



HP FlexFabric Reference Architecture Overview

Building modern, virtualization-optimized data center networks

Introduction.....	2
Business and technology drivers for new data center architectures.....	2
Data center consolidation	2
Increasingly powerful and scalable compute	2
New application deployment and delivery models.....	2
Virtualization and cloud mega trends	3
Current legacy data center design challenges	3
HP FlexFabric Reference Architecture—building a cloud-optimized network.....	5
HP FlexFabric—for an agile, cost-effective network from core to the edge.....	6
HP Data Center Ethernet Switches	7
Simplifying the data center network architecture	8
HP Intelligent Resilient Framework	8
HP and F5 Networks application delivery solutions.....	9
Interconnecting geographically dispersed data centers.....	10
HP FlexFabric Interconnect solution	10
I/O consolidation and network convergence	11
Multi-hop FCoE across the data center	11
Securing the virtual data center	12
Securing the FlexFabric: HP TippingPoint Secure Virtualization Framework.....	13
Unifying physical and virtual network management	14
Delivering the cloud with Virtual Application Networks.....	15
Summary	16

Introduction

New application architectures and software deployment models are fundamentally transforming the data center. Server virtualization, cloud computing, and everything-as-a-service (XaaS) imperatives are altering data center traffic flows, escalating bandwidth and performance demands, and introducing new security and service orchestration requirements. Legacy data center networks are simply too complex, costly, and rigid to meet the needs of the new on-demand world.

Tomorrow's virtualized data center demands more agile, efficient, and scalable networking solutions. HP believes customers should take a long-term view towards data center modernization and pursue strategies that protect and extend existing investments and minimize disruptions. Customers should seek approaches that maintain continuity with existing storage, server, and network management practices, and choose standards-based solutions that gracefully evolve to meet emerging scalability and performance needs.

HP offers a flexible, virtualization-optimized data center network architecture that requires far fewer devices, interconnections, layers, and discrete appliances. HP Networking solutions streamline network operations and accelerate application and service delivery; reduce space, power, cooling, and capital requirements; and protect investments while providing a solid foundation for the future.

This white paper reviews data center trends and describes HP solutions for building cost-effective, advanced data center networks that meet the evolving performance, reliability, and agility demands of the 21st century.

Business and technology drivers for new data center architectures

Data center consolidation

For many enterprise customers, the data center *is* the business. With mission-critical applications and services deployed to provide the foundation for day-to-day operations and delivery of end-customer services, the data center must deliver exceptional availability and meet stringent service-level agreements. Exploiting server virtualization and low-cost computing power, customers are deploying more and more sophisticated applications on a larger scale. To reduce the sheer complexity and improve operations of these deployments, customers are seeking to consolidate fragmented, dispersed facilities into fewer, centralized locations. These new 'mega' data centers are fundamentally challenging how networks must be built. Today's networks must be designed to deliver much higher levels of performance, scalability, and availability than previously expected in order to meet service-level agreements (SLAs) 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 help ensure continued performance and minimize service disruptions.

Increasingly powerful and scalable compute

The adoption of increasingly powerful multicore processor servers, higher-bandwidth interfaces and blade servers is dramatically increasing the scale of data center deployments. Now, thousands of virtual machines can be 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 are introducing high-volume machine-to-machine traffic flows and impacting existing administrative practices, creating a new 'virtual edge' that 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 mashups, 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, virtual machine-to-virtual machine 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.

Virtualization and cloud mega trends

Enterprises are turning to the cloud to improve business agility, reduce expenses, and accelerate business innovation. Cloud computing redefines the way business deploys and consumes IT assets and architects and manages data center networks. Conventional hierarchical data center networks built to support traditional siloed IT architectures can't meet the security, agility, and price/performance requirements of virtualized cloud computing environments. Public cloud service providers and enterprises deploying private clouds must implement flatter, simpler data center networks to support the bandwidth-intensive, delay-sensitive server-to-server traffic flows that accompany cloud computing. Enterprises must also evolve and adopt new management systems and practices to administer and secure virtual resources and orchestrate on-demand services.

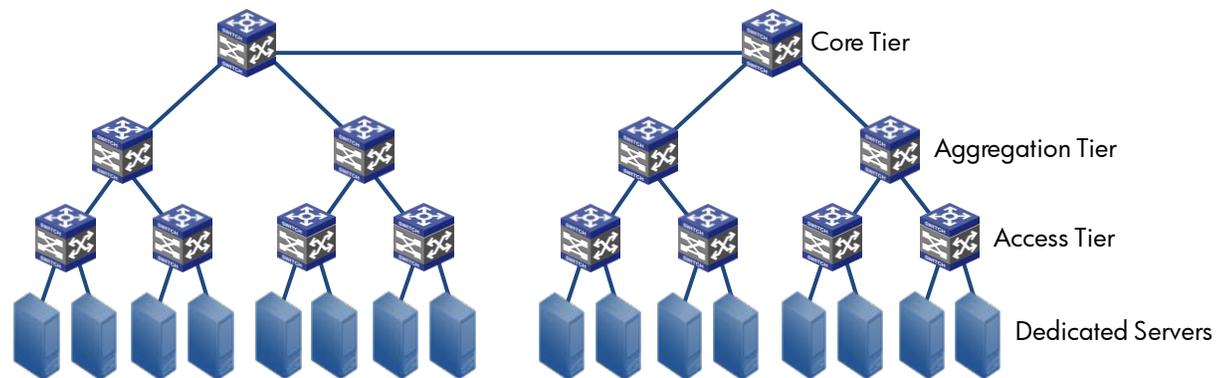
"The consequence of changing application deployment and increased use of VM migration will mean that traffic patterns in the data center network are changing from being predominantly client/server (north-south) to a significant level of server-to-server (east-west) flows. By 2014, network planners should expect more than 80 percent of traffic in the data center network to be between servers."¹

Gartner, April 2011

Current legacy data center design challenges

Today's data center networks were designed to support conventional siloed IT architectures in which servers are dedicated to specific functions or organizations and the vast majority of traffic flows in and out of the data center. Most data center networks are based on hierarchical designs comprised of an access tier, an aggregation tier, and a core tier (figure 1).

Figure 1. Legacy three-tier data center network architecture



The access tier is made up of Ethernet switches connecting rack servers and IP-based storage devices (typically 10/100Mbps or 1GbE connections). The access switches are connected via Ethernet to a set of aggregation switches (typically 1/10GbE connections), which in turn are connected to a layer of core switches or routers that forward traffic to an intranet, the Internet, and between aggregation switches. Layer 2 VLANs are typically implemented across the access tier and aggregation tier, and Layer 3 routing is implemented in the core.

Bandwidth is typically over-provisioned in the access tier, and to a lesser extent in the aggregation tier. The server infrastructure and the networking infrastructure are typically administered independently, by separate teams using distinct toolsets. Each server is dedicated to a specific function (e.g., Web server, application server, database server) and can be reasonably well protected using conventional security solutions such as intrusion prevention systems.

