. In the event of failure, DCI provides the means to recover data from remote data centers and ensures business continuity

Key considerations for DCI design

Current network conditions

The network resources of a user between data centers will determine the solution to use, as follows:

- **Ethernet Virtual Interconnect (EVI):**

  This IP-based solution option is extremely useful in simplifying DCI. Transmission between datacenters over DWDM or MPLS can be complex to manage and is often highly dependent upon costly dedicated and rigid service provider's infrastructures. In contrast, EVI runs over IP infrastructure so it can be deployed without requiring changes to an existing infrastructure. This characteristic simplifies deployment by allowing Layer 2 connectivity across the network without having to deal with Layer 3 networking dependencies. EVI is currently supported by the HP 12500 series switch

- **Ethernet LAN extension:**

  This option extends Ethernet natively over a dark fiber or DWDM optical transport. As such, this solution mostly applies to point-to-point deployments, where the sites are connected via dedicated dark fiber links or DWDM optical circuits. HP FlexFabric Networking products that support this Ethernet LAN Extension option include the HP 12500 series switches and HP 8800, 6600 series routers

- **MPLS Point-to-Point or Multipoint using MPLS or VPLS:**

  This option uses MPLS technologies to provide L2 connectivity services over a L3 network service. Depending on the nature of the transport infrastructure between data center sites and the number of data center sites to be interconnected, different technologies can address the connectivity requirements. HP FlexFabric Networking products that support MPLS/VPLS include the HP 12500 series switches, and HP 8800, 6600 series routers

High availability (HA)

In the perspective of VM resource scheduling and remote cluster access, multiple interconnected data centers can be considered as a logical large-size data center. The links interconnecting data centers can be considered as the DCI backbone links of the large-size data center. More importantly, the backbone links interconnecting data centers transmit control signaling, in addition to vMotion or data packets. Therefore, once the links interconnecting the data centers fail, the large-size data center fails to work properly, and causes service interruption for users.

Therefore, a key consideration of Layer 2 DCI is to improve availability. The best way of improving HA is to design backup DCI links and backup nodes (DCI devices). To improve HA and increase the interconnecting bandwidth at the same time, you can design load-sharing interconnection links, so that you can increase the bandwidth and enable the services to rapidly converge when the system encounters errors, thus improving HA.

For example, whenever possible, HP recommends using IRF technology on the DCI devices as well as the CE core devices. Aggregating two or more links not only simplifies the DCI topology but provides HA and achieves load sharing. Use LACP to aggregate two or more links between the dedicated DCI devices into a logical link, so that the DCI topology is greatly simplified. At the same time, the bandwidth of the two HA links is optimized, and load sharing is achieved.

Figure 38 IRF + LACP load sharing design

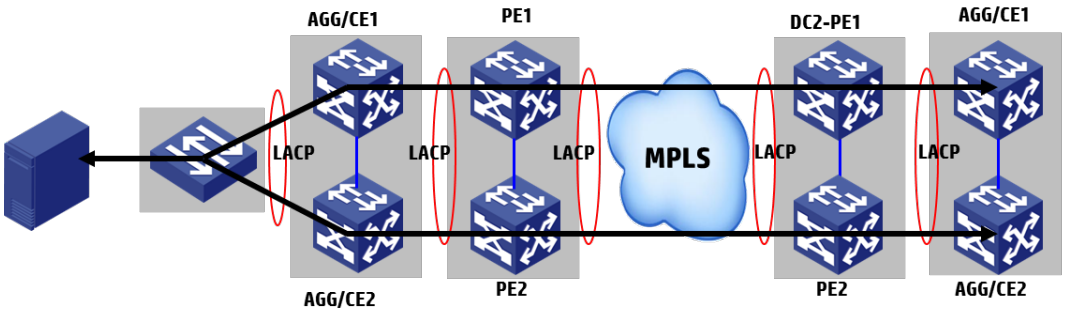

Because many routers do not support IRF, there may still be an end-to-end loop between DCI devices of each DC, thus smart link should be used to obtain HA and loop avoidance. As shown in the figure below, if the PE or link between PE and CE devices fails, the flow can achieve fast convergence by using smart link and quickly switch the traffic path to another one.

Figure 39 Smart link design

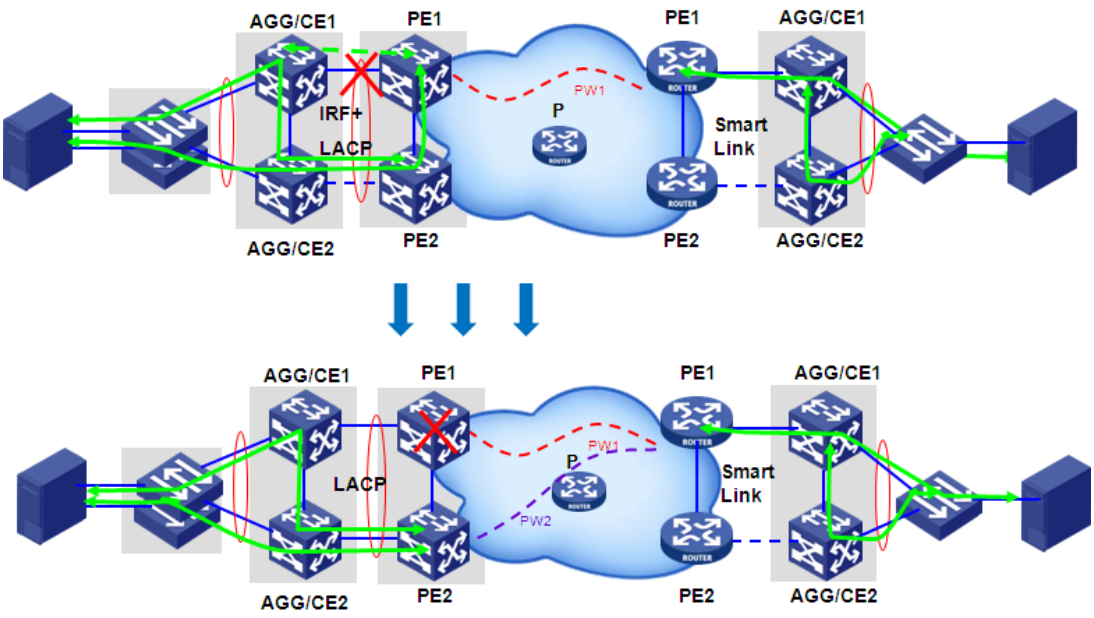

Layer 2 topology management (STP Domain Isolation & loop management)

When you deploy STP in a Layer 2 DCI network, HP recommends that you isolate STP domains to each data center. On each port connected to the DCI link, enable BPDU drop (Edge port) and disable STP.

Additionally, end-to-end loop management needs to be addressed to manage the loops across multiple interconnected Layer 2 domains. Since STP domains will be isolated at each data center you will need to ensure no physical loops exist. Of course this could be done by utilizing only one single link between data centers, but a better solution would be to use IRF & LACP, or smart link if the devices do not support IRF.

Figure 40 End-to-end loop management design

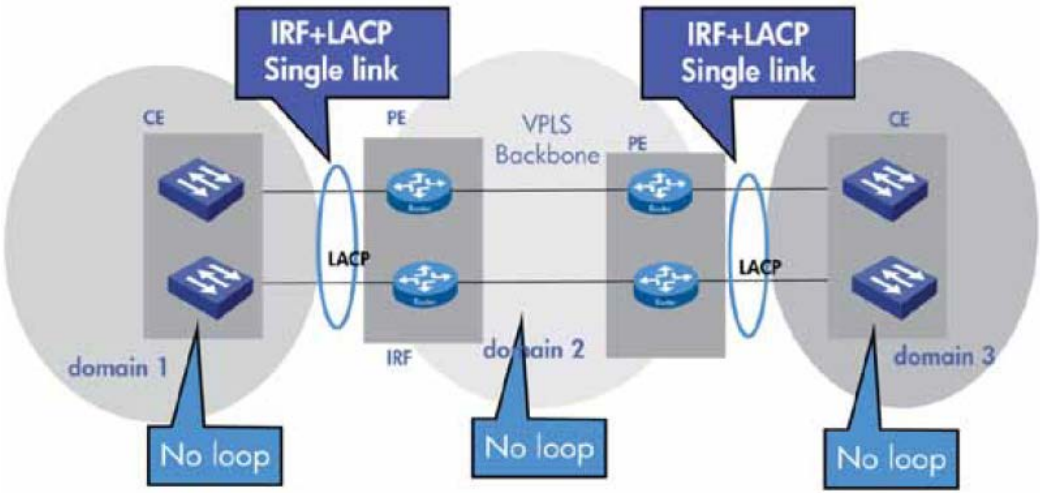

Path optimization for DCI

Typical traffic models of data centers show that the incoming traffic and the outgoing traffic are seriously asymmetrical: the request traffic entering the data center is usually very light and the response traffic from the servers to clients is usually very heavy.

If a large volume of data traffic passes through the DCI core network links, the data traffic consumes heavy bandwidth and degrades the quality of transmitted control data. This causes a loss in control data packets and affects the migration or disaster recovery process.

Therefore, in the event of a DC failure, you must prevent the response data traffic from passing back through the DCI core network link. The key to data path optimization is drawing the response traffic of servers away from the original DCI core network path and towards the data center local gateway.

To draw the response traffic of servers away from the original incoming DCI core network link, you can configure a VRRP group on each data center gateway, as shown in the figure below.

More specifically:

- Configure the same VRRP group (the same virtual IP address and priority) on the gateways for data center A and data center B.
- Configure an ACL on the outgoing interface that connects the gateway to the DCI core network link to prevent the VRRP packets from entering the other data center through the DCI link.

Figure 41 Data path optimization for data centers

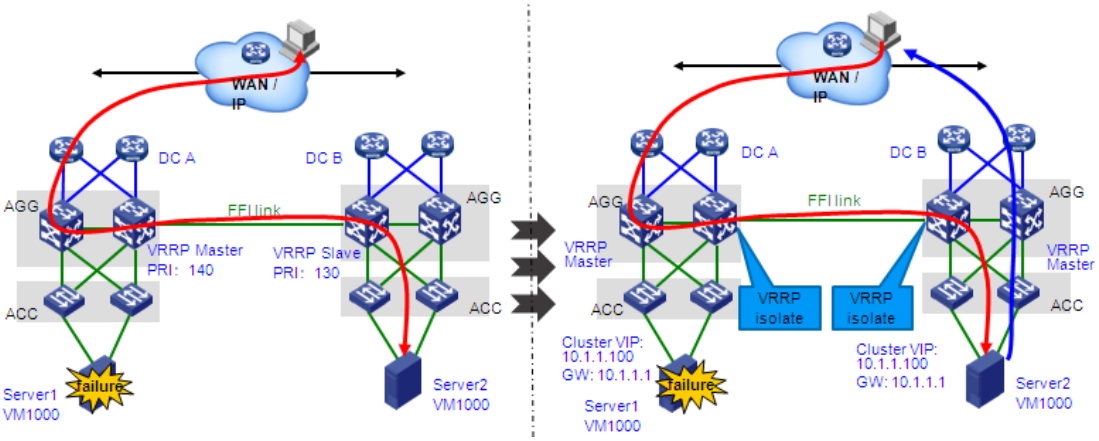

With these configurations performed, when a VM is moved from data center A to data center B, its IP address does not change, and data center B is configured with the same VRRP gateway. As a result, the response traffic of servers can be directly sent to the clients through the data center B gateway. This prevents a large volume of response traffic from passing back through the DCI link.

Ethernet Virtual Interconnect (EVI)

Along with existing DCI solutions based on DWDM and MPLS/VPLS, HP has developed data center interconnect solution with multiple technologies to help organizations quickly and easily extend layer 2 networking to multiple data centers.

Ethernet Virtual Interconnect

EVI runs over Internet Protocol (IP) transport and extends layer 2 domains across a WAN network, typically between data centers. By virtualizing and automating the link layer domain across data centers, EVI delivers the elements necessary to enable a Software Defined Networking (SDN) data center infrastructure. It enables several data centers to work as one that is more responsive, with higher efficiency and solid high availability for business resiliency.

With EVI, enterprises are able to:

- Accelerate the delivery workload mobility with remote vMotion
- Increase applications performance with multipathing and load balancing
- Allow organizations to scale to up to 8 geographically dispersed data centers without requiring them to change the underlying network
- Simplify the Layer2 connection by encapsulating traffic over GRE and automatically isolate Spanning Tree Protocol
- Achieve optimum degrees of high availability and disaster recovery for valuable data
- Allows clients to have a simple set of Layer 2 routing extensions that can provide data interconnectivity in minutes rather than the months of legacy approaches like VPLS

Figure 42 HP Ethernet Virtual Interconnect

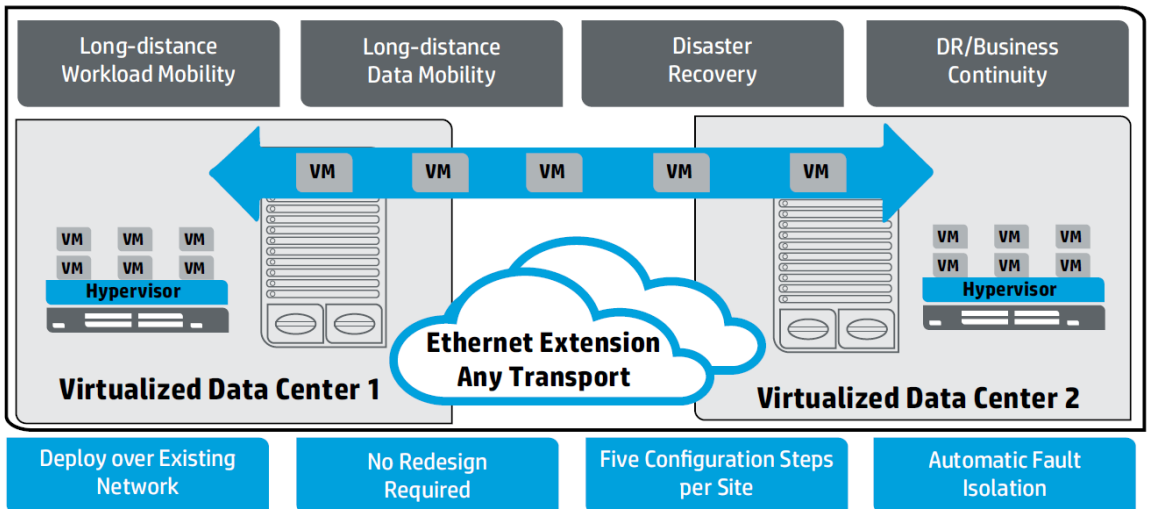

When used along with the IRF switch virtualization technology, EVI delivers greatly enhanced reliability, resilience and faster remote vMotion capabilities. The combination of EVI and MDC brings multi-tenancy to cloud-ready and remotely connected data centers.

EVI is currently supported on the HP 12500 series switches. All technology references and design considerations focus on positioning these switches in the data centers to establish Layer 2 connectivity between locations. EVI is a feature supported within Comware 7.

EVI Basics

EVI is a layer 2 routing technology that uses EVI Links and GRE tunnels to extend VLANs across up to 8 locations. Each EVI network has a unique network ID, extends a unique list of VLANs, and has separate control and forwarding planes. When used in conjunction with MDC, each MDC can support 32 EVI

networks with each MDC running up to 4K VLANs.

EVI Control Plane

The EVI control plane is responsible for the discovery of new EVI nodes, the connection of these nodes, and maintaining MAC learning and advertisement.

EVI Neighbor Discovery Protocol (ENDP)

ENDP is the protocol responsible for setting up and maintaining the EVI Links between end devices. There are two types of entities used with ENDP:

- **ENDP Server (ENDS):**

  An ENDS is responsible for receiving registration requests, maintaining the ENDC database and propagating member information.

- **ENDC:**

  These are the locations that communicate with the ENDS. Requests are sent from ENDCs and are received by the ENDS.

Once EVI has been configured the discovery process is autonomous. The high level process is described below:

1. The main data center switch(es) have EVI server (ENDS) enabled
2. When a new location is added, the switches need to be configured with EVI client enabled
3. The new location sends a registration request to ENDS
4. The main data center sends a response back
5. The two locations exchange EVI member information

Figure 43 HP EVI discovery process

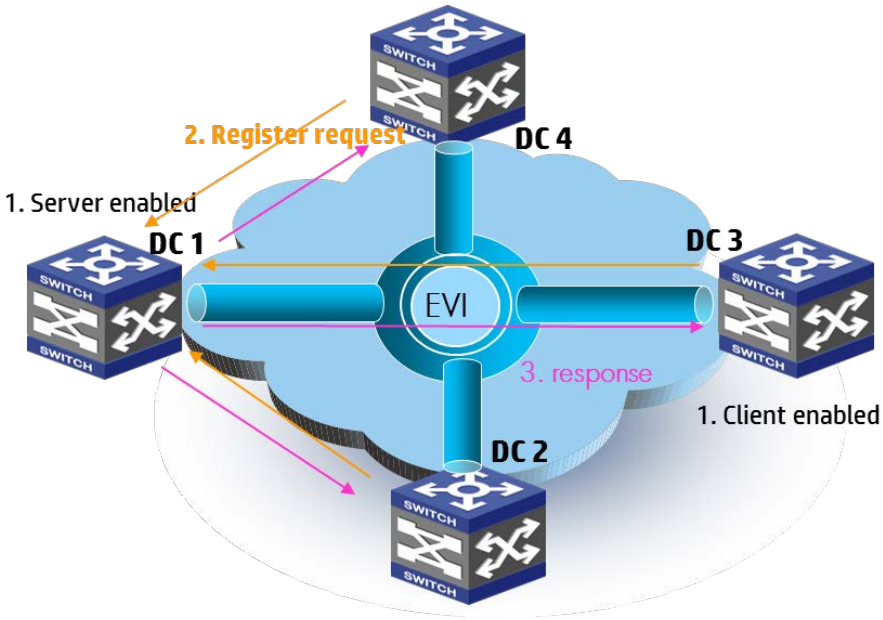

EVI MAC Address learning and advertisement

The standard MAC learning behavior is not affected by the configuration of EVI. In fact, EVI adds enhanced functionality by advertising known MAC addresses to other EVI enabled devices as well as other features that are highlighted in the EVI features section below. EVI uses IS-IS to propagate MAC addresses reachability information. The EVI IS-IS packets are transported across the underlying carrier network.

The figure below demonstrates how a MAC address is learnt at Site 1 (IP1) and sent to the other locations.

1. Site IP1 learns about new MAC entries mac1 and mac2 to VLAN 100
2. The EVI IS-IS process creates an LSP update that includes the mac addresses and VLAN info
3. The Edge Device (ED) at Site IP1 sends the LSP to its neighbors
4. Each neighbor reads the LSP and delivers the information to the local EVI IS-IS process. Locally the MAC address looks like it was learnt by the EVI tunnel interface. When a packet is destined for either MAC at any location other than IP1, it is the EVI IS-IS process that routes this back to IP1
5. The LSP updates can also include MAC addresses that have been aged out

Figure 44 HP EVI MAC learning

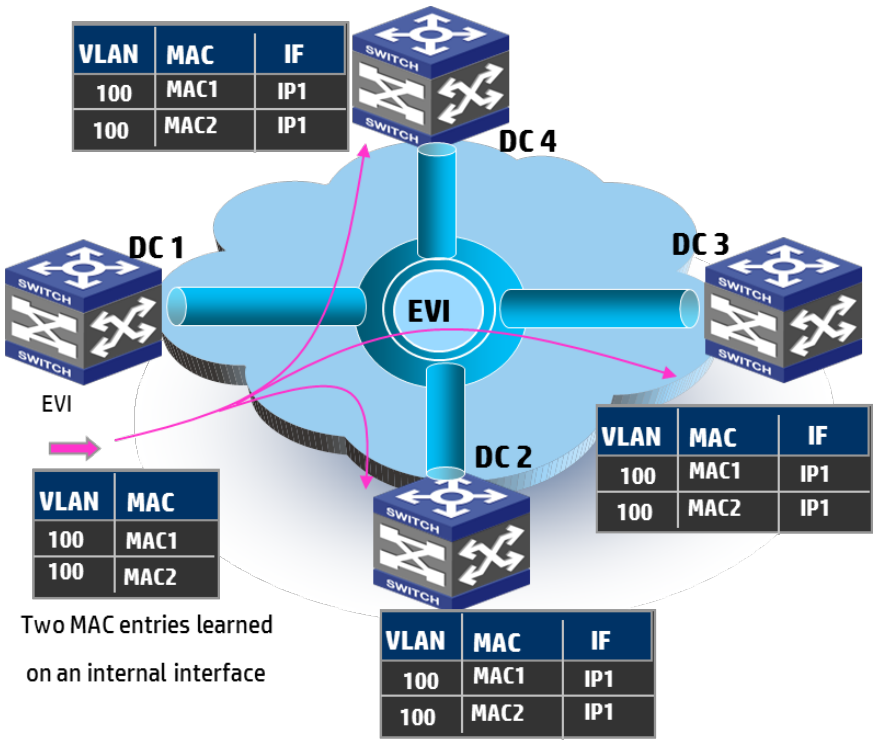

EVI Data Plane

The EVI data plane is separate from the EVI control plane and is used to forward unicast, multicast, and broadcast traffic across the EVI network. This section covers the forwarding of traffic in an EVI network in more detail.

Local Forwarding

Local packet flow is unchanged and not influenced by EVI.

EVI Data Plane Unicast Forwarding

When it is determined that a packet is destined for a device at another location, a lookup is done and the correct EVI interfaces are identified. The packet is encapsulated in a GRE tunnel and flows across the EVI network from the source ED to the destination ED. The destination ED removes the encapsulation and forwards the Ethernet packet to the end device as if the original sender were local.

