

Build a Smarter Data Center with Juniper Networks QFabric



Redguides
for Business Leaders



Bill White
Hoi T Hon
Stephen Sauer
Raymund Schuenke

- Understand the limitations of traditional data center networks
- Exploit leading edge technology to build a smarter network
- Optimize network performance, simplify implementation, and reduce costs



Executive summary

To respond to the forces of change in today's business environment, organizations need flexible and cost-effective IT infrastructures. Many are taking a new approach by consolidating and virtualizing their IT resources—servers, storage, applications, networks, and even desktops. Virtualization decouples hardware and software resources from their physical implementation and thus allows organizations to make IT resources available to applications as needed, rather than requiring a dedicated server for each application. For example, applications can call on additional server processing power or storage capacity to meet changing demands. Adding advanced, automated provisioning capabilities to this virtualized environment allows resources to be applied dynamically, which provides organizations with an even more efficient, responsive, and flexible infrastructure.

Networking plays an essential role in enabling the dynamic infrastructures that are the foundation for smarter data centers. In dynamic environments with virtualized IT resources, the network must do more than just carry traffic and support the provisioning of new IT services. It must also have the built-in flexibility and capability to adapt quickly while maintaining comprehensive security, visibility, and management. The data center network (DCN) is the key enabler for smarter data centers.

This IBM® Redguide™ publication highlights the key requirements for a smarter data center and shows how the characteristics of the data center fabric, a new switching architecture, provide the performance, scalability, flexibility, and manageability that is required.

We take a close look at Juniper Networks' QFabric, an innovative DCN fabric product, and describe how its characteristics and functions provide real business value in the key areas of rapid service deployment, cost efficient service delivery, energy efficiency (QFabric has 68% to 89% lower power consumption), and business resiliency and security. We also discuss the key network innovations in QFabric.

We then examine Juniper Networks' QFabric design, product software, hardware, and deployment options, and illustrate how QFabric can drastically improve your DCN while reducing your business costs. Recent STAC benchmark testing confirms that QFabric performs with ultra-low, consistent latency connecting thousands of devices. QFabric is the first network architecture where adding more devices does not increase latency and degrade performance.

To demonstrate the value of the Juniper Networks' QFabric design, functions, and characteristics, we examine three common QFabric network use cases. These are based on real life enterprise network requirements, namely optimized application delivery control, secure isolation provisioning of a multi-tenant environment, and support of business continuity. The use cases highlight fundamental changes in DCN architecture and demonstrate how QFabric addresses business challenges and adds value to your network.

IBM understands that the first step to transform the network infrastructure is to develop a sound enterprise network architecture, one that takes the business and IT environments, security and privacy policies, service priorities, and growth plans into account. With this in mind, this guide describes how to migrate to a smarter data center network with QFabric, and considers the organizational aspects involved in the migration.

IBM and Juniper Networks have built a strong partnership that offers leading-edge network products and technologies. This partnership can help your company plan and implement its own innovative data center network design.



Challenges with the data center

CIOs and IT managers face many challenges in improving and expanding the information services they supply to their respective businesses. Practitioners face accelerated change, increasing complexity, and the need to quickly and efficiently handle ever-growing volumes of data while furnishing the business with the information it needs to be competitive in the marketplace. Issues such as business intelligence and analytics, virtualization, mobility, risk management, and compliance arise daily and often affect business requirements while creating new challenges.

Management must handle these dynamic issues as they arise, and keep an eye on the bottom line, that is, capital (CapEx) and operating expenses (OpEx), while also staying competitive, encouraging innovation and growth, and of course, remaining profitable. Businesses are continually expanding, requiring escalating amounts of data that needs to be secured, processed, stored, and delivered at near real time rates.

Most of the top CIO priorities¹ are closely related to IT standardization and consolidation because both these approaches efficiently leverage existing IT components. To respond to these requirements, businesses are building more flexible and cost-effective IT infrastructures, beginning with consolidating and virtualizing their IT resources, such as servers, applications, storage, networks, and devices.

Many IT infrastructures were not built to support the explosive growth in computing capacity and information that we see today. Data centers are often highly distributed and somewhat fragmented. As a result, they are limited in their ability to change quickly and support the integration of new types of technologies, or to easily scale to power the business as needed. IT service delivery requires change to move beyond today's operational challenges to a data center model that is more efficient, service-oriented, responsive to business needs, and that offers improved levels of economy, rapid service delivery, data security, business resilience, and tighter alignment with business goals.

Planning a data center network to support the growing, diverse set of requirements can no longer focus on merely buying larger switches, more routers, and application devices that respond to short-term performance issues. Rather, you need to plan and architect the network with the flexibility to support future, as yet unknown, business and technology

¹ IBM conducted a study with 3,018 CIOs, spanning 71 countries and 18 industries. The results can be found at: <http://www.ibm.com/ciostudy>

requirements. In other words, a DCN that is cost-effective, saving money in both capital and operating costs; one that furnishes monitoring and management features; one that provides adaptability, high performance, scalability, and energy efficiency.

All these efforts define a move from a static legacy design towards a dynamic and scalable design in a smarter data center, as illustrated in Figure 1.

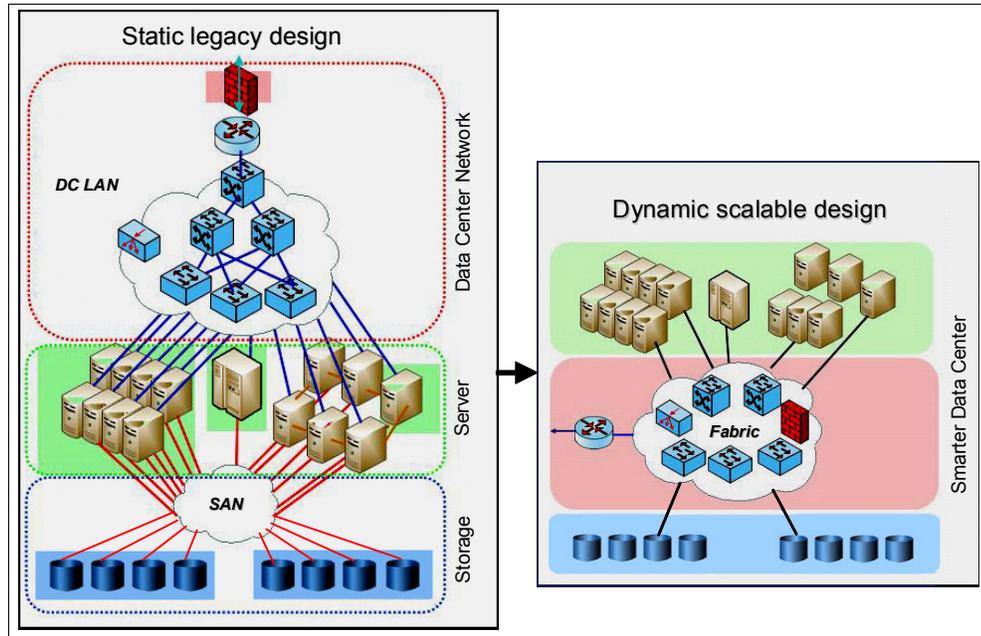


Figure 1 Evolving data center network architectures and technologies

Modern DCNs must be flexible, responsive, and capable of being managed with the rest of the IT infrastructure. If this does not accurately describe your DCN, your organization will find the network limiting your ability to leverage information resources to competitive advantage. To make matters worse, a rigid and unresponsive network could be exacting a high cost from the business.

To help you construct a next generation DCN, Juniper Networks offers a strategic DCN solution. That solution is QFabric, an innovative network technology that allows you to build flexible, scalable, and easily managed DCNs that deliver high performance.

Building a better network with Juniper QFabric

QFabric is a packet switched networking technology specifically designed to create highly efficient, cost-effective, dynamic, and easily managed data centers.

QFabric is a scalable product that lets you use a wide range of off-the-shelf devices. Those devices can connect to QFabric through standard network interfaces such as Ethernet and Fibre Channel, allowing QFabric to play an integral role in a virtualized, cloud-ready DCN.

QFabric is a platform that is immediately ready for deployment in the data center. It is best represented as a circle with N identical interfaces, each of which is available to attach to a server, storage device, appliance, or network device, as shown in Figure 2 on page 5.

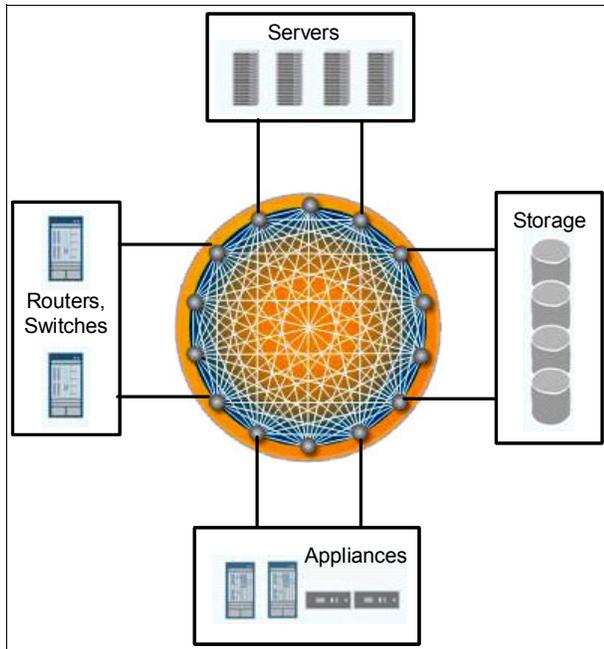


Figure 2 The QFabric platform

The circle symbolizes that all QFabric interfaces are equal in terms of latency, bandwidth, and connectivity. The result is a flattened, ultra-low latency yet highly efficient network, a trending architecture that is replacing the static legacy design in today's data centers.

QFabric improves the DCN

QFabric features and advantages address your business needs, regardless of the industry, and will help reduce costs, deploy services quickly, and add deep flexibility to your DCN. QFabric tackles today's and tomorrow's challenges in the following ways:

- ▶ QFabric is a more scalable network because the fabric acts as a single logical device, thus reducing the complexity commonly associated with large scale DCNs.
- ▶ QFabric reduces the time to deploy services by flattening the network and providing non-blocking any-to-any connectivity. This type of connectivity is critical for pooling the computing and storage resources in a data center.
- ▶ QFabric provides cost-efficient resiliency and availability because all inter-switch links are running in an active/active mode and data packets are never dropped.
- ▶ All QFabric ports are one hop away. This feature provides excellent application performance and decreases latency.
- ▶ QFabric is built using a set of modular, distributed components; the modules are kept independent and can be added or removed without affecting operation.
- ▶ DCNs are increasingly multi-tenant, with requirements for application security and reliability. QFabric allows you to partition resources yet communicate with them securely whenever necessary.
- ▶ QFabric lowers operational expenses (OpEx) because it acts as a single logical device, reducing complexity and simplifying operations. QFabric significantly lowers capital expenses (CapEx) through scaling due to the reduced number of chassis, racks, and cabling. And QFabric also increases TCO benefits by reducing power consumption, cooling, and rack and floor space.



A fabric-enabled DCN brings value to the business

If an IT organization spends most of its time mired in day-to-day operations, it is difficult to evaluate and exploit new technologies that could streamline IT operations and keep the company competitive and profitable. As noted in “The Enterprise of the Future: Implications for the CEO”¹, a paper based on a 2008 IBM CEO Study, the IT environment design must accomplish the following:

- ▶ Provide a flexible, resilient, highly scalable IT infrastructure
- ▶ Enable collaboration and turn information into business insight
- ▶ Facilitate global integration with a shared service model
- ▶ Support evolving business models and rapid integration of acquisitions and mergers
- ▶ Provide support for company-wide green initiatives

In the previous section we discussed how QFabric can address your business needs, regardless of the industry. Now we discuss the business benefits you can realize with the QFabric DCN solution, for example:

- ▶ Rapid service deployment
- ▶ Cost efficient services delivery
- ▶ Green efficiency
- ▶ Business resiliency and security
- ▶ Networking innovations

¹ Global CEO Study: The Enterprise of the Future:
<http://www.ibm.com/ibm/ideasfromibm/us/ceo/20080505/>

Rapid service deployment

The ability to deliver quality service is critical to businesses of all sizes. Service management enables visibility, control, and automation to deliver quality service at any scale. Maintaining stakeholder satisfaction through cost efficiency and rapid return on investment depends on the ability to see the business (visibility), to manage the business (control), and to leverage automation (automate), to drive efficiency and operational agility.

In today's data center network, application and service deployment performance is often dependent on the physical locations of servers that are trying to communicate, or, where each resource resides in the tree hierarchy relative to other resources. The more hops required to complete a transaction, the more that transaction is subjected to additional latency, contributing to unpredictable performance. This can also affect the delivery of differentiated service levels within shared or multi-tenant environments.

Figure 3 shows a typical tree configuration where server and storage resources are contained in localized bubbles. The resources are as close together as possible, which works well when few applications are operating. But scalability becomes an issue as applications grow in tree structures.

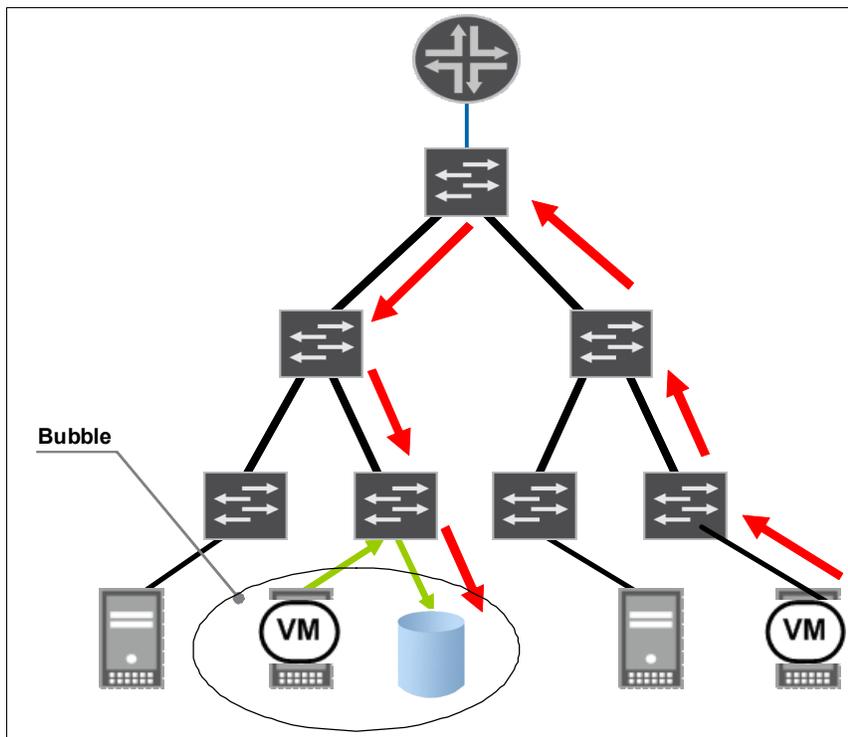


Figure 3 Typical tree configuration

So why are bubbles interesting? When provisioning or dynamically migrating a virtual machine instance, the VM and its external data sources would ideally reside within the same bubble. However, if the VM is instantiated outside this bubble, it might end up five or more network hops away from the data source, thus slowing the application. Therefore, tree structures force you to proactively manage VM locations in the physical environment in order to maintain predictable application behavior.