¹ Source: "Your Data Center Network Is Heading for Traffic Chaos, Bjarne Munch," 27 April 2011/ID Number: G00210674, Gartner

Hierarchical networks aren't well suited for server-to-server communications. Server-to-server traffic must traverse multiple layers of switches. Each hop adds delay.

"This means that network design must change. Traditional three- and four-tier data center network architectures focus on aggregating traffic flows from servers to users. This is not an optimum design for traffic flow between servers, because it may have to move through several switches to the core to get to other servers. Thus, the traffic would encounter a number of switch hops, as well as move through high user-traffic aggregation points. Instead, networks must be designed to support arbitrary traffic flows, which means a meshed topology, instead of the traditional tree topology. This requires changes to the physical and logical topology design."²

Gartner, April 2011

Contemporary data center networks designed to support siloed IT architectures simply can't meet the performance, security, availability, and agility requirements of the cloud. Public cloud service providers and enterprises deploying private clouds must implement simpler and more efficient networks that support the bandwidth-intensive, delay-sensitive server-to-server traffic flows and stringent SLA and security demands that accompany cloud computing. They must also adopt new management systems and practices to orchestrate on-demand services and administer and isolate virtual resources. Requirements for the new cloud-optimized data center network include the following:

- **Low-latency server-to-server connections:** Today's three-tier hierarchical networks aren't well suited for high-volume server-to-server communications. Inter-server traffic is forced to traverse multiple layers of switches, and each switch adds latency to the connection (figure 1). Enterprises and service providers must implement flat, low-latency data center networks to accommodate the delay-sensitive, volume-intensive server-to-server traffic flows that accompany cloud computing models.
- **Greater performance and resiliency:** Hierarchical data center networks typically rely on some variant of the spanning tree protocol (STP) for resiliency. STP is 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 protocol automatically selects a backup path. STP can take several seconds to recover from link failures and is not well suited for delay-sensitive applications. Cloud computing requires more efficient and resilient network designs that make full use of networking resources (no idle backup paths) and recover from failures in milliseconds to meet stringent resiliency requirements.
- **Large Layer 2 domains:** VM migration (vMotion/Live Migration)—the ability to seamlessly move VMs from one physical server to another without impacting applications or users—is driving the requirement for large-scale Layer 2 domains that offer high throughput and low latency. VM migration is critical for executing routine maintenance, business continuity, and disaster recovery functions in the cloud.
- **Higher bandwidth at the server edge:** Blade servers and server virtualization pack more and more computational power into smaller and smaller form factors—increasing bandwidth demands at the server edge, driving requirements for new switching solutions that offer better performance and greater port densities.
- **Reliable application delivery across public or private cloud:** Providing maximum application performance and availability. Organizations responsible for private cloud services must deliver predictable and reliable network services to support mission-critical business applications and address the objectives of internal service-level commitments.
- **High availability, Business Continuity (BC)/Disaster Recovery (DR):** With hundreds or even thousands of virtualized applications now in play across multiple, consolidated data centers, network resiliency and high availability (HA) take on a new, heightened level of importance. Network platforms and designs must be able to recover quickly from hardware and software faults to maintain continuity of service and business continuity.
Enabling distributed workloads and replicating data and applications across multiple, geographically dispersed data centers are a challenge. Conventional Layer 3-oriented WAN solutions can't meet the stringent performance and latency requirements—and server virtualization technologies require contiguous network domains. Customers wishing to extend and connect Layer 2 networks across data centers require connectivity and technologies that stretch those networks across multiple physical sites.
- **Converged network infrastructure:** With the proliferation of virtualization and the escalating demand for communication and storage performance, I/O sprawl has IT at the breaking point. The traditional model of completely parallel, autonomous data and storage networks with dedicated interface cards, switches, and cabling

² Source: "Your Data Center Network Is Heading for Traffic Chaos, Bjarne Munch," 27 April 2011/ID Number: G00210674, Gartner

plants can be costly and inefficient. Enterprises are looking to consolidate server and storage connectivity to reduce equipment and operations expenses; eliminate clutter and complexity; and make efficient use of shared networking resources while ensuring continuity of service.

- **Unified management:** Server virtualization creates a new virtual edge that blurs the traditional boundaries between network and server administration and introduces a variety of operational challenges. New tools are required for efficiently administering virtual switches, servers, and connections; for orchestrating on-demand applications and services; and for ensuring SLAs and enforcing service policies as VMs migrate across the data center.
- **Virtualization-aware security:** The virtual edge is beyond the scope of existing security systems and practices. In contemporary data centers, distinct workloads (database, application, web-hosting) and tenants are deployed on discrete physical servers. Workload-to-workload communications always occur over physical connections and can be secured using conventional intrusion prevention tools. With server virtualization, workloads can communicate over virtual connections within the same server in a manner transparent to existing network-based intrusion prevention systems. New 'virtualization-aware' security solutions to police intra-server communications flows, protect virtual resources, and partition multitenant environments must be considered.

Legacy Design vs. Emerging Cloud Requirements

Traditional Legacy Network Design	Emerging Cloud Network Requirements
80% Traffic North to South	80% Traffic East to West
Inter-rack latency 100+us	Inter-rack latency 15+us
Up to 64:1 oversubscription	1:1 oversubscription
10/100M and GbE attached rack servers	10/40 GbE attached Blade Servers
Siloed, serial application/network provisioning model, CLI	Agile, flexible control plan provisioning
Designed for peaks (over-provisioned)	Right-sized design burst as needed
High cost and complexity	Lower (shared) cost and complexity

The bottom line is that contemporary data center networks designed to support siloed IT architectures simply aren't designed meet the performance, security, availability, and agility cloud requirements.