Figure 45 HP EVI data plane forwarding

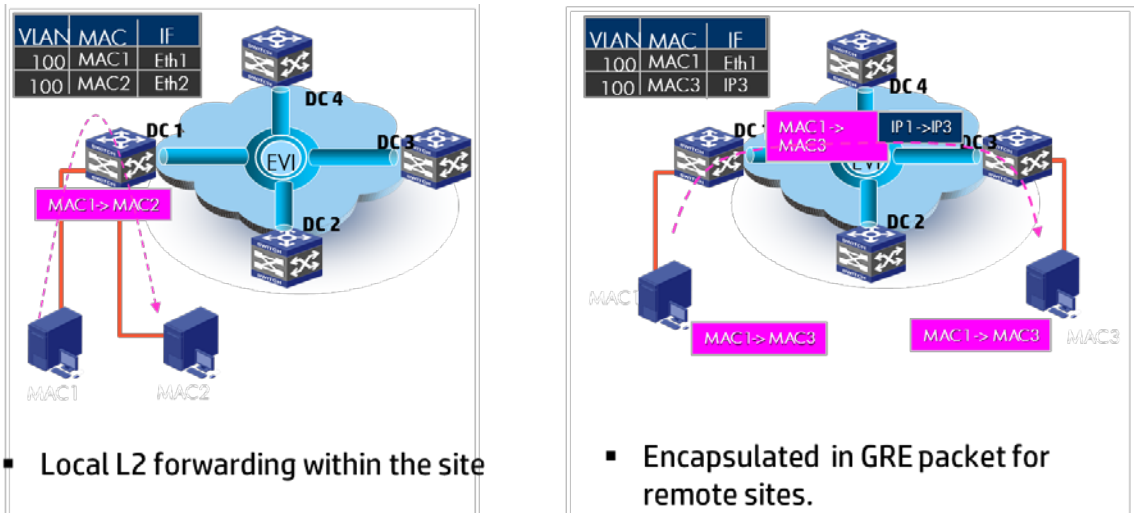

EVI Data Plane Multicast Forwarding

In the current release, multicast is supported. The EVI ED closest to the multicast traffic source will unicast the frames across the EVI Links. The local site will propagate the multicast packet based on the EVI control plane learning mentioned previously.

In the future, the end device will map the multicast frame to the multicast tree on the transport. The transport then replicates the frame for multicast member sites. This will require the devices to support Asynchronous Source Multicast (ASM) and Source Specific Multicast (SSM).

EVI Data Plane Broadcast Forwarding

By default, broadcast traffic is allowed across the EVI network. Unknown Unicast and Multicast packets are dropped at the ED.

EVI Features

The following features are used to optimize the control plane traffic thus increasing efficiency.

Selective MAC Routing

This feature stops unknown unicast and multicast frames from flooding the EVI Links. The internal interfaces are capable of flooding to the internal interfaces while the EVI-Tunnel interfaces will drop these frames. Similar to selective flooding, it is possible to permit or deny a MAC route.

Automatic Loop Prevention

EVI has the following loop prevention mechanisms built in and enabled by default on both the data and control planes:

- **EVI-Split Horizon:**

  This is enabled on the data plane to prevent frames received from EVI tunnels from being forwarded to the transport layer (EVI Links). Its primary function is to prevent loops among EDs
- **STP Domain Boundary:**

  This disables STP on the EVI Links. STP domains and BPDUs are not extended across sites keeping topology changes local. BPDUs are blocked from one location to another so that STP changes are contained within a site. This also allows each site to run different versions of STP. The following versions of STP are supported 802.1d, 802.1s and 802.1w
- **Selective Flooding:**

  By default unknown unicast and multicast are dropped at the EVI Links. If an application uses a special MAC address for traffic identification it would break. Selective flooding enables the ED to flood frames with a certain unknown destination MAC to an EVI tunnel interface

  An example of this would be Microsoft NLB (Network Load Balancer) that uses a special MAC address (cluster MAC) to identify a cluster. If cluster members are located in multiple sites, selective flooding would be used to propagate the cluster MAC address on the EVI-Tunnel interface(s)

Automatic VRRP Isolation

To provide routing path optimization, EVI enables VRRP isolation by default. This allows each DC to have an active layer 3 gateway that leverages the VRRP protocol to hide the specific HW details (IP and MAC addresses). All data centers will run separate sets of VRRP instances. VRRP isolation stops the population of VRRP keep-a-lives of each VLAN over EVI Links, so that each data center always has active/active gateways.

ARP Flooding Suppression

This feature will reduce the number of broadcasts that traverse the EVI network. When an ARP request is made and initially flooded across the EVI network, the EVI ED listens to ARP responses on the EVI Link. These ARPs are cached for the remote MACs so that subsequent ARP requests can be handled directly by the ED.

Figure 46 HP EVI ARP Flooding Suppression

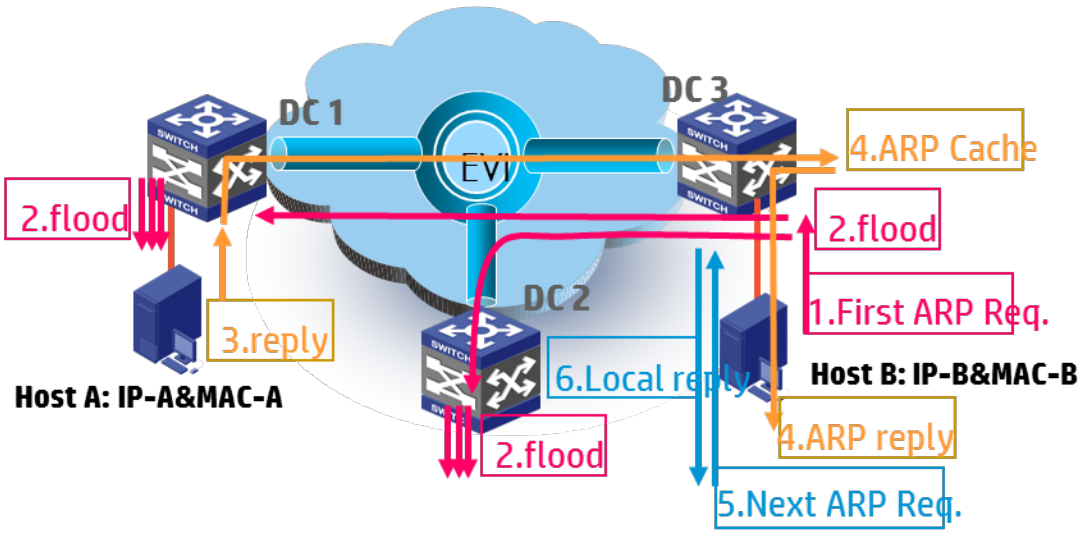

In the figure above:

1. Host B sends an ARP for Host A's MAC
2. The ED at Site IP3 floods the request out all port, including the EVI-Tunnel interfaces (vLinks). The remote EDs decapsulate the packet and flood at their respective sites
3. Host A responds to the ARP request
4. Site IP3 caches the entry for Host A and forwards the response to Host B
5. Any subsequent ARPs for Host A are now handled locally by the RD at Site IP3 rather than repeating the process

EVI with IRF

At each location, High Availability (HA) is achieved by configuring IRF on the HP 12500 series switches to simplify the network topology and increase uptime. IRF can also be configured in conjunction with MDC.

Figure 47 HP EVI combined with IRF

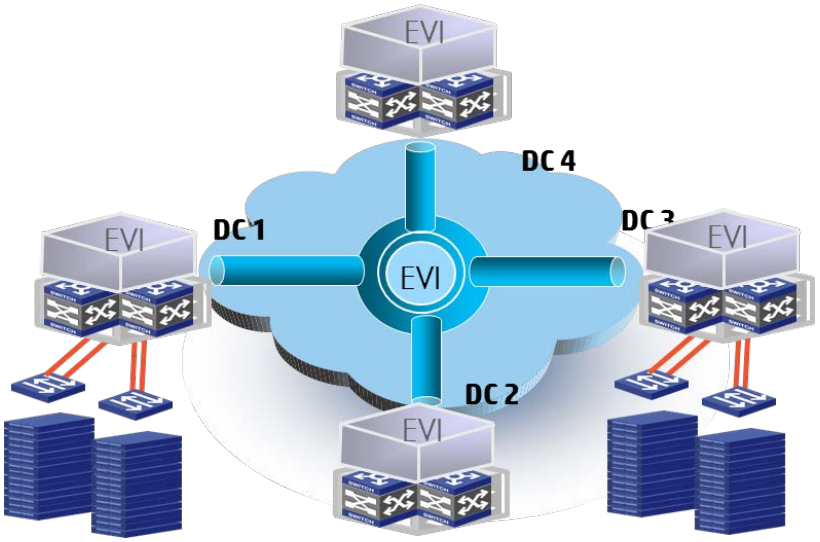

EVI and MDC

MDCs can be configured with IRF in an EVI environment to create completely secure isolation between tenants. Up to 32 EVI networks can be deployed in each MDC and each MDC also has a completely functional L2 and layer 3 environment.

Figure 48 HP EVI and MDC key takeaways

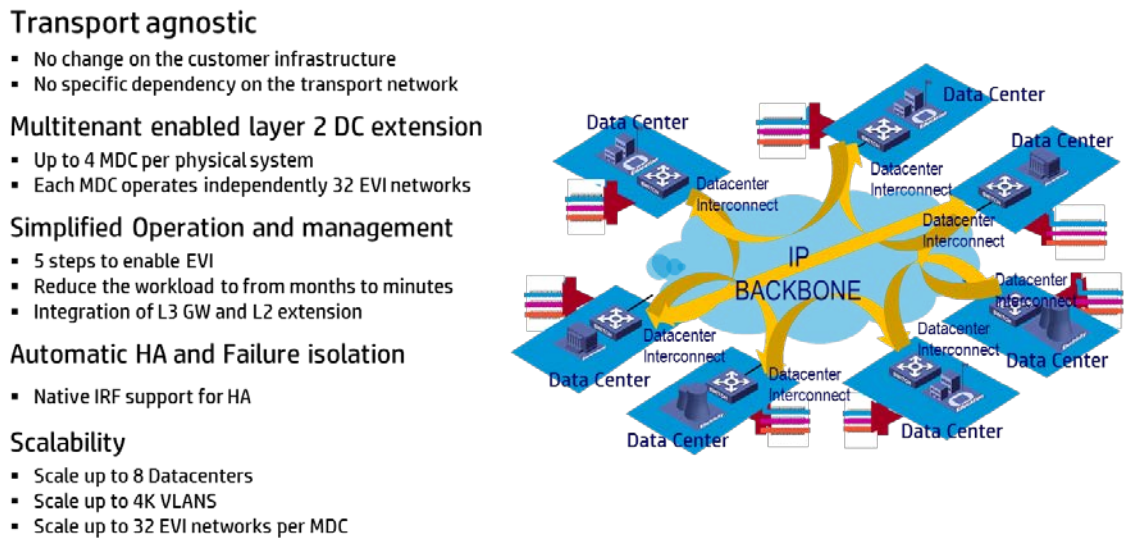

Other DCI options

For those scenarios where HP 12500s and EVI cannot be used, customers can still leverage a variety of other HP products which support other types of DCI deployment models.

The following DCI deployment models are supported on a variety of HP FlexFabric devices such as HP

12500 switches and 8800, 6600, and HSR routers.

DCI design guidelines for Dark fiber or DWDM

Dark fiber or DWDM DCI: Point to point for two data centers

When two data centers are directly connected through dark fiber or DWDM you can use pairs of dedicated DCI devices at each data center utilizing IRF and LACP to create the DCI connection. To facilitate expansion to more than two data centers you can use alternate solutions listed later in this document.

This solution utilizes dual dedicated DCI devices at each data center (can be considered customer-owned PE devices) to create the DCI connection. The dedicated DCI devices need to support virtualization and should support IRF and cross-device link aggregation. Note that this does not mean you will create a cross data center IRF, instead an IRF fabric is formed between the dedicated DCI devices at each data center, and then the two IRF fabrics are connected through an aggregate link. By using IRF to combine the two dedicated DCI devices at each data center to form an IRF fabric, you can decrease the failure convergence time from dozens of seconds to several milliseconds, improving the HA performance by almost two magnitudes. As discussed, you need to isolate the STP domains at each data center by disabling STP on the DCI interfaces.

If IRF-based switches are currently deployed in each of the data centers, it is also viable to leverage the existing IRF devices at each site to help minimize the extra investment for DCI devices.

Figure 49 Point-to-point dark fiber/DWDM for two data centers

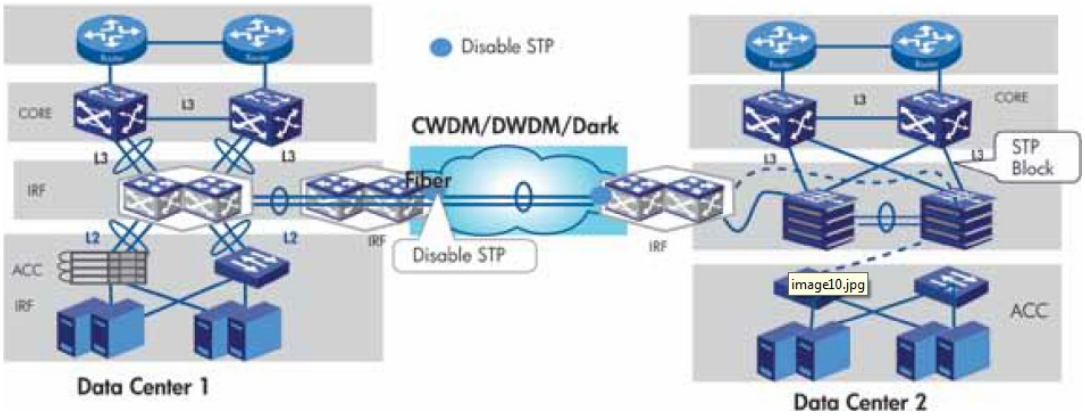

Dark fiber or DWDM DCI: RRPP solution for multiple data centers

RRPP is a convenient solution that can be used to interconnect up to four data centers. You can use RRPP to connect the aggregation layers of multiple data centers to form a DCI core network.

The major purpose of RRPP is to connect multiple DCs and to avoid a core node failure, which can split and isolate the data centers. In an RRPP network, all nodes (the aggregation layer nodes of data centers) are equal. No matter which node fails, the network connectivity can be implemented through backup links. All backup nodes and links are deployed in distributed mode.

Figure 50 RRPP solution

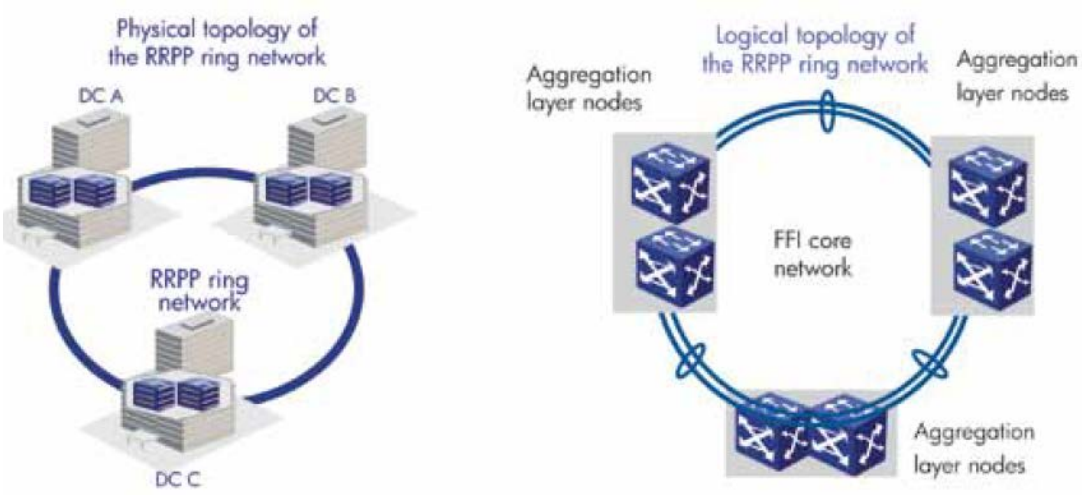

To further improve the HA performance for the RRPP network, you can deploy IRF + LACP in the RRPP ring network. More specifically, you can deploy IRF on the aggregation layer of each data center and configure LACP to aggregate links for interconnecting aggregation layer nodes. This solution improves the HA performance for each node and links in the RRPP ring network.

The combination of RRPP, IRF, and LACP increases the HA performance to millisecond-level convergence. At the same time, link aggregation increases the interconnecting bandwidth for the DCI core network, improves its performance, and enhances the service quality for users.

The benefits of RRPP solution are:

- More robust network architecture, and reasonable risk factor distribution
- Very high HA
- Easy, simple cabling

The RRPP solution has the following disadvantages:

- Supports a maximum of 4 data centers within the same city. Over 4 data centers and across too large a distance will cause the forwarding path in the ring network to be too long. This will increase the network delay and decreasing the service quality

Dark fiber or DWDM DCI: Hub-and-spoke solution for multiple data centers

You can use the dark fiber or DWDM and RRPP solution to interconnect up to 4 data centers nodes. To facilitate expanding to more data center nodes, you must use the hub-and-spoke model, where a core node is connected to the aggregation layers of multiple data centers. Logically, the multiple data centers and the core node form a hub-and-spoke star topology, where the core node is the hub, and the aggregation layer of each data center is a spoke.

Figure 51 Hub-and-spoke solution for multiple data centers

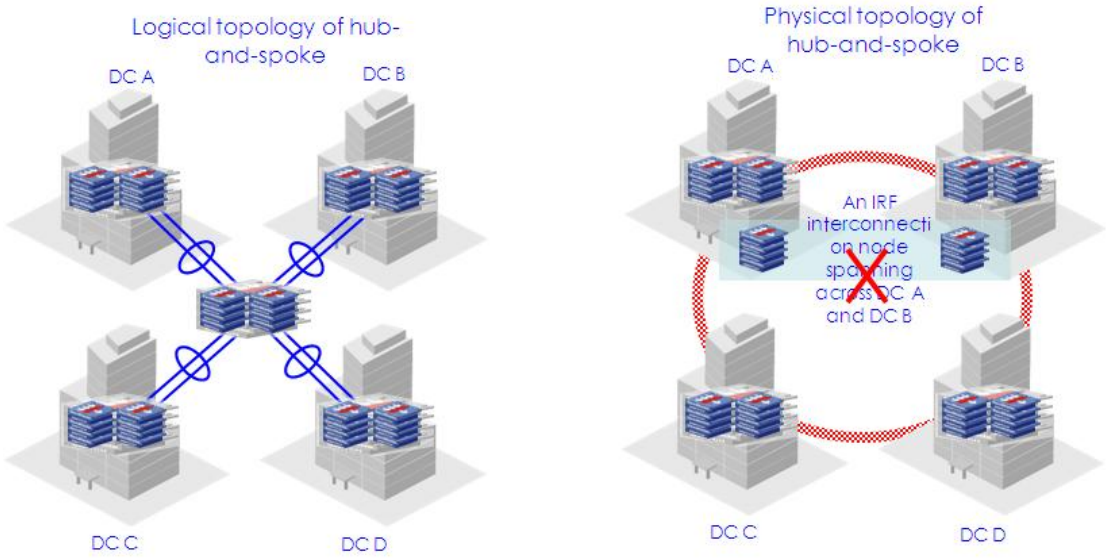

In a hub-and-spoke network, the core node is crucial, it determines whether the whole network can work properly and is a key factor for guaranteeing HA for multiple data centers. By using IRF to combine the two devices to form an IRF fabric, you can decrease the failure convergence time from dozens of seconds to several milliseconds, improving the HA performance by almost two magnitudes. The HA performance improvement is paramount for the hub-and-spoke solution.

Note that, HP recommends putting the two-core node IRF member devices of the IRF fabric in the same data center.

As shown on the right in the figure above, the two core devices have been deployed in two separate data centers (**not recommended**), connected through DWDM, and combined to form an IRF fabric as a logical core node. However, the two data centers are connected through more than one link. The nodes at the aggregation/access layer of each data center are connected to the core node through an aggregate link, improving the HA performance of the links. Note that, a data center usually has multiple nodes at the aggregation/access layer, and if you use a flat Layer 2 network solution (including only the core layer and access layer), you may need lots of aggregate links which connect to the other data center. This will multiply the investment costs for DCI.

Note that, this example with core devices in separate data centers forming an IRF fabric also greatly lengthens the control signaling link of the IRF fabric and thus increases the probability of IRF partitioning. Once the control signaling link failures cause IRF partitioning packets can be lost in the DCI network, thus seriously affecting the availability of the services. Therefore, this example solution brings high risks, is given as a use cases scenario, and is **not recommended**.

In all, the hub-and-spoke solution using IRF delivers the following benefits:

- Least number of logical nodes, and simple architecture
- Scalability

However, the hub-and-spoke solution has the following disadvantages:

- When creating IRF fabrics across data centers:
  - The solution greatly lengthens the control signaling link of the IRF fabric and thus increases the probability of IRF partitioning
  - Adds to the number of links from aggregate/access devices which would cross the DCI link – increasing the investment costs of the DCI

RRPP + hub-and-spoke solution

The RRPP solution and the hub-and-spoke solution each have disadvantages: the RRPP solution can cover only a small scope but provide high HA performance, and the hub-and-spoke solution provides a simple architecture, covers a large scope, delivers high scalability, but provides relatively low HA performance. However, you may find scenarios where you need to combine the two solutions to form an RRPP + hub-and-spoke solution. More specifically, you can use RRPP to connect multiple data centers in the same MAN and use the hub-and-spoke solution to interconnect the RRPP ring networks.

Figure 52 RRPP + hub-and-spoke solution

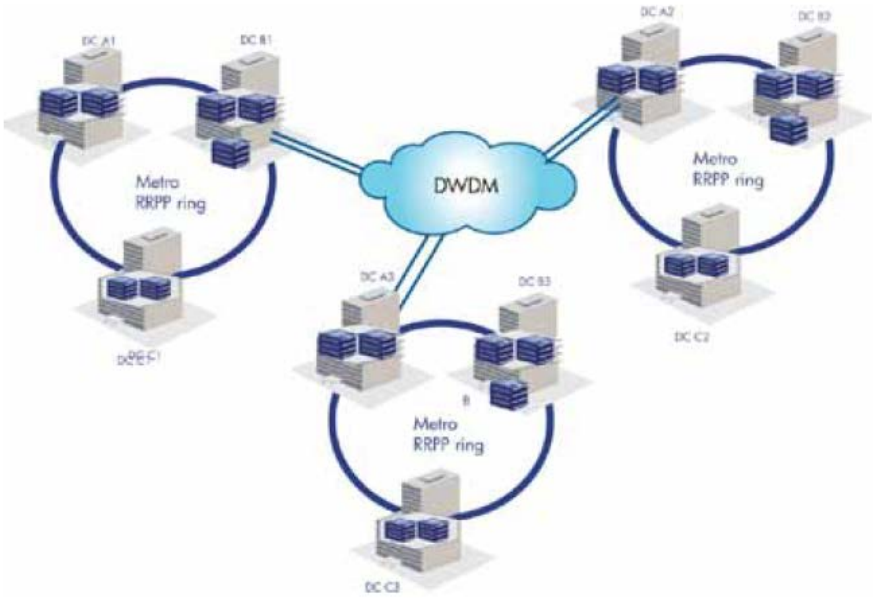

DCI design guidelines for VPLS

VPLS overview

VPLS controls packet forwarding by using two layers of labels and can implement point-to-multipoint DCI connections. With VPLS, you can simulate an Ethernet switch among multiple data centers in the MPLS network to make inter-Layer 2 port forwarding decisions based on MAC address or the combination of MAC address + VLAN ID. A VPLS instance for implementing Layer 2 DCI contains multiple data centers, which are connected to multiple DCI devices (could be viewed as customer-owned PE devices). Data center aggregation layer switches directly communicate with the other aggregation layer switches associated with the VPLS instance.