QFabric not only allows services to be delivered rapidly, but also reduces the time to deploy services by flattening the network and providing non-blocking connectivity. With QFabric, all devices are one hop away from each other, removing the requirement to proactively manage

the location of the VM and its external data source, thus ensuring fast, consistent communications. This flat connectivity also simplifies the writing of applications because developers no longer need to worry about the performance hierarchy of communication paths inside the data center. Additionally, operations staff are freed from concerns about the affinity of application components as they work to provide good performance.

To reduce the problems with legacy DCN designs, a new architectural approach is required, one that is transformational and innovative, not incremental. The ideal next-generation network architecture for modern DCNs would directly connect all processing and storage elements in a flat, any-to-any network fabric. Optimized for performance and simplicity, this next-generation architecture would address the latency requirements of today's applications; support virtualization, cloud computing, convergence, and other data center trends; scale elegantly; and eliminate much of the operational expense and complexity of today's hierarchical architecture.

Juniper Networks' data center solution is not a one-size-fits-all approach. Rather, it is about addressing clients' diverse data center needs with Juniper Networks' innovative data center technologies. EX Series switches with Virtual Chassis technology simplify the data center network by collapsing three tiers into two tiers. QFabric takes it to the next level, introducing a single tier, high-performance data center network, enabling cutting-edge cloud and low-latency networks to truly take off. Figure 4 illustrates this strategy.

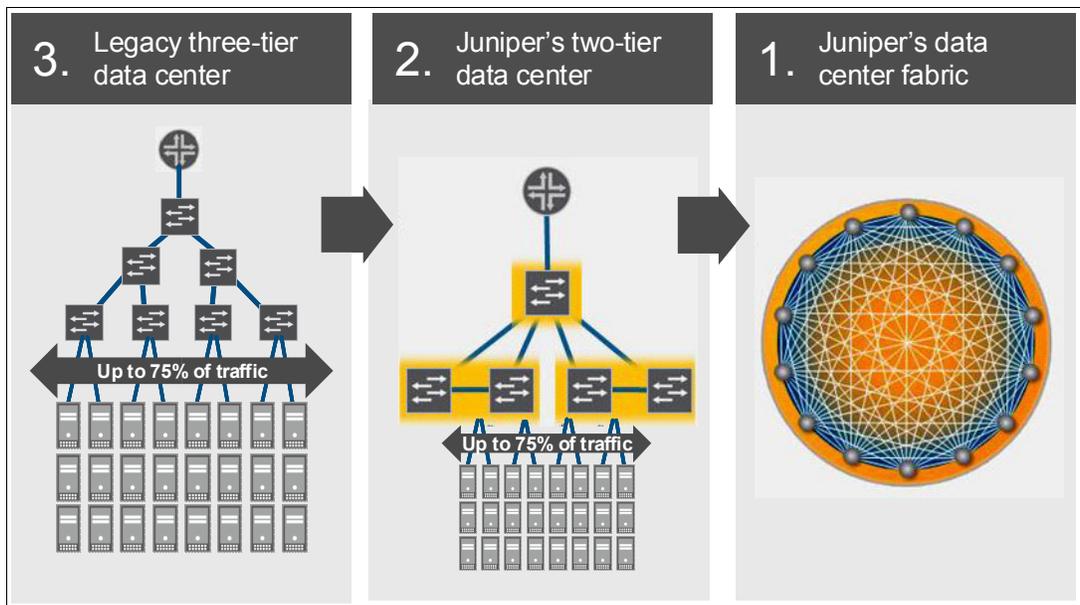


Figure 4 Juniper's transformative Data Center Network architecture

The resulting new DCN architecture allows you to migrate from the inefficient three-tier network to a flattened one-tier network, depending on non-functional requirements such as serviceability, performance, and scalability. Figure 5 on page 10 shows an example of a flattened network architecture.

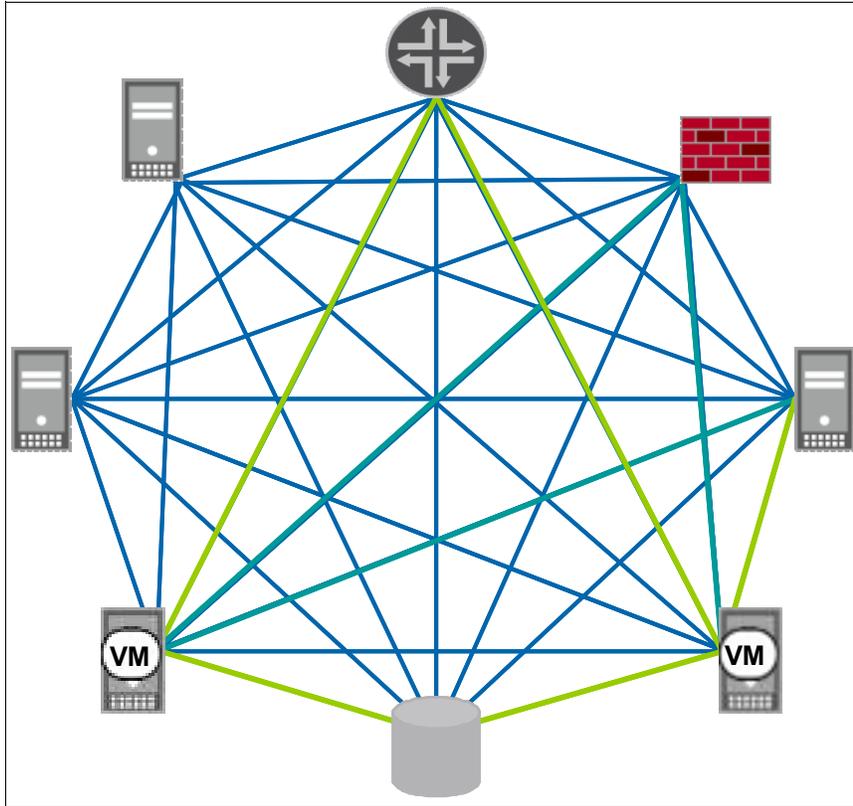


Figure 5 Single-tier flattened architecture

Non-blocking connectivity is critical to pooling the computing and storage resources in a data center. These resource pools are an integral part of a dynamic, cloud-enabled infrastructure in a highly virtualized data center environment. By supporting this “ensemble” approach of resource pools, QFabric can provide significant business benefits, such as:

- ▶ Reducing management complexity by combining stand-alone resources into virtualized pools for higher availability and more flexibility.
- ▶ Lowering IT costs by improving the efficiency of systems, energy, people, and space.
- ▶ Improving the quality of IT service by allowing new services to be provisioned easily, reducing downtime, and aligning IT with business goals by dynamically adjusting the amount of resources.
- ▶ Preparing you to get in front of new risks that can arise in a connected and collaborative world.

Cost efficient service delivery

The daily expense of managing systems and networks is increasing, and the cost of skilled labor continues to go up while the size of the available talent pool diminishes. Meanwhile, there is an explosion in the volume of data and information that must be managed, stored, and shared in a secure way. These pressing issues result in growing difficulty for IT departments to deploy new applications and services. Here are some facts to consider:

- ▶ Data volumes and network bandwidth consumed double every 18 months; the number of devices accessing data over networks doubles every 2.5 years.

- ▶ The number of physical servers is flattening out but the number of virtual servers has been growing exponentially. The cost of managing virtual networks has increased operational expenses drastically.

We now look at the results of a total cost of ownership (TCO) analysis and how QFabric supports cost reduction by lowering operational and capital expenses when compared to the network architectures of the market share leader.

QFabric supports cost reduction

QFabric lowers operational expenses (OpEx) because it acts as a single logical device, reducing complexity and simplifying operations. This device characteristic is important because it provides a network model that permits resource pools and management applications to be as simple as possible. QFabric also significantly lowers capital expenses (CapEx) through scaling due to the reduced number of chassis and racks. QFabric also increases TCO benefits by reducing power consumption, cooling, and rack and floor space.

ACG Research conducted a TCO comparison of QFabric versus the market share leader's multitiered network architecture for a mid scale to large scale 10 gigabit Ethernet (GbE) data center². This TCO analysis considered the following configuration parameters and assumptions:

- ▶ A range of 1,000 to 5,000 server ports.
- ▶ Oversubscription ratios of 3:1 and 6:1.
- ▶ Operational expenses include costs for maintenance, power, cooling, rack space, and network administration.

Compared to the market share leader's network architecture, QFabric produced the following benefits:

- ▶ Overall 58% to 76% lower TCO.
- ▶ 58% to 75% lower CapEx, with better scaling because of the lower number of devices and reduced rack space.
- ▶ 53% to 77% lower OpEx due to simplified management and better support for automating network operation tasks.
- ▶ 68% to 89% lower power consumption because of fewer chassis and less floor and rack space. See "Green efficiency" on page 14 for more details.

Next, we look at how QFabric's operational management model can support TCO benefits to your business. QFabric management is organized in layers, providing the network administration teams with the tools and capabilities to manage and automate the virtualized network infrastructure and services. This management model consists of three functional layers based on Junos and Junos space software deployments, as shown in Figure 6. The management operational layers include:

- ▶ Built-in automation for QFabric management tasks internally
- ▶ Standard application program interfaces (APIs) to integrate with leading management systems
- ▶ Support of application automation and integration with systems management solutions

² ACG Research Juniper Networks' QFabric: Scaling for the Modern Data Center:
<http://acgresearch.net/UserFiles/File/Juniper%20Documents/Juniper%20QFabric%20TCO%20Whitepaper%281%29.pdf>

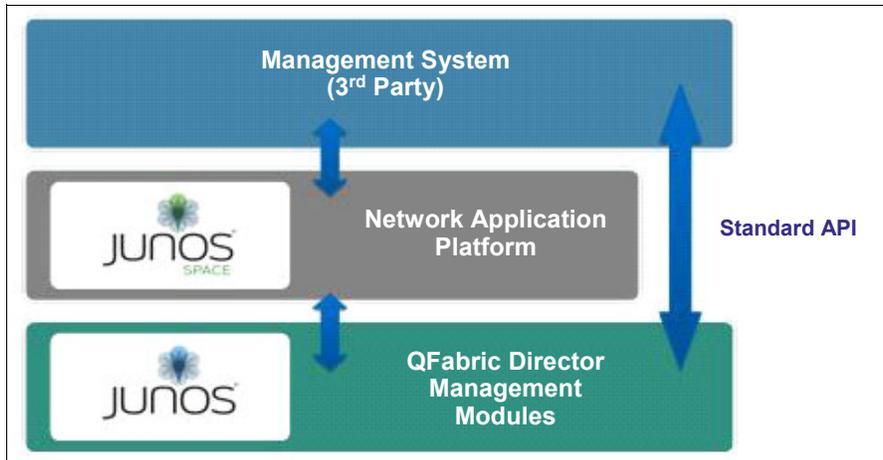


Figure 6 High level view of QFabric management architecture

QFabric provides built-in automation and standard interfaces

The standard Junos Operation System command line interface (CLI) and Simple Network Management Protocol (SNMP) module offers the same granular management capabilities and scripting parameters found in any router or switch powered by Junos OS. Performance, configuration, and fault data for the QFabric architecture can also be exported to leading third-party management systems such as HP OpenView, IBM Tivoli®, and Computer Associates Unicenter software, providing a complete, consolidated view of network operations. The QF Director contains the application logic that performs functions to simplify the operation, such as built-in automation tools, centralized image management, and customized monitoring views.

Keep in mind that QFabric is one logical device with one operating system, Junos, which means fewer managed devices and far fewer device interactions, resulting in dramatically simplified management. This management allows network managers to focus on strategic initiatives and to spend less time managing day-to-day network operations. Consequently, the QFabric management architecture and operational model provides significant business benefits, including the following cost and productivity features:

- ▶ Seamless integration into existing DC environments
- ▶ Single logical device and Junos abstraction which provides operational simplicity
- ▶ Built-in automation scripts for operational efficiency and increased productivity
- ▶ Image management and flexible upgrade options to minimize upgrade downtime
- ▶ Monitoring and troubleshooting tools for visibility and for reducing overhead

In September 2011, Juniper Networks commissioned Forrester Consulting to examine the total economic impact and potential return on investment (ROI) enterprises might realize by deploying Junos in an enterprise network environment.³ In conducting in-depth interviews with four enterprises from various industries, Forrester found that these companies, through the use of Junos and Juniper switches and routers, achieved a significant reduction in operations support costs as well as reductions in frequency and duration of unplanned events that cause network downtime.

QFabric integrates with system management

Managing virtualized environments creates new and unique possibilities. When you can dynamically resize virtual machine memory or processor power, it brings new capabilities that

³ The Total Economic Impact Of Juniper Networks' Junos Network Software
<http://forums.juniper.net/jnet/attachments/jnet/architectingthenetwork/452/1/Forrester%20Study%20Junos%20TEI.pdf>

you can exploit to deliver higher efficiencies. And when you can move virtual machines from physical host to physical host while the virtual machines remain operational, it adds another compelling capability to your business.

Businesses are embracing virtualization because it brings value and enhances capabilities for business continuity and disaster recovery. The ability to use business policy-based process automation for orchestrating, provisioning, workload management, and service level management in line with business goals will drive higher levels of virtualization adoption.

In dynamic infrastructure and cloud computing environments, automated provisioning of servers, storage, and networks is key in supporting services in a highly-virtualized infrastructure. Managed provisioning is the interlock between service management, such as that supported by Tivoli's Service Automation Manager, and the resources management supported by single platform applications, such as Systems Director or Tivoli Network Manager.

The following features highlight QFabric integration levels, which adhere to the requirements of data center automation and dynamic provisioning:

- ▶ Event monitoring allows Juniper network devices to send alerts to monitoring systems, directly to OMNibus or through Tivoli Monitoring.
- ▶ Network management allows Juniper network devices to be discovered by network infrastructure management applications like Tivoli Network Manager and Tivoli Application Discovery Dependency Manager (TADDM) to discover application interdependencies.
- ▶ Provisioning management allows the provisioning system, such as Tivoli Provisioning Manager, to automate the provisioning network resources and configurations on Juniper network devices.

QFabric supports virtualized infrastructure management by exploiting the virtualized network management capabilities of Junos Space Virtual Control along with systems management integration. This is illustrated in Figure 7.