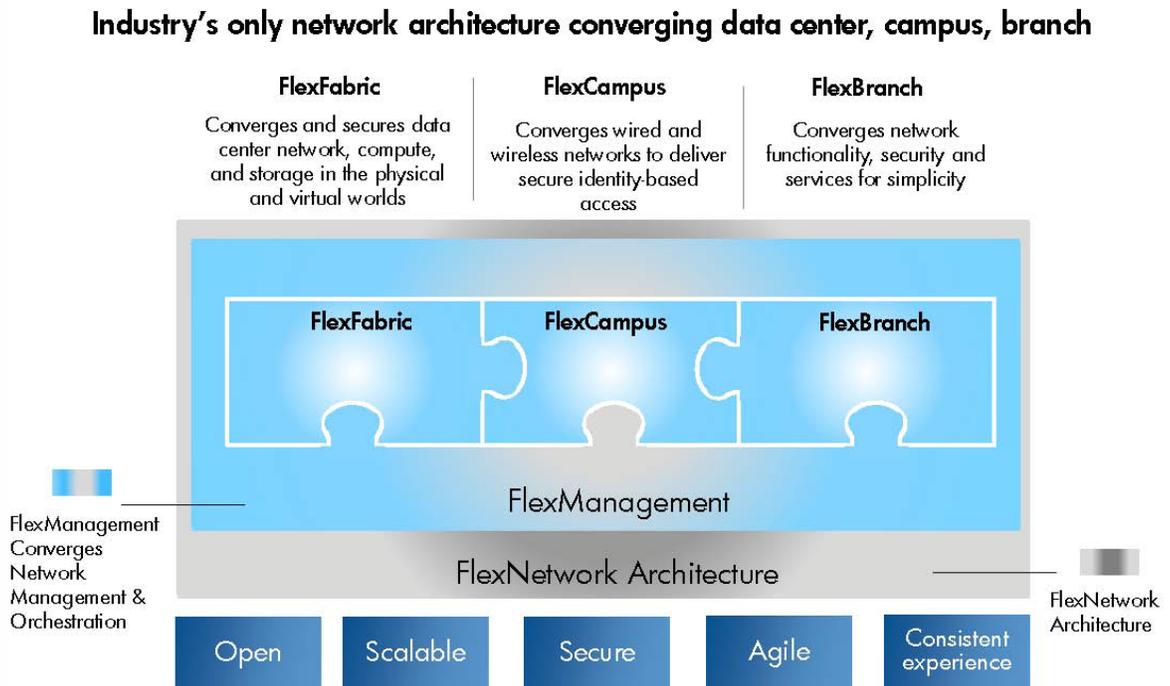
HP FlexFabric Reference Architecture—building a cloud-optimized network

HP FlexFabric Reference Architecture is an integral part of HP FlexNetwork Architecture, which is HP's blueprint for cloud-optimized networking which enables enterprises to align their networks with their business needs. With FlexNetwork, enterprises can segment their networks into the four interrelated modular building blocks: FlexFabric, FlexCampus, FlexBranch and FlexManagement (figure 2).

HP FlexFabric and **HP FlexCampus** enable the construction of flat, low-latency data center and campus networks with fewer layers, less equipment and cabling, and greater port densities. **HP FlexBranch** includes comprehensive WAN optimization and routing solutions for delivering dynamic cloud-based services to geographically distributed enterprises. **HP FlexManagement** provides a unified view into the virtual and physical network infrastructure that accelerates application and service delivery, simplifies operations and management, and boosts network availability.

HP FlexNetwork allows enterprises to securely deploy and centrally orchestrate cloud-optimized architectures that scale from the data center to the network edge. It enables the construction of flatter, simpler data center networks to support the bandwidth-intensive, delay-sensitive server-to-server virtual machine and workload traffic flows that accompany cloud computing, and it provides rich management tools for administering and securing virtual resources as well as orchestrating on-demand services.

Figure 2. HP FlexNetwork lets enterprises align networks with business needs.



HP FlexFabric—for an agile, cost-effective network from core to the edge

Flexible network designs are at the core of HP's approach to building data center networking solutions. HP Networking platforms are built using open-standards technologies and built to interoperate with the entire range of third-party server interfaces and standards-based switches and routers across Layer 2, Layer 3, IPv4, IPv6, MPLS, and VPLS protocol deployments. This compatibility helps ensure cohesion with existing network infrastructures and provides flexibility to integrate best-in-class third-party capabilities.

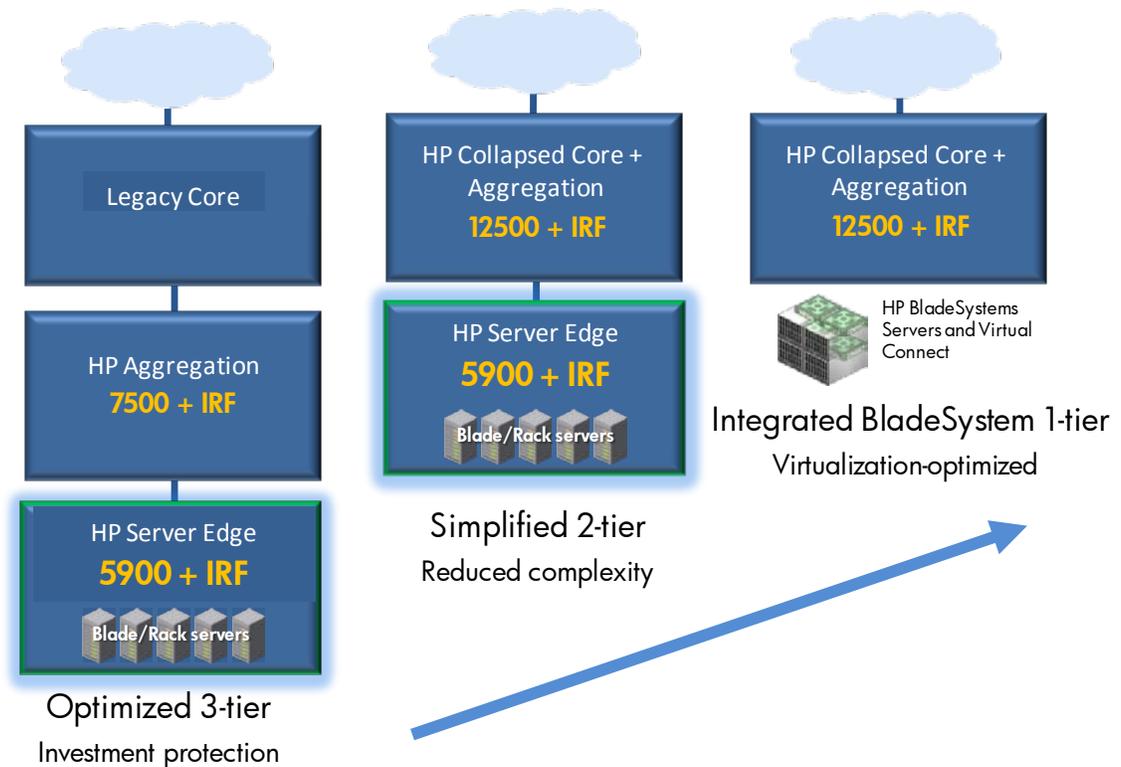
Customers looking to protect investments in legacy core infrastructure can implement a three-tier traditional network design and deploy cost-effective HP top-of-rack server edge and aggregation platforms that interoperate with their existing core switches. This approach allows customers to preserve existing assets and gradually migrate to a more agile network design over time while enjoying the benefits of Intelligent Resilient Framework (IRF) switch virtualization and cost-effective, energy-efficient HP enterprise data center switches in the server edge and aggregation layers.

Customers looking to reduce cost and complexity can implement a two-tier collapsed network design that completely eliminates a dedicated aggregation layer. These designs leverage HP Virtual Connect or HP 58x0/59x0 series switches at the server edge along with highly scalable HP12500 series core switches as a collapsed core/aggregation layer. These flat network designs help ensure direct-flight server-to-server performance while dramatically reducing network port counts. A two-tier collapsed design also simplifies and streamlines network management, and reduces capital expense and energy consumption.