Such a design is relevant in these two scenarios:

- Enterprise customer owns or manages their own Layer 1 and VPLS network
- Enterprise customers acquire a Layer 1 or Layer 2 type of service from a provider, and VPLS is run between the enterprise PE devices at the edge of the provider's cloud

Figure 53 Layer 2 DCI solution for interconnecting two data centers by using MPLS and VPLS

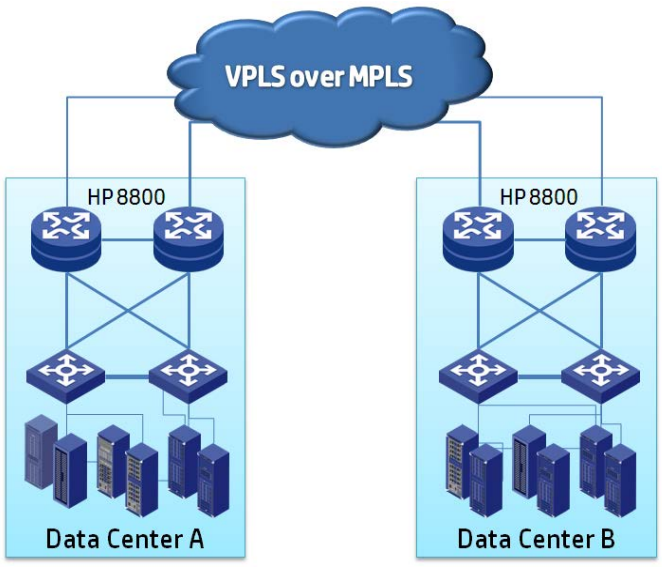

Introduction to smart link: Router-based solution

To implement dual-homing from the aggregation layer to routers, which do not support IRF, you can use the smart link feature to connect to the routers through the links of a smart link group. If a router or an aggregation layer to router link fails, smart link can implement 50-millisecond convergence, and rapidly switch the traffic to the other link.

Figure 54 Smart link solution for dual-homing CEs

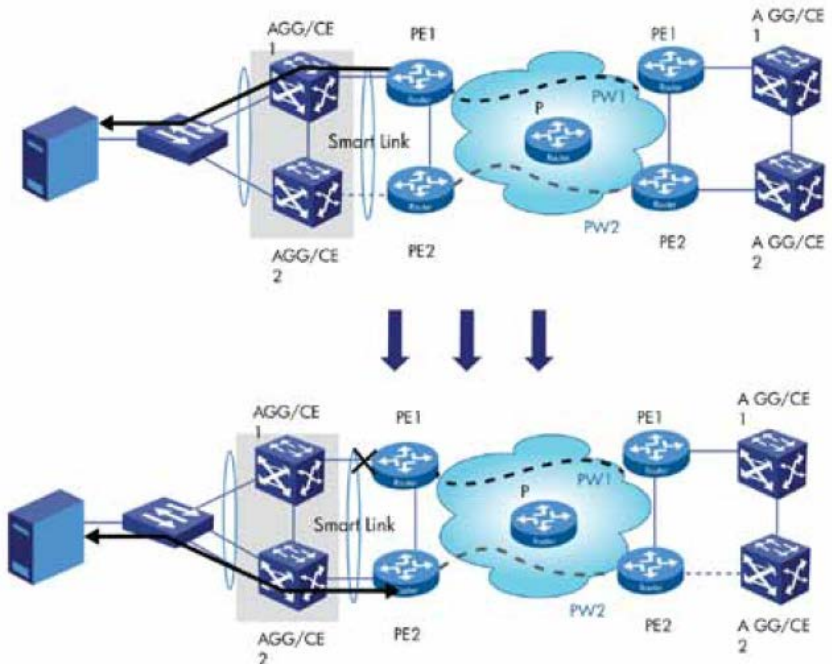

Interconnecting multiple data centers using VPLS + smart Link

In the VPLS network, consider the following two factors with caution:

- Some MPLS networks do not support Ethernet interfaces. Therefore, in those scenarios you must use routers to connect to the network.
- To improve the HA performance, you must implement dual-homing for the routers.

Considering the two factors, HP recommends that you use the VPLS + smart link solution.

Figure 55 VPLS + smart link solution

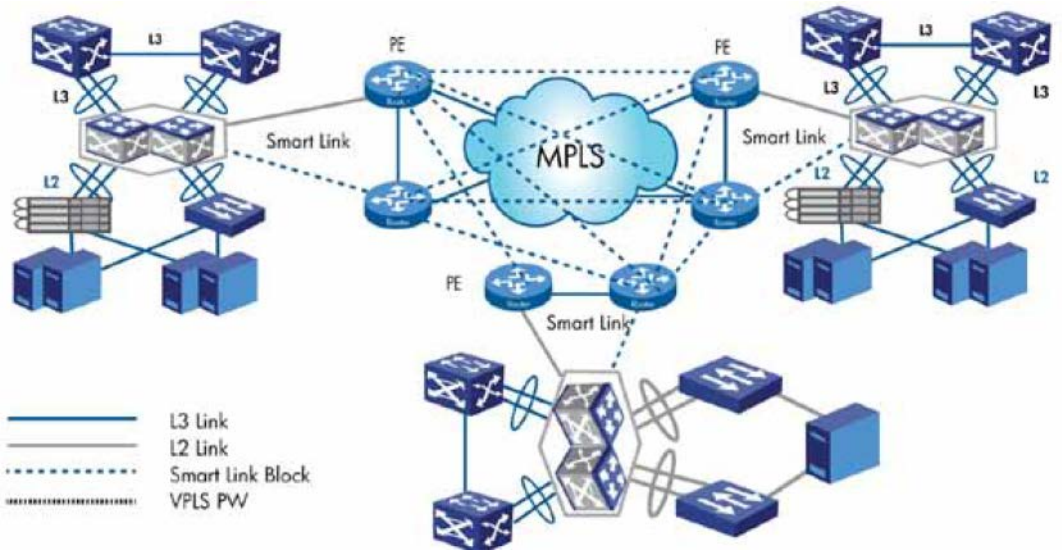

In the VPLS + smart link solution, use the existing MPLS network as the core network, configure two routers (could be viewed as customer-owned PE devices) for each data center, and configure the aggregation layer devices (IRF-capable devices) of each data center. Configure smart link on the aggregation layer to implement 50-millisecond convergence for the two links connecting to the routers.

Set up full-mesh connections for the routers. Note that, the HA performance of the VPLS network depends on the HA design of the MPLS network. Therefore, if the customer owns the MPLS network, HP recommends that you simplify the architecture to facilitate the traffic path planning.

The VPLS + smart link solution delivers the following benefits:

- High standardization level, interoperability and compatibility with most MPLS networks
- Fast switchover of smart link, which improves the HA performance of the system

The VPLS + smart link solution has the following disadvantages:

- The routers do not support IRF. As a result, there are many routers and a large volume of broadcast traffic in the network
- Smart link cannot sense the link failures at the outbound sides of the router PEs and may result in black hole routes

Interconnecting multiple data centers using VPLS + IRF + LACP

Suppose a VPLS network provides the following conditions:

- The MPLS network supports Ethernet interfaces. Therefore, you can use dedicated IRF-enabled

switches (could be viewed as customer-owned PE devices) for connecting to the MPLS network

- The users need dual-homing aggregation layer devices to improve the HA performance

In this case, HP recommends that you use the VPLS + IRF+ LACP solution.

Figure 56 VPLS + IRF + LACP solution

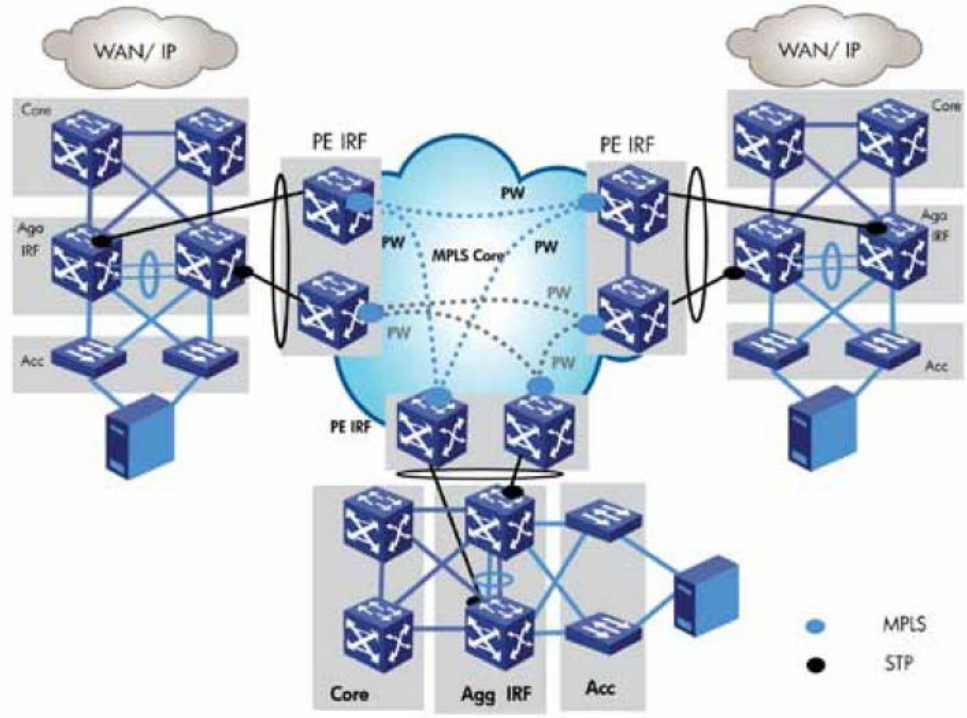

In the VPLS + IRF + LACP solution, use the existing MPLS network as the core network, configure two IRF-capable switches acting as the dedicated DCI devices for each data center (could be viewed as customer-owned PE devices), and configure the IRF-capable aggregation layer devices of each data center. Configure link aggregation between the aggregation layer and dedicated DCI devices to implement millisecond-level convergence for the failure of any node or any link between the layers.

Set up full-mesh connections for the dedicated DCI devices. Note that, the HA performance of the VPLS network depends on the HA design of the MPLS network. Therefore, if the customer owns the MPLS network, HP recommends that you simplify the architecture to facilitate the traffic path planning.

The VPLS + IRF solution delivers the following benefits:

- Per-layer IRF design, which has a simple architecture and is easy to maintain
- Per-layer IRF design, which provides high HA performance and implements millisecond-level convergence for the failure of any node or any link
- Implementing IRF on the dedicated DCI devices decreases the number of nodes in the MPLS network and the number of broadcast packets
- Per-layer IRF design implements load sharing for all links, improves the bandwidth utilization, and improves the performance of the whole network

DCI design guidelines for IP-based solution

VPLSoGRE overview

When only IP networks are used for interconnecting data centers, you can use both EVI or VPLS over GRE (VPLSoGRE) to implement Layer 2 DCI. This VPLSoGRE design solution is only relevant in those scenarios, where an enterprise acquires an IP type of service from a provider and VPLSoGRE is run between the enterprise PE devices at the edge of the provider's cloud.

The VPLSoGRE design requires a tunnel that can carry multiple pseudo wires (PWs) between the various dedicated DCI devices at each data center. The tunnel can be an MPLS tunnel or a GRE tunnel. For the MPLS tunnel implementation of VLL or VPLS, the outer label of a packet is an MPLS label. The GRE tunnel implementation of VPLS is VPLS + GRE. Therefore, the network topology and Layer 2 loop prevention of the VPLSoGRE solution are similar to those of the VPLSoMPLS solution.

Figure 57 The VPLSoGRE solution

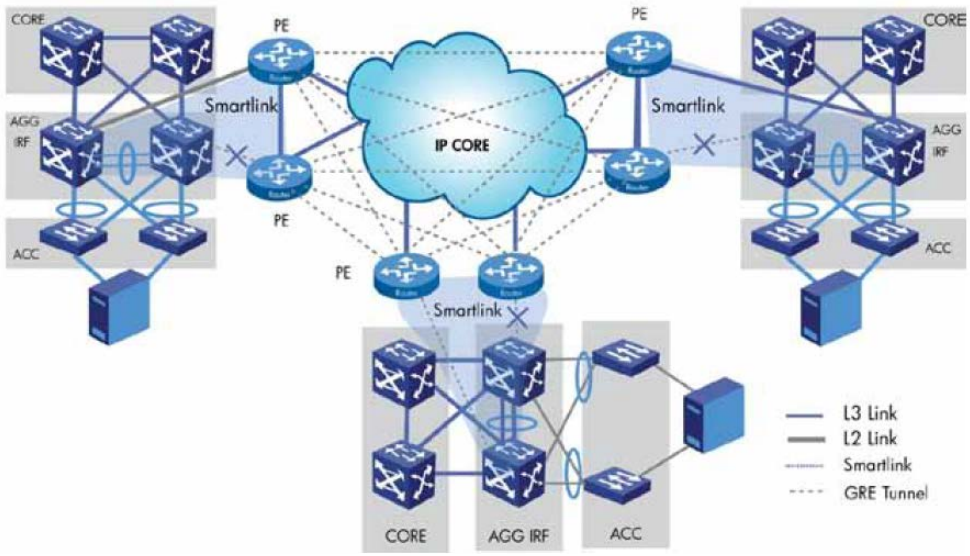

In the VPLSoGRE solution, use the existing IP network as the core network, configure two routers as the dedicated DCI devices (could be viewed as customer-owned PE devices) at each data center, and configure the IRF-capable aggregation layer devices of each data center. Configure smart link on the aggregation layer devices to implement 50-millisecond convergence for the two links connecting to the dedicated DCI devices.

Set up full-mesh GRE tunnel connections for the dedicated DCI devices (can be viewed as customer-owned PE devices).

The VPLSoGRE solution delivers the following benefits:

- High standardization level, powerful compatibility, and compatibility with all IP networks
- Fast switchover of smart link, which improves the HA performance of the system

The VPLSoGRE solution has the following disadvantages:

- The dedicated DCI devices do not support IRF. As a result, there are many dedicated DCI devices and the PW configuration is complicated

The IP network provides low-service quality and you must configure network-wide end-to-end QoS, which can be very difficult.

Summary

EVI is the preferred DCI method from HP and stands out when compared to other vendors like Brocade, Cisco, and Juniper.

Table 3 HP EVI – Unmatched in the market

DCI Feature	Brocade	Juniper	Cisco	HP
Layer 2 Routing Extension	No	No	Yes – OTV	Yes – EVI
Layer 2 Routing Extension Scaling	N/A	N/A	512 VLANS 6 Data centers	4K VLANS 8 Data centers
Layer 2 Routing Extension with Multitenancy	No	No	No	YES EVI & MDC
Multiple MPLS/VPLS-based & Layer 2 DCI on one switch	No	No	No	Yes 12500

For data center clients who have multiple data centers, Cisco and HP are the only vendors that offer a L2 DCI that doesn't rely on MPLS and allows for multiple network instances.

Table 4 HP EVI Advantages vs. Cisco OTV

DCI Feature	VPLS	Cisco OTV	HP EVI
End-to-End Loop-free without STP	√	√	√
Failure Domain Isolation	√	√	√
Multi-pathing and Load-balancing			√
Independent of Infrastructure	MPLS required	Multicast required by Default	√
Multi-DC Interconnectivity	√	√ Up to 6	√ Up to 8
Active/Active Physical Redundancy			√
Large Number of VLANs	256	√ 512	√ 4K
Number of Network Instances		10 per Nexus 7000	32 per MDC

When utilizing DCI solutions other than EVI, select the proper solution by considering the performance, HA, Layer 2 management, and security factors.

Table 5 Traditional DCI performance, HA, Layer 2 management, and security factors

Your network resource	Solution	Performance	HA	Layer 2 management	Broadcast control	Security
Dark fiber or DWDM	RRPP	High	High	High	High	Medium
	Hub-and-spoke	High	Medium	High	High	Medium
MPLS core network	VPLS + IRF	Medium	Medium	Medium	Medium	Low
	VPLS + smart Link	Medium	Medium	Medium	Low	Low
IP core network	VPLSoGRE	Medium	Low	Medium	Low	Low

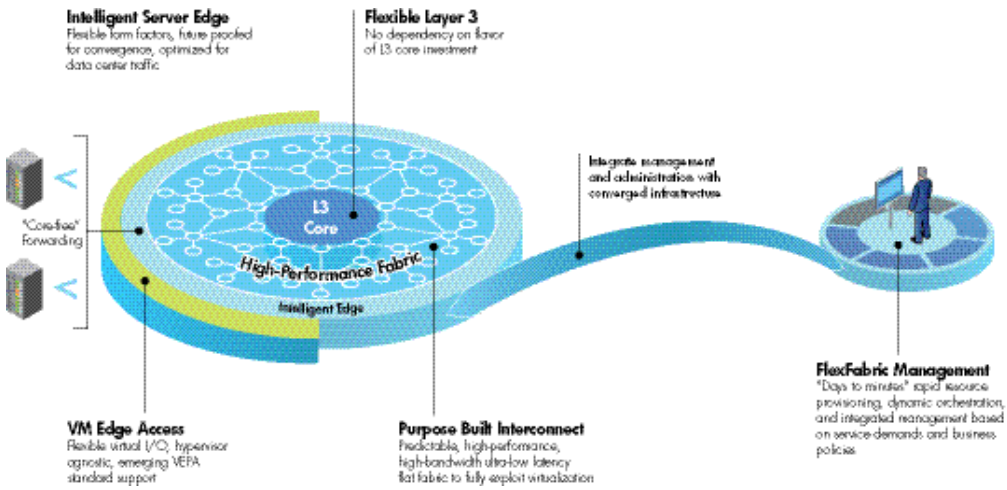

HP Data Center Networking Portfolio

Data center solutions portfolio overview

- **HP 12500 Switch Series:**

  Large core/data center switching platforms with future-proof backplane scalability and the ability to deliver more than 13.32 terabits of high-performance switching capacity and to aggregate up to 576 10GbE or 864 GbE ports

- **HP 11900 Switch Series**

  The FlexFabric 11900 Series are data center switches focused on aggregation and small core deployments with high density 10GbE and 40GbE interfaces and next generation data center features including TRILL, DCB/FCoE along with Comware V7 innovations.

- **HP 5900AF and 5920AF Switch Series:**

  Family of high-density 10GbE and ultra-low latency ToR switches which are ideally suited for deployment at the server access layer of large enterprise data centers. These switches support TRILL, and are Virtual Ethernet Port Aggregator (VEPA), and FCoE ready for virtualized networks and data center convergence

- **HP 5830AF Switch Series:**

  High-density, GbE and 10GbE (uplinks) data center ToR switches with deep packet buffers to eliminate the network congestion at the I/O with heavy use of server virtualization

- **HP 5800 and 5820 Switch Series:**

  High-density, low-latency GbE and 10GbE, FCoE-capable Ethernet access layer/ToR switches that deliver industry-leading price/ performance value

- **HP HSR6800 and 8800 Router Series:**

  Data center WAN access and Layer 3 routers that support IRF, BGP, MPLS, security, QoS, and optional network service module integration

- **HP Virtual Connect:**

  VC modules provide a better way for IT to work together and offer benefits to all parties involved, without the traditional compromises. It is simply the best way to connect servers to network LANs and SANs for the lowest costs and least amount of power. HP continues to expand this technology and its capabilities across ProLiant, Integrity and Storage product lines. VC can simplify and converge your server edge connections, integrate into any standards based

networking infrastructure and reduce complexity while cutting your costs

- **HP 612x Blade Switch Family:**

  Designed for the HP BladeSystem c-Class enclosure, the HP 612x Blade Switches provide 16 1GbE or 10GbE server downlinks and various options for uplinks, along with 10GbE cross-connects. A robust set of industry-standard Layer 2+ switching functions, QoS metering, security and high-availability features round out this extremely capable blade switch family

  With a variety of connection interfaces, the 612x Blade Switch Family offers excellent investment protection, flexibility, and scalability, as well as ease of deployment and reduced operational expense

  The new 6125 Blade Switch Family has added features which include support of IRF, Comware OS, powerful QoS, and IMC management

- **HP S-Series Security:**

  An enterprise security suite that provides proven, scalable, high-performance appliance-based Intrusion Prevention Systems, security management, and security subscription services based on HP TippingPoint technology that unifies physical and virtualization security in a common high-performance framework

- **HP IMC:**

  A comprehensive platform that integrates management of physical and virtual network resources and provides full FCAPS management functionality for IT infrastructures

Core/distribution switches

HP 12500 Switch Series

Figure 58 HP 12500 Switch Series

These data center core/distribution switches provide unprecedented levels of performance,

scalability, HA, density, and flexible deployment options validated by independent testing. They drive down data center operations costs while enabling new service levels and delivering the resiliency and low latency required for mission-critical networking.

The 12500 Switch Series is ideal for organizations contemplating large-scale data center or campus consolidations, business continuity and disaster recovery sites, MAN deployments, and other applications requiring a robust, high-performance switching platform.

Key features and benefits

- 13.32 Tbps of high-performance switching capacity with more than 4.3 billion packets per second of forwarding performance
- Support for up to 576 10GbE or 864 GbE ports
- A future-proof design able to accommodate 40/100 GbE to support emerging unified network requirements such as FCoE
- High-availability, non-blocking design for zero service interruption; control plane features 1+1 redundancy
- Modern, energy-efficient architecture that dramatically reduces space, power, and cooling requirements

More information regarding the HP 12500 Switch Series can be found at the HP 12500 product page.

HP 11900 Switch Series

Figure 59 HP 11900 Switch Series

The HP FlexFabric 11900 switch series is a high performance data center aggregation switch with line-rate, high density 10/ 40 GbE. This switch is SDN capable and has full layer 2 and layer 3 features. The initial offering includes an 8 slot chassis and supports advanced features like IRF, SPB, and TRILL to build large, resilient switching fabrics.