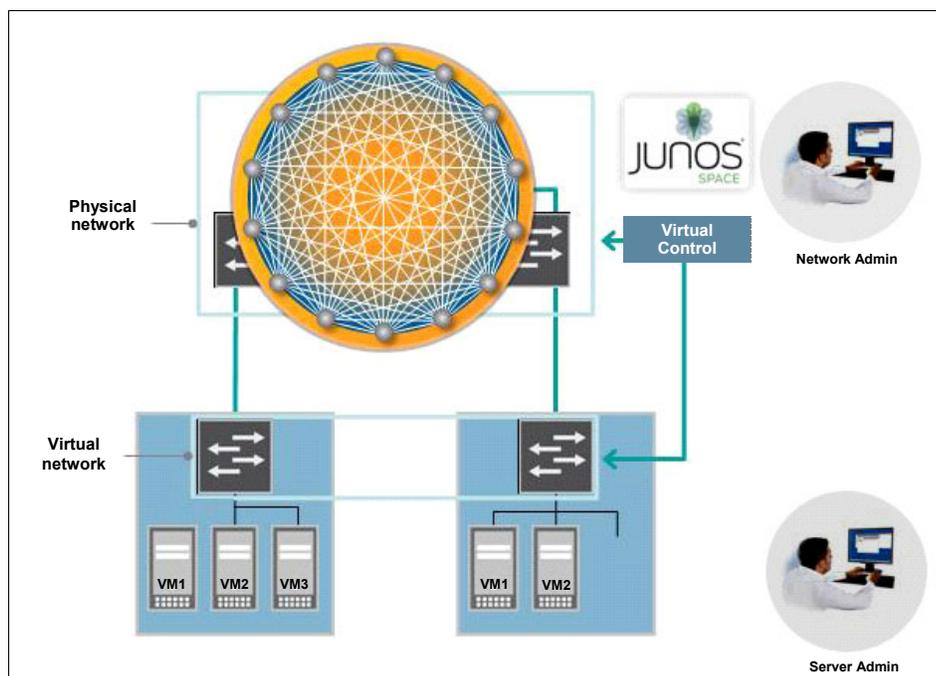


Figure 7 Virtualization with Junos Space Virtual Control

This integrated approach can optimize your virtualized management by:

- ▶ Creating a clear understanding of roles and responsibilities.
- ▶ Automating the orchestration, or resource pooling, of physical and virtual networks.
- ▶ Supporting standards of virtualized network edge management, such as Virtual Ethernet Port Aggregation (VEPA), which allows the virtualized network to communicate directly with virtual network interfaces in the server systems.

Green efficiency

The need for power and cooling increases as a company grows. At the same time, the cost of energy continues to rise; in fact, power and cooling costs grew eight-fold between 1996 and 2008. With power at a premium, and even capped in some areas, organizations must become more energy efficient. Technology groups are tasked with controlling energy costs while developing a flexible foundation from which to scale.

QFabric architecture is environmentally conscious, allowing your business to optimize every facet of the DCN while consuming less power, requiring less cooling and producing a fraction of the carbon footprint when compared to multitier data center networks.

Figure 8 illustrates how annual power consumption (KWH) of QFabric compares with that of multitier architecture. For the purpose of this example the power consumption includes power to computer equipment and for cooling. Each column in the figures shows the annual power consumption for a specific data center size. The data is based on a 3:1 oversubscription ratio, which is the most resource-intensive option analyzed in this study.

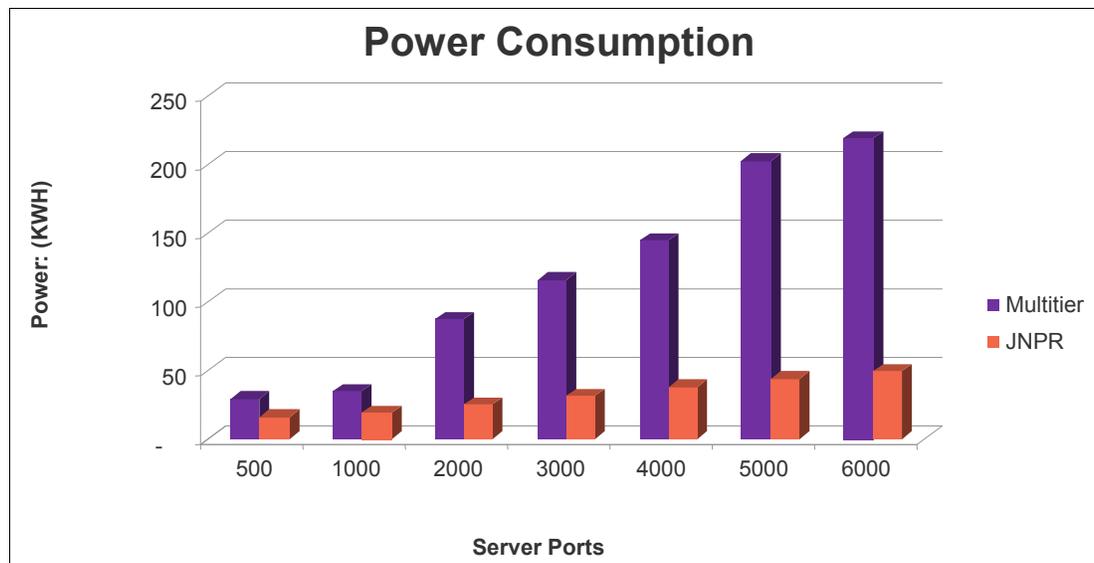


Figure 8 Power consumption multitier versus Juniper

QFabric uses power more efficiently than multitier architecture because it dramatically reduces the number of chassis and interconnections required to provide any-to-any connectivity. Compared to multitiered network architecture, QFabric has 68% to 89% lower power consumption due to fewer chassis and reduced rack space, and can potentially save up to \$350,000 per year⁴ in power costs.

⁴ 6,000 server ports, power / KWH \$0.11, HVAC and lighting / KWH \$0.13, rack space / year \$9,000

Business resiliency and security

Businesses of all sizes and various industries now realize that global expansion, emerging technologies, and the rising number and sophistication of new threats have greatly increased security needs and demand for resiliency measures. These issues are so important that enterprise risk management is now integrated into corporate ratings delivered by such organizations as Fitch, Moody's and Standard & Poor's. Users require real-time access to confidential, critical data while companies demand that both internal and external users have instantaneous access to this information, putting extra, and often conflicting, pressure on the business to provide improved availability, security, and resilience in the IT environment.

As IT enhances security and resiliency, it can be more responsive and better prepared to meet business needs, as shown in Table 1.

Table 1 IT responses to business needs

Business need	IT response
Infrastructure security and resilience	Protect against evolving threats while enabling accelerated innovation, agility, and reduced operational costs through improved security, disaster recovery, and continuity efforts.
Information and data protection	Ensure that data is accessed by only authorized users, remains available and accessible during disruption, and that it is protected both at rest and in flight.
Regulatory compliance	Plan for and respond to regulatory requirements associated with security and business resiliency, such as the Health Insurance Portability and Accountability Act (HIPAA), Sarbanes-Oxley (SOX), and Payment Card Industry Standards (PCI).

QFabric delivers a highly available and resilient infrastructure

QFabric is built on a set of modular, distributed components and the modules are kept independent. This use of physical or logical separation has two important benefits:

1. Failures in either the hardware or software do not compromise the entire system.
2. You can increase or decrease the system size while it is running so network changes can be done without disruption. In a multi-tenant environment, changes to a single tenant will not impact other tenants, as is often the case in "tree-structured" networks.

QFabric achieves cost-efficient resiliency and availability because all inter-switch links are running in an active/active mode and also because QFabric does not drop packets under congestion. Not dropping packets during congestion is critical to efficiently transporting bursty server-to-disk traffic. This is an important key in supporting multi-tenancy and multiple applications within a tenant.

QFabric builds resiliency into many levels: hardware, data plane, control plane, and management plane, as shown in Figure 9 on page 16.

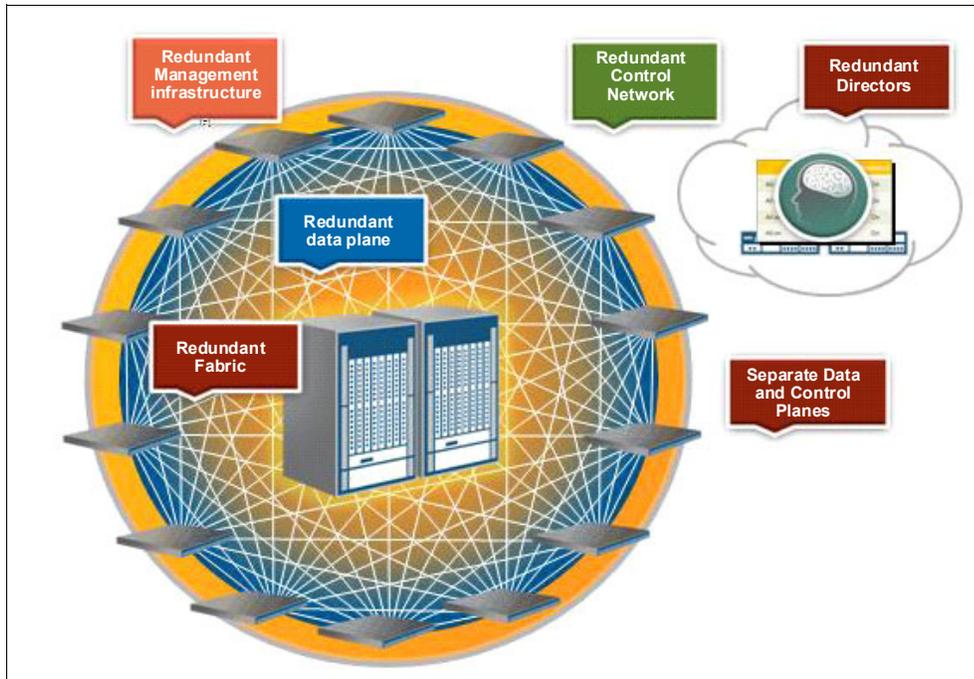


Figure 9 QFabric resiliency

Service availability is supported by the following aspects of QFabric design:

- ▶ Hardware resource resilience, such as redundant fabric, interconnect, and directors
- ▶ Data and control plane resilience, such as redundant fabric links, load balanced routes, and redundant control ports
- ▶ Highly available management infrastructure, such as redundant management connections and control networks

We now continue to the server and network nodes located at the intelligent edge and show how QFabric supports resilience there with higher availability for business and management applications.

QFabric scalability supports services access on demand

QFabric can scale to 6144 ports and offers predictable, consistent latency of less than 5 microseconds. QFabric was designed for virtualization; it behaves and is managed as a single switch with Junos delivering the benefits of a single operating system. QFabric's non-blocking and lossless nature supports storage connections for Fibre Channel over Ethernet (FCoE) as a transit switch. This scale-out design benefits your business by translating into a linear build-as-you-grow model that matches with sourcing and service models of cloud computing. These key benefits also address a cost-effective approach because:

- ▶ QFabric operates as a single logical device, so scaling the network is as simple as adding ports to an existing switch.
- ▶ You can seamlessly integrate new services into the fabric.
- ▶ IT can easily scale the DCN with minimal additional management and operations overhead.

Aligning network security with IT security

The increasing use of cloud computing creates new security challenges for organizations. In the near future, cloud computing will shift much data and many controls to cloud providers. It might be difficult to maintain control and accountability when the operational responsibility for data centers falls to one or more third parties.

As illustrated in Figure 10, Juniper offers the following integrated security solutions for virtualized data center and cloud environments⁵:

- ▶ Protecting physical workloads by using external security services gateways.

Juniper's SRX Series Services Gateways are high-performance security, routing, and network solutions for the enterprise and service provider. The SRX Series platform provides high port density, advanced security, and flexible connectivity in a single, easily managed platform that supports fast, secure, and highly available data center operations. SRX gateway is a zone-based firewall that allows you to create security zones to segment and isolate traffic among physical workloads.

- ▶ Protecting virtualized workloads by deploying virtual security gateways into the server world.

Juniper's vGW Virtual Gateway Series technology is based on a three-tiered architecture consisting of a hypervisor-based module, a security virtual machine, a management server, and a web interface. The hypervisor-based module resides in the hypervisor of each virtual machine host and performs security functions, including packet inspection and security policy enforcement. The security VM communicates with the vGW management server, where security policy information and VM details are stored, and the hypervisor module. The vGW management server stays in constant communication with the VMware vCenter so that as VM changes occur, they are synchronized to the vGW management server.

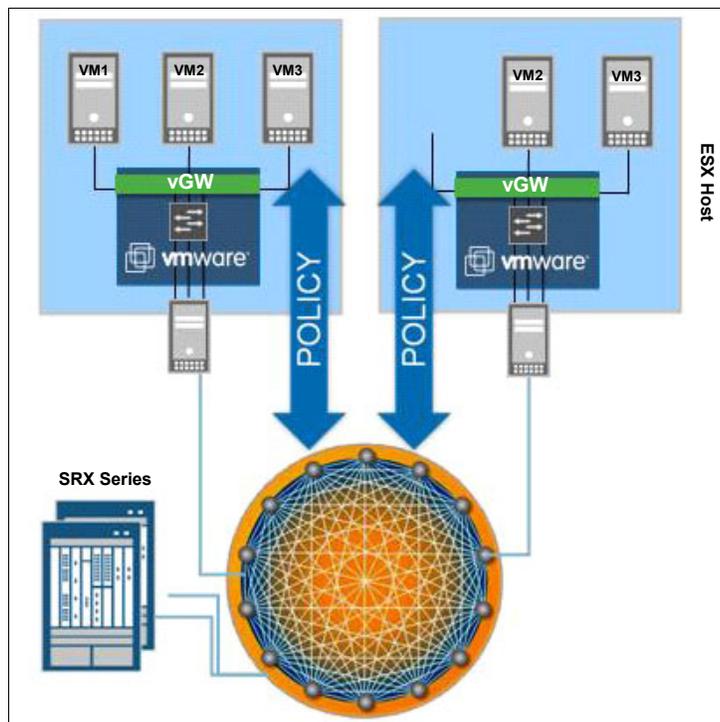


Figure 10 Integrated security zone enforcement

⁵ Juniper White Paper - An Integrated Security Solution for the Virtual Data Center and Cloud
<http://www.juniper.net/us/en/local/pdf/whitepapers/2000431-en.pdf>

By using this integrated zone enforcement approach, security is applied in a collaborative way on physical and virtualized workloads. The design ensures that security policies applied to workloads are consistent with their logical use, regardless of the platform on which they are deployed. The SRX Series zone concept integrates with the vGW VM enforcement engine such that the zone information is synchronized to and used by vGW within the virtualized environment.