Customers can reduce capital and operating expenditures even further by implementing an optimized one-tier collapsed network design using HP BladeSystem Servers and HP Virtual Connect, which provides the ultimate in simplification, agility, and cost optimization.

HP Enterprise Ethernet Switches and innovative IRF technology enable flat, low-latency network designs to support highly virtualized data centers. In addition, enterprise switches offer industry-leading server edge port density to meet the increased bandwidth demands that accompany the implementation of server virtualization and blade server technology.

Figure 3. FlexFabric supports a range of network designs to support diverse customer requirements



HP Data Center Ethernet Switches

HP 12500 series core switches leverage the latest generation of ASICs and a fully non-blocking design based on a CLOS architecture to deliver ultimate performance, density, and scalability. The product family delivers 13Tbps performance and offers very high port density today (512 10GbE per chassis) with support for 40GbE and 100GbE connections in the future.

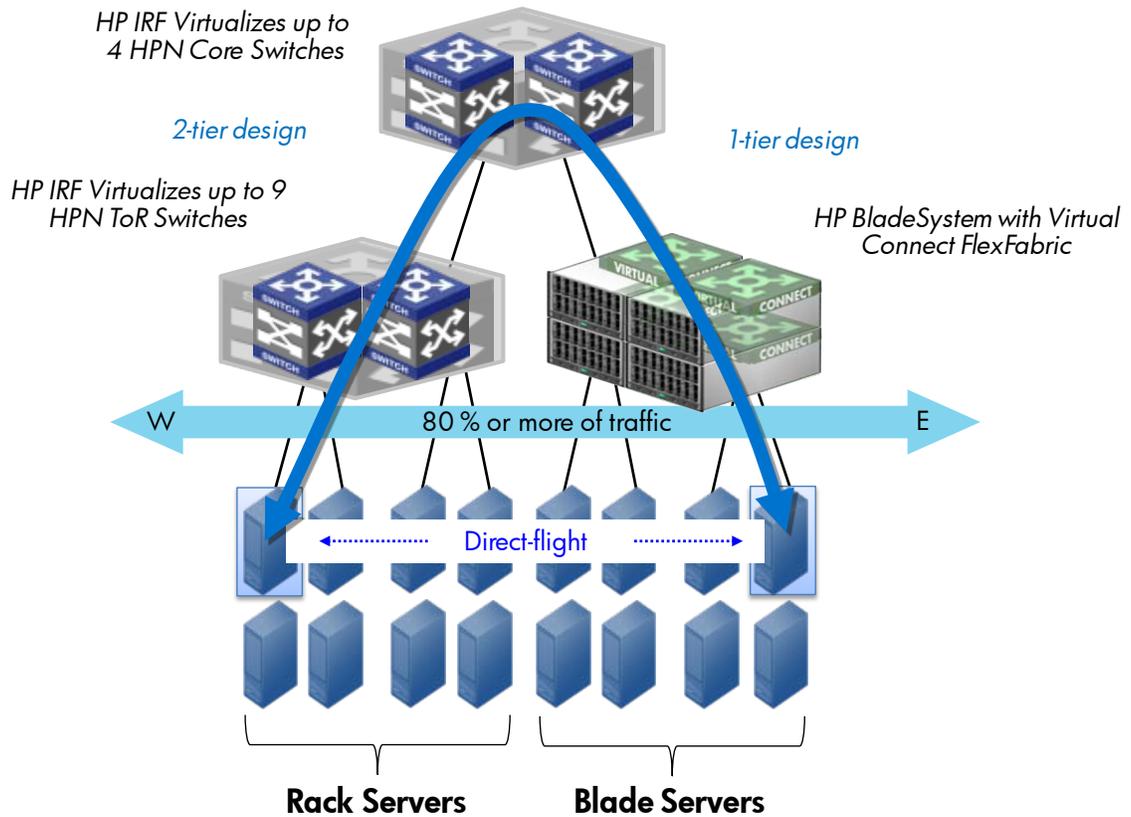
HP 58x0 and 59x0 Top of Rack (ToR) switches leverage ultra deep packet buffers, cut-through switching technology and a high-availability architecture to deliver line-rate, low-latency performance and outstanding reliability at the server edge. The product family's high port density (up to 48 10GbE and 4 40GbE ports per unit) meets escalating bandwidth demands at the server edge.

HP provides flexible solutions for delivering high-performance server-to-server connectivity at the server edge. HP solutions can directly interconnect hundreds of virtual machines at the edge of the network, eliminating unnecessary network hops, reducing latency, and optimizing performance for high-volume server-to-server traffic flows.

For traditional, top-of-rack server edge installations, HP 58x0/59x0 ToR switches can be deployed with IRF virtualization technology to provide high-throughput, low-latency server-to-server connectivity at the server edge. With IRF, multiple switches can be virtualized and logically combined to enable low-latency, ultra-resilient, virtual switching fabrics comprising hundreds or even thousands of 1GbE or 10GbE switch ports—all managed via a single IP address.

For BladeSystem deployments, HP Virtual Connect delivers direct server-to-server connectivity within the rack, enabling wire-speed, machine-to-machine communications for delay-sensitive, bandwidth-intensive traffic. In addition, HP Virtual Connect Flex-10 and FlexFabric modules can be leveraged to dynamically fine-tune application-specific performance across server and storage networks to improve scale and make best use of shared connectivity resources.

Figure 4. High-performance server-to-server connectivity: HP data center ToR switches with IRF for rack servers; HP Virtual Connect for deployment in HP BladeSystem servers



Simplifying the data center network architecture

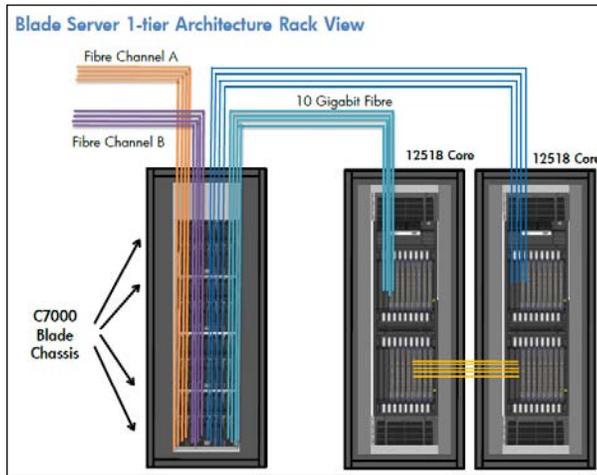
In the core of the network, HP 12500 switches can be deployed in conjunction with IRF to completely eliminate the aggregation layer found in conventional three-tier data center networks. IRF overcomes the limitations of legacy spanning tree networks by fully leveraging all network connectivity (no inactive backup paths) and by providing rapid failover to dramatically improve network utilization and performance in the network core.