The 11900 switch enables network and storage convergence, and has very low latency and low energy consumption. The 11900 Switch Series is a key addition to the FlexFabric portfolio that is ideal for aggregation and cost effective end of row and small core deployments.

Key features and benefits

- 7.7 Tbps of high-performance switching capacity with up to 5.7 billion packets per second forwarding performance
- Support for up to 384 non-blocking 10GbE ports, 192 non-blocking 10GbE Base-T ports, or 64 non-blocking 40GbE ports
- 8 slot chassis with front to back airflow optimized for hot aisle cold aisle environments
- Simplify the data center complexity with flatter layer 2 networks and convergence of network and storage infrastructure and enables you to simplify data center architectures, and deliver reduced footprint, power and cost
- Provides the best of both worlds by adding mature virtualization technology like IRF with standards based protocols like SPB and TRILL
- With front-to-back airflow, redundant and hot swappable management, power, fabric and cooling modules, the 11900 is energy efficient and delivers latency as low as 3 us

More information regarding the HP 11900 Switch Series can be found at the HP switch product page.

Access layer/edge switches

HP 5900AF and 5920AF Switch Series

Figure 60 HP 5900/5920 Switch Series

With the increase in virtualized applications and server-to-server traffic, customers now require ToR switch innovations that will meet their needs for higher-performance server connectivity, convergence of Ethernet and storage traffic, capability to handle virtual environments, and ultra-low latency all in a single device. The two new models, 5900AF and 5920AF switches, are ideally suited for deployment at the server access layer in large virtualized enterprise data centers. They are also designed for deployment at the core layer of data centers at medium-size enterprises.

The HP 5900AF and 5920AF series usher in the arrival of the world's first virtualized server access network. These 10GbE ToR switches are built on open industry standards and set new benchmarks for performance, low latency, reliability, scalability, and greener data centers with a simpler network architecture. These switches deliver high 10GbE port density, deep packet buffers, and ultra-low latency (~1 microsecond) performance, with a choice of front to back (port side to power side) or back to front (power side to port side) airflow. These switches also future-proof network investments by providing support for full Layer 2 and Layer 3, IPv4 and IPv6 dual-stack support. The HP 5900AF and 5920AF ToR series are fully ready for virtualized data center and cloud computing environments.

Key features and benefits

- Industry-leading HP IRF technology radically simplifies the architecture of server access networks and enables massive scalability—this provides up to 300% higher scalability as

compared to other ToR products in the market

- Industry's only support for multiple use cases in a single ToR switch—up to 50 percent device reduction ––Ultra-low-latency (~1 microsecond) IP switching with ALL features enabled ––Server-edge storage + Ethernet convergence with DCB today and FCoE (future)
- Industry's only ToR switch in its class with IPv6 routing and IPv4/IPv6 dual-stack support for advanced future networks and future-proof investment
- TRILL supported – combines the simplicity and flexibility of Layer 2 switching with the stability, scalability, and rapid convergence capability of Layer 3 routing
- VEPA ready for virtualized networks and data center convergence
- 48 port 10GbE/1 Gbps ToR options with 40GbE uplinks for high performance and scalable networking with full L2/L3 features
- New 48 port 10GbE Base-T model with 40GbE uplinks for high performance and scalable networking with full L2/L3 features
- Lower OpEx and greener data centers with reversible air flow and advanced chassis power management
- Full Layer 2 and Layer 3, IPv4 and IPv6 dual-stack support
- Choice of front to back (port side to power side) or back to front (power side to port side) airflow with dual fan trays and redundant internal power supplies

More information regarding the HP 5900AF and 5920AF Switch Series can be found at the HP 5900 and 5920 product pages.

HP 5830AF Switch Series

Figure 61 HP 5830AF Switch Series

The HP 5830AF Switch Series offers high-density, GbE and 10GbE (uplinks) data center ToR, as well as campus, and service provider aggregation or access deployments. The HP 5830AF-48G Switch w/1 Interface Slot provides 48 GbE, and up to four 10GbE ports. The HP 5830AF-96G Switch delivers 96 GbE ports and ten 10GbE uplinks. The HP 5830AF Switch Series adopts ultra-deep packet buffering (1GB and 3GB for 1RU and 2RU respectively), to eliminate the network congestion at the I/O associated with heavy use of server virtualization, as well as bursty multimedia, storage applications and other critical services. In addition, the HP 5830AF Switch Series is also data center optimized design, supporting redundant hot-swappable power supplies and fans to ensure hardware availability, front-to-back or back-to-front airflow hot and cold aisle isolation requirements, and low power consumption to optimize operating power and reduce operating expenses.

Key features and benefits

- Industry-leading HP IRF technology radically simplifies the architecture of server access networks and enables massive scalability

- Stackable 96 port GbE model with ten 10GbE ports in just 2RU
- Ultra deep packet buffer (1GB/3GB) for virtualized environments that are loss sensitive, or with bursty traffic or just demanding applications in multimedia rich and storage environments
- Choice of front to back (port side to power side) or back to front (power side to port side) airflow with dual fan trays and redundant internal power supplies
- Support for full L2 & L3 features and IPv4 & IPv6 dual stack

More information regarding the HP 5830AF Switch Series can be found at the HP 5830AF product page.

HP 5800 and 5820 Switch Series

Figure 62 HP 5800 Switch Series

Figure 63 HP 5820 Switch Series

These unique flex-chassis switches can function as a modular chassis, as well as a fixed-form-factor stackable switch, providing the flexibility, scalability, and reliability of a modular platform and the ease of use of a stackable solution. The 5800 and 5820 Switch Series deliver line-rate connectivity with support for a combination of 10GbE, 10GbE FCoE, GbE, and 8 Gbps Fibre Channel ports which facilitates investment-protection.

The high-performance switching and high-density capacity of the HP 5800 and 5820 Switch Series

provide support for a variety of flexible deployment options. The switches can be used for ToR, building or department core, Layer 3 10GbE aggregation switches in a campus network, at the network access layer as a GbE PoE switches with 10 GbE uplinks, or as FCoE-enabled data center server access to a 12500 core switch.

Key features and benefits

- Wire-speed, line-rate performance on all ports for IPv4 and IPv6 traffic; Layer 2/3 routing capabilities (IPv6)
- High-performance, high-density Gigabit and 10GbE connectivity; two front-facing expansion slots (2RU 5800-48G/5820-14XG) can dramatically increase 10GbE or GbE port density for even greater deployment flexibility
- Flexible, FCoE module that provides cost-effective Fibre Channel (FC) storage and server network I/O consolidation with better interoperability with existing FC SANs and 1G/10G dual-speed support to enable cost-effective migration
- Next-generation traffic prioritization for converged traffic, including advanced policy-based CoS/QoS, eight priority queues per port, committed access rates, bandwidth limiting, and filtering
- High-availability architecture based on IRF technology
- Redundant, fully hot-swappable power supplies and fans typically found in modular core platforms

More information regarding the HP 5800 and 5820 Switch Series can be found at the HP 5800 product page, the 5820 product page.

HP data center routing

HP HSR6800 Router Series

Figure 64 HP HSR6800 Router Series

The HP HSR6800 Router Series is a portfolio of high-performance WAN services routers, ideal for large-scale data center and campus WAN networks.

The HP HSR6800 features an independently-developed high-performance communication processor (Apollo), and an advanced multi-core distributed processing architecture that scales up to 420 Mpps forwarding and up to 2 Tbps switching capacity. They deliver robust routing, security, full layer 2 switching, traffic analysis capabilities, continuous bandwidth provisioning capability, carrier-level high availability, and-high density 10 GbE (and 40/100 GbE-ready) WAN interface options, all integrated in a single high-performance routing platform.

In addition, the HSR6800 Router Series are the first aggregation routers in the industry to support system virtualization by taking advantage of HPs innovative Intelligent Resilient Framework (IRF) technology (available - Q4 2013).

Key features and benefits

- **Multi-core High Performance Advanced Distributed Service Architecture:**

  The HP HSR6800 adopts multi-core CPUs and fully distributed independent routing and service engines. All engines have separate control and service planes to avoid interference and to ensure service continuity during an active/standby switchover. Each engine can independently process NAT, IPsec, GRE, NetStream services in a distributed way, improving service availability and system processing capability
- **Intelligent Resilient Framework (IRF) (available - Q4 2013):**

  The HP HSR6800 uses IRF to implement WAN aggregation virtualization. IRF enables two

HSR6800 routers to form a single device. This technology improves device performance, speeds up whole-network convergence, simplifies configuration, reduces O&M costs, and improves network availability.

- **Industry-leading design:**

  The HP HSR6800 provides a wide range of line cards including FIPs and SAPs to meet various network requirements.

- **Large capacity aggregation:**

  The HP HSR6800 can provide large-capacity, high-density narrowband aggregation and access. For example, the HSR6808 can provide up to 32 channelized 155 M POS interfaces that can also be channelized to E3/T3 and E1/T1/DS0 interfaces. The HSR6808 can suport up to1512 wire-speed E1s, or 2016 wire-speed T1s

  It supports the PPPoE server. Each line-card allows for the access of 8K PPPoE broadband users, so the HSR6808 chassis can support 64K PPPoE broadband users

  It supports high-density GE or 10GE ports. A single SAP card can provide 48 GE ports or four 10GE ports, fully satisfying MSTP link aggregation requirements from different users

- **Secure cloud access:**

  The HP HSR6800 supports multiple secure access technologies. You can use GRE and IPsec to connect branches to the headquarters, use L2TP and IPsec to connect SOHO users, and use IPsec to secure MSTP links. In addition, you can use DVPN, GRE, and IPsec RRI to build dynamic VPN networks. DVPN solves challenges facing traditional VPN technologies, such as dynamic address access, node configuration simplification, full-mesh, and user authentication.

  It also provides data encryption/decryption, high-density concurrent access, simplified routing deployment, branch NAT traversal, dynamic addressing, MPLS VPN integration, and hub/VAM server backup and integration functions to ensure secure cloud access

- **Hardware PPP multi-link bundling:**

  The HP HSR6800 supports two types of high-speed CPOS interface modules (HIM-CL2P and HIM-CL1P) that support hardware Multilink PPP (MP) bundling. When the HP HSR6800 serves as an aggregation node of a WAN, you can implement PPP multi-link bundling on the downlink E1 or T1 through the CPOS interfaces. You can also implement reassembly and fragmentation of MP packets without affecting traffic forwarding

  Each FIP-210 line-card can implement 10 groups of 12E1 MP bundling and 14 groups of 12T1s MP bundling at wire-speed, thus providing sufficient bandwidth for narrowband aggregation networks

- **Abundant interface types. The HP HSR6800 supports:**

  10 GbE, GbE, 100 Mbps Ethernet ports, and WAN ports such as POS OC-48/OC-12/OC-3, cPOS OC-3 (channelized to E3/T3 or E1/T1), ATM OC-3, E3/T3, E1/T1, and serial

  Serves as both aggregation devices for WANs and access devices for LANs, allowing for building a flat network and simplifying network topology

- **All-round security protection. The HP HSR6800 provides in depth security protection functions:**

  All-round firewall functions: Packet filtering and stateful firewalls, filtering attack packets and generating filtering logs. The HSR6800 supports multiple attack protection mechanisms: ARP attack protection technologies, Single-packet attack protection function, Scanning attack protection function, Flood attack protection function, Blacklist function, Traffic statistics function, URL filtering

- **Carrier-class high availability:**

  The HP HSR6800 supports dual-MPUs, redundant power modules, distributed service

architecture. With two MPUs equipped, the system allows you to perform an active-standby switchover, without interrupting the ongoing data forwarding or service processing on the line-cards, thus ensuring reliability. All the line-cards and modules on the HP HSR6800 are hot-swappable, without interrupting one another. The HSR6802, HSR6804 and adopts “1+1” power module redundancy and the HSR6808 provides four power modules that support multiple redundancy modes

The HP HSR6800 supports the following software high availability features: Hotfix, □ OSPF/IS IS/BGP/LDP NSR, BFD & NQA, MPLS TE FRR, OSPF/IS-IS IP FRR, VRRP and VRRP load balancing mode, OSPF/IS-IS/BGP/MPLS LDP/MPLS RSVP-TE graceful restart (GR), IGP fast routing convergence, RRPP

More information regarding the HP HSR6800 Routers can be found at the HP router product page.

HP 8800 Router Series

Figure 65 HP 8800 Routers Series

The HP 8800 routers series are components of the FlexFabric and FlexCampus modules of the FlexNetwork Architecture. It features a distributed high-performance network processor as well as high-capacity crossbar nonblocking switching technology that delivers high performance and flexibility. A distributed QoS control unit provides end-to-end service with granular control. The routers' distributed operation, administration, and maintenance detection engines implement fault detection within 30 ms to provide uninterrupted core services. These innovative technologies, paired with the QoS control mechanism, deliver smooth operation and HA of multiple services within HP FlexNetwork. The 8800 routers are commonly deployed in IP backbone networks, IP MANs, the core or convergence layers of large IP networks. Offering high forwarding performance and abundant services, the 8800 series delivers on routing performance and flexibility.

Key features and benefits

- **Innovative Architecture:**

The 8800 integrates the crossbar non-blocking switching technology into its 10G NP platform, thus satisfying users' requirements on application processing performance and capacity.

The distributed forwarding engine enables the use of hardware in processing applications on the Line Processing Unit (LPUs)

The HP 8800's LPUs accommodate NP engine, QoS engine and table search engine at the same time

- **Dedicated QoS Engine:**

  The built-in QoS engine supports hierarchical QoS (HQoS), enabling QoS scheduling based on ports, individual users, user groups, and user applications

  QoS engine provides a variety of functions including multiple queuing mechanisms (priority queuing, low latency queuing, custom queuing, weighted fair queuing, class-based weighted fair queuing), congestion avoidance, traffic policing, traffic shaping and priority marking, ensuring differentiated services (bandwidth, delay and jitter) for different users and applications

  QoS engine also provides a powerful packet buffer capability of 200 ms, which can effectively solve the packet loss problem caused by burst traffic on the network.

  The Crossbar switching fabric on the HP 8800 supports virtual output queuing (VOQ) and end-to-end flow control technologies, which can effectively avoid head of line (HOL) blocking and implement differentiated services on the Crossbar

- **Comprehensive IPv6 Support:**

  The 8800 supports the complete IPv6 protocol suite, including IPv6 static routing, RIPng, OSPFv3, IS-ISv6 and BGP4+. It also supports the IPv4-to-IPv6 transition technologies, such as IPv6 manual tunnels, 6to4 tunnels, ISATAP tunnels, GRE tunnels, and IPv4-compatible automatic tunnels

- **MPLS VPN:**

  HP 8800 supports distributed Layer 3 MPLS VPN, VLL and VPLS/H-VPLS services, and high-performance P/PE applications. It provides a high-quality and multi-layer MPLS VPN solution

- **NAT support:**

  HP 8800 supports distributed and centralized NAT, and supports load sharing over multiple NAT interface card to enhance the entire NAT performance. The NAT interface card for the 8800 delivers abundant features, including NAT, NAPT, internal servers, application layer gateway (ALG), blacklist, and multi-instance NAT

- **Security Features:**

  HP 8800 supports a variety of security cards, such as firewall card, IPS card, SSL VPN card, application control gateway (ACG) card, and LB card

- **Dedicated OAM Engine:**

  Monitors network operation and ensures non-interrupted operation with the capability of 30ms failure detection and 20ms switchover. The detection plane, which is independent of the control plane and the forwarding plane, ensures carrier-class reliability for users

More information regarding the HP 8800 Routers can be found at the HP router product page.

HP Virtual Connect for BladeSystem —the simplest, most flexible way to connect servers to any network

Figure 66 HP Virtual Connect Module

HP VC portfolio is a collection of Ethernet, Fibre Channel and converged fabric interconnect modules and firmware that provide an ideal alternative to traditional switches and patch panels installed in the interconnect bays of HP BladeSystem c-Class enclosures.

VC simplifies change by streamlining the provisioning and maintenance of blade infrastructure. It allows 'silo'-ed data center teams to more work together more efficiently and ultimately enable more agile, responsive IT services. It ultimately allows physical compute resources to be added, moved or changed in minutes instead of days or weeks. Server teams can independently manage these changes with just a few mouse clicks or commands. They no longer need to choreograph every add, move, or change to servers with the network and storage administrators. Without VC, they WAIT for everyone to finish the pieces they are responsible for – with VC, they simply move forward with connectivity.

VC simplifies connections by converging and integrating networking resources better than any other product available. With converged adapters built-into HP server blades and fully converged and integrated interconnect modules, VC eliminates unmatched amounts of traditional networking components required to connect servers to data and storage networks.

HP's innovative wire-once connection management is the key technology that simplifies server connectivity and the ability to move, add and change servers in minutes instead of hours or days. It sharply contrasts with traditional switch environments where connectivity changes are hampered by inflexible, hard-wired connections requiring extensive coordination among data center server, network and storage teams. Wire-once technology allows 'silo'-ed data center teams to efficiently "bulk-provision" enclosure connectivity as pooled connection resources.

Server teams are able to create server profiles containing virtual network and storage addresses and connection information and assign them to each server bay in an HP BladeSystem enclosure. This gives server admins the ability to independently move, add or change servers HP server blades as they automatically assume the virtual identification and connection parameters of their assigned bay to connect to the network resource pool. With VC, you wire once, and then have the flexibility to change HP ProLiant and Integrity server blades without re-cabling or reconfiguring upstream network infrastructure. It's all transparent to the upstream networks and greatly simplifies network management. This technology also provides the essential foundation for fully automating network services which are essential to any cloud-ready environment.

VC helps to rapidly deploy servers. The server profiles are set up in VC by pre-populating the LAN and SAN connectivity information. These profiles are assigned to enclosure bays – whether they are populated or empty. A server may be installed at any time and connect to data and storage networks with the pre-assigned profile information.

VC also assists with rapid system recovery, or workload reassignment. If for any reason, you need

to move an entire server workload to a different server, VC allows you to do that transparently to the LAN and SAN. The original server's entire network identity, boot from SAN definitions and other connection parameters can be moved to any server bay within a VC domain. In a stacked, multi-enclosure environment, you have the flexibility to move the workloads anywhere within a VC domain of up to 4 enclosures with 64 server bays.

When you replace a physical server, the network identity assigned to the server bay over-writes factory worldwide name (WWN) and MAC of the new server so the replacement is completely transparent to the network. Once removed, the replaced server defaults back to its factory assigned identity.

VC Flex-10 has led the industry in reducing I/O sprawl and simplifying the way servers are connected to data and storage networks. Flex-10 allows you to consolidate Ethernet connections and converge data and storage fabrics onto 10Gb server connections. It allows you to carve up 10Gb bandwidth into 4 separate virtual connections resulting in significant cable, adapter and switch port reduction. Flex-10 also supports QoS on each of those connections by allowing you to dynamically adjust speed in 100Mb increments ensuring that each connected application receives the bandwidth it requires. You no longer need to over-provision or under-provision bandwidth.

Figure 67 Flex-10 device

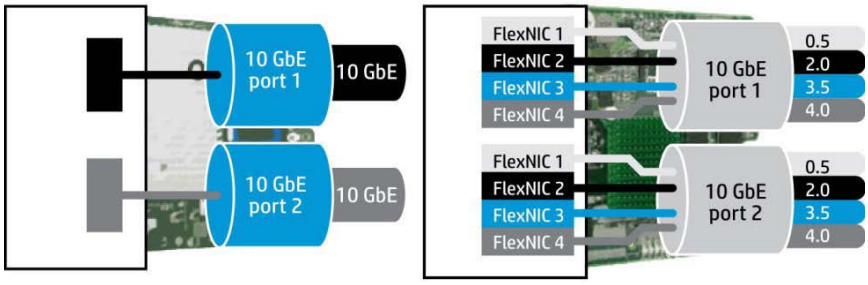

HP is also building storage network protocols into its VC technology that will combine the existing Flex-10 capabilities with DCB and FCoE technologies. This will allow customers to use a single VC server connection to access both storage and server networks. HP FlexFabric uses industry standards to converge storage, server, and network resources. It provides an orchestration and management layer to provide virtualization-enabled networking and rapid provisioning. FlexFabric offers a wire-once data center that can respond to application and workload mobility. Network connections can move with workloads across or even between data centers.

HP VC technology offers significant advantages over other 10GbE devices that provide large bandwidth, but do not provide segmentation. The ability to adjust transmit bandwidth by partitioning data flow makes 10GbE more cost effective and easier to manage. It is easier to aggregate multiple 1GbE data flows and fully utilize the 10GbE bandwidth. The fact that VC is hardware-based means that multiple FlexNICs are added without the additional processor overhead or latency associated with server virtualization. Significant infrastructure savings can also be realized because additional server network mezzanine cards and associated interconnect modules may not be required.

There are several advanced features available when using the VC FlexFabric module and adapter technology when enabling FCoE. One of these advanced features is support for Flex-FC. This feature allows for higher-level software tools such as Insight Dynamics and Orchestration to assign up to 16 WWNs to an individual FCoE Physical Function (Each 10GbE port has four Flex-10 physical functions [PFs]) and map all those WWNs to the VC SAN Fabric. In addition, it allows up to eight SAN boot targets to each FCoE PF.

Figure 68 Flex-10 with Fibre Channel or accelerated iSCSI

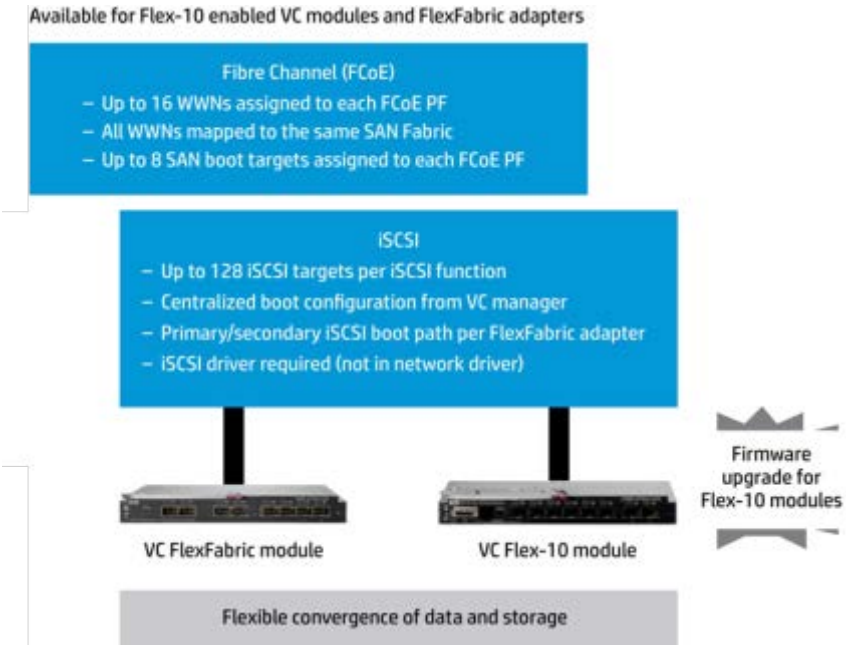

Also, iSCSI boot and offload support on the FlexFabric adapters enables TCP/IP and iSCSI processing offload as well as iSCSI boot functions of these adapters when connecting to either Flex-10 or FlexFabric modules. One advanced iSCSI feature is support for up to 128 iSCSI boot targets per iSCSI function. Volume Connect Manager (VCM) provides centralized boot configuration management of iSCSI boot parameters, and provides primary and secondary boot path management. Although it requires a separate iSCSI driver, it does not require an NIC driver and will drastically reduce CPU utilization by offloading processing to the adapter.