This integrated security approach gives you the following benefits:

- ▶ It meets strategic requirements of network security alignment with IT security:
 - Security policies applied to workloads are consistent from the data center perimeter to the server VM with their logical use, regardless of the platform on which they are deployed.
 - Zone synchronization provides an automated way to link the virtual server security layer with physical device and network security; this also supports the alignment of network security with IT server security.
- ▶ It optimizes operational efforts and expenses by
 - Optimizing the operational efforts of performing of security changes.
 - Allowing you to centralize the policy management, lowering operational expenses.
 - Automating and verifying that VM connectivity do not violate security zone policies.
 - Automating migration of security policies and port profiles when a VM migrates to a new server port.

Harnessing networking innovations

The increasing speed and availability of network bandwidth creates new opportunities to deliver services across the web and to integrate distributed IT resources. Easier access to trusted information and real-time data and analytics will soon become basic expectations. QFabric introduces technology innovations that address these expectations by offering features such as:

- ▶ Separate data and control
- ▶ Moving toward a single, flat data plane
- ▶ Moving toward a distributed control plane
- ▶ Moving toward a single management plane

Separating data and control planes in a DCN switching environment is a current top trend. Figure 11 on page 19 shows this data transformation and control plane separation.

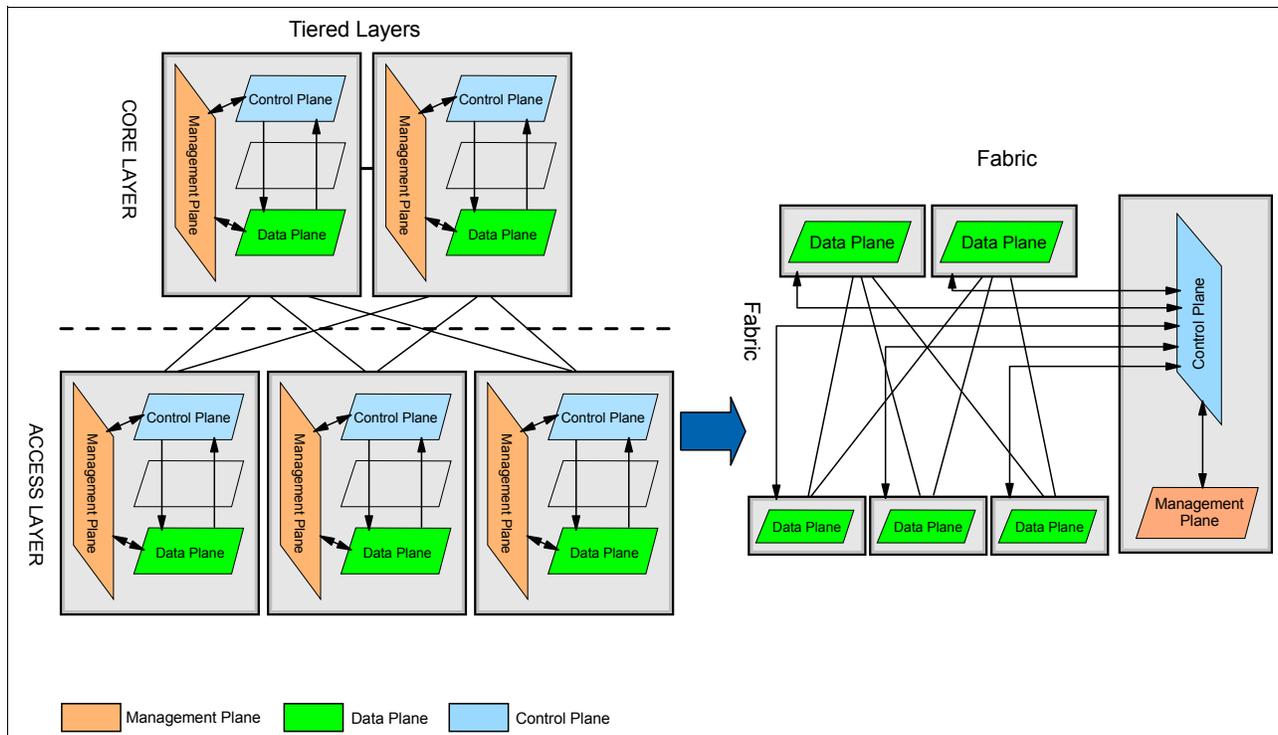


Figure 11 Data transformation and control plane separation

This paradigm implies that two parallel networks are needed: one to forward data and the other to control traffic via an out-of-band signalling infrastructure. QFabric created an advanced industry approach geared toward this paradigm, which also provides additional business benefits:

- ▶ Optimized operational costs
 - The single control plane node has a complete picture of the network, not a partial one.
 - A centralized point of control removes some of the boundaries related to node interdependence and can simplify automated provisioning in the DCN.
 - Leveraging a centralized control plane eases security policy enforcement.
- ▶ Enhanced support of business continuity
 - Suboptimal traffic paths are avoided because the control plane node makes traffic path selection decisions based on link utilization, availability, and other parameters.
 - A centralized control plane eliminates the need to use the spanning tree protocol to exchange topology information among network nodes.
 - Using this approach eases network state mobility orchestration due to the centralized collection point that enforces network configuration stickiness after a VM movement inside the data center.
- ▶ Secure support of multi-tenant services
 - Virtual networks support is important because data centers are increasingly multi-tenant and have requirements for application security, performance, and reliability.
 - Virtual networks provide the basic tools that allow resources to be partitioned and yet allow them to communicate securely whenever necessary. These abstractions are essential to decoupling the applications from the infrastructure.
 - Full mobility of virtual machines from any interface to any other interface is supported. Support of virtual networks does not compromise any of the other properties.

Ultimately, all these innovations have a tremendous effect on improving communication and service, reducing barriers for market entry, and enhancing how organizations conduct business.

STAC benchmark results

The Securities Technology Analysis Center (STAC) is a vendor-neutral specialist in high-performance capital markets technology. STAC facilitates the STAC Benchmark Council, which creates a set of STAC Benchmarks to help vendors and customers understand the performance of workloads in the financial industry using different products in combination.

Juniper Networks, IBM, and Mellanox Communications have completed an audited STAC-M2 Benchmark⁶, which tests the ability of a given “stack” of switches, middleware, NICs, and servers to handle real-time market data under a variety of conditions. The tests provide key performance metrics such as latency, throughput, power efficiency, and CPU/memory consumption under several scenarios, including undisturbed flow and exception conditions like slow consumers.

This was the first STAC-M2 benchmark to utilize a data center network fabric as the interconnect, versus a single switch. The benchmark results confirm that QFabric can create a large-scale, high-performance, ultra-low-latency data center fabric with the simplicity of a single switch.

The STAC-M2 benchmark revealed an end-to-end mean latency of 10 microseconds, which includes the application layers. The QFabric latency was 5 microseconds. Figure 12 show the STAC-M2 benchmark results from a small scale QFX3500 network compared to a large scale QFabric network.

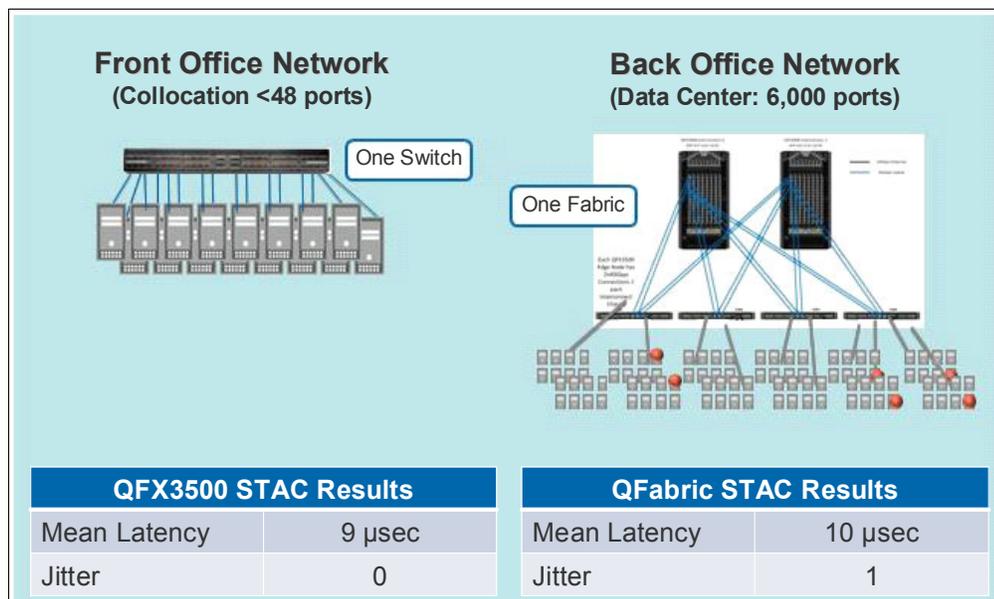


Figure 12 STAC test results with QFX3500 and QFabric

⁶ For an audited STAC-M2 benchmark report, go to:
<http://www.juniper.net/us/en/local/pdf/industry-reports/2000459-en.pdf>
 For an Executive Summary of the STAC-M2 test report, see:
<http://www.juniper.net/us/en/local/pdf/industry-reports/3000082-en.pdf>



Why QFabric is innovative

We have discussed how a fabric network is an integral part of a smarter data center. In this chapter we examine Juniper Networks' QFabric design, product set (software and hardware), and deployment models.

QFabric is built using modular hardware and software components that are highly reliable and scalable; together they provide the ability to interconnect and manage all interfaces in the network like a single logical switch. This new design flattens the network to a single logical layer and does away with the conventional "tree-structured" network.

Tree-structured networks are typically complex and costly to maintain and operate. In addition, tree-structured networks rely on indirect data paths and redundant uplinks that consume more than half of the data center's available device ports. This can in turn produce traffic patterns that result in significant performance degradation. With QFabric, any servers, storage devices, or network devices are just one hop away from any others, regardless of their location in the fabric; thereby creating an any-to-any high-performance network.

The chassis switch: the foundation of QFabric

The QFabric architecture adopted the three basic components of a self-contained switch chassis: line cards, backplane, and routing engines, and transformed them into independent, stand-alone components consisting of the QFabric Nodes, the QFabric Interconnects, and the QFabric Directors, respectively. The QFabric architecture is also composed of three distinct planes: the data plane, the control plane, and the management plane. QFabric behaves like any classical switch with distributed components, but it is managed as a single logical switch. Figure 13 illustrates this concept.

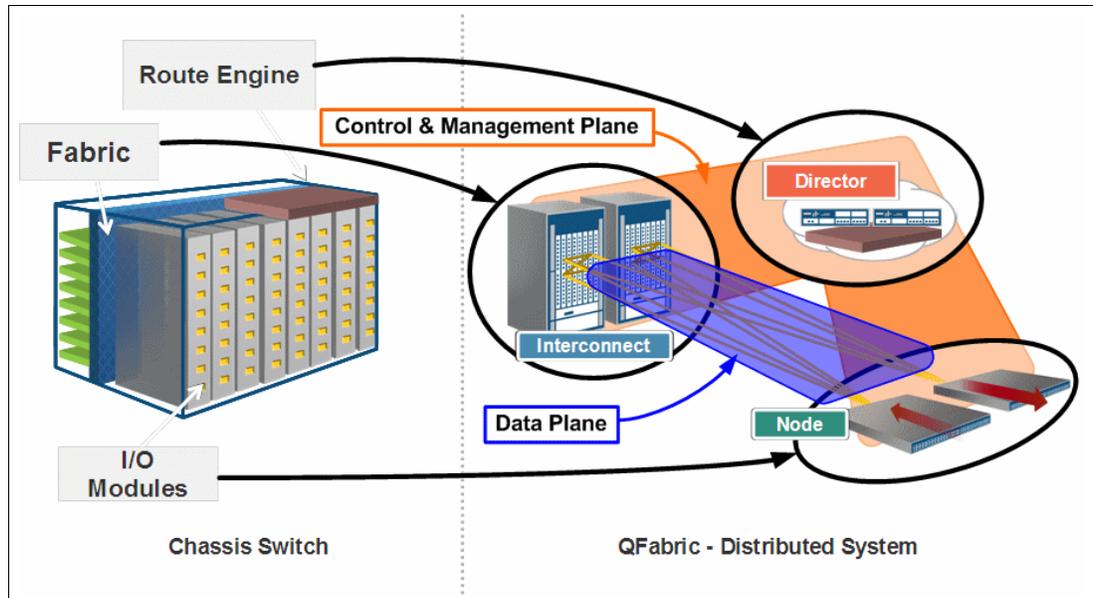


Figure 13 QFabric design concept

Although QFabric is based on the operational model of the single switch, its scaling model has been modified to suit several thousand network ports. This modification flattens the tree structure, which allows you to add capacity without requiring a new branch to the tree as in today's static legacy design. Hence, the architecture evolves from a traditional chassis to a distributed chassis, allowing smooth scaling to meet developing business needs.

QFabric architecture

When designing QFabric, Juniper took a physical fabric design approach as opposed to a protocol fabric design approach. The physical fabric design involved re-architecting the physical layers of the DCN and applying key principles, such as:

- ▶ Creating a lossless, loop-free, converged Ethernet network with a single switch view
- ▶ Reducing the number of network tiers to optimize data traffic
- ▶ Separating the data and control plane inside the data center network
- ▶ Moving toward a flat data plane that provides any-to-any non-blocking connectivity
- ▶ Enabling full utilization of the connectivity by eliminating old approaches
- ▶ Moving toward a new control plane that enables multi-pathing
- ▶ Simplifying management by decreasing the number of discrete devices

- ▶ Centralizing management and provisioning to support automation
- ▶ Moving toward a single management plane to manage the network as though it were one big switch

In short, QFabric collapses the traditional multitiered data center model into a single-tier model where all edge devices are directly interconnected via a large scale fabric backplane. With this design approach, QFabric functions like a single logical switch, and as such, switch-to-switch interactions and network protocols like Spanning Tree Protocol (STP), Shortest Path Bridging (SPB), and Transparent Interconnect of Lots of Links (TRILL) are not needed. This, in turn, eliminates the significant overhead typically found in multitiered data center models.

QFabric components

QFabric is a switching architecture designed to dramatically transform the DCN while improving performance and simplifying operations. At the core of this new technology is a product set capable of supporting up to 6,144 ports in a single logical switch. QFabric is composed of three major components:

- ▶ QFabric Node is equivalent to line cards offering ports and route engine functions.
- ▶ QFabric Interconnect is related to the backplane, especially to the data plane, allowing quicker forwarding of packets inside the QFabric.
- ▶ QFabric Director is the CPU of QFabric with route engine functions.

QFabric Node

A QFabric Node is an ultra-low latency, high port density, fixed configuration 1 U top-of-rack device that provides access in and out of the network. The QFabric Node can be compared to the line card in a legacy switch chassis. QFabric Nodes are located at the edge of the QFabric and are connected to the QFabric Interconnects. Servers, appliances, storage devices, or external networks can connect to a QF Node using standard protocols. In addition to its role at the edge of QFabric architecture, QFabric Node also acts as a high performance stand-alone converged access switch.