A collapsed, two-tier data center network architecture enables direct-flight server-to-server performance, requires significantly fewer connections and port counts (no aggregation switches), streamlines provisioning and network management, and reduces capital expense and energy consumption. In addition, these two-tier networks provide large Layer 2 domains to enable VM migration across the data center (move workloads from one server to another server in the same VLAN/IP subnet.)

HP Intelligent Resilient Framework

By deploying IRF in conjunction with high-performance HPN switches, enterprises can directly interconnect hundreds of virtual machines at the edge of the network, eliminating unnecessary network hops, reducing latency, and improving performance for large intra-data center workloads. In the core of the network, HP leverages highly scalable modular platforms and the same IRF technologies deployed at the server edge to significantly improve scalability and provide unified, high-performance, multi-data center network extension. IRF overcomes the limitations of legacy spanning tree networks providing rapid failover for delay-sensitive, mission-critical applications and dramatically improving network utilization and performance in the network core.

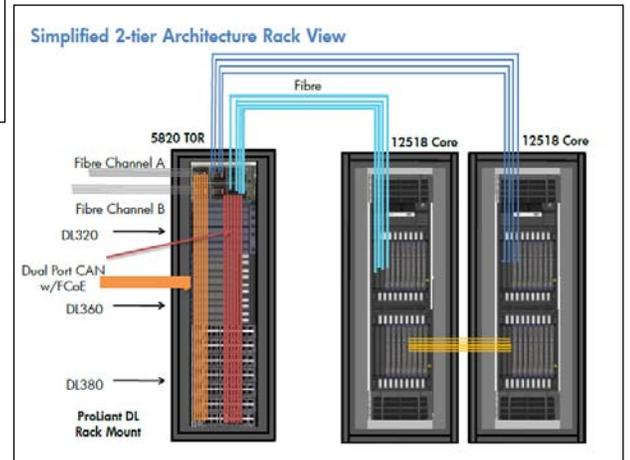
Figure 5. HP Networking 1 and 2 tier architectural designs (rack view)



To visualize the layers in a real environment, consider this 1-tier design. Logically, the network is divided into one core layer with an HP Networking 12500-series switch connected directly to an HP BladeSystem via Virtual Connect FlexFabric interconnects providing a simplified, high performance, low latency design.

Logically, the simplified 2-tier network is divided into two layers. The access layer consists of HPN 5800- and 5900-series LAN switches. The 12500-series switches are set in the core layer to provide a highly available, high-performance core in the network.

Both the 5800- and 5900-series switches in this design are configured as L2 switches, although they can also handle full L3 switches and routing to ensure ultralow latency, optimized server to server application traffic flow.



“By leveraging IRF, operators can dramatically simplify both initial network design and ongoing network operations— as much as 85 percent or more—by consolidating configuration files. The IRF approach allows each IRF domain (as many as 9 devices) to operate from a single configuration file. This reduction slashes administrative load, while also drastically reducing the potential for introduction of function-impacting manual configuration errors.”³

Enterprise Management Associates, Inc. (EMA)

By deploying IRF in conjunction with highly scalable HPN modular switches, customers can completely eliminate the requirement for a dedicated aggregation layer as they scale-out data centers. Plus, they can enjoy the benefits of large Layer 2 domains with increased network uptime and simplified management. To support inter-data center virtual machine mobility, HP offers options for flexibly and reliably extending large-scale Layer 2 domains across geographically dispersed sites.

HP and F5 Networks application delivery solutions

Application delivery is a suite of technologies that comprises application security, application acceleration, and network availability. It ensures that applications are always secure, fast, and available across any network.

An Application Delivery Controller (ADC) is a device that is typically placed in a data center between the firewall and one or more application servers (an area known as the DMZ). First-generation application delivery controllers primarily performed application acceleration and handled load balancing between servers.

The newest generation of application delivery controllers can handle a wider variety of functions, including global server load balancing, support for real-time protocol manipulation (for example, Transmission Control Protocol [TCP]

³ Source: "HP IRF Simplifies Networks and Operations," Enterprise Management Associates, Inc, 2011. Please see <http://h20195.www2.hp.com/v2/GetPDF.aspx/4AA3-6079ENW.pdf>

connection management and HTML optimization), data compression, network address translation, traffic shaping, Quality of Service (QoS), DOS protection, SSL offload, and resiliency/failover.

HP's partnership with F5 Networks (the undisputed ADC market share leader) provides customers with an industry-leading value proposition focused on addressing mission-critical application performance demands, virtualized data centers, and cloud services.

This partnership delivers a host of differentiated customer solutions from cloud-enabled data centers to L4-7 application optimization and server load balancing to business continuity (BC) and disaster recovery (DR)—all backed by fully tested reference architectures, management integration, and global professional services.

HP and F5 Networks Solutions include:

- **Application/VM, Server Optimization**
Optimize application traffic using load balancing, acceleration, yielding improvements in performance/availability
- **Microsoft Exchange Server 2010**
Quickly and securely replicate mailboxes, increase Exchange server capacity, simplify Web access, consolidate infrastructure, reduce download times for end users, and reduce spam
- **Business Continuity/ Disaster Recovery**
Control global traffic to ensure maximum availability and performance of applications running across multiple data centers
- **Cloud Optimized Networking**
Deliver reliable application services across public or private cloud with maximum application performance and availability
- **Live Migration between data centers**
Enable vMotion migration between data centers without downtime or user disruption—automate, encrypt, accelerate

HP and F5 solutions are an integral part of HP's FlexFabric Reference Architecture. These solutions are designed to help customers create a dynamic and responsive data center infrastructure that helps organizations align IT functions to constantly changing business needs—ensuring their applications stay secure, fast, and available.

Interconnecting geographically dispersed data centers

HP FlexFabric Interconnect solution

Business continuity, performance through workload mobility and SAN extension

The HP FlexFabric Interconnect is an innovative solution that enables the interconnection of geographically dispersed data centers. As enterprises expand their data centers to meet ever-growing customers' needs, they are adopting virtualized environments and cloud computing. This in turn drives the need to deliver Infrastructure as a Service (IaaS) and to connect geographically dispersed data centers to meet rising customer expectations.

Unfortunately, current interconnect methods suffer from limitations, including transport dependency, complexity, and lack of resiliency. HP FlexFabric Interconnect is designed to address these limitations by delivering responsive, efficient, and resilient data center interconnect solution. The HP FlexFabric Interconnect solution is transport agnostic and extends Layer 2 connectivity across the networks of the various data centers. It enables several data centers to work as one that is more responsive, with higher efficiency and solid high availability for business resiliency. With FlexFabric Interconnect, enterprises are able to deliver workload mobility with remote vMotion, increase applications performance with load balancing, and achieve optimum degree of HA and disaster recovery for valuable data. When used along with HP's IRF switch virtualization technology, FlexFabric Interconnect delivers greatly enhanced reliability, resilience, and faster remote vMotion capabilities.