FlexFabric helps modify the tasks of provisioning SANs

The old way:

- Connect servers individually to LAN and SAN
- Require LAN and SAN coordination to add, move, or replace a server
- Wait for everyone to finish their pieces

The new way:

- Pool server connections to LAN and SAN
- System administrator can add, move, or replace servers transparent to LAN and SAN

HP VC offers Direct-Attach Fibre Channel for 3PAR Storage Systems – using Flat SAN technology to eliminate traditionally complex multi-tier SAN infrastructure and flattening it. HP VC FlexFabric Modules can be directly connected using a Fibre cable to a 3PAR SAN Storage System which reduces latency and the cost of the SAN Fabric. This solution removes complexity between compute and storage and simplifies the entire Fibre Channel SAN fabric layer by eliminating Fibre Channel switches and associated cables. With a reduction of latency by 55% and time to provision SAN infrastructure , and a cost reduction of up to 50%, there is less equipment to manage and maintain, and less time is needed to make changes to the infrastructure.

Current Virtual Connect modules include:

- **Virtual Connect FlexFabric 10Gb/24-Port Module:**

  Each VC FlexFabric module provides 16 adjustable 10GbE downlink connections to server blade NICs and FlexFabric adaptors, and eight uplinks for connections to upstream Ethernet and Fibre Channel switches (2 of the uplink ports can also be used as 10GbE cross connect ports to other VC FlexFabric module). Each 10GbE downlink server connections supports up to 3 adjustable FlexNICs and 1 adjustable FlexHBA or 4 adjustable FlexNICs. Each connection can be configured from 100 Mb up to 10 Gb, allowing just the right amount of network bandwidth based on your application needs
- **Virtual Connect Flex-10 10Gb Ethernet Module:**

  Each VC Flex-10 10Gb Ethernet module provides 16 adjustable 10GbE downlink connections to server blade NICs and FlexFabric adaptors, and eight uplinks for connections to upstream Ethernet switches (2 of the uplink ports can also be used as 10GbE cross connect ports to other VC Flex-10 modules). The HP Virtual Connect Flex-10 10Gb Ethernet Module allows administrators to fine-tune network bandwidth at the server edge by allocating each 10 Gb server port into four independent physical FlexNIC server connections. Each FlexNIC can be configured from 100 Mb up to 10 Gb, allowing just the right amount of network bandwidth based on your application needs
- **Virtual Connect 10Gb-F Ethernet Module:**

  Each VC 10Gb-F Ethernet module provides 16 1GbE downlink connections to server blade NICs and FlexFabric adaptors, 4 10/100/1000 Base-T uplink ports, 2 1GbE SFP uplink ports, 2 10GbE XFP uplinks, and 1 10GbE cross connect port
- **Virtual Connect 20-Port Fibre Channel Module:**

  This Fibre Channel module provides 16 internal 8Gb server downlinks presented as F-Ports, and 4 external 8Gb uplinks presented as N-Ports. The Virtual Connect 20-Port Fibre Channel Module supports NPIV for VMs which provides “any-to-any” provisioning between storage and VMs on server blades and provides the best cost per port of the VC Fibre Channel modules
- **Virtual Connect 24-Port Fibre Channel Module:**

  This Fibre Channel module provides 16 internal 8Gb server downlinks presented as F-Ports, and 8 external 8Gb uplinks presented as N-Ports. The Virtual Connect 24-Port Fibre Channel Module supports NPIV for VMs which provides “any-to-any” provisioning between storage and VMs on server blades and provides the greatest port density of the VC Fibre Channel modules

More information regarding the HP Virtual Connect Modules can be found at the HP VC product page.

HP 612x Blade Switch Family

The HP 6120 Blade Switch Series provides customers using the HP c3000 and c7000 enclosures with new options for connecting server and storage blades to the enterprise network. There are currently two models available:

- HP 6120G/XG Blade Switch
- HP 6120XG Blade Switch

These blade switches are managed switches that are easily configured like so many other HP fixed port and modular switches. In general, a managed switch is a network device that can be typically configured and monitored using a number of interfaces. For instance, the 6120 Blade Switch Series can be managed using the Command Line Interface (CLI) through a console port, telnet, or SSH session. These switches can also be managed using the web browser UI, or using SNMP applications such as PCM.

The 6120 Blade Switch Series allows the network administrator to easily deploy and provision network access to server and storage blades that are co-resident with the blade switch. Despite the small footprint, the 6120 Blade Switch Series can support network access for a variety of applications, even those requiring high bandwidth such as video streaming. A robust set of industry-standard Layer 2+ switching functions, QoS metering, security and high-availability features are supported on these extremely capable blade switches.

Key features and benefits

- **Integrated networking solutions:**

  The 6120 Blade Switch Series are integrated solutions with robust Layer 2+ switching capabilities. The blade switches offer 1GbE and 10GbE external port connectivity for uplinks. Along with the c3000 or c7000 enclosure, blade systems allow the data center manager to simplify cable deployments and reduce the overall cost of ownership

- **Investment protection:**

  The external ports of the 6120G/XG Blade Switch provides flexible 1GbE and10GbE connectivity options over copper and fiber links. Similarly, the 6120XG Blade Switch provides all 10GbE connectivity options over copper and fiber links. The lifetime warranty of these blade switches, just as with many HP switches and routers, lowers the data center's long term operational costs

- **Flexible network architecture:**

  The 6120 Blade Switch Series reduces data center setup complexity and provide choice and flexibility when deploying either blade switch and its accompanying enclosure using a top-of-rack or end-of-row server connectivity solution

- **Enhanced security:**

  The 6120 Blade Switch Series supports end-user security through any of three popular authentication methods; 802.1X, Web, and MAC authentication. These authentication methods are deployed using a centralized user directory system that is accessed using the RADIUS or TACACS+ protocol

  Another popular security feature supported by the 6120 Blade Switch Series is port security. On a per-port basis, you can configure security measures to block unauthorized devices, and to send notice of security violations. Once port security is configured, you can then monitor the network for security violations. For instance, you can configure specific MAC addresses of upstream switches that are allowed to connect to the external ports of the blade switches

  Management access security, which limits access by switch administrators, can be implemented using a variety of methods including SSL for secure web browser access to the web UI, SSH for secure access to the CLI, and the Authorized Managers list to control the source computers that can manage a given blade switch

- **Converged application support:**

  The 6120 Blade Switch Series supports the data-driven Internet Group Management Protocol (IGMP) used for managing multicast video applications, and the Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) protocol commonly used in voice deployments

- **Automated server provisioning:**

  The 6120 Blade Switch Series integrates with the HP Data center Connection Manager (DCM) to simplify server and network provisioning. The DCM application runs on the HP DCM appliance or a specialized 5400zl/8200zl module and assists the data center manager by allowing physical and virtualized servers to be more easily provisioned as they come online in the network

- **Comprehensive network management and diagnostics:**

  The 6120 Blade Switch Series are interconnect modules supported by the c-Class BladeSystem enclosures. These interconnect modules can be managed and monitored using the HP

BladeSystem Onboard Administrator application just like any other server, storage, and interconnect module supported by the enclosures

The typical method for performing the initial network and ongoing network configuration is through the switch web browser or CLI user interfaces. The web browser interface can be launched from the HP BladeSystem Onboard Administrator. The switch CLI can be accessed in a number of ways, This includes using a console port connection to the blade switch, a console port connection to the HP BladeSystem Onboard Administrator's CLI, and through a telnet or SSH session

These blade switches can also be monitored and managed by PCM, a network device management application, and through SNMP-based applications like a MIB browser. Like many HP switches the blade switches support traffic monitoring using the port mirroring feature. Other diagnostic support capabilities include Remote Network Monitoring (RMON) and UniDirectional Link Detection (UDLD), to name a few

- **Converged Fabric support:**

  The 6120XG supports the full suite of Converged Ethernet protocols including DCBX, PFC, ETS and FIP snooping with the optional CEE license. When connected to a HP Blade server configured with a Converged Networking Adapter, the 6120XG will support Ethernet, FCoE and iSCSI data between the c-Class enclosure and an upstream switch. Combining Fibre Channel and Ethernet traffic on a single link can dramatically reduce networking expenditures and reduce network complexity

HP 6120G/XG Blade Switch

Figure 69 HP 6120G/XG Blade Switch

External Network Ports

The 6120G/XG Blade Switch has nine external Ethernet ports for which the connectors or transceiver slots are accessible from the front panel. The external ports function as uplinks that can be used to connect to one or more upstream switches. For example, an upstream switch could be an HP 5500 Series Switch installed in the same rack with a c3000/c7000 enclosure containing the 6120G/XG Blade Switch. The upstream switch could just as easily be located in another rack or require a cable run to a wiring closet where the upstream switch may be located.

The external Ethernet ports consist of the following:

- 1x 10GbE CX4
- 2x 10GbE XFP ports
- 2x 1GbE SFP ports
- 4x 10/100/1000 RJ-45 ports

Internal Network Ports

The 6120G/XG also has “internal” Ethernet ports that function as downlinks to devices in the enclosure. The internals ports are essentially communication interfaces used to transfer data to and from server blades and even another 6120G/XG (if it is installed in an adjacent bay of the enclosure). The internal ports use the enclosure’s midplane circuitry to communicate with the other c-Class enclosure components.

The internal ports consist of the following:

- 16x 1GbE ports for communications with server blades
- **1x 10GbE port for communications with another 6120G/XG Blade Switch**. Note that the 6120G/XG Blade Switches must be installed in adjacent interconnect bays (for instance, bays 1 and 2, or 3 and 4, and so forth) to be able to communicate using this internal port

HP 6120XG Blade Switch

Figure 70 HP 6120XG Blade Switch

External Network Ports

The 6120XG Blade Switch has nine external Ethernet ports for which the connectors or transceiver slots are accessible from the front panel. Several of these ports have shared functionality or can be thought as being a form of dual personality ports.

The external Ethernet ports consist of the following:

- **1x 10GbE CX4 or 1x 10GbE SFP+:** The port labeled 17 can be used for either 10GbE CX4 or 10GbE SFP+ connectivity
- 5x 10GbE SFP+ ports
- **2x 10GbE SFP+ ports**: individually, the two ports labeled 23 and 24 can be used as either external 10GbE SFP+ ports or internal 10GbE switch-to-switch (cross-link) ports. For SFP+ connectivity, these ports support the same fiber optic transceivers as the other SFP+ ports

In summary, you can have a maximum of eight 10GbE external ports operating concurrently. One scenario you could deploy is 8 SFP+ ports, or 7 SFP+ ports and 1 CX4 port. As a second scenario, you could have up to seven 10GbE external ports operating concurrently (7 SFP+, or 6 SFP+ and 1 CX4) and one of the right-most external ports (port 23 or 24) operating as a10GbE internal switch-to-switch port. Lastly, as a third scenario, you could have up to six 10GbE external ports operating concurrently (6 SFP+, or 5 SFP+ and 1 CX4) and both of the right-most external ports (ports 23 and 24) operating as 10GbE internal switch-to-switch ports.

Internal Network Ports

Like the 6120G/XG Blade Switch, the 6120XG also has “internal” Ethernet ports that function as downlinks. The internals ports are essentially communication interfaces used to transfer data to and from server blades and even another 6120G/XG (if it is installed in an adjacent bay of the enclosure).

Internal ports of the 6120XG operate as 10GbE ports or 1GbE ports depending on the NIC or adapter they are connected to, and are autonegotiating.

The internal ports consist of the following:

- 16x 1GbE/10GbE ports for communications with servers
- **0x, 1x, or 2x 10GbE ports for communications with another 6120XG Blade Switch**. The 6120XG Blade Switches must be installed in adjacent interconnect bays to be able to communicate using these internal ports. You have the option of individually using ports 23 and 24 as SFP+ ports or internal switch-to-switch ports. Therefore you can have no switch-to-switch ports active, or only one, or both

HP 6125 Blade Switch Family

The HP 6125 Blade Switches are the next generation versions of the HP 6120s. The devices listed below are due to be released September 24th 2012.

The upcoming HP 6125 Blade Switches will deliver the following added new added features:

- **Intelligent Resilient Framework (IRF):**

  IRF is a technology that enhances ordinary Ethernet switching designs, allowing substantial improvements in the way Ethernet switches communicate. HP IRF provides the ability to flatten data center and campus networks, eliminating the need for multiple tiers of aggregation switches and unutilized data paths. IRF provides a framework that enhances Ethernet and provides better link management, utilization, and redundancy
- **Comware OS:**

  Comware is HP Networking's common operating system for blade, top of rack and core Ethernet switches. The use of a common OS means that today's demanding data centers can be managed and configured from edge to edge under a single stream of FW using common configuration scripts, troubleshooting procedures and upgrade policies
- **HP Intelligent Management Center (IMC) management:**

  HP IMC is a next-generation management software which provides the data center operations team with a comprehensive platform that integrates network technologies and provides full fault, configuration, accounting, performance, and security management functionality. Built from the ground up to support the Information Technology Infrastructure Library (ITIL) operational center of excellence IT practices model, IMC's single-pane-of-glass management paradigm enables efficient end-to-end business management to address the stringent demands of today's mission-critical enterprise IT operations
- **Powerful QoS feature:**

  Creates traffic classes based on access control lists (ACLs), IEEE 802.1p precedence, IP, DSCP or Type of Service (ToS) precedence; supports filter, redirect, mirror, or remark; supports the following congestion actions: strict priority (SP) queuing, weighted round robin (WRR), weighted fair queuing (WFQ), weighted random early discard (WRED), weighted deficit round robin (WDRR), and SP+WDRR

HP 6125G Blade Switch

Figure 71 HP 6125G Blade Switch

The HP 6125G is the first member of the 6125 family of switches. HP understands that there is still a need for a 1GbE switch with exceptional performance for customers with basic data center needs, remote locations and also cluster applications with lots of 1GbE links. The 6125G is designed to fit that niche. The 6125G may be low on cost but is as fully featured as a switch can be.

Running the same Comware OS as the HP Networking ToR switches, the 6125G supports full Layer 3, is IPv6 compatible, and has more memory and processing power than competing blade switches.

External Network Ports

The 6125G Blade Switch has eight external Ethernet ports for which the connectors or transceiver slots are accessible from the front panel.

The external Ethernet ports consist of the following:

- 2 x 1GbE SFP ports
- 4 x 1GbE RJ45 ports
- **2 x 1GbE/10GbE SFP/SFP+ ports**: These ports support 1GbE SFP modules and can also be configured as 10GbE IRF switching links. Multiple switches can be combined into a single virtual switch using IRF, and any combination of 6125G or 6125G/XG switches can be IRF'd together

Internal Network Ports

Like the 6120G/XG Blade Switch, the 6125G also has "internal" Ethernet ports that function as downlinks. The internals ports are essentially communication interfaces used to transfer data to and from server blades and even another 6125G/XG (if it is installed in an adjacent bay of the enclosure).

The internal ports consist of the following:

- 16 x 1GbE ports for communications with server blades
- 1 x 10GbE port for cross connect communications to a switch in the adjacent interconnect bay

HP 6125G/XG Blade Switch

Figure 72 HP 6125G/XG Blade Switch

A cost effective 1GbE/10GbE hybrid, the HP 6125G/XG switch is designed for customers requiring low cost aggregation at the edge along with virtualized environments and scale out applications where customers desire lots of 1GbE downlinks, cross server traffic and consolidated uplinks to “leaf layer switches for any server to any server communications.

Running the same Comware OS as the HP Networking ToR switches, the 6125G/XG supports full Layer 3, is IPv6 compatible, and has more memory and processing power than other 1/10GbE hybrid blade switches.

External Network Ports

The 6125G/XG switch has four RJ45 copper 1GbE uplinks and up to four 10GbE uplinks. The Four SFP based uplinks support 1Gb SFP modules for Fiber or additional copper uplinks, four 10GbE SFP+ optical uplinks or any combination of uplink or 10Gb IRF switching links (up to four). When combined with the internal 10GbE cross connect, ample cross sectional and redundant bandwidth is available for almost any switching application.

The external Ethernet ports consist of the following:

- 4 x 1GbE RJ45 ports
- 4 x 1GbE/10GbE SFP/SFP+ ports: These ports can also be configured as 10GbE IRF switching links. Multiple switches can be combined into a single virtual switch using IRF, and any combination of 6125G or 6125G/XG switches can be IRF’d together

Internal Network Ports

Like the 6120G/XG Blade Switch, the 6125G/XG switch also has “internal” Ethernet ports that function as downlinks. The internals ports are essentially communication interfaces used to transfer data to and from server blades and even another 6125G/XG (if it is installed in an adjacent bay of the enclosure).

The internal ports consist of the following:

- 16 x 1GbE ports for communications with server and storage blades
- 1 x 10GbE port for cross connect communications to a switch in the adjacent interconnect bay

More information regarding the HP 612x Blade Switch Family can be found at the HP product page.

HP data center and VM security

Figure 73 HP TippingPoint Security

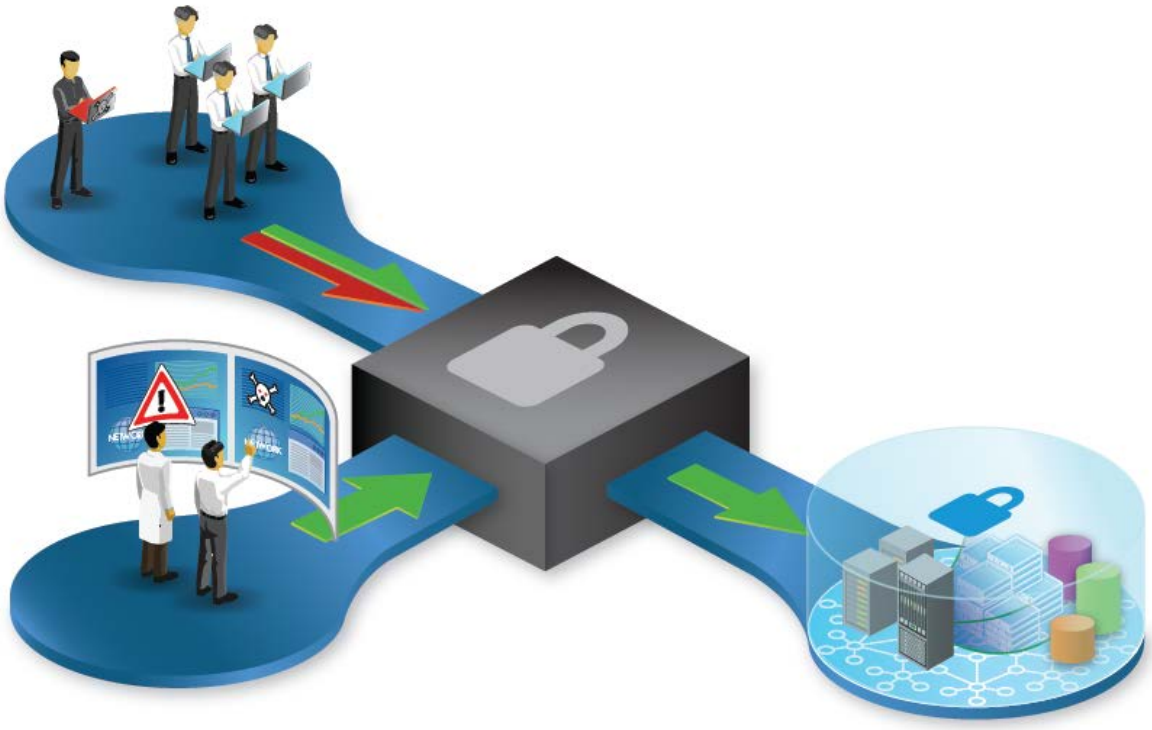

All of the traditional security duties must be accomplished to secure any data center. On top of the standard duties, data center security teams need to secure the virtual platforms. HP TippingPoint has a suite of solutions to support a converged network infrastructure and VMs.

The HP TippingPoint Secure Virtualization Framework (SVF) is a suite of products designed to help prevent network threats from impacting virtualized environments. Because several VMs are hosted on a single physical server, a security breach on a VM can impact the other VMs on that server. We have earlier discussed the components of virtualization and the challenges introduced by virtualization. In this section, we take a look at the components of the SVF and how they help address the challenges.