The QF Node is a QFX3500¹ and has the following features:

- ▶ 48 x 10 GbE ports in total:
 - 12 ports (10 GbE or Fibre Channel (FC) 2/4/8 Gbps)
 - 36 ports (10 GbE or 1 GbE)
- ▶ 4 x 40 GbE fabric uplink ports to the QF Interconnects
- ▶ All ports have Layer 2 and Layer 3 support
- ▶ All ports have FCoE and Data Center Bridging (DCB) capabilities
- ▶ Redundant AC power supply
 - Power (nominal): 295 watts
- ▶ Front-to-back air flow
 - Dissipation (nominal): 1.01kBTU

QFabric Interconnect

QFabric Interconnect acts as the primary fabric for data plane traffic traversing the switch between QF Nodes. It is an ultra-low latency, 21 U, eight slot chassis with sixteen 40 GbE QFabric Node facing ports per slot. QFabric Nodes exchange traffic via the QFabric

¹ See this link for QFX3500 product details:

http://www.juniper.net/techpubs/en_US/release-independent/junos/topics/concept/qfx3500-hardware-overview.html

Interconnect by forming a full mesh topology, using standard high-speed optics. To compare it with a legacy switch chassis configuration, QFabric Interconnect functions like the backplane.

QFabric Interconnect directs traffic between QFabric Nodes using single tag lookups instead of the full Ethernet lookups used in legacy switches. This is one of the features that makes QFabric dramatically different from legacy switch components.

The backplane of a single switch becomes the QFabric Interconnect, which connects all QFabric Nodes in a full mesh topology. The QFabric Interconnect is a QFX3008²; it contains the following features:

- ▶ 128 QSFP (40 GbE) ports
- ▶ Eight fabric cards (10.24 Tbps/chassis)
- ▶ Dual redundant control board
- ▶ Redundant AC power supply
 - Power (nominal): 3,000 watts
- ▶ Front to back air flow
 - Dissipation (nominal): 10.24kBTU

QFabric Director

The QFabric Director is a QFX3100³ that provides control and management services to the QFabric, acting as a central point to manage and control all network components as a single logical switch. The Director can be compared with the legacy switch supervisor component and route engine. The QFabric Director communicates directly with all the QF Nodes and QF Interconnects to build a global view of the entire topology. QFabric Director is a 2RU device based on x86 architecture that has GbE ports to connect to the QF Nodes and QF Interconnects.

Command Line Interface

QF Director software contains the application logic for implementing the command line interface (CLI) for the distributed system. A single Junos CLI command is scattered across all the individual federated systems via the device interfaces it has to all the Junos instances in the system. After results are successfully executed locally, they are returned to gather application logic. The final response is returned to the user. The QF Director software ties together all the aspects of presenting a single Junos CLI to the QFabric administrator.

Software management

QFabric Director contains functions that further manage the network, that is, the Fabric Manager, the Fabric Administrator, and the Fabric Viewer. These components manage and control all aspects of QFabric processes with the following features:

- ▶ Fabric Manager performs provisioning using an active/backup model.
- ▶ Fabric Administrator provides the single logical switch abstraction, CLI, user session, using an active/active model.
- ▶ Fabric View is the management GUI application.
- ▶ Fabric Control synchronizes all fabric communication and tables using an active/active model.
- ▶ Network Node Group RE is the routing engine for network node groups, running all Layer 2 and Layer 3 protocols. It uses an active/backup model.

² See this link for QFX3008 product details:

http://www.juniper.net/techpubs/en_US/release-independent/junos/topics/concept/qfx3008-hardware-overview.html

³ For more information about QFX3100, see:

http://www.juniper.net/techpubs/en_US/release-independent/junos/topics/concept/qfxctr1-hardware-overview.html

Route engine

The route engine is a Juniper-specific processing entity that implements QFabric switch control plane functions, routing protocols, system management, and user access.

Route engines reside in the QF Nodes and the QF Directors and work together to provide Layer 2 learning of local hardware, the propagating of information and status, and the routing of data through the optimum path.

Because QFabric only runs Ethernet protocols at the QF Nodes, a simple transport protocol is used to move data upstream to the QF Interconnects and downstream to the QF Nodes. Routing within the fabric is based on a MAC reachability table. The MAC reachability table has an entry for each QF Node and enumerates all the possible paths for reaching every QF Node in the fabric through the QF Interconnects. Note that the QF Interconnect only looks into the fabric header and does not do a MAC address lookup.

QFabric connectivity

Each QFabric Node is connected to all QFabric Interconnects in the fabric, establishing redundancy for high availability and multiple paths for ultra-low latency. There are no direct connections between the QFabric Interconnects. Because of the way the QF Nodes are interconnected in the fabric, additional routing protocols outside of the route engine algorithms are not needed.

Multi-chassis link aggregation groups (MC-LAGs) are used to allow the QF Interconnects to present all QF Nodes as a single logical switch. Figure 14 shows the connectivity that builds the fabric between the QF Nodes and QF Interconnects⁴, as well as the QF Directors.

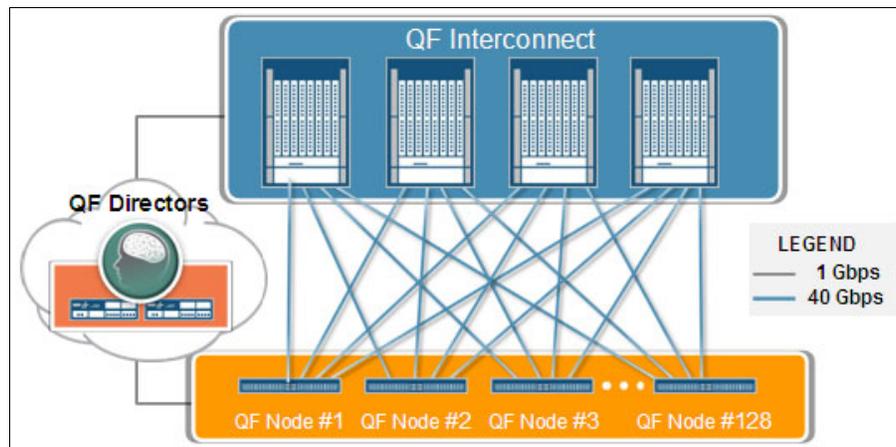


Figure 14 QFabric connectivity

QFabric planes

Ensuring high reliability is a priority for Juniper, so separating the data plane from the control plane has always been an important design principle. This is shown in Figure 15 on page 26, where data and control traffic are carried over two separate networks. The distributed nature of the control plane is pivotal to QFabric scalability and simplicity and also eliminates a single point of failure.

QFabric architecture consists of three planes: the data plane, the control plane, and the management plane.

⁴ The number of QFabric Interconnects can be two or four depending on the number of QFabric Nodes and oversubscription required.

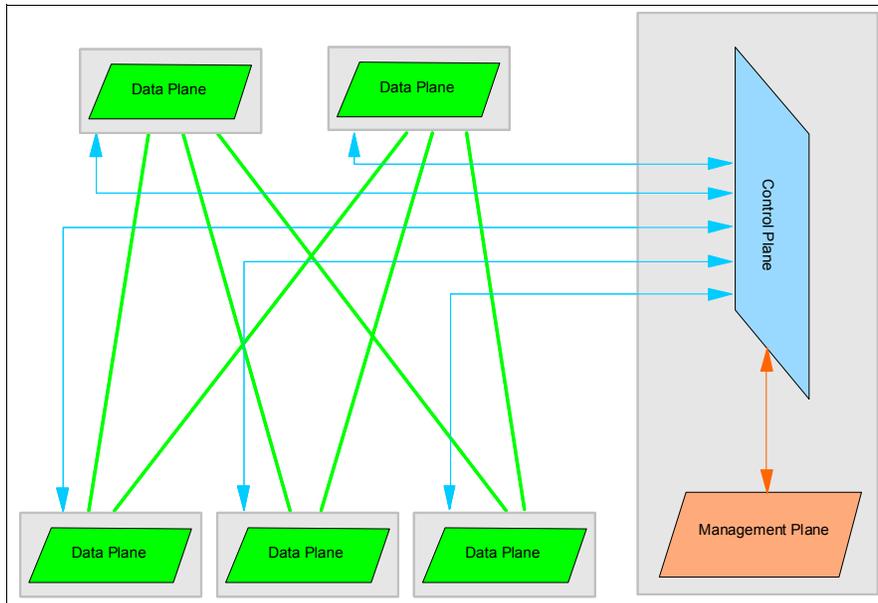


Figure 15 QFabric planes

Data plane

The data plane⁵ in the QFabric switch rapidly and efficiently transfers application traffic between QFabric switch components. The data plane uses QSFP+ interfaces and fiber-optic cabling to connect QFabric switch components at speeds of 40 Gbps. By creating a redundant set of connections between the QFabric Nodes and the backplane-like QFabric Interconnects, the data plane enables the QFabric Nodes to appear as though they are directly connected to one another in a single tier network.

Control plane

The QFabric control plane⁶ Ethernet network is built on EX4200 switches acting as a Virtual Chassis. This out-of-band network provides a separate control plane within the QFabric switch to handle management traffic. This design enables the data plane network to focus on efficient, low-latency delivery of data, voice, and video traffic. The control plane generally uses two sets of four EX4200 switches configured as a pair of Virtual Chassis to connect all components within the QFabric switch. The dual Virtual Chassis architecture provides redundancy and high availability to ensure reliable QFabric switch operation for the QFabric Director group, the QFabric Interconnects, and the QFabric Nodes.

Management plane

The management plane function is under the control of the QFabric Director, which provides all management services for the QFabric architecture. The QFabric Director communicates directly with all QFabric Nodes and QFabric Interconnects and builds a global view of the entire network, which might consist of thousands of server-facing ports. This provides a single point of visibility, control, and management for the entire data center fabric. The QFabric Director also interfaces with the network management ecosystem via standards-based protocols such as XML/NETCONF, SNMP, or CLI.

⁵ For a data plane description, see:

http://www.juniper.net/techpubs/en_US/junos11.3/topics/concept/data-plane-qfabric-understanding.html

⁶ For a control plane description, see:

http://www.juniper.net/techpubs/en_US/junos11.3/topics/concept/control-plane-qfabric-understanding.htm

1

Junoscript automation and Juniper Networks Junos SDK provide a rich automation framework that enables network customization and tuning as required, ensuring the QFabric architecture deployment fits the existing ecosystem without special tools.

Logical components

QFabric Nodes can run in different modes depending on the function needed, for example, Routing/Spanning tree activated, single node, or dual nodes in a virtual edge. The QFabric Node should be configured in one of the following node group types⁷.

Server nodes group (SNG)

The server nodes connect servers whenever cross-node resiliency is not needed. The SNG has the following characteristics:

- ▶ A single QFabric Node comprises a single logical edge entity in the QFabric distributed system. The server node group connects server and storage endpoints to the QFabric system. Members of a link aggregated group (LAG) on a server are connected to the server node group to provide a redundant connection between the server and the QFabric system.
- ▶ Server node groups do not run network protocols like STP, PIM, and OSPF. However, they contain mechanisms like BPDU guard and Storm control that are ON by default to detect and disable loops across their ports. You can connect network devices, for example, firewalls and load-balancers, to server node groups using LAG connections.
- ▶ The local CPU in the QFabric Node in SNG performs Routing Engine⁸ (RE) and Packet Forwarding Engine (PFE⁹) functions. The Forwarding functions are local to the SNG, which is the default mode.

Redundant server node group (RSNG)

RSNG spans a maximum of two QFabric Nodes and contains the same protocol restrictions as SNG. The RSNG uses the same approach as the Virtual Chassis Model. A pair of QFabric Nodes in a RSNG has the same architectural model as that of Juniper's Virtual Chassis configuration. However, a key difference is that there is no explicit cable dedicated for this virtual chassis connection and that all connectivity is only through QFabric Interconnect.

Multi-chassis LAG (MC-LAG) members on a server are distributed across the RSNG to provide a redundant connection between the server and the QFabric system. RSNG configuration is recommended for use cases where redundancy is not built into the software application running on the server. Local CPU of each QFabric Node in RSNG performs PFE functions. RE functions are active on one QFabric Node and backup on the other, just as in Virtual Chassis.

Network node group (NNG)

NNG is a modular switch model where QFabric Node CPUs perform PFE function and QFabric Directors perform RE function with an Active/Backup mechanism. The NNG connects to external network devices and runs network facing L2 and L3 protocols, for example, xSTP, OSPF, BGP, and so forth. One NNG is used per QFabric network. Inside a NNG, all the QFabric Nodes, maximum of 8, are seen as one device.

⁷ For more information about the node group types, see:

http://www.juniper.net/techpubs/en_US/junos11.3/topics/concept/node-groups-qfabric-understanding.html

⁸ For more information about RE, see:

http://www.juniper.net/techpubs/en_US/junos11.3/topics/concept/routing-engine-qfabric-understanding.html

⁹ For more information about PFE, see:

<http://juniper.cluepon.net/index.php/PFE>

QFabric deployment models

Pivotal to understanding the QFabric solution is the description of various deployment options, or types, namely, small, medium, and large. The number of QFabric Node device ports determines the number of QFabric Interconnect devices that you will need. QFabric Nodes can also transport FCoE and FC traffic, so FCoE is yet another type of deployment.

You can use various QFabric deployment types based on your changing business demands, another facet of QFabric flexibility. Starting with a small deployment model, you can increase deployment types as your business grows by adding QFabric Nodes as necessary. From a management point of view, the process is similar to adding a linecard in a chassis.

Small deployment

A small deployment type consists of two QFabric Interconnects, each with two 16-port (40 GbE ports) line cards, two QFabric Directors, and up to 16 QFabric Nodes providing up to 768 ports with an oversubscription of 3:1.

Medium deployment

The medium deployment type can be used when a certain level of expansion is required, adding QFabric Nodes and line cards inside the QFabric Interconnects as the DCN grows.

The starting configuration is based on two QFabric Interconnects, each with four 16-port (40 GbE ports) line cards, two QF Directors, and up to 32 QF Nodes that can provide up to 1536 ports with an oversubscription of 3:1.

The medium deployment type can be expanded to four QFabric Interconnects, each with four 16-port (40 GbE ports) line cards, two QF Directors, and up to 64 QFabric Nodes that can provide up to 3072 ports with an oversubscription of 3:1.

Large deployment

A large deployment type consists of four QFabric Interconnects, each with eight 16-port (40 GbE ports) line cards, two QFabric Directors, and up to 128 QFabric Nodes that can provide up to 6144 ports with an oversubscription of 3:1.

Deploy as your business grows

You can use various QFabric deployment types, from small to large, based on your changing business demands. You can move from the minimum to maximum with an oversubscription of 1:1 (keeping the total bandwidth under or equal to 160 Gbps for each QF Node) to an oversubscription of 6:1.