Key highlights:

- Innovative transport-agnostic solution that can be deployed without requiring changes in existing infrastructure, which simplifies deployment and allows HP FlexFabric Interconnect to be added to your data center seamlessly and without disruptions.

- HP FlexFabric Interconnect extends Layer 2 connectivity across the network, eliminating the need to deal with Layer 3 interconnect dependencies; it also makes loop isolation implementation and undesirable failure prevention possible. HP FlexFabric Interconnect enables reliable vMotion support for mobile workloads that are necessary for virtualized and cloud computing environments. When coupled with HP IRF, the combination gives enterprises a reliable means for remote vMotion and increased available network bandwidth.
- HP FlexFabric Interconnect solution offers greater scalability as it supports four data centers that can be expanded to eight sites in the future.

I/O consolidation and network convergence

With the proliferation of virtualization and the escalating demand for communication and storage performance, I/O sprawl has IT at the breaking point. The traditional model of completely parallel, autonomous data and storage networks with dedicated interface cards, switches, and cabling plants can be costly and inefficient. Enterprises are looking to consolidate server and storage connectivity to reduce equipment and operations expenses, decrease clutter and complexity, and make more efficient use of shared networking resources while ensuring continuity of service.

HP is playing a key role in driving (Fibre Channel over Ethernet) FCoE and the set of emerging Data Center Bridging (DCB) standards in the IEEE and T11. By committing to delivering proven, standards-based solutions, HP Networking products give customers a seamless path to data center network consolidation that delivers convergence within the rack today, and provides an orderly migration to a fully converged LAN/SAN infrastructure as FCoE and DCB gain maturity in the market and become commercially available across the full spectrum of data center products.

In addition to having the potential to reduce complexity at the network edge, FCoE also promises to reduce equipment costs beyond the access layer by enabling convergence in the extended network. However, given the critical nature of storage networking and its central importance to the integrity of the end-to-end data center architecture, customers should carefully evaluate the implications of these solutions when considering more extensive FCoE deployments.

Multi-hop FCoE across the data center

FCoE is a strong edge technology for native Fibre Channel networks, with 1-hop FCoE proven and 2-hop FCoE viable in selected cases. Today, end-to-end FCoE is viable only in single-vendor proprietary networks for the foreseeable future. Any deployment of end-to-end FCoE should be carefully scrutinized for cost, benefit, risk, and ROI.

A completely unified, end-to-end shared infrastructure built to transport server and storage traffic over a common Ethernet-based network also requires much more stringent maintenance and upgrade practices, as these activities now have the potential to affect ALL traffic in the data center. This inherent risk will be mitigated over time with the advent of mature, tested solutions based on ratified industry standards that yield robust, fault-tolerant network designs and assured operations.

In the interim, a 'common but separate' approach—in which customers continue to deploy separate extended server and storage networks using common, FCoE-enabled networking building blocks—will allow customers to enjoy the incremental cost savings of Ethernet-based technologies deployed uniformly across both network infrastructures.

Complementing other widely implemented Ethernet-based networking technologies such as Network-attach storage (NAS), Direct-attached storage (DAS) and iSCSI, HP Networking offers FCoE-based solutions that provide customers even greater flexibility to meet varying storage networking demands.

I/O Consolidation with FCoE/DCB: With HP ToR data center switches and Virtual Connect FlexFabric modules, enterprises can consolidate LAN and SAN I/O connectivity onto a common FCoE server edge fabric within the rack while preserving investments in extended LAN and SAN infrastructures. This approach allows customers to remove expense and complexity, and improve network performance without impacting the installed LAN/SAN infrastructure or disrupting existing management practices.

This approach can also simplify operations by enabling common equipment sparing and unified Ethernet network-based management tools. In a common but separate model, storage networks can be independently designed and tuned to address unique storage traffic patterns, giving storage network designers more freedom to manage network congestion compared to a completely shared network infrastructure. The end result is reduced risk and management complexity in the short term with a solid foundation for transitioning from dedicated Fibre Channel-based networking to Ethernet-based networking over the long term.

As FCoE-based solutions continue to evolve, so too will a wider array of Ethernet-based storage networking technologies. Customers should carefully characterize consolidation/simplification opportunities and consider FCoE alongside NAS, DAS, and iSCSI as potential viable approaches. For those instances in which FCoE is a good choice, HP recommends that customers start at the server edge—where the real-world savings are most substantial—and then extend the solution over time. An evolutionary approach reduces cost and complexity in the near term while maintaining continuity of operations and maximizing investments over the long haul.

Figure 6. Converged Network Candidates

	FCoE on DCB	iSCSI	InfiniBand	iSCSI on DCB
Directory, security, other SAN services	Fibre Channel switch software (feature rich, single-vendor network)	iSNS (open source)	(minimal)	iSNS (open source)
Flow control (one hop)	Per priority pause (PFC) (part of DCB)	TCP	Per priority pause	PFC
Flow control (end to end)	QCN (part of DCB)	TCP	Manual tuning of workload	iSER or QCN
How would a cluster connection be added here?	RoCE	iWARP (RDMA over TCP)	InfiniBand is the lowest latency, fastest cluster network today	RoCE
Comments	Best coexistence with and transition plan for Fibre Channel	Success in smaller environments and not enterprise	Success only in low latency and supercomputer environments including storage system internals	Emerging limited vendor support

HP recommends exploring storage connection alternatives and estimating what percent of your storage needs are best met by block storage on a SAN (Fibre Channel, FCoE, iSCSI), what percent are best met by file storage (NAS, and at what scale), and what percent are best met by DAS.

Securing the virtual data center

Server virtualization introduces a new ‘virtual edge’ that significantly impacts traditional network and server security systems and practices. Previously, database, application, and web-hosting workloads were deployed on discrete physical servers. Traffic flows between workloads were wholly contained within a single physical server, allowing straightforward, physical network-based threat management using traditional intrusion prevention tools. Server virtualization introduced the ability to host multiple workloads on a single physical server and initiated the concept of virtual switches (vSwitches) to facilitate intra-server communications. VMs can communicate directly with each other and pass traffic in a manner that is transparent to conventional security systems and never traverses the physical network fabric. Practically speaking, the virtualized nature of these intra-server communications makes applying security policies or monitoring the network very difficult. Existing security tools and practices built around physical servers and physical switches are unaware of these traffic flows. Specifically, virtualization introduces a wide variety of new security challenges:

- **Hypervisor security**
New security procedures are needed to safeguard the VM hypervisor.
- **VM-to-VM threats**
Traffic moving from one virtual machine to another inside the same physical host is not visible to the external network, meaning the traffic cannot be subjected to security, QoS, management, or mirroring policies.
- **VM mobility**
Security policies must be preserved as VMs migrate from host to host within the data center.
- **Host-to-host threats**
Virtualization increases host-to-host communications and requires cost-effective solutions to inspect and control server-to-server traffic. Deploying an IPS in front of every server is impractical.