Figure 74 Application VMs

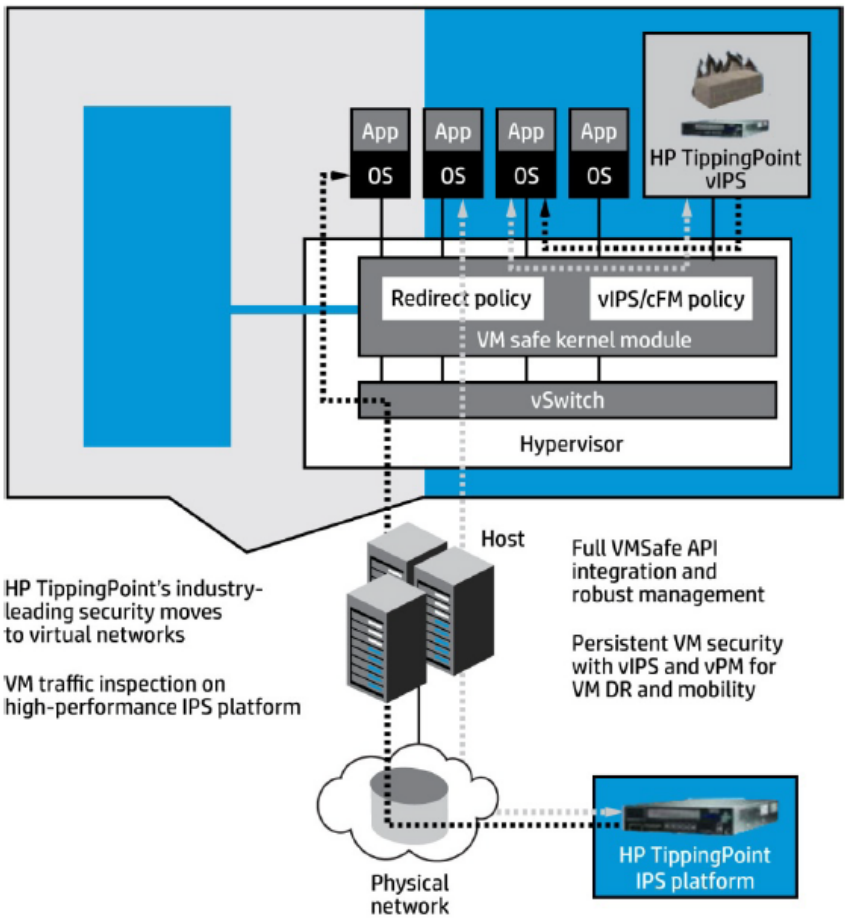

SVF is comprised of the following components:

- **IPS platform:**

  The IPS platform is an in-line device that examines packets as they flow through it and blocks malicious packets. It also includes HP DVLabs, which is the security research organization that supplies the security filters to customer IPS platforms to keep the IPS devices updated with the latest security vaccines. The IPS solution delivers automated, in-line traffic inspection designed to protect the most demanding 10 Gbps network environments

- **Virtual controller (vController):**

  The HP TippingPoint vController works with an HP TippingPoint N-Platform IPS to provide high-performance intrusion prevention for a virtualized server. Software that is easily installed in a virtualized server, the vController extends security protection from physical to virtual networks by routing traffic through an HP TippingPoint N-Platform IPS appliance. The vController prevents security attacks by inspecting all VM traffic as it moves through the network, either between VMs or from VMs to traditional servers. It offers low-cost ownership because vController leverages investments in the same N-Platform that the customers use to protect **the enterprise network**

- Virtual IPS (vIPS):

  The vIPS is a virtual appliance that provides the same IPS capabilities as the IPS platform. This leverages the resources on the host system and can provide security in cloud environments or

added security in virtualized environments in combination with the IPS platform

- **Security management system (SMS):**

  The HP TippingPoint security management system is a hardened appliance responsible for discovering, monitoring, configuring, diagnosing, and reporting for multiple HP TippingPoint systems. It features a secure Java client interface that enables analysis with trending reports, correlation and real-time graphs on traffic statistics, filtered attacks, network hosts and services, as well as IPS inventory and health

- **Virtual management center (VMC):**

  The VMC is used to automatically discover every VM in the data center and deploy vController and the virtual firewall on each virtualized host. This helps ensure appropriate security policies are dynamically applied to and enforced by vController and the IPS platform for all deployed/discovered VMs

- **HP Secure Network Fabric:**

  The unified, core-to-edge network protection provided by the HP Secure Network Fabric addresses the limitations of perimeter-based security and the increasing number and frequency of security attacks. Security virtualization and segmentation, industry-leading performance, and comprehensive security management enhance protection of business-critical data, applications, and infrastructure

- **HP TippingPoint SVF:**

  SVF provides consistent, unified security across virtualized and physical data center network infrastructures. SVF extends to virtualized servers the same strong HP TippingPoint security provided to physical infrastructure. It offers simplicity and low-cost operations because it is centrally managed by the HP TippingPoint Security Management System, and enables infrastructure-wide policies to be consistently applied to VMs

- **HP S1500 SSL appliance:**

  The HP S1500 SSL appliance provides hardware accelerated SSL offloading and bridging to enable high-performance IPS inspection of SSL encrypted traffic. This delivers increased security coverage in next-generation data centers, prevents encrypted attacks, and helps enterprises address compliance requirements without impacting the performance or availability of the network

More information regarding the HP TippingPoint can be found at the HP security product page.

HP data center management

Figure 75 HP data center management

HP IMC is a core building block of the FFRA. It allows IT administrators to gain new efficiencies and higher levels of control by converging network management and orchestration. Instead of turning to a myriad of discrete network management tools, IT staff can use HP IMC for a single pane-of-glass management across heterogeneous networks. IMC provides full FCAPS management and scales easily from small to very large deployments.

HP IMC is a powerful and flexible platform that saves enterprises time and money in deploying, managing, and monitoring mission-critical converged networks. By bringing virtual and physical network management and monitoring of services, resources, and users together in a single, integrated, and modular platform, HP solutions deliver greater value, helping enterprises align their networks with their business objectives.

HP IMC provides scalability by supporting distributed and hierarchical system architectures, as well as through additional operating system and database support to meet the requirements of complex networks. IMC uses a SOA model to provide full resource, service and user management. Its modular design enables the integration of traditionally separate management tools. IMC enables enterprises to expand their infrastructure management in scale and to seamlessly accommodate new technologies at the same time.

HP IMC Base Platform Features

HP IMC consists of a base platform and service modules that offer additional functionalities. The base platform provides administrators and operators with the basic and advanced functionality needed to manage IMC and the devices managed by IMC. The IMC base platform provides the following functions:

- Administrative controls for managing IMC and access to it. This includes granting or restricting operator access to IMC features through operator and operator group management. The base platform also includes features for the system-wide management of device data collection and information shared by all IMC modules including the creation and maintenance of device, user, and service groups, and device vendor, series, and device model information. It also includes SNMP MIB management and other system-wide settings and functions
- A broad feature set for network device management, from the ability to manage SNMP, Telnet, and SSH configurations on a device to configuring Spanning Tree and PoE energy management for managed switches and much more

- Management of the configuration and system software files on devices managed by IMC. This includes storing, backing up, baselining, comparing, and deploying configuration and software files
- Real time management of events and the translation of events into faults and alarms in IMC. This includes creating, managing, and maintaining alarm lists, trap and Syslog filters and definitions, and configurations for notifications of alarms
- Monitoring, reporting, and alarming on the performance of the network and the devices that comprise it. This includes managing global and device specific monitors and thresholds as well as creating views and reports for displaying performance information
- ACL management. This includes creating and maintaining ACL templates, resources, and rule sets and deploying ACL rule sets to devices managed by IMC. It also includes monitoring and leveraging ACLs that exist on devices for deployment to other network devices
- Monitoring and managing security attacks and the alarms they generate
- Global management of VLANs on all devices that support VLANs, managed by IMC

HP IMC service modules

HP IMC's modular and scalable SOA architecture supports extension of IMC's scope of coverage beyond the functionality of the base platform. The following optional service modules are available:

- **Virtual Application Network Manager Module (VAN):**

  VAN is an IMC module, which delivers a template based approach for managing network configuration policies, yielding consistency and reliability across network infrastructure. Administrators define the policies in a template, which will be applied as a configuration policy to the edge switch associated with the virtual machine of interest. Virtual machine network connectivity is automated and orchestrated by IMC VAN Manager which lends to the acceleration of application deployment and service rollout and greatly reduces risk during virtual machine migration. The module includes a plugin into the VMware hypervisor manager which enables the connection policies defined in IMC to be applied to the virtual machine
- **Extended API Software (eAPIs):**

  The eAPIs are an extension of IMC's open and extensible platform and can be used to share this SOA platform with an organization's homegrown and in-house applications. By integrating with IMC, developers can ensure their applications will work with all the aggregated network data collected by IMC. Developers can write their programs only once to interface with IMC, instead of many times to integrate with the operating system of each third-party device on their network. IMC is built upon an open and extensible architectural platform that leverages REST-style web services as one of the three web services in IMC. These REST-style web services enable 3rd party developers to create applications which interface and leverage IMC services. IMC Extended APIs includes over 200 APIs that provide access to core platform services. The Extended APIs are included with the Enterprise Platform and are an optional license upgrade for the Standard Platform
- **Application Performance Manager (APM):**

  APM is an IMC module that allows administrators to visualize and measure the health of critical business applications and their impact on network performance. With the available data, you can easily determine which business process is affected and which application issues to prioritize—all leading to quick and effective troubleshooting. The comprehensive monitoring and management that APM provides includes fault management, and performance monitoring of application servers, servers, and databases. Applications can easily be discovered by APM, and administrators can be informed of application issues through generated alarms. As with many of IMC modules, APM provides comprehensive reporting features

- **Service Health Manager (SHM):**

  IMC SHM is an IMC module that provides end-to-end service monitoring and service assurance through the visualization of infrastructure or network variance/factors that are in the service path. SHM leverages data derived from other IMC components to yield critical performance metrics. SHM then aggregates key performance indicators (KPIs) to generate key quality indicator (KQI) metrics. KQIs can be modeled to provide a visual representation of service-level agreement (SLA) obligations. With SHM, administrators can visually determine the level of quality for defined services and take proactive measures to maintain SLAs

- **Remote Site Manager (RSM):**

  The RSM securely extends the IMC core platform capability to remote sites by deploying remote agents. These remote agents manage and monitor the remote network, and apply policies and configurations to the remote network devices on behalf of the central IMC server. The use of agents allows IMC to securely manage remote networks, even in a firewalled environment. Additionally, these local distributed agents increase polling efficiency, allowing you to monitor your network with higher granularity, which provides more accurate and real-time visibility

- **TACACS+ Authentication Manager (TAM):**

  TAM is an IMC module that provides basic authentication, authorization, and accounting functions for network device or IT users in order to deliver network device management security. TAM, which utilizes the TACACS+ protocol, can assign users with different privileges, monitor login and command execution operations, and simplify user management. TAM works with devices that support the TACACS+ protocol

- **Intelligent Analysis Reporter (IAR):**

  IAR extends the reporting capabilities within IMC to include customized reporting. These extended reporting capabilities enable network administrators to perform proper analysis and make informed decisions. IAR makes customized reporting easy by including a report designer, which can save designs into report templates. Report outputs include a variety of formats, including charts. Reports can be automatically generated at specified intervals and distributed to key stakeholders

- **Wireless Service Manager (WSM):**

  With IMC's WSM module, operators can perform wireless LAN WLAN device configuration, topology, performance monitoring, RF coverage and planning, WLAN intrusion detection and defense, and generate WLAN service reports from the same platform they use to manage wired networks. In addition, IMC provides operators with historical reports for monitoring how wireless network usage, performance, and roaming patterns have changed over months or years

- **Voice Service Manager (VSM):**

  IMC coupled with the VSM module offers operators an integrated management solution that provides a comprehensive set of tools for managing converged voice and data networks easily and efficiently. IMC's VSM offers a single pane for voice resource and service management for 3Com and H3C voice infrastructures, including VCX® Connect platforms, Media Gateway, and IP phones. IMC VSM also provides comprehensive management and notification of any issues that may impact service quality. VSM monitors the network using built-in rules, diagnoses problems, tracks changes to IP phone status, and tracks inventory of communications devices and IP phones. It also provides tools to facilitate rapid troubleshooting and fault isolation; service-level, real-time alerting, and reporting are built in

- **User Access Manager (UAM):**

  The UAM module provides IMC users with authentication and authorization services for endpoints accessing the network edge. As a component of the IMC management platform, UAM

supports access policies across devices such as Ethernet switches, routers, broadband access servers, and VPN access gateways to centrally manage access for wired, wireless, and remote users. UAM, together with the base IMC platform and other IMC modules, provides network operators with integrated management of users, resources, and services

- **Endpoint Admission Defense (EAD):**

  The EAD module supports operators in reducing network vulnerabilities by integrating security policy management and endpoint posture assessment for identifying and evaluating, alerting on, and isolating risks at the network edge. NAC solutions have typically involved the integration of several functions that were usually deployed, configured, managed, and audited as independent systems. The HP IMC management platform provides all of these functions in a single platform, eliminating the complexity of managing multiple systems. With EAD, IMC integrates security threat evaluation, identification, location, security event awareness, and the execution of protective measures into a centrally managed and monitored platform. IMC reduces implementation costs and complexity while increasing overall network security

- **User Behavior Auditor (UBA):**

  The UBA module is a log auditing tool that enables operators to view user and network access information. UBA is designed to process large, complex log files and present the information in a simplified tabular format. UBA provides auditing of NAT, Web, and FTP site visits and more

- **QoS Manager (QoSM):**

  The QoSM is the core component of IMC's QoS solution. QoSM provides operators with a common set of QoS device and configuration management features for easily managing QoS for different device types. IMC's QoSM straightforward implementation of QoS management enables operators to focus on most critical aspects of QoS management - service planning

- **Network Traffic Analyzer (NTA):**

  The NTA provides operators with real time traffic analysis. NTA is a graphical network monitoring tool that leverages industry standard sources of network traffic data to generate real-time displays of TopN users and applications. Routers and switches that support NetFlow provide NTA with the data that feeds NTA reports. NTA analysis and reports support operators in understanding how network bandwidth and resources are being used as well as with information on which hosts and uses are consuming network resources. NTA also supports operators in identifying network bottlenecks, with support in taking corrective measures. The information provided by NTA supports mission critical network management activities such as network planning, monitoring, optimization, and troubleshooting

- **MPLS VPN Manager (MVM):**

  The MVM module brings operators VPN monitoring and management features for MPLS networks to the IMC suite of network management applications

- **Service Operation Management (SOM):**

  SOM focuses on operations and management flow to provide full IT lifecycle management. It allows IT organizations to adhere to ITIL v3.0, including IT services such as policy design, operation, and improvement. Through flow management, SOM software provides controls, measures, and audit capabilities for configuration changes, fault identification, and recovery. Based on a unified configuration management database (CMDB), SOM software provides configurable flows and options for self-service, as well as management of asset configuration, change, fault events, problem recognition, and auto-generation of a knowledge base. Self-service reduces IT involvement by allowing end users to recognize known network issues as well as to create and track service requests. SOM software integrates with the HP IMC platform to correlate information about network performance, traffic flows, and user controls

- **IPSec/VPN Manager (IVM):**

IVM provides features for all aspects of IPSec VPN management. This includes real-time and historic status, performance monitors, problem recognition, and resolution. In addition, IVM expedites IPSec VPN deployments and displays a graphical VPN topology, VPN channel status, and other configurable monitors. It is scalable and can configure and monitor multiple devices at once

- **Branch Intelligent Management System (BIMS):**

  BIMS uses an intelligent component-based architecture to provide powerful support for service operations, delivering high reliability, scalability, flexibility and IP investment returns. Based on the TR-069 protocol, IMC BIMS offers resource, configuration, service, alarm, group, and privilege management. It allows the remote management of customer premise equipment (CPE) in the WAN

More information regarding the HP IMC can be found at the HP Network Management product page.

Partnering architecture

Network virtualization and F5

The benefits of virtualization with VMware are clear. F5, an HP AllianceONE partner, provides virtualization optimized solutions that can be placed at various locations in the FFRA data center architecture to help support the virtualization scale-out requirements previously discussed.

Specifically F5 has worked closely with VMware to integrate solutions such as:

- vSphere
- vCenter Server
- vCloud Director
- vCenter Site Recovery Manager
- VMware View
- Local Traffic Manager™ (LTM)
- Local Traffic Manager™ Virtual Edition (LTM VE)
- Global Traffic Manager™ (GTM)
- WAN Optimization Manager™

Application/VM, server optimization for virtualized/cloud data centers

F5 BIG-IP LTM optimizes connections, routes traffic, and balances loads across VMs (VMs), offloading many of the functions that create CPU and memory strain. Offloading these functions to a purpose-built appliance can free up to 40 percent[3] of server CPU and memory resources, creating a 60 percent increase in VM guest density on the same host. This enables servers to operate more efficiently and

[3] F5 tests performed on Microsoft Exchange and vSphere 4 yielded a 40% reduction in CPU utilization from BIG-IP LTM, thus an equivalent 60% increase in VM guest density on the same host. See www.F5.com/vmware a detailed overview of the test plan and results.

frees up server capacity.

F5 BIG-IP LTM also integrates with vCenter via the F5 iControl® API to receive instructions that enable it to adjust network traffic in response to changing application conditions. When new VMs are provisioned by vCenter, BIG-IP LTM can automatically add those new servers to its load balancing pool and can direct traffic to them. At the same time, BIG-IP GTM knows when VMs or entire data centers are overloaded or unavailable and reroutes traffic accordingly. Both BIG-IP devices can respond to fluctuating traffic without the need for manual intervention.

Additionally, organizations that rely heavily on Web-based applications can take advantage of F5 BIG-IP® WebAccelerator™. BIG-IP WebAccelerator speeds up access and decreases webpage load time. It also drastically reduces the CPU load on Web application servers in virtualized environments through a combination of intelligent caching, connection pipelining, and exploitation of browser behavior.

The F5 Management Plug-In for VMware vSphere allows virtualization administrators to more easily manage their BIG-IP Application Delivery Networking policies as they relate to VMware-virtualized applications. The F5 Management Plug-In for VMware vSphere eliminates manual synchronization of information between BIG-IP devices and the vSphere consoles. It also helps automate common networking tasks involved in routine VM maintenance and administration. Finally, it can automatically apply Application Delivery Networking policies to newly provisioned VMs and ease the process of de- provisioning VMs. Overall, these features simplify and automate many of the networking tasks common to VMs, thereby improving the agility of the overall infrastructure.

Business continuity and disaster recovery

F5 and VMware have developed a complete solution for running vMotion and Storage vMotion events together, between vSphere environments and over long distances. The solution components enable vMotion migration between data centers without downtime or user disruption. Key solution components include:

- Encryption and compression of vMotion traffic between sites using BIG-IP LTM iSessions feature
- Byte-level data deduplication of vMotion traffic between sites using BIG-IP WAN Optimization Manager
- Client traffic management with BIG-IP LTM to direct user traffic to the correct VM
- Data center traffic management with BIG-IP GTM

One example is a Windows Server guest vMotion event across a 622 Mbps link with 40 milliseconds of round-trip time and zero packet loss, which would normally take more than five minutes to complete. With BIG-IP WAN Optimization Manager it takes less than 30 seconds. The worse the WAN conditions, the greater the potential for improvement. When the vMotion event acceleration is combined with dynamic global and traffic management, newly migrated VMs are recognized quickly, without disrupting existing user sessions.

The integration of BIG-IP GTM and VMware Site Recovery Manager (SRM) provides a complete solution for automated disaster recovery between two data centers, or to the cloud. In the event of disaster, SRM automatically orchestrates the failover of VM guests and virtual infrastructure between the two sites, while BIG-IP GTM redirects all incoming client application traffic to the secondary site. BIG-IP GTM and SRM are easily integrated via the F5 iControl API.

In addition, F5 BIG-IP WAN Optimization Manager improves the transfer of data over the WAN during a failover. This module enables large volumes of data to be transferred from a source to a target

data center quickly using compression and deduplication. BIG-IP WAN Optimization Manager encrypts traffic before transmission and decreases bandwidth requirements.

F5 also provides a number of solutions that enable organizations to leverage public or private cloud solutions from VMware easily, securely, and with maximum application performance and availability.

- BIG-IP GTM is used to direct traffic between multiple data centers in cases where the application may be running in more than one location at times (for example, cloudbursting).
- BIG-IP LTM enables organizations to retain authentication and authorization locally, when running applications in the cloud, by redirecting incoming authentication requests to the home data center.
- BIG-IP LTM Virtual Edition enables clouds to provide full BIG-IP LTM services as VMs, which can be provisioned and configured on-demand.
- BIG-IP® Application Security Manager™ can provide application firewall security to a wide variety of applications running in the cloud.

Microsoft Exchange Server 2011 optimization

F5's application-ready solution for Microsoft Exchange Server 2011 helps ensure a secure, fast, and available deployment, providing the following benefits to organizations and their end users.

- F5 increases Microsoft Exchange Server 2011 efficiency. F5 optimizations dramatically reduce the number of objects the client access servers have to deliver to the clients, allowing those servers to spend more processing power on the delivery of actual mail
- F5 helps eliminate spam. F5 provides a reputation-based, perimeter anti-spam solution that is integrated into the application delivery control network. This allows F5 to extend security for message applications to the edge of the corporate network, eliminating up to 70 percent of unwanted email. By eliminating 70 percent of unwanted email before it even reaches the Exchange Servers, F5 greatly reduces the chance that an unwanted and potentially dangerous email gets through to the Exchange 2011 servers
- F5 provides reliable, real-time availability of globally dispersed edge transport servers (SMTP). If one data center goes down, F5 immediately recognizes that it is unavailable and seamlessly re-routes incoming email to the available data center

Figure 76 F5 in the data center

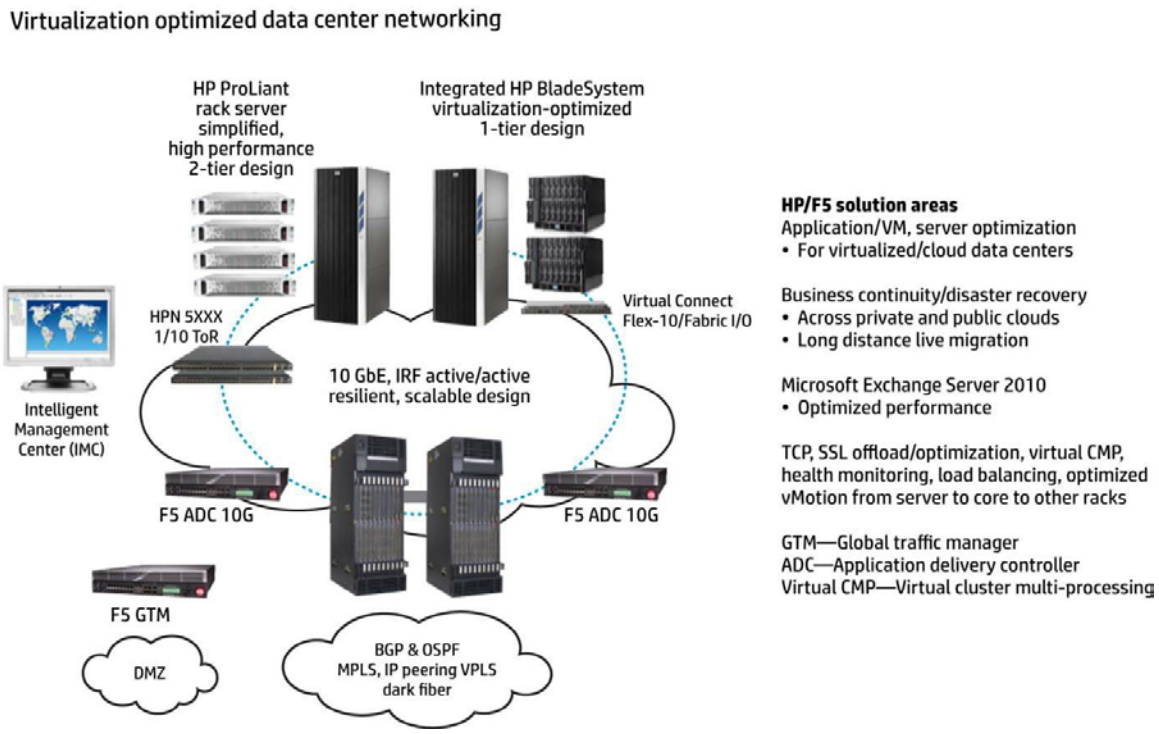

Enhancing application security with F5

Providing security specific to an application deployment is fast becoming an essential component of launching and maintaining a new application. Security personnel must work closely with the network and application teams to help ensure the successful and secure deployment of an application, especially one like Microsoft Exchange which is often used by all employees, everyday. F5 has a number of ways to help increase the security of Exchange 2011 deployments.