This flexibility is a key characteristic of QFabric. Starting with a small deployment model as described in the previous section, you can expand QFabric as your business grows by adding QFabric Nodes when necessary. From a management perspective, it is similar to adding line cards to a chassis. QFabric is a complete solution, not merely a product. View it as your core network infrastructure, flexible enough to seamlessly grow from 768 ports to 6144 ports.

FCoE deployment

QFabric Nodes can also transport FCoE and FC traffic.

FCoE transit switch

The QFX3500 offers a full featured DCB implementation that provides strong monitoring capabilities on the top-of-rack switch. In addition, FC Initiation Protocol (FIP) snooping provides perimeter protection, ensuring that the presence of an Ethernet layer does not impact existing SAN security policies. To transport FCoE traffic¹⁰ with QFabric, use QFabric as a Transit Switch L2 Forwarding with FCOE-FC Gateway connected to QFabric.

FCoE-FC gateway

In FCoE-FC gateway mode, the QFX3500 eliminates the need for FCoE enablement in the SAN backbone. You can add a converged access layer and interoperate with existing SANs without disrupting the network. The QFX3500 allows up to 12 ports to be converted to Fibre Channel without requiring additional switch hardware modules.

DC interconnection considerations

Because QFabric is designed to be the heart of the data center, deploying QFabric is an intra data center solution. Data center interconnections can be enabled through components like Juniper's MX and current technologies, such as MPLS L3VPN, MPLS L2VPN, VPLS, or DWDM over dark fiber.

DC interconnection issues

DCs are becoming more important network elements as data is consolidated to a smaller number of distributed DCs for multiple organizations. Networks stretch between DCs, using some of the following:

- ▶ L2 stretch
- ▶ MPLS/VPN
- ▶ Internet Edge
- ▶ WAN aggregation
- ▶ WAN Core

Consolidating DCNs improves business agility, but it can also increase traffic on the WAN. As DCs continue providing class of services (CoS) on the interconnects to various types of clients, it becomes increasingly important to do so while providing consistent Quality of Experience (QoE) to users¹¹ accessing the applications.

Study and consider interconnection needs carefully, especially when Layer 2 stretching between DCs is required. If you use the Layer 2 extension to move VMs from one site to the other, pay attention to available bandwidth, the latency between the two sites, and the number of VMs that could be moved.

If you use a Layer 2 extension for clustering, a dangerous situation could occur if the WAN links go down. In this case, a possible split brain in the cluster could result in serious database synchronization issues. Layer 2 extension should only be used between DCs when no other

¹⁰ FCoE Deployment: Opportunities and Challenges with the Convergence of Data Center Networks:
<http://www.juniper.net/us/en/local/pdf/whitepapers/2000315-en.pdf>

¹¹ Learn more about the many possible levels of quality of experience for users:
http://en.wikipedia.org/wiki/Quality_of_experience

mechanisms are available. Make sure that you study the configuration carefully, while considering the associated costs, operations, and technical constraints.

DC interconnection solutions

QFabric views the interconnection between DCs as an external connection through its connections to routers. The DC Interconnection¹² services are performed by routers such as MX.

You can also build a Layer 2 extension with VPLS¹³. As we have noted, consider the issues involved with clustering or VM Move. Whenever possible, use dark fiber and DWDM (owned or leased) to do the Layer 2 extension. DWDM allows you to securely transport SAN (FC or FCoE) traffic. Currently, the best practice is to use MPLS L3VPN for the DC interconnection¹⁴.

You can also use QFabric to build an overlay network to offer services to servers, applications, and databases.

¹² For a description of MPLS DC Interconnection for Disaster Recovery, see:

<http://www.juniper.net/us/en/local/pdf/whitepapers/2000407-en.pdf>

¹³ A description of VPLS implementation for DCs Interconnection is available here:

<http://www.juniper.net/us/en/local/pdf/implementation-guides/8010050-en.pdf>

¹⁴ For a description of MPLS VPN for Secure DCs interconnection, see:

<http://www.juniper.net/us/en/local/pdf/implementation-guides/8010079-en.pdf>



Use cases: Practical QFabric scenarios

In this chapter, we examine three common network uses based on our actual project experiences with clients:

- ▶ Optimized application delivery control
- ▶ Secure isolation provisioning of multi-tenant environments
- ▶ Support for business continuity and recovery

The three use cases highlight the need for fundamental changes in the data center network architecture and demonstrate how Juniper Networks' QFabric can address those challenges while adding value to your network.

Use case 1: Optimized application delivery control

Application Delivery Controller (ADC) capabilities vary from vendor to vendor. Most appliance ADCs and software-based ADC products provide wide area network/application acceleration, load balance between servers, Global Server Load Balancing (GSLB), SSL offload, and TCP header compressions. In this use case, we use the term ADC to cover both appliance controller and software-based products.

Figure 16 illustrates a typical multi-tier application architecture deployment, such as for SAP, ERP, or industry-specific transaction-based applications. Based on a typical security zone concept in a smarter data center, this architecture consists of several services:

- ▶ Business-to-business (B2B) front end services in a controlled sub-zone is defined to provide connectivity for specific services for limited trusted partners and Internet untrusted connections.
- ▶ Restricted front end server services that provides the front end web layer for internal web applications.
- ▶ Restricted mid server services based upon IBM WebSphere® Application Servers that provide a second tier, offering application services to web servers located in the controlled or restricted front end server zones.
- ▶ Back end server services that offer a second and third tier, for example, mainframe-based database services to the restricted mid server zone.

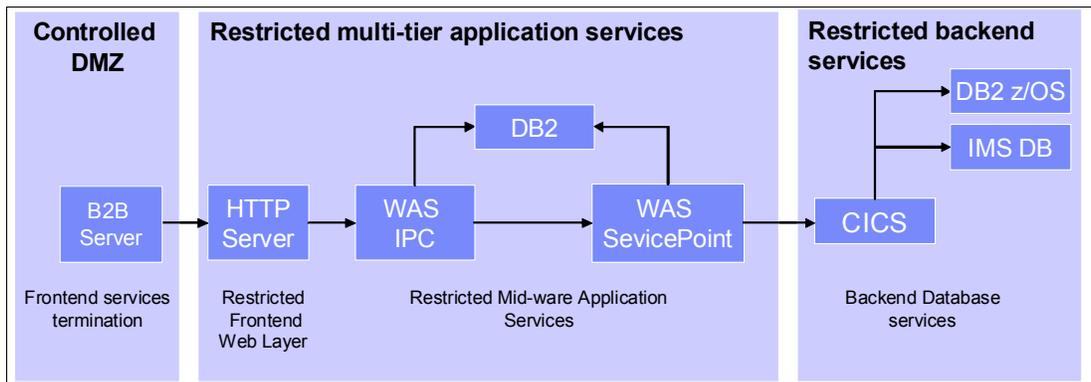


Figure 16 High level multi-tier application architecture

Application delivery and security, performance, and infrastructure security are top concerns for businesses today. DCN complexity has increased drastically in logical and physical domains in order to support a variety of application deployments. Application delivery control increases the importance of delivering application efficiency, performance, and web security to a wide range of applications and speed-driven web users. This is especially true as cloud providers produce more offerings and as more enterprises move their applications to a cloud environment.

In response to this demand, enterprises and cloud providers are aggressively shifting their simple load balancing deployment strategy to more application friendly and feature-rich ADC platforms. These platforms support ERP systems, SAP, Web 2.0 on IBM WebSphere, business applications on VM vCenter, and vSphere on virtualized and hybrid cloud environments. In the past, traditional data center infrastructure was deployed in hierarchical core-distribution-access and multitiered topologies. Application, server, and storage traffic was routed and switched through unnecessary intermediate paths to arrive at destinations. But recently, data center agility has been challenged by smarter SOA and Web 2.0

application traffic. In this scenario we show how Juniper Networks' QFabric offers an easy and innovative alternative to address these issues.

Preconditions

The environment must meet the following preconditions to deploy a successful optimized ADC in a QFabric network:

- ▶ Requirements for application deployment
- ▶ Hardware and software requirements for the QFabric environment
- ▶ Secure internet access
- ▶ Secure DMZ¹⁵ provisioning that adheres to company security rules and policies
- ▶ Set up the Server Load Balancer and ADC to provide or accelerate applications either on the internal side or facing the public internet
- ▶ Define perimeter security policy on edge security devices, for example, firewalls and intrusion detection and prevention systems (IDPS)
- ▶ Storage requirements

Scenario

Juniper QFabric brings a whole new approach to optimizing traffic. When application traffic enters the data center, it passes through a series of firewalls or security checkpoints for packet inspections, and through switches and routers to depacketize, packetize, and redirect. A hardened and secure DMZ tier is imperative to protect the internal network from the insecure internet, where threats such as virus, SYN flood, distributed denial-of-service (DDoS), spoofing, and man-in-the-middle attacks lurk. However, within a secure DC environment, we can reduce all the redundant South-North and North-South traffic by connecting the ADC devices, web and application servers, and FCoE data storages via 10 GbE switch fabric.

For example, many enterprises and cloud providers currently maintain traditional multitiered and top-down hierarchical architecture. These tiered and hierarchical gigabit architectures are complex to operate, require top technical skill sets, and are costly to maintain. Over time, the design principles for the traditional approach deviate from its original intent and benefits are steadily diminished. For example, servers that need high speed access to storage and other resources have to hop through multiple switches within a DC to reach the resources and SAN storage (see Figure 17). Also, 3-4 tiers defense-in-depth DC networking architecture needs to be re-evaluated because advanced and more sophisticated security devices offer better protection. In some large tier-3 or tier-4 DC designs and deployments, physical cabling can easily contribute 5%¹⁶ to the cost of the total budget. Therefore, reducing infrastructure cabling cost and minimizing DCN complexity are top priorities for IT management.

¹⁵ Demilitarized zone (DMZ); also known as a perimeter network or screened subnetwork.

¹⁶ Gartner Inc: Best Practices for Data Center Costs and Design (Doc: G00213184)

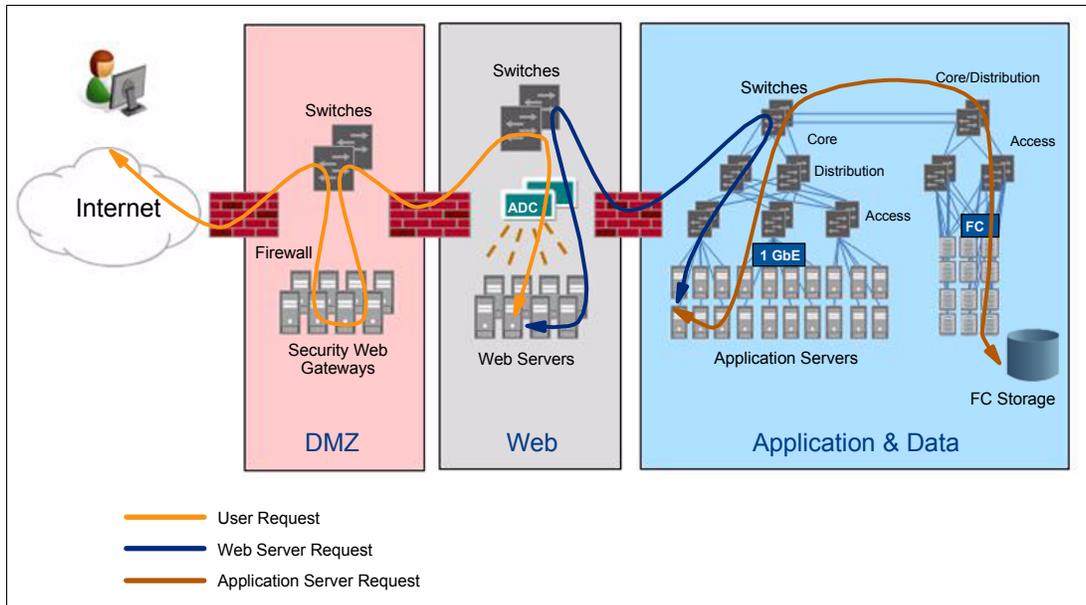


Figure 17 Traditional large scale data center multi-tier architecture

QFabric design collapses core-distribution-access into a simple switch fabric and innovative high speed any-to-any architecture to devices supporting 10 GbE and FCoE. QFabric operates as a single logical switch to reduce the number of tiers, thus reducing complexity, deployment effort, CapEx and OpEx costs. You can easily scale and modify the QFabric network to support several thousand ports when DC devices increase. These modifications flatten the tree structure, which means that adding capacity does not require adding a new branch to the tree as required in the traditional hierarchical model. Currently each Juniper QFX Node supports the following:

- ▶ 48 10 GbE ports in total
 - 12 ports (10 GbE or FC 2/4/8 Gbps)
 - 36 ports (10 GbE or 1 GbE)
- ▶ 4 x 40 GbE fabric uplink ports (each 40 GbE port can be broken out to 4 x 10 GbE)
- ▶ All ports have FCoE and DCB capabilities

A single QFabric infrastructure can provide up to 6,144 ports. All ports have Layer 2 and Layer 3 support. ADC, web servers, application servers, and SAN storage are always one hop away from each other. This simplifies ADC deployment regardless of where devices are connected in the QFabric environment. QFabric improves application performance and reduces the time to deploy services by flattening the hierarchical architecture and creating a non-blocking fabric network for any-to-any connectivity.

Figure 18 on page 35 depicts the QFabric design with the new network layout.

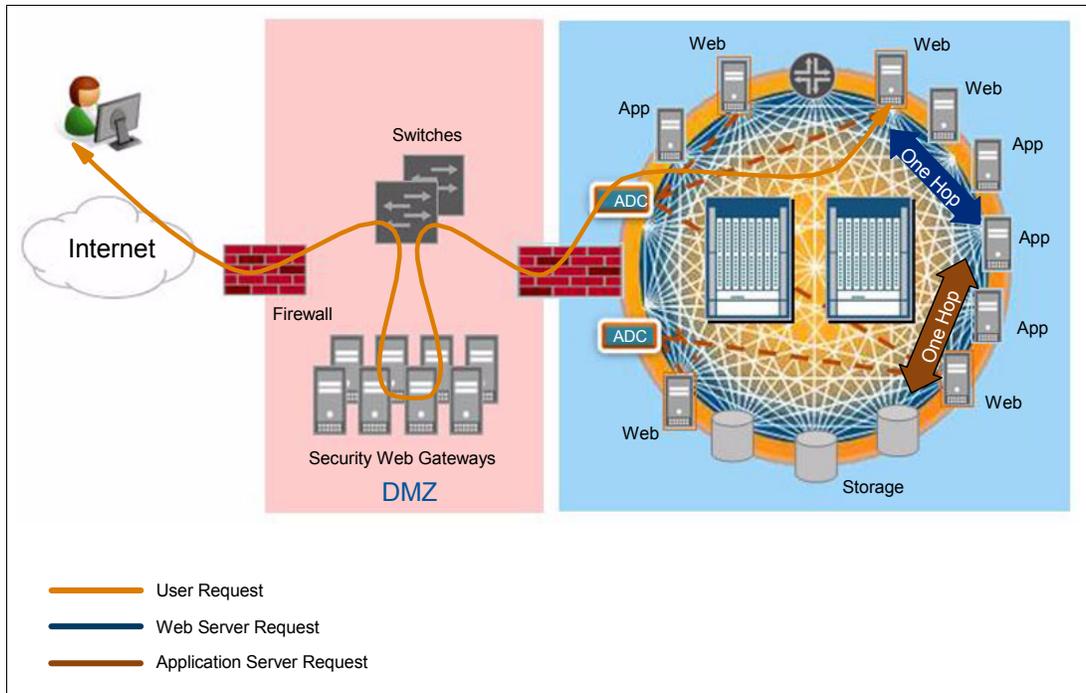


Figure 18 Simplified data center multitiered architecture using QFabric

You can create a network configuration where the DMZ is physically separated with specific hardware, or, depending on your security policies, you can integrate DMZ components into QFabric. Current best practices recommend that the DMZ be kept separate; however, future innovations and enhancements in technology could allow for other configuration options.