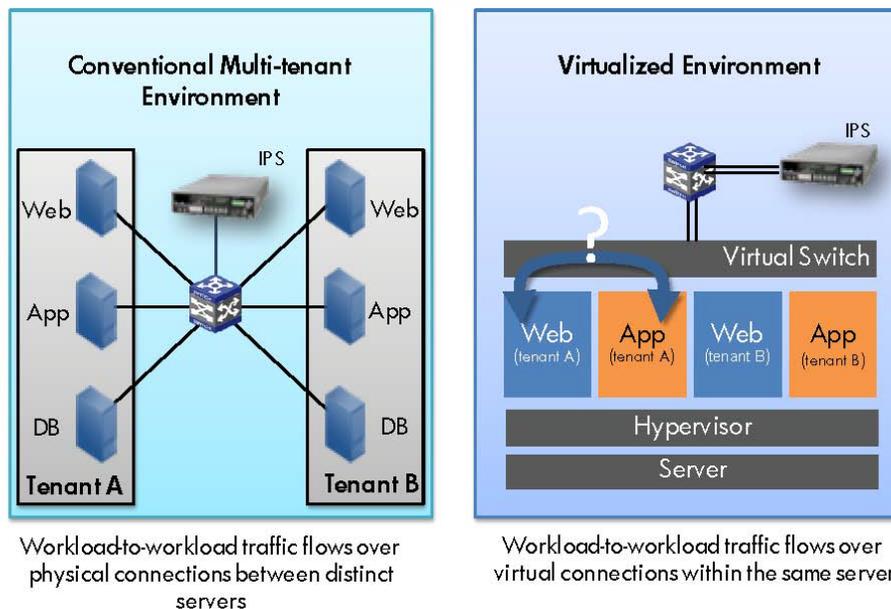
Securing the FlexFabric: HP TippingPoint Secure Virtualization Framework

HP TippingPoint Secure Virtualization Framework (SVF) enables unified security across virtualized and physical domains, safeguarding VM-to-VM as well as inter-server and inter-network traffic in a common platform. The framework streamlines administration and reduces operations expenses by centralizing and automating security management functions. Administrators define rich, infrastructure-wide security policies, which are implemented across virtual machines and virtual switches in a transparent fashion. The solution helps cloud service providers and enterprises implementing public clouds address the unique challenges associated with partitioning distinct user communities and securing multi-tenant environments.

SVF brings best-of-breed TippingPoint intrusion prevention, threat mitigation, and security management features to the virtual edge, safeguarding IT assets and optimizing service availability. HP TippingPoint vController—an integral SVF component—works with an HP TippingPoint N-Platform IPS to provide high-performance intrusion prevention for a virtualized server. A software-based solution that is easily installed in a virtualized server, vController directs virtual machine traffic to an N-Platform where robust intrusion protection services are applied with line-rate performance (figure 7). The solution segregates virtual resources and inspects and polices intra-server traffic flows, providing consistent, unified security across virtualized and physical data center network infrastructures.

HP's high-performance security solutions allow enterprises to support dynamic traffic demands and service flows without adding devices and complexity or compromising security. Enterprise private cloud data centers may 'burst' to public clouds from time to time in order to accommodate demand spikes; therefore, data centers must ensure real-time protection of both their private cloud and the virtual machines delivering burst capacity in the public cloud. HP TippingPoint Secure Virtualization Framework delivers real-time security as VMs are provisioned in the public cloud, helping to ensure seamless, secure, and context-aware protection, enabling enterprises to dynamically shift and re-allocate services without sacrificing security.

Figure 7. Conventional multi-tenant environment vs. virtualized environment



With flexible deployment options, vController + Firewall integrate with high-availability network fabrics and automated vSphere deployment provides streamlined installation and configuration of the virtual environment. The HP TippingPoint Virtualization Management Center (VMC) is integrated within the Secure Virtualization Framework and included with vController + Firewall. VMC allows security teams to monitor and enforce security policies throughout the virtual environment, which enables integration of virtualization security within organizations' existing roles, responsibilities, and best practices. Tightly integrated with vSphere, VMC provides real-time visualization of all vCenter managed virtual environments. With the ability to scale across multiple vCenter servers, VMC provides a

single unified environment for efficiently managing security across the entire virtual environment, offering the unique capability to define virtual security zones and policies using VMware attributes in addition to traditional Layer 2 and Layer 3 traffic attributes. This environment enables security teams to develop highly automated zones and policy definitions that adapt to the ever-changing virtual environment in real time—management solutions that enable organizations to build security zones and policies based on VM and/or infrastructure attributes in order to enhance security and implement best practices for separating duties. Attribute-based definitions provide protection from accidental or malicious configuration changes that may otherwise create security vulnerabilities. And, as a VMware Global Technology Alliance Partner, HP works closely with VMware to validate product integration and interoperability.

Unifying physical and virtual network management

As described in the previous section, server virtualization introduces a new virtual edge that blurs the traditional demarcation between network and server administration. Previously, the network edge was defined as the point where the server connected to the switch, with servers statically deployed with a single operating system and a set of interfaces. With virtualization and the insertion of vSwitches in the server, system administrators now have greater ability to control, configure, and manage server connectivity. Most importantly, since administrators can easily migrate virtual servers between physical servers, the tools for managing connections to workloads must evolve to be more agile and accommodating of the dynamic nature of the environment.

Managing an increasingly virtual data center has become a daunting task for data center managers. Plus, managing the assignment and allocation of highly dynamic and mobile virtual servers across physical and virtual networks has added tremendous complexity to overall data center network operations and administration. The configuration of servers, virtual machines, and physical and virtual networks (vSwitches) can often be complex and difficult to coordinate across IT staff. Server adds, moves, and changes can be time consuming and error prone. The lack of a single-pane view of the virtual and physical network infrastructure makes troubleshooting difficult, if not impossible. In summary, virtualization introduces a host of new network management challenges and requirements such as:

- Configuring servers, virtual machines, and physical and virtual switches is a complex, time-consuming undertaking requiring coordination between network and server teams.
- Implementing adds, moves, and changes is a manual, error-prone process involving multiple applications and data center teams.
- Isolating and resolving problems is an arduous process involving multiple management systems with overlapping functionality.