F5's message security offering provides an additional layer of protection for Exchange 2011 deployments. Spam email can contain virus attachments and other malicious content, like phishing attempts and Trojan attacks. The F5 solution leverages reputation data from the McAfee® TrustedSource™ multi-identity reputation engine to accurately filter email. By eliminating 70 percent of unwanted email before it even reaches the Exchange Servers, F5 greatly reduces the chance that an unwanted and potentially dangerous email gets through to the Exchange 2011 servers.

All data can now be symmetrically encrypted between local and remote F5 devices, providing a new way to help ensure site- to-site data security by preventing clear text from being passed on the wire. This secure connection, or tunnel, also improves transfer rates, reduces bandwidth, and offloads applications for more efficient WAN communication. As mentioned previously, F5 can perform DAG replication across data centers inside this encrypted tunnel for secure mailbox replication for the entire mailbox store.

For remote users who might be trying to access Microsoft Office Outlook or Outlook Web App from an airport kiosk or other unknown device, F5's comprehensive Endpoint Security provides the best possible protection for remote users. F5 technology prevents infected PCs, hosts, or users from connecting to your network and the applications inside, and delivers a secure virtual workspace, pre-login endpoint integrity checks, and endpoint trust management.

And when the remote user has finished their session with Outlook or Outlook Web App, F5's post-logon security protects against sensitive information being left on the client. F5 can impose a cache-

cleaner to eliminate any user residue such as browser history, forms, cookies, auto-complete information, and more. Post-logon security can also be configured to close desktop search applications so nothing is indexed during the session. Post-logon actions are especially important when allowing non-trusted machines access without wanting them to take any data with them after the session.

F5 security devices report previously unknown threats—such as brute force attacks and zero-day Web application attacks—and mitigate Web application threats, shielding the organization from data breaches. Our full inspection and event-based policies deliver a greatly enhanced ability to search for, detect, and apply numerous rules to block known L7 attacks.

F5 makes security compliance easy and saves valuable IT time by enabling the exporting of policies for use by offsite auditors. Auditors working remotely can view, select, review, and test policies without requiring critical time and support from the Web application security administrator.

Not only does F5 provide comprehensive application security, but produce extremely secure devices. We help make sure your Microsoft Exchange Server deployment, and the information it contains, remains secure.

DCI and the Alcatel-Lucent 1830 Photonic Service Switch

Today's enterprises expect anywhere, anytime access to storage and computing from their service providers. Providers are therefore deploying virtual and shared resources across geographically diverse data centers — all connected with a high-speed, redundant mesh of links. These metro and long-haul links must deliver a high-quality experience to enterprises over providers' data center connect (DCC) networks.

Additionally, today's enterprises recognize that their data centers are continuously at risk from internal and external security threats. More than simply deploying antivirus and firewall defenses, enterprises must establish comprehensive IT security programs that protect a virtual and distributed environment of computing and storage resources.

For these distributed resources to work effectively and meet diverse end-user requirements, applications require secure, low-latency real-time communications with guaranteed QoS. The data center interconnect (DCI) infrastructure and optical fibers used for DCI transport are key components of a holistic and systematic IT security program.

Alcatel-Lucent offers a versatile, scalable, secure and easy to use photonic service fabric for the DCI and plays an important role in enabling cloud computing based services. The Alcatel-Lucent 1830 Photonics Switching System (PSS) is the networking product of choice for enterprise and communication service provider DCI solutions. Through versatile, scalable and secure configurations, it allows enterprises to realize lower IT costs while improving network reliability, performance and security.

With a unique ability to seamlessly "virtualize" data center and networking resources that facilitate a cloud services model, HP and Alcatel-Lucent are enabling enterprises and communication service providers to reduce overall IT costs, improve business continuity and agility while reducing space, power consumption and management costs. The 1830PSS is an integral part of the HP and Alcatel-Lucent alliance for DCI solutions.

Note: Photonic sometimes also referred to as Optical or WDM (Wavelength De-multiplexing). WDM virtualizes the physical fiber infrastructure from one into many, and in case of DWDM up to 88 virtual fibers per physical fiber.

Figure 77 Converged MAN & WAN with ALU 1830PSS

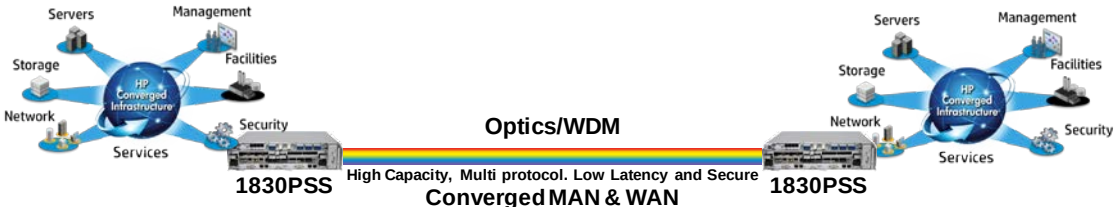

Key business problem solved

Enterprises face major inter-data center networking challenges as bandwidth demands for storage applications continue to grow. There are four important market trends that shape the future of DCI and its challenges to enterprises:

- **Continued improvement of Business Continuity and Disaster Recovery (BCDR) capabilities (solution needs to be Latency Optimized):**

  Improving BCDR capabilities continue to be one of the top priorities for IT decision-makers. Enterprises are also moving towards Cloud Infrastructure services to realize capital and

operational cost savings. Some enterprises will pursue a strategy of implementing Cloud Infrastructure services delivered by communication service providers for increased elasticity. Therefore although enterprises are moving towards public Cloud Infrastructure services, they continue to prefer dedicated infrastructures to protect their mission critical data. Enterprises will continue to be challenged in meeting the performance, reliability and bandwidth requirements of BCDR for mission critical data using their private data center inter-connect. As public cloud services mature, they could become a viable low cost alternative for enterprise less critical applications and systems (increasing the elasticity)

- **Growth of mission critical data (solution needs to provide Dynamic and Flexible Bandwidth):**

  Enterprises recognize that more and more applications are critical. Enterprises consider email and collaborative applications like Sharepoint critical applications because applications ecosystems have become more complex and developed interdependencies that force less critical application data to have the same degree of availability as mission critical application data

  As storage data will continue to grow and the percentage of data considered critical increases, the growth of critical data will grow at a higher rate. The growth of mission critical data increases the bandwidth required by enterprises to replicate the information synchronously within the boundary of a Metropolitan Area Network (MAN)

- **Secure mission critical data (solution needs to support in-flight encryption and intrusion monitoring):**

  In a world filled with electronic security threats, companies are recognizing that their data centers are at continuous risk. Security threats to the data center arise not just from traditional malware or hobbyist hackers, but increasingly from criminal organizations that are directly targeting the enterprise. Multi-site data center security is not just about technical countermeasures such as antivirus and firewalls used inside the data center, but a much more systematic and holistic approach to enterprise-wide security. Enterprises must establish a comprehensive and coordinated counter-attack, implementing solutions that provide detection and mitigation of security threats

  Detection of security breaches requires the deployment of embedded security monitoring technology in network devices that can detect intrusions even when the traffic traversing the network is undisturbed

  Encryption of data is required to mitigate the impact of security breaches. It is the most effective method of mitigation as it renders stolen data useless to the intruder. Encryption is the algorithmic process of transforming data into unreadable cryptographic text. Encryption is no longer an exotic mechanism whose use is limited to secret organizations: it is a common tool used as part of normal business workflow for security

- **Consolidation of data centers (solution needs to support MAN and WAN Distance Optimization):**

  Data center consolidation is one of the most fundamental ways of lowering the cost of IT operations. Larger data centers are simply more cost-effective on a per-unit basis. Enterprises are consolidating their existing data centers to new strategic locations outside major cities where real estate costs are lower

  Enterprises are also searching for proximity to greener energy sources (hydro dams, improved solar or wind locations) to realize lower energy costs. These larger data centers hold larger volumes of storage data that need to be shared with remote data center sites geographically apart

  Enterprises cannot look at the benefits of data center consolidation without considering the networking challenges they bring. The need to communicate effectively over the MAN and WAN

becomes a critical factor in achieving the lower costs of IT operations promised by data center consolidation

- **Introduction of Federated Storage technology (solution needs to support meshed networks):**

  Federated storage is the collection of autonomous storage resources that form a scale-out cluster governed by common software that manages the rules of the federation: how disk capacity is shared, joined and presented to the hosts

  Storage federation is implemented in the storage arrays or specialized appliances that sit in the data path between hosts and storage ports. With this technology enterprises can scale-out storage resources over distance to create networks of virtually unlimited capacity, improve BCDR, eliminate disruptive migrations and allow applications to share a storage volume over local, metro or global distances

  Federated Storage solutions have similar inter-connect requirements to BCDR solutions in terms of latency and multi-protocol support. In addition data migration applications introduce the challenge to dynamically allocate bandwidth for the duration of the data migration process

Use cases and potential customer profile

The joint HP and Alcatel-Lucent target market will typically consist of larger enterprises and communication service providers looking to HP and Alcatel-Lucent leadership to offer both intra- and inter-data center networking solutions. Cost reduction isn't their only driver for data center consolidation. They want to create competitive advantage by moving to a Cloud services model and will be looking for single/minimal suppliers to achieve that aim. The following use cases, among others, have been identified:

Data center inter-connect

- Enterprise
  - Own and Manage your Private Transport
  - Own and Manage your Secure Private Transport
- Communication Service Provider
  - Managed Private Transport Services (in case Enterprise only wishes to Own the Private Transport Network)
  - Managed Transport Services (in some cases known as Leased Lamda Services or Managed Wavelength Services)
  - Managed Secúre Private Transport Services (in case Enterprise only wishes to Own the Secure Private Transport Network)
  - Managed Secure Transport Services

Converged MAN & WAN

The Converged MAN & WAN is essentially a broader set of use cases not necessarily relating to pure DCI. Example in the utility / vertical space:

- Power utility in need for reliable mission critical and multi-protocol transport (such as the Alcatel-Lucent reference case Statnett in Norway)
- Railway utility with separated physical fiber infrastructure (could be different company) in need of flexible mission critical and multi-protocol transport
- Road utility in need for Converged IP and WDM low latency transport (such as the Alcatel-

Lucent reference case Highways Agency in England)

Figure 78 ALU 1830PSS use cases

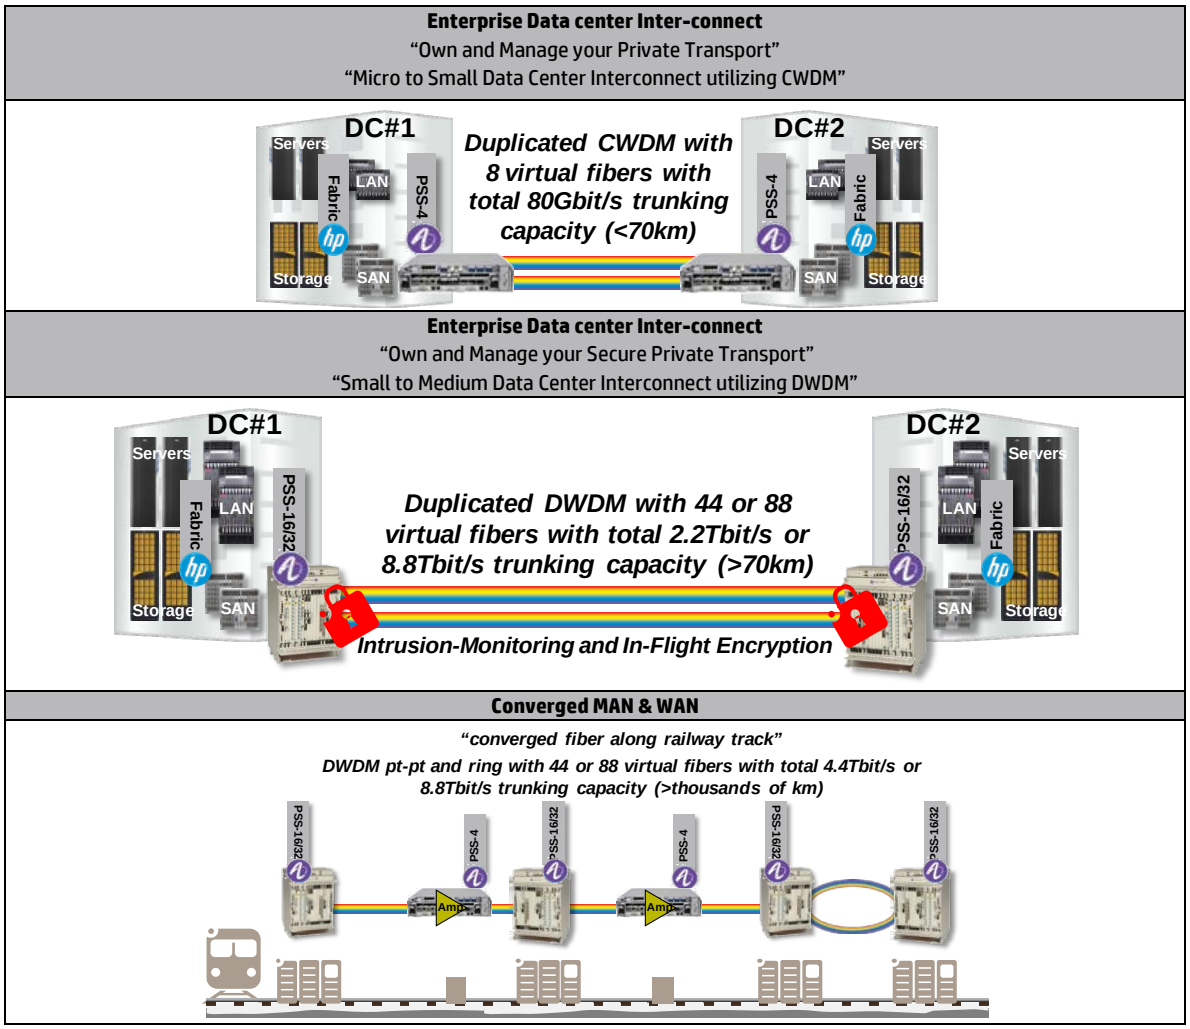

What is an Alcatel-Lucent 1830 Photonic Service Switch?

The Alcatel-Lucent 1830PSS is a versatile, scalable, secure and easy to use photonic service fabric, but how does the configuration provide for this?

Figure 79 Alcatel-Lucent 1830 Photonic Service Switch

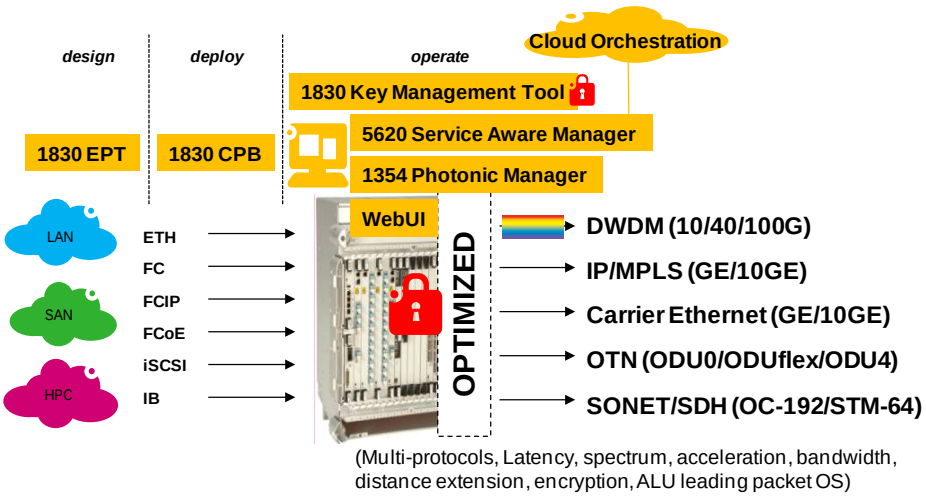

The client (SAN, LAN, HPC) signal is connected to a universal Transponder/Muxsponder that converts and aggregates the multi-protocol client signals to a CWDM or DWDM Wavelength. The wavelength is connected to a Wavelength Multiplexer (Filter)/Tuneable ROADM port that multiplexes all wavelengths onto a line fiber connected the other data center. In case of longer distances all wavelengths are amplified and typically compensated for dispersion. All connections between the discrete modules are provided by fiber optics patches supporting an unmatched scalability and flexibility in terms of configuration.

Figure 80 Alcatel-Lucent 1830 Photonic Service Switch deployment options

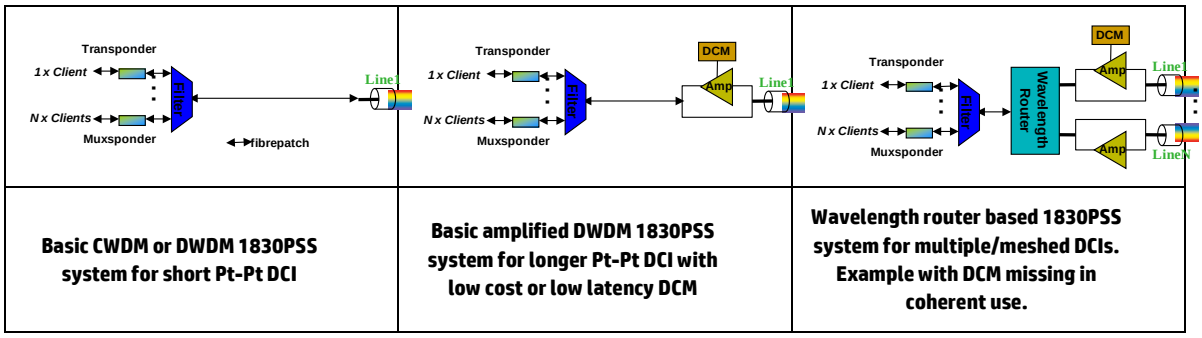

Why Alcatel-Lucent 1830 Photonic Service Switch?

Versatile:

- Multi-layer cloud service fabric agility facilitates ease of response to unpredictability
- Low latency optimized offers enabling application monetization
- Protocol and bit-rate agnostic for a variety of LAN, SAN & HPC DCI applications
- Unconstrained topology support including Pt-Pt, ring and mesh configurations
- De-risked deployments with certified partner solutions
- Contentionless, colorless, directionless multi-degree optical switching (Wavelength Router) supporting any-to-any data center inter-connectivity
- Enable seamless mobility and high availability of storage data across data centers

- Hands-free operation based on autonomous programmability of the network elements using controller-based administration and management
- Evolutionary Cloud modes of operation offering investment protection
- Optimized TCO through 1.5X density and 2/3X power at 100G

Scalable:

- Support for 10G to 40G to 100G with a path to 400G with the Photonic Service Engine (PSE) for growth
- Flexibility of choice in networking capacity with C/DWDM, ROADM, flexible grids, OTN, GMPLS & Packet technologies
- Right-sized platform commonality across any premise from branch to HQ to DC
- Cost optimized Cloud service levels enabled by flexible configuration of increasing network availability
- Cost optimized Metro, Regional and Global network implementations supported
- High performance fiber capacity up to 8.8 Tb (17.6 Tb @ 400G)
- Lower cost interconnection through transparent and statistical multiplexing
- Mixed 10G/40G/100G wavelength compatibility within same fiber pair
- Enhanced application performance enabled by Ethernet fabric extension into the packet optical transport

Secure:

- In-flight security mitigation through AES 256 encryption
- Security breach detection through Optical Intrusion Detection (OID)
- Security attack prevention through secure configuration operational mode
- Assurance of specification & proven implementation via CC / FIPS certifications
- Assured secure networking achieved with fully disassociated user roles
- Customer controllable key management enabling differentiated services

Easy to use:

- Plug & play, self-restoring managed photonics
- Lower TCO due to FCAPS/FAB commonality derived from extended packet OS
- Elastic environment delivering next-generation cloud services employing pre-integrated orchestration northbound interfaces for flow-through provisioning
- Common practices derived from a unified network management platform
- Measurement, monitoring, recording and reporting of SLA attributes through continuous monitoring on Quality of Service (QoS) enabling a Cloud-ready service fabric
- Simplified engineering and planning resulting from equipment supporting multi-protocols and multiple bit-rates in addition to streamlining equipment and sparing costs

Which Alcatel-Lucent 1830PSS product to position?