Results

This use case illustrates the functional requirements of optimized application deployment and provisioning. In this scenario, QFabric supports the business requirements by applying some key characteristics at packet transport system level. The goal is to make sure that application delivery to the end users is optimized. A successful outcome is achieved when the application is reached with the desired level of performance.

This use case applies whenever you need to deploy a new application in the DC. The application might be available in a repository and can be packaged with an operating system for rapid VM provisioning.

Note the following network considerations:

- ▶ Is the OS image available from a storage device? If so, is there enough or dedicated bandwidth to achieve rapid service deployment?
- ▶ If the operating system and the application have to be deployed in separate steps, additional delay might negatively affect application delivery time. What is the business criticality of the application that is being deployed? If it is business-critical, it is essential that network resources are aware of this to provide differentiated levels of service in the most effective way from an enterprise-wide standpoint.

We now demonstrate how a typical QFabric deployment model successfully supports this use case. Performance requirements and the reduced complexity of the network architecture in this scenario address those aspects of the supporting DCN infrastructure.

Reduced complexity

QFabric reduces the number of tiers from the traditional multitiered and top-down hierarchical architecture. QFabric directly connects all processing and storage elements in a flat, any-to-any network fabric. Optimized for performance and simplicity, this next-generation architecture would address the latency requirements of today's applications; support virtualization, cloud computing, convergence, and other data center trends; scale elegantly; and eliminate much of the operational expense and complexity of today's hierarchical architecture.

Highest performance

The following features of QFabric architecture provide low latency to boost application performance:

- ▶ Web servers, application servers and SAN storage are always one hop away from each other. This simplifies ADC deployment regardless of where devices are connected in the QFabric environment. QFabric has interface-to-interface latency on the order of 2 microseconds at the small scale, growing slowly to about 5 microseconds at the largest scale. Latency also grows slowly with offered traffic load. QFabric also provides very low jitter.
- ▶ Full non-blocking connectivity. This is critical to pooling the computing and storage resources in a DC to support rapid access to business applications. These resource pools are an integrated part of a dynamic, cloud-enabled infrastructure in a highly virtualized data center environment.

Use case 2: Secure isolation provisioning of multi-tenant environments

Dynamic provisioning in a multi-tenant environment requires network, server, and storage resources to be virtualized as needed and then returned to the main pool when no longer needed. This concept of rapid and dynamic provision of hardware, software, and services can be a challenge when dealing with multiple tenants in one physical network. For example, a provider network might have a mixture of financial services, oil, gas, and biotech clients, all residing on the same physical network but needing to be securely separated into multiple logical networks.

Preconditions

These are the preconditions to automate and dynamically provision a QFabric network:

- ▶ Tenant perimeter security requirements
- ▶ VLAN requirements
- ▶ Virtual machines requirements
- ▶ Number of VM instances per tenant in the deployment
- ▶ Orchestration and automation requirements
- ▶ Storage space requirements
- ▶ High availability and resiliency NFR requirements

Each client business and operational models can vary and have different requirements, as shown in Table 2 on page 37.

Table 2 Sample multi-tenant technical requirements

Tenant requirements	Tenant A	Tenant B	Tenant C
Security Zone	<ul style="list-style-type: none"> ▶ Internet DMZ ▶ Internal network #1 ▶ Admin network ▶ VM network ▶ Storage network 	<ul style="list-style-type: none"> ▶ Internet DMZ ▶ Internal network #1 ▶ Internal network #2 ▶ Admin network ▶ VM network #1 ▶ VM network #2 ▶ Storage network 	<ul style="list-style-type: none"> ▶ Internet DMZ ▶ Vendor DMZ ▶ Internal network #1 ▶ Internal network #2 ▶ Admin network ▶ VM network #1 ▶ VM network #2 ▶ Storage network #1 ▶ Storage network #2
VLANs	100-199	200-399	500-999
VM instances	100	250	3000
Storage capacity	1 TB	25 TB	2000 TB

Scenario

In this use case, the goal is to build multiple virtual infrastructures over the physical infrastructure and to offer a secured and isolated environment for each tenant. Each infrastructure component provides functions to achieve the goal. The components consist of the following:

- ▶ Network components
 - QFabric with Nodes, Interconnect, Director
 - MX routers
 - Embedded switches in the Blade Chassis
 - Virtual switches to interconnect the VMs
- ▶ Security components
 - SRX Firewalls
 - SA for secure access
 - Other security components can be added, for example, IDS/IPS
- ▶ Servers
 - Blade Chassis
 - Virtual Machines and Hypervisor
- ▶ Storage
 - NAS equipment in this use case providing NFS access

The physical components are connected to QFabric through two or more 1 GbE or 10 GbE ports. To add redundancy, the blade chassis, the MX, and the SRX are connected through a LAG to a RSNG in the QFabric. See “Logical components” on page 27 for RSNG details.

QFabric is responsible for the interconnection of all the components to build the virtual infrastructure at layer 2 of the OSI model. This isolation is created using VLANs. MX provides the isolation of the routing through the Virtual Router, but you can also do this using the NNG inside QFabric. For security reasons, this function has been put outside of the QFabric.

SRX provides the isolation between the zones inside a tenant with a context defined for each tenant. This context is provisioned with a set of rules to permit or deny traffic from one zone to another. All these rules, context, and zones are depicted under the term policy in the diagram.

To better explain a secured and an isolated tenant, we consider a flow from a user to the application inside the tenant infrastructure as shown in Figure 19. The user connects to the environment through the internet. At this level, the user is not authenticated or authorized to access the environment. The flow comes from the internet to the MX routers connected to it. The MX routers' main function is to route traffic from the users to the access zone defined in a context in the SRX firewall. This context contains another zone, a global one for all users, where the MX routers are connected.

The access zone is a DMZ containing the Secure Access (SA) devices.¹⁷ The role of these devices is to secure each user in their own environment through an SSL VPN in its outside interfaces; this gives access to its tenant through a specific VLAN in the inside interfaces. Here the traffic is isolated between the users depending on their authentication. They have access to their tenant and only their tenant.

Each VLAN from the inside interface of the SA is connected to an external zone in a specific tenant context in the firewall. This tenant context has multiple zones. So we have an external zone and several internal zones. These internal zones are linked to specific VLANs created inside QFabric. We link the virtual switches and the VMs to this specific tenant on these internal tenant VLANs. From the storage view, a virtual environment is built through vFILTER and linked to a virtual routing instance for each tenant.

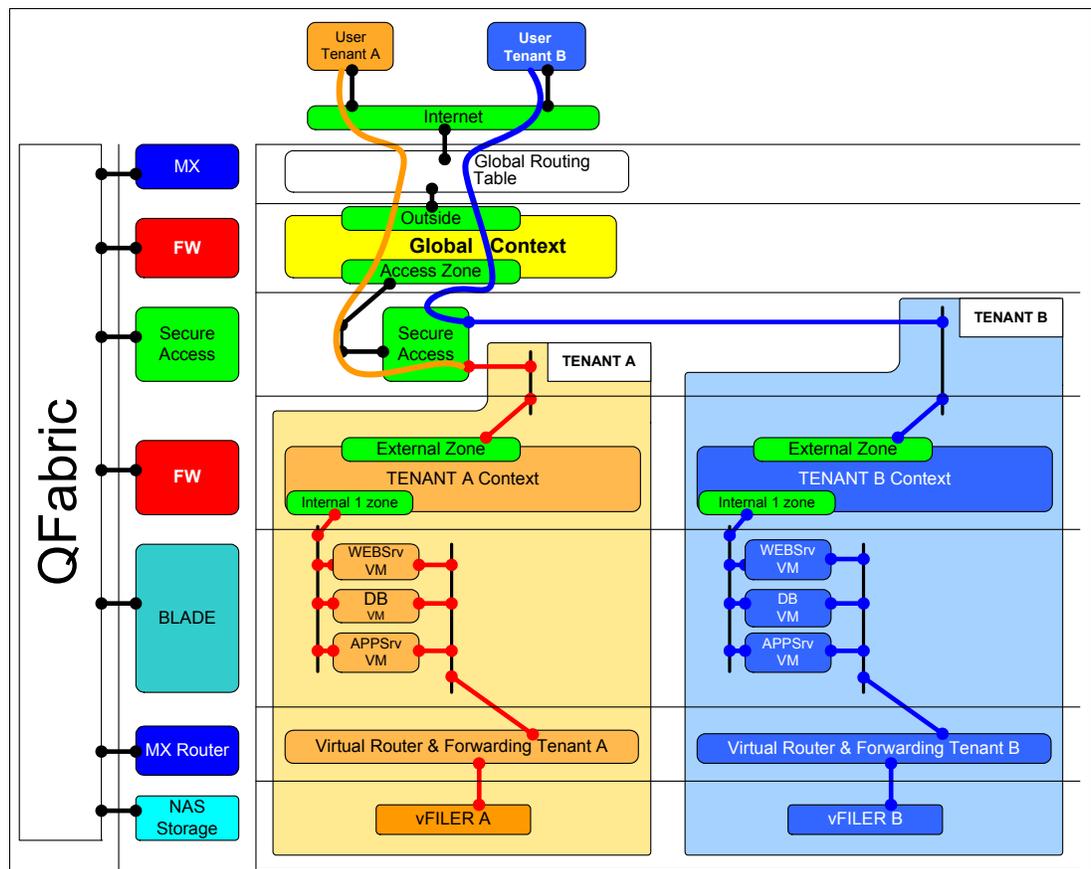


Figure 19 Multi-tenant user flow

Dynamic multi-tenant provisioning can leverage the IBM Cloud Service¹⁸ Provider Platform (CSP²)¹⁹ built on Tivoli. The IBM Cloud Service Provider Platform accelerates and simplifies

¹⁷ Secure Access to the Virtual Data Center
<http://www.juniper.net/us/en/local/pdf/solutionbriefs/3510352-en.pdf>

deployment of a complete Public Cloud services environment that rapidly and cost-effectively enables Cloud Service Providers to:

- ▶ Create partner-enabled cloud services by harnessing the power of the ecosystem to differentiate their brand and drive profitable growth
- ▶ Manage and deliver cloud services in a highly secure and automated way with an operation that scales while maintaining a low cost structure
- ▶ Monetize networks, systems, and other resources further by offering a diverse and compelling portfolio of attractively priced cloud services

The manage pillar is the heart of the CSP² offering as shown in Figure 20. The core management component of the solution is based on IBM Service Delivery Manager. ISDM enables creation, delivery, and management of cloud services. Besides the core service automation and management capabilities provided by IBM Service Delivery Manager, the manage pillar of CSP² includes enhanced options. The enhanced options are extensions for security management, network management, storage management, and advanced monitoring and service level management.

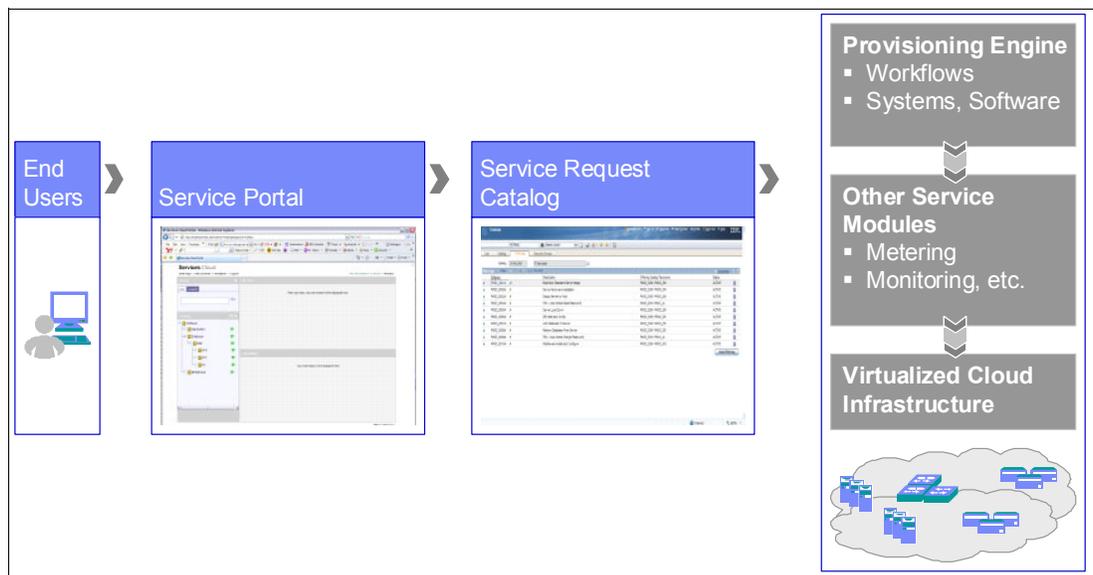


Figure 20 Key components of the managed CSP² platform

IBM has a complete set of software for managing virtualized compute, storage, and network resources in a secure multi-tenant, highly scalable, carrier-grade environment to securely provision hardware, network, and application services at a needed level. Among these, CSP² uses Tivoli TSAM, TPM, and TNCM²⁰, which offers easy provision to Juniper EX series, MX series, SRX series, and QFabric on secure VLAN, VRF, security policy, NAS, SAN storage and external API to vSphere for VMWare configurations. In short, TSAM handles workflow creation; TPM provisions Juniper devices and Juniper APIs; and TNCM maintains and pushes configurations to Juniper devices.

CSP² provides reusable and systematic automatic provisioning. CSP² provides flexibility for administrators to perform moves, adds, and changes without affecting other tenants residing

¹⁸ Gartner Inc: Vendor Focus for IBM Global Services: Consulting Services for Cloud Computing: http://www-935.ibm.com/services/us/gbs/bus/pdf/vendor_focus_for_ibm_global_174046.pdf

¹⁹ IBM Whitepaper: Becoming a cloud service provider - <ftp://public.dhe.ibm.com/common/ssi/ecm/en/tlw03009usen/TLW03009USEN.PDF>

²⁰ TSAM - Tivoli Service Automation, TPM - Tivoli Provisioning Manager, TNCM - Tivoli Netcool® Configuration Manager

on the same physical hardware. The CSP² network provision and automation procedure has the intelligence to reflect the tenant network requirements to create their environment. Dynamic multi-tenant provisioning is shown in Figure 21.