HP Intelligent Management Center

HP Intelligent Management Center (IMC) unifies physical and virtual network management and helps IT overcome the challenges of administering the new virtual server edge. The solution provides a unified view into the virtual and physical network infrastructure that accelerates application and service delivery, simplifies operations and management and boosts network availability. Capabilities include the following:

- Automatic discovery of virtual machines, virtual switches, and their relationships with the physical network
- VM and virtual switch resource management, including creation of virtual switches and port groups
- Automatic and transparent configuration of virtual and physical network infrastructure
- Unified performance and alarm monitoring of hosts, workloads, and virtual switches
- Topology views and status indicators for networks, workloads, and virtual switches
- Automatic reconfiguration of network policies as workloads migrate across the data center

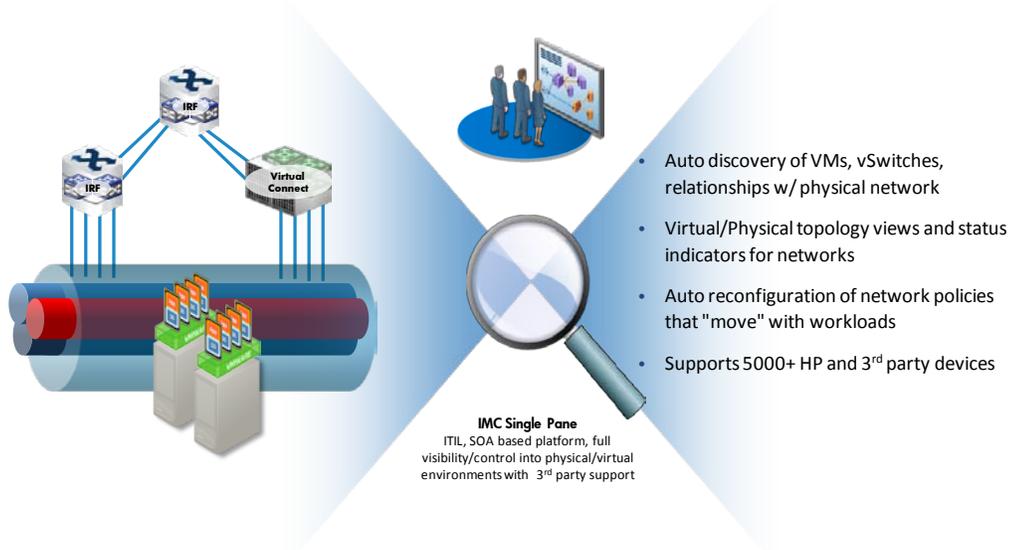
IMC offers a new proactive, dynamic, application-aware provisioning model with comprehensive solution integration to align with end-to-end IT operations. This new paradigm shifts businesses to a more agile model by eliminating unnecessary steps in virtualizing business environments. Time to deployment has been accelerated through the upfront definition of profiles with virtual machine connectivity characteristics which are filed in a library, rather than the use of an iterative manual process for defining network connectivity characteristics that cannot be leveraged and

repurposed. These ready-to-use profiles allow for rapid deployment and follow the workload if it is moved, paused, and/or resumed.

Figure 8. HP Intelligent Management Center

HP Intelligent Management Center (IMC)

Unify management for virtualized and physical resources



HP IMC can help eliminate service interruptions caused by virtual/physical network configuration errors; reduce administration and troubleshooting by providing unified management of physical and virtual network infrastructure through a single pane of glass; and accelerate the delivery of new applications and services by automating configuration of virtual and physical network infrastructure.

*"HP's IMC solution represents a mature, integrated approach for managing complex, multivendor, heavily virtualized infrastructures across multiple functional areas. IMC's integral support for security management sets it apart from the many others in the marketplace, as does its approach to covering the vast majority of management tasks within a single product."*⁴

Enterprise Management Associates, Inc. (EMA)

Delivering the cloud with Virtual Application Networks

Today, networks are managed as a series of individual components that are supervised by individual groups. The network is largely unaware of the applications. Network configuration and management is too labor-intensive, with administrators relying heavily on the command-line interface and writing scripts.

IT needs a way to shift to a more strategic command-and-control model, so IT staff can administer and secure network resources based on an awareness of the applications themselves, and more easily orchestrate on-demand services for global deployment. To separate network configuration and orchestration, enterprises need a way to separate the logical and physical provisioning of resources: enterprise network virtualization technologies. Just as server virtualization allowed IT to separate the allocation and management of physical and logical server resources, the same efficiencies can be brought to bear on the network with network virtualization in order to provision services as needed.

⁴ Source: "Seven Priorities for Integrated Network Management: How HP Intelligent Management Center Delivers an Enterprise-class Solution," March 2011, Enterprise Management Associates (EMA)

With HP Networking, enterprises can build a scalable, agile, and secure network that streamlines operations and enables a better operating model for IT. With innovative solutions from HP Networking, enterprises can overcome the limitation of legacy networks and move applications to the cloud with confidence.

HP Enterprise Network Virtualization allows businesses to simplify and automate network design, configuration, and monitoring. It allows IT to virtualize network resources from the user to the data center or cloud. HP Enterprise Network Virtualization runs across the HP FlexNetwork architecture—FlexFabric, FlexCampus, FlexBranch, and FlexManagement—so businesses can gain efficiencies across all functional areas of their organization.

Think of HP Enterprise Network Virtualization technologies as a control plane for the network—analogue to a hypervisor for server virtualization. IT can logically provision physical network resources into purpose-built VANs for each enterprise service or set of services. By separating infrastructure provisioning from infrastructure management, IT can focus on connecting users to applications—and the subsequent quality of experience—rather than the details of configuring the networks.

Summary

Enterprises are turning to the cloud to accelerate business innovation, improve business agility, and contain costs. Cloud computing reshapes the way applications are deployed and consumed, and influences data center network designs. HP helps enterprises build unified, virtualization-optimized data center networks that meet the rigorous performance, scalability, availability, and agility demands of the cloud. HP FlexNetwork networking solutions deliver:

- Flatter and more efficient data center networks with fewer layers, less equipment and cabling, and greater port densities
- High-performance, low-latency intra-data center connectivity for VM migration and bandwidth-intensive server-to-server communications
- Virtualization-aware security to partition multi-tenant environments and isolate virtual resources and intra-server communications flows
- Optimal WAN performance for the highest quality end-user and application experiences and most efficient use of WAN resources
- Unified administration and service orchestration to accelerate the delivery of cloud-based applications and services
- Multi-site, multi-vendor management to connect and control thousands of physical and virtual resources from a single pane of glass

To learn more about how HP can help you build a cloud-optimized data center network, please contact your HP account manager or reseller.

For additional information, please visit:

HP Networking Solutions

HP Networking Solutions home page: hp.com/go/networking

HP Cloud Solutions

HP Cloud Solutions home page: hp.com/go/cloud

HP CloudSystem home page: hp.com/go/cloudsystem

HP Data Center Switches

HP Switches data sheets and product details:

hp.com/us/en/products/switches/index.aspx?tab=tab_A-Series

HP Intelligent Resilient Framework

HP IRF White Paper – “Reducing network complexity, boosting performance with HP IRF technology”:

hp.com/docs/irf/irf.pdf

HP Virtual Connect

HP Virtual Connect data sheets and videos: hp.com/go/virtualconnect

HP TippingPoint Security

HP TippingPoint data sheets and product details: hp.com/us/en/index.aspx?banner=security

HP Intelligent Management Center

HP IMC data sheets and product details: hp.com/us/en/products/network-management/index.aspx

Get connected

hp.com/go/getconnected

Current HP driver, support, and security alerts
delivered directly to your desktop

© Copyright 2012 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

4AA4-0789ENW, Created April 2012