Right-sized platform (Scalable)

- 1830PSS-4: “Micro to Small DCI utilizing CWDM”:

  The 1830 PSS-4 offers 4 half-height slots or 2 full-height slots that support C/DWDM Fixed OADM point-to-point applications. It supports AC or DC power options. The PSS-4 shelf can be

stacked in multi-shelf applications for increased low-end scalability. It can also be used as a small Amplifier in case of a longer distance Converged WAN

Figure 81 Micro to Small DCI utilizing CWDM

Identify Clients (from intra Data center)	Select C/DWDM Filter and non/Amplified line	Select PSS-4 Chassis
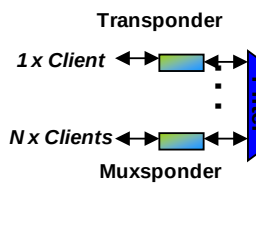	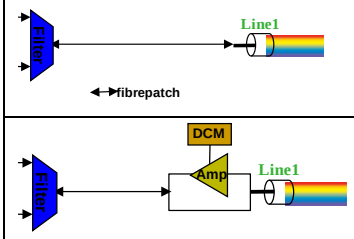	

- 1830PSS-16: “Small to Medium DCI utilizing DWDM”:

  The 1830 PSS-16 platforms offer 16 half-height slots or 8 full-height slots. It typically is used for DWDM FOADM ring or pt-to-pt applications. It can also be used as a large Amplifier in case of a longer distance Converged WAN

Figure 82 Small to Medium DCI utilizing DWDM

Identify Clients (from intra Data center)	Select (C/)DWDM Filter (typically 44 channels) and non/Amplified line	Select PSS-16 Chassis
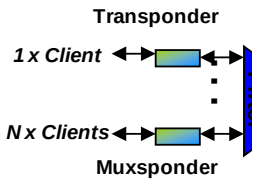	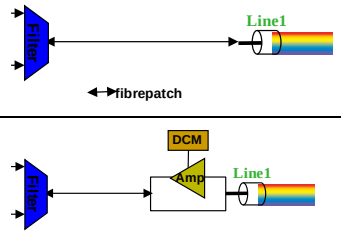	

- 1830PSS-32: “Medium to Large DCI utilizing DWDM”

  The 1830 PSS-32 platforms offer 32 half-height slots or 16 full-height slots. It typically is used for DWDM Reconfigurable OADM ring, mesh or pt-pt applications

Figure 83 Medium to Large DCI utilizing DWDM

Identify Clients (from intra Data center)	Select (C/)DWDM Filter (typically 44 or 88 channels) &or Wavelength Router and Amplified line	Select PSS-16 Chassis
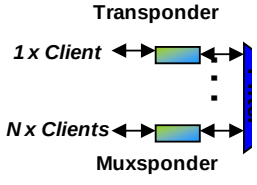	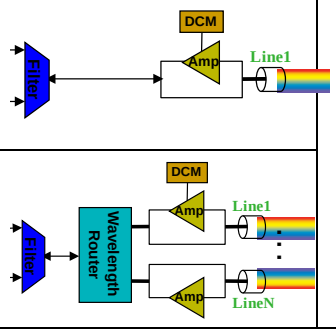	

Note: For more information on the 1830PSS chassis and variants please find datasheets on www.alcatel-lucent.com by searching for “1830 PSS”.

HP Networking Services

HP offers a comprehensive set of network services to help design, deploy, integrate, manage, and support your next- generation connectivity and communications environment. HP Services for Data Center Transformation can help you simplify, optimize, and integrate your existing facilities and implement a next-generation data center that is efficient, virtualized, and highly available. Contact your HP account manager or reseller to find out how HP Networking Services can help you implement an automated, high-performance data center network.

Support, services, and partners

The core foundation to the solution (HP and its technology partners):

- HP is #10 on the list of Fortune 500 companies
- HP has a global reach, so your data center solution will be supported internationally
- HP has industry-leading support and professional services to support your network
- HP products can be integrated with other vendor solutions through the use of standards-based technologies
- HP can provide a complete end-to-end virtual computing solution encompassing racks, servers, storage, networking technologies, security, and management

More information regarding the HP Services can be found at the HP Business Services product page.

Glossary of acronyms

ACL: An access control list (ACL) is a set of rules (or permit or deny statements) for identifying traffic based on criteria such as source IP address, destination IP address, and port number. ACLs are primarily used for packet filtering. A packet filter drops packets that match a deny rule and permits packets that match a permit rule. ACLs are also used by many modules, for example, QoS and IP routing, for traffic identification.

ALG: An application layer gateway (ALG) is a type of security device that can provide security based on applications, protecting the servers and clients from malicious traffic.

BASEL II: The international accord on banking operations (BASEL II) designed to provide international standard for banking regulators.

BEB: Within a PBB network, the Backbone Edge Bridge (BEB) maps frames and performs MAC header encapsulation and decapsulation.

BPDU: STP uses bridge protocol data units (BPDUs), also known as configuration messages, as its protocol packets. STP-enabled network devices exchange BPDUs to establish a spanning tree. BPDUs contain sufficient information for the network devices to complete spanning tree calculation.

BGP: Border Gateway Protocol (BGP) is a dynamic routing inter-Autonomous System (inter-AS) Exterior Gateway Protocol.

BIMS: HP IMC Branch Intelligent Management System (BIMS) uses an intelligent component-based architecture to provide powerful support for service operations, delivering high reliability, scalability, flexibility and IP investment returns.

DCB: Data center bridging (DCB) is a series of enhancements to the IEEE 802.1 standard to provide extensions to Ethernet for support for converged technologies such as Fibre Channel over Ethernet (FCoE).

CEE: Converged enhanced Ethernet (CEE) is an enhanced Ethernet that enables the convergence of various applications in data centers (LAN, SAN, and HPC) onto a single interconnect technology.

CIFS: The common internet file system (Microsoft CIFS), an enhanced version of Microsoft Server Message Block (SMB), is the standard way that Windows computer users share files across intranets.

CoS: Class of Service (CoS) is one type of the techniques or methods used to deliver Quality of Service (QoS) in a network.

C-VID: The QinQ 802.1Q standard enables edge devices on a service provider network to tag Ethernet frames from customer networks (private networks) with an outer VLAN tag so that the Ethernet frames can travel across the service provider network (public network) with double VLAN tags. QinQ enables a service provider to use a single service-provider VLAN (S-VIDs) to carry multiple customer VLANs (C-VIDs).

DCI: Data center interconnection (DCI) is a method to connect geographically dispersed data centers.

DoS: A Denial of Service (DoS) attack is an occasion in which a legitimate user or a group of users is prevented from accessing the services and information of network resources they would normally receive.

DVPN: Dynamic Virtual Private Network (DVPN) is a hub and spoke VPN architecture that allows multiple remote branch and regional offices (spoke) to establish site-to-site IPsec VPN tunnels to secure connectivity to the headquarters or data centers (hub).

DWDM: Dense wavelength division multiplexing (DWDM) is a technology which combines data from different sources on to an optical fiber where each signal is signaled via its own wavelength.

EAD: HP IMC Endpoint Admission Defense (EAD) module supports operators in reducing network vulnerabilities by integrating security policy management and endpoint posture assessment for identifying and evaluating, alerting on, and isolating risks at the network edge.

EoR: End-of-row topologies, which rely on larger switches placed on the end of each row for server connectivity.

FCAPS: Fault, configuration, accounting, performance, and security **-** An extension of the popular network management conceptual frameworks called telecommunication management network (TMN), FCAPS describes network management in four layers. Each TMN layer needs to perform some or all FCAPS functions in certain ways.

FCIP: Fibre Channel over TCP/IP (FCIP) describes mechanisms that allow the interconnection of islands of Fibre Channel storage area networks over IP-based networks to form a unified storage area network in a single Fibre Channel fabric.

FCoE: Fibre Channel over Ethernet (FCoE) is an encapsulation of Fibre Channel frames over Ethernet networks. This allows Fibre Channel to use 10 Gigabit Ethernet networks (or higher speeds) while preserving the Fibre Channel protocol.

FC: Fibre Channel, a Gigabit-speed network technology primarily used for storage networking.

FEC: MPLS is a forwarding technology based on classification. It groups packets to be forwarded in the same manner into forwarding equivalence class (FEC). Packets in an FEC are handled in the same way on an MPLS network.

FFI : HP FlexFabric Interconnect (FFI) is an HP method to connect geographically dispersed data centers.

FFRA: FlexFabric Reference Architecture (FFRA) is an overall discussion of the use of current HP technologies to fulfill our vision of architecting data centers.

FISMA: Federal Information Security Management Act 2002 (FISMA) Is a federal program which details specific responsibilities to various federal agencies, designed to strengthen information system security.

GLBA: The Gramm-Leach-Bliley Act of 1999 (GLBA) requires financial institutions explain their information-sharing practices to their customers and to safeguard sensitive data.

GRE: Generic Routing Encapsulation (GRE) is a protocol designed for encapsulation at network layer protocol over other network layer protocol.

HA: High availability (HA) refers to system designs and architectures that are able to provide a reasonable assurance of uptime and availability. HA systems may consist of redundant line cards and management modules, redundant paths, or a combination of those examples, and others.

HBA: Host bus adapters (HBA) are integrated circuit adapters that provide input/output processing and physical connectivity between servers and storage devices.

HIPAA: Health Insurance Portability and Accountability Act (HIPAA) applies national standards to health plans, health care clearinghouses, and those health care providers who electronically conduct certain financial and administrative transactions that are subject to the transactions standards adopted by the Department of Health and Human Services.

HOL: Head of line (HOL) blocking occurs when a stream of packets are held-up by the first packet.

HQoS: Hierarchical QoS (HQoS) uniformly manages traffic and hierarchically schedules traffic by user,

network service, and application. It provides more granular traffic control and quality assurance services than traditional QoS.

IMC: HP Intelligent Management Center (IMC) delivers next-generation, integrated and modular network management capabilities that efficiently meet the end-to-end management needs of advanced, heterogeneous enterprise networks.

IDS: An intrusion detection system is a device or software application that monitors network and/or system activities for malicious activities or policy violations and produces reports to a management station.

IETF: The Internet Engineering Task Force (IETF) is a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet.

IPS: An intrusion prevention systems (IPS) is a device or software application that monitors networks for malicious activity, and then actively blocks the malicious activity.

IPSec: IP Security (IPSec) is a security framework defined by the IETF for securing IP communications.

IRF: Intelligent Resilient Framework (IRF) is a software virtualization technology developed by H3C (3COM). Its core idea is to connect multiple devices through physical IRF ports and perform necessary configurations, and then these devices are virtualized into a distributed device.

iSCSI: The Internet SCSI (iSCSI) is a TCP/IP-based protocol for establishing and managing connections between IP-based storage devices, hosts, and clients, called the storage area network (SAN).

ISSU: In-Service-Software-Upgrade provides a mechanism to update device software without service interruption.

ITIL: Information Technology Infrastructure Library (ITIL) provides a framework for identifying, planning, delivering and supporting IT services.

IVM: HP IMC IPSec/VPN Manager (IVM) provides features for all aspects of IPSec VPN management.

Jumbo Frames: Jumbo frames often mean 9,216 bytes for Gigabit Ethernet, but can refer to anything over 1,500 bytes.

LACP: Link Aggregation Control Protocol (LACP) provides a mechanism to control the aggregating of multiple physical ports together to form a single logical port.

LAN: Local area networks (LAN) are computer networks which interconnects computers and other devices in an area such as a home or office building.

LB: A load balancer (LB) is a device that can distribute network or application traffic across a number of servers. Load balancers can increase capacity and reliability of applications.

LSP: label switched path (LSP) is the path along which a forwarding equivalence class (FEC) travels through an MPLS network.

MAD: Multi-active detection (MAD) mechanism helps increase the usability of the IRF system by decreasing the impact of an IRF split on services.

MAN: Metropolitan Area Networks (MAN) are a networks that interconnects clients in geographic area that are larger than usually covered by typical LANs, but not as large as areas covered by WANs.

MMF: Multi-mode fiber (MMF) is a type of optical fiber used for short distance communication, normally used within a building or single room.

MPLS: Multiprotocol Label Switching (MPLS) is a tunneling technology and a routing and switching technology platform that combines label switching with Layer 3 routing.

MSTP: The multiple spanning tree (MST) protocol carries the concept of the IEEE 802.1w rapid spanning tree protocol (RSTP) a leap forward by allowing the user to group and associate VLANs to multiple spanning tree instances (forwarding paths) over link aggregation groups (LAGs).

MVM: HP IMC MPLS VPN Manager (MVM) module brings operators VPN monitoring and management features for MPLS networks to the IMC suite of network management applications.

NAS: Network Attached Storage (NAS) are network connected devices dedicated to file storage that operates in a client/server mode.

NAT: Network address translation (NAT) is a service which can be used to translate data sent between IP-heterogeneous nodes.

NBAD: Network Behavior Anomaly Detection (NBAD) is the continuous monitoring of a network for unusual events or trends.

NERC: The North American Electric Reliability Corporation (NERC) provides standards and guidelines which are established for the national electric grid and its operation.

NFS: Network File System (NFS) is a file system protocol developed to allow clients to access files over the network in a manner similar to how local storage can be accessed.

NIC: Network interface cards (NIC) are adapters attached to a computer (or other network device such as a printer) to provide the connection between the computer and the network.

NMS: The network management system (NMS) is a combination of hardware and software used to monitor and administer a network.

NTA: HP IMC Network Traffic Analyzer (NTA) is a graphical network monitoring tool that utilizes industry-supported flow standards to provide real-time information about the top users and applications consuming network bandwidth.

OOBM: Out-of-band management (OOBM) is a devices management method which operates on the management plane, rather than the data plane which is used by data traffic and by in-band management traffic.

PBB: IEEE 802.1ah-2008, Provider Backbone Bridging (PBB) or MAC-in-MAC , extends the work achieved in 802.1ah by providing a hierarchical network infrastructure that completely abstracts the Service Provider Backbone from the Customer Network. PBB networks are deployed to aggregate existing PBB networks.

PCI-DSS: Payment Card Industry-Data Security Standard (PCI-DSS) is a security standard which describes requirements for security management, network architecture, software design and other critical protective measures. The standard is intended to help organizations proactively protect customer account data.

Port mirroring: Port mirroring is used on a network switch to send a copy of network packets seen on one switch port (or an entire VLAN) to a network monitoring connection on another switch port.

PWE3: Pseudo Wire Emulation Edge-to-Edge (PWE3) is a mechanism that emulates the essential attributes of a service such as ATM, Frame Relay or Ethernet over a Packet Switched Network (PSN).

PXE: Pre-execution environment (PXE) is an environment to boot computers using a network interface independently of data storage devices (like hard disks) or installed operating systems.

QoS: Quality of Service (QoS) is the ability of a network to provide differentiated service guarantees for diversified traffic regarding bandwidth, delay, jitter, and drop rate.

QoSM: HP IMC QoS Manager (QoSM) is the core component of IMC's QoS solution. QoSM provides operators with a common set of QoS device and configuration management features for easily

managing QoS for different device types.

RRPP: Rapid Ring Protection Protocol (RRPP) is a link layer protocol designed for Ethernet rings. RRPP can prevent broadcast storms caused by data loops when an Ethernet ring is healthy, and rapidly restore the communication paths between the nodes in the event that a link is disconnected on the ring.

RSTP: The rapid spanning tree protocol (RSTP IEEE 802.1w) can be seen as an evolution of the IEEE 802.1d standard more than as a revolution. IEEE 802.1w is also capable of reverting back to IEEE 802.1d in order to interoperate with legacy bridges (thus dropping the benefits it introduces) on a per-port basis.

RTT: Round trip time (RTT) refers to the length of time takes for signals to be sent, plus the time it takes to receive and acknowledgment.

SAN: A storage area network (SAN) is a high-speed special-purpose network (or subnetwork) that interconnects different kinds of data storage devices with associated data servers on behalf of a larger network of users.

SATA: Serial Advanced Technology Attachment (SATA) is an interface used to connect hard drives to a computer's motherboard.

SCSI: The small computer system interface (SCSI), an ANSI standard, is a parallel interface standard used by Apple Macintosh computers, PCs, and many UNIX systems for attaching peripheral devices to computers.

SDH: Synchronous digital hierarchy (SDH) is a technology for utilizing synchronous data transmission over optical media.

SIEM: Security Information and Event Management (SIEM) tools provide real-time analysis of security alerts which can be generated by network hardware and applications.

SIM: HP System Insight Manager (SIM) is HP's server-storage management tool which helps maximize IT staff efficiency and hardware platform availability for small and large server deployments alike. It is designed for end-user setup, and its modular architecture enables system administrators to plug in additional functionality as needed.

SMB: The Server Message Block (SMB) protocol is an IBM protocol for sharing files, printers, serial ports, etc. between computers.

SMS: The HP TippingPoint Management Security Management System (SMS) is a hardened appliance that provides global vision and control for multiple HP TippingPoint IPSs.

SNMP: The simple network management protocol (SNMP) is used by network management systems to communicate with network elements.

SOA: Service-oriented architecture (SOA) is a set of principles used for designing and developing software. HP IMC uses a service-oriented architecture (SOA) model to provide full resource, service and user management.

SOM: HP IMC Service Operation Management (SOM) focuses on operations and management flow to provide full IT lifecycle management. It allows IT organizations to adhere to ITIL v3.0, including IT services such as policy design, operation, and improvement. Through flow management, SOM software provides controls, measures, and audit capabilities for configuration changes, fault identification, and recovery.

SOX: The U.S. Sarbanes-Oxley Act of 2002 (SOX) requires CEOs and CFOs to attest to the accuracy of corporate financial documents, as well as provide IT and networking controls and their audit as per Section 404 of the Act.

SPAN: Switched Port Analyzer (SPAN), also known as port mirroring, is a method of monitoring network traffic.

SPB: Shortest path bridging (SPB) provides logical Ethernet networks on native Ethernet infrastructures using a link state protocol to advertise both topology and logical network membership.

SRM: VMware Site Recovery Manager (SRM) ensures the simplest and most reliable disaster protection for all virtualized applications. Site Recovery Manager leverages cost-efficient VMware vSphere® Replication or third-party storage-based replication to provide centralized management of recovery plans, enable nondisruptive testing, and automate site recovery and migration processes.

SSL: Secure Sockets Layer (SSL) is a protocol used for securing transmissions on a network.

STP: The spanning tree protocol (STP) is an L2 protocol designed to run on bridges and switches. The main purpose of the spanning tree is to prevent loops from forming in a bridged network.

S-VID: The QinQ 802.1Q standard enables edge devices on a service provider network to tag Ethernet frames from customer networks (private networks) with an outer VLAN tag so that the Ethernet frames can travel across the service provider network (public network) with double VLAN tags. QinQ enables a service provider to use a single service-provider VLAN (S-VIDs) to carry multiple customer VLANs (C-VIDs).

SVF: A Secure Virtualization Framework is an HP TippingPoint technology designed specifically for implementing best-of- breed threat protection for the virtualized infrastructure.

ToR: Top-of-rack utilizes a switch at the top of each rack (or close to it).

TRILL: Transparent Interconnect of Lots of Links is a specification that enables multipathing in the data center.

UAM: HP IMC User Access Manager (UAM) module provides IMC users with authentication and authorization services for endpoints accessing the network edge.

UBA: HP IMC User Behavior Auditor (UBA) module is a log auditing tool that enables operators to view user and network access information.

VAN: HP Virtual Application Network Manager (VAN) is designed to characterize applications, virtualize the network to align the resources needed for delivering the application, and automate the orchestration.

VC: HP Virtual Connect (VC) portfolio is a collection of Ethernet, Fibre Channel and converged fabric interconnect modules and firmware that provide an ideal alternative to traditional switches and patch panels installed in the interconnect bays of HP BladeSystem c-Class enclosures.

VCM: Embedded on Virtual Connect (VC) modules, the HP Virtual Connect Manager (VCM) provides a built-in Web GUI interface and a fully scriptable Command Line Interface (CLI) designed to manage single Virtual Connect domains.

VCEM: HP Virtual Connect Manager (VCEM) is a software application that centralizes server connectivity and workload management for hundreds of Virtual Connect domains and thousands of servers from a single console.

VEPA: A standard being led by HP for providing consistent network control and monitoring for VMs (of any type).

vIPS: Virtual IPS (vIPS) is a virtual appliance that provides the same IPS capabilities as the IPS platform. This leverages the resources on the host system and can provide security in cloud environments or added security in virtualized environments in combination with the IPS platform.

VLL: Virtual Leased Line (VLL) is a way to provide Ethernet-based point to point communication over IP/MPLS networks.

VM: A virtual machine (VM) is a system that enables multiple operating systems to concurrently run on a single physical server, providing much more effective utilization of the underlying hardware.

VMC: Virtual management center (VMC) is used to automatically discover every VM in the data center and deploy vController and the virtual firewall on each virtualized host. This helps ensure appropriate security policies are dynamically applied to and enforced by vController and the IPS platform for all deployed/discovered VMs.

VLANs: Virtual LANs (VLANs) provide the capability to overlay the physical network with multiple virtual networks. VLANs allow you to isolate network traffic between virtual networks and reduce the size of administrative and broadcast domains.

VPLS: Virtual Private LAN Service (VPLS), also called Transparent LAN Service (TLS) or virtual private switched network service, can deliver a point-to-multipoint Layer 2 VPN service over public networks.

VPN: Virtual Private Networks (VPN) are primarily private networks that interconnect remote networks through primarily public infrastructures such as the Internet.

VSM: HP IMC Voice Service Manager (VSM) module offers operators an integrated management solution that provides a comprehensive set of tools for managing converged voice and data networks easily and efficiently. IMC's VSM offers a single pane for voice resource and service management for 3Com and H3C voice infrastructures, including VCX® Connect platforms, Media Gateway, and IP phones.

WAN: Wide Area Network (WAN) typically refers to a network that covers a broad area.

WSM: With IMC's Wireless Service Manager (WSM) module, operators can perform wireless LAN WLAN device configuration, topology, performance monitoring, RF coverage and planning, WLAN intrusion detection and defense, and generate WLAN service reports from the same platform they use to manage wired networks.

WWN: A world wide name or world wide identifier (WWID) is a unique identifier which identifies a particular Fibre Channel, advanced technology attachment (ATA) or serial attached SCSI (SAS) target.

For more information

HP Converged Infrastructure

HP Converged Infrastructure white papers and videos

hp.com/go/ci

HP Converged Infrastructure Reference Architecture Guide

http://h20195.www2.hp.com/v2/GetPDF.aspx/4AA2-6453ENW.pdf

HP FlexFabric

HP FlexFabric white papers and videos

hp.com/go/flexfabric

HP Intelligent Resilient Framework

HP IRF white paper–reducing network complexity, boosting performance with HP IRF technology

http://h10144.www1.hp.com/docs/irf/irf.pdf

HP Virtual Connect

HP Virtual Connect technology information

http://h18004.www1.hp.com/products/blades/virtualconnect/index.html

HP Intelligent Management Center

HP IMC data sheets and product details

http://h17007.www1.hp.com/us/en/products/network-management/index.aspx

HP TippingPoint Security

HP TippingPoint data sheets and product details

http://h17007.www1.hp.com/us/en/products/network-security/index.aspx

HP Networking Services

HP Networking Services brochures and videos

http://www8.hp.com/us/en/business-services/index.html#/page=1&/sort=csdateweblaunch|DESC&/cc=us&/lang=en

Get connected

hp.com/go/getconnected

Current HP driver, support, and security alerts

delivered directly to your desktop

Created April 2011; Updated April 2013, Rev. 3.0