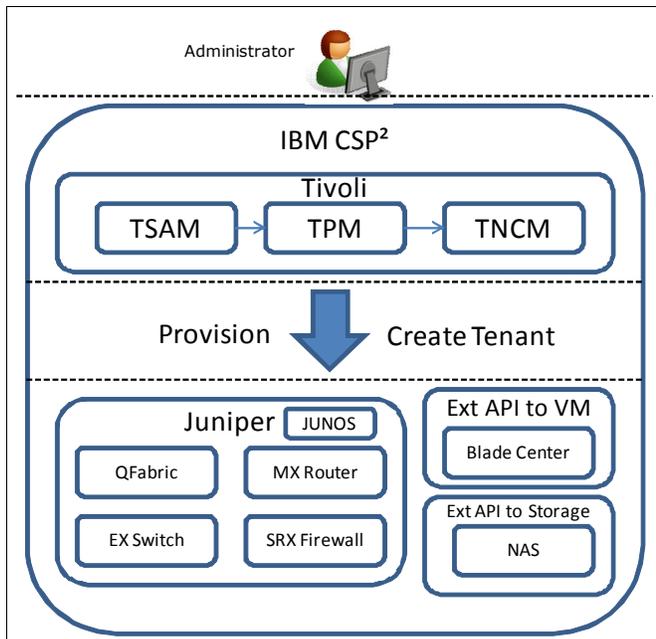


Figure 21 Secure provisioning of multi-tenant environments

Results

This use case applies whenever network resources must be dynamically provisioned (IMAC – install, move, add, change) to provide a certain service in a secure multi-tenant dynamic, shared infrastructure.

By applying this use case we try to answer the following questions:

- ▶ Is the QFabric network just supposed to deliver automation inside the data center, or can it also be involved in the automation process?
- ▶ To what extent can the QFabric be dynamically provisioned today?

One of the conclusions is that the QFabric data center network has to be seen as the key enabler of automation in the data center because it provides access to remote and distributed resources, making it possible to speed up repetitive tasks without having to gain physical access to distinct systems.

Now we demonstrate how an IBM CSP² /Juniper QFabric deployment model can support this use case successfully based on key requirements for cost efficient service delivery:

- ▶ CSP² delivers reusable and systematic automatic provisioning.
- ▶ CSP² provides flexibility for the administrator to perform IMAC services without affecting other tenants residing on the same physical hardware.
- ▶ CSP² network provision and automation procedure has the intelligence to reflect the tenant network requirements to create their environment.
- ▶ Providing safety functions such as fallback and isolated provisioning for multiple tenants

- ▶ Tenants might have control on their server and storage resources and also have secure access to manage and monitor their secure, scalable, and highly resilient network.
- ▶ Juniper network products (QFabric, EX, MX, SRX series) have an advantage from an implementation effort standpoint because their devices run the same OS (JUNOS) across different families of products, for switching, routing and security.
- ▶ QFabric supports a high level of integration following the scenario used in
 - Event Monitoring: The integration allows Juniper network devices to send alerts to OMNIbus directly or through Tivoli Monitoring.
 - Network Management: The integration allows Juniper network devices to be discovered by ITNM and also by TADDM to discover application interdependencies.
 - Provisioning Management: The integration allows TPM to automate the provision of network resources and configurations on Juniper QFabric and other devices.

This scenario results in considerations from various points of view that need to take into account when services will be end-to-end provisioned efficiently:

- ▶ Provisioning considerations such as topology and configuration dependencies
- ▶ Automation tools such as standard, OS-level scripting and systems management integration
- ▶ Process and cross-organization considerations
- ▶ Infrastructure management integration into automation and provisioning

Use case 3: Support of business continuity and recovery

According to the Federal Emergency Management Agency (FEMA), 25% of Nashville businesses did not reopen after the 2010 flood. Imagine how your business could be affected in a natural disaster like an earthquake, flood, or major storm. When disaster hits, it is critical that you have a resilient, redundant, and flexible network, along with a solid Business Continuity and Recovery Plan (BCRP), to help your business survive and navigate during a difficult time.

Business resilience and recovery have become top priorities for both executives and stakeholders²¹. Areas of concern that organizations need to focus on include enterprise and work area risk, availability of critical data and business applications, IT stability, recoverability of IT infrastructure, data backup, and disaster recovery. They also need to consider the constantly changing governmental and industry regulations associated with business continuity that apply to them, and have a plan in place to manage compliance.

Your business must be resilient and able to recover from all types of disruptive events, as shown in Figure 22 on page 42.

²¹ IBM Whitepaper “Business continuity and resiliency services from IBM”:
<http://public.dhe.ibm.com/common/ssi/ecm/en/buw03007usen/BUW03007USEN.PDF>

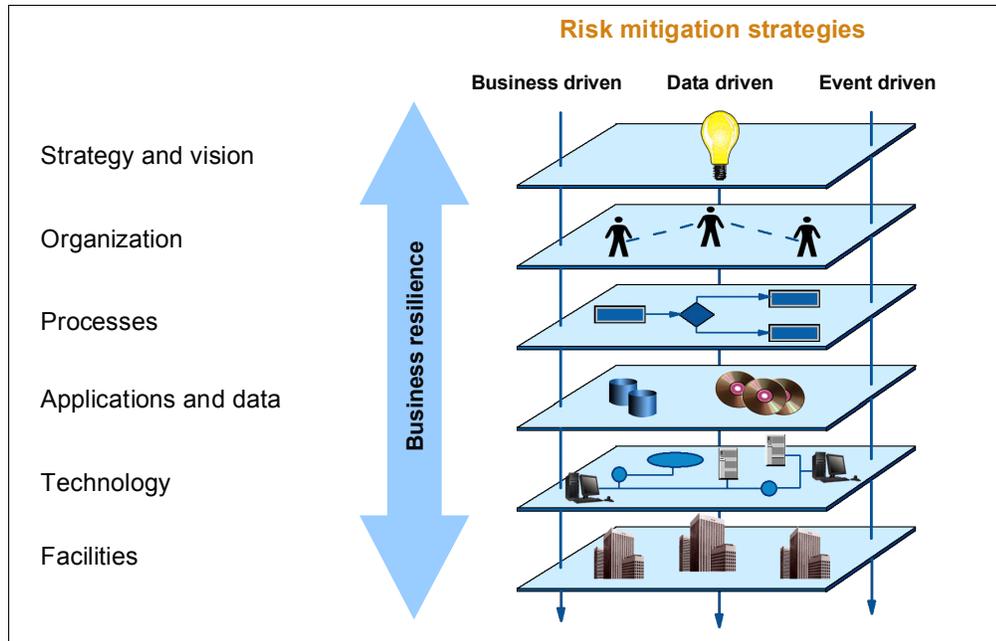


Figure 22 IBM business resilience framework

To help you keep your business continuously operating and simplify management of compliance with industry and government regulations, Juniper Networks' QFabric offers a business resilient infrastructure that allows you to mitigate business-driven risks. To mitigate event-driven risks caused by power outages, natural disasters, fires, and other IT disruptions, QFabric can help your organization at the technology level to be able to distribute operations beyond the area of immediate impact as well as implement an effective disaster recovery and crisis management plan.

Now we demonstrate how QFabric can support your business resilience and recovery strategy and operation.

Preconditions

Here are the critical elements for a comprehensive and resilient BCRP, along with DCN interconnect for rapid recovery:

- ▶ High availability requirements
- ▶ Resilience requirements
- ▶ Flexibility requirements in a DC
- ▶ Easy recovery, making services available across DCNs in the event of major disaster.
- ▶ Business continuity and recovery plan.

Scenarios

We describe two interlocking scenarios in this use case:

1. An intra-DC design
2. An inter-DC design

The intra-DC scenario is more focused on business continuity (BC), whereas the inter-DC scenario is geared more to both business continuity and business recovery (BR).

Intra-DC design

The intra-DC scenario (shown in Figure 23 on page 44), represents a highly available (HA), redundant and resilient network. It is deployed to satisfy business continuity and easy recovery requirements through the use of QFabric and other high performance, highly reliable Juniper networking components:

- ▶ Two Juniper MX series routers are used to provide redundant WAN connectivity and critical VRRP configuration to eliminate a single point of WAN failure.

Virtual Router Redundancy Protocol (VRRP) is deployed to increase the availability and reliability using multiple paths to the WAN connections. Each QFX node physically connects to two MX series routers acting as a group. The default gateway of the blade server is assigned to the virtual router instead of a physical router. If the MX router where the virtual router is running fails, the redundant MX router is selected to automatically replace it to pick up the packet forwarding function.

- ▶ Two separate QFabric networks are deployed to provide any-to-any high-speed core fabric switching services. In this case, both QFabric networks can be used as active switches passing data simultaneously, or one QFabric network can be used as an active switch, while the other is used as a passive switch. The passive switch will become active in case of a crash of the active QFabric network. Several QFabric features are also deployed in the QFabric to meet the HA and resilience requirements:
 - Redundant Server Node Group (RSNG) will be deployed in this design. In addition, LAG connections from a server are distributed across the RSNG to provide a redundant connection between the server and the QFabric system. RSNG supports IBM BladeCenter® and OSA connectivity to IBM System z®.
 - Network Node Group (NNG) is also configured to connect the QFabric Nodes to both MX series routers that run network-facing L2 and L3 protocols, such as xSTP, PIM, OSPF, BGP, to provide physical network level redundancy.
- ▶ From a management standpoint, the QFabric Director provides a single point for controlling and managing all QFabric components as a single logical switch. The QFabric Director is the routing engine embedded within a switch and externalized in the QFabric architecture.
- ▶ Redundant Juniper SRX series Services Gateways are deployed to offer high performance advanced security functions. They are dual-homed to two QFabric networks and can be easily managed and configured as zone-based firewalls to isolate and secure networks.
- ▶ Redundant ADCs are also dual-homed to HA QFabric network to provide wide area network/application acceleration and load balancing between servers.
- ▶ vMotion is used to provide high performance replication of virtual machines (VM) between BladeCenter servers within a data center spanning across two instances of QFabric to increase flexibility and availability for your business needs.
- ▶ Redundant FC/FCoE gateways are deployed to bridge Fibre Channel over Ethernet (FCoE) connectivity between servers and legacy SAN storage. However, enterprise systems such as System z will connect to the SAN storage via FC or IBM FICON® connections. Multiple Open Systems Adapters (OSA) in System z provide high bandwidth data throughput, network availability, reliability, and recovery.

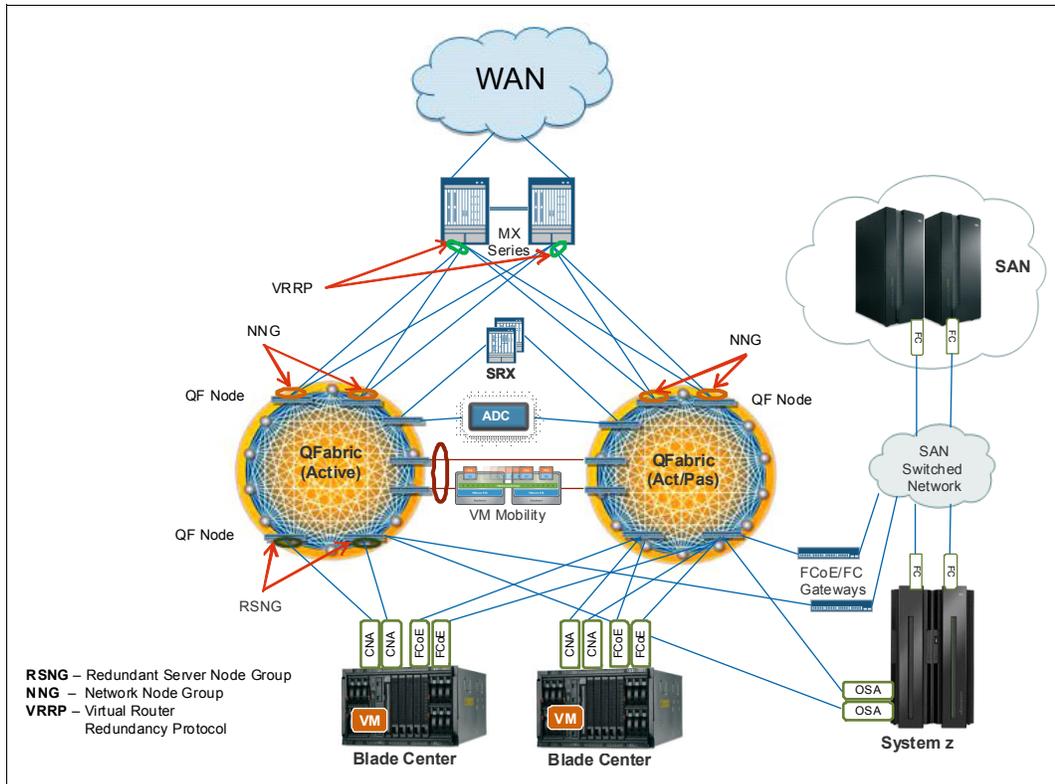


Figure 23 HA, resilient, flexible intra-data center scenario

Inter-DC design

The inter-DC design is more important in today's network disaster recovery. Depending on the applications that data centers need to support, the network can run on L3 or L2 stretches between two DCs using MPLS/VPLS²² over DWDM²³. See "DC interconnection issues" on page 29 for details.

Interconnecting data centers with MPLS/VPLS over DWDM increases high availability with expedited application replication service across data centers and improves time-to-recovery in the event of a disaster.

In the inter-DC scenario shown in Figure 24 on page 45, the attributes and benefits of using RSNG, NNG, VRRP, and multiple FC/FCoE gateways are similar to those for the intra-DC configuration.

Redundant MX series routers are deployed to provide VRRP features and for internet or internal network routing. The current QFabric supports 12 K routes, which might not be enough to hold an internet routing table. Most large data centers have a full BGP feed, which requires MX series routers.

Enabling vMotion within a data center can improve business continuity; however, across data centers over distance this could be very challenging using today's technology. For example, the following time delay limits are difficult to meet:

- ▶ Less than 10 ms delay for vSphere version 5
- ▶ Less than 5 ms delay for vSphere version 4
- ▶ Less than 1 ms delay for vMotion Fault Tolerance deployment

²² MPLS/VPLS - Multiprotocol Label Switching/Virtual Private LAN Services

²³ DWDM - Dense Wavelength Division Multiplexing

In addition, FCoE is not supported over WAN inter-site connections. FCoE requires Lossless Ethernet for WAN network connections, which is currently not DCB compliant.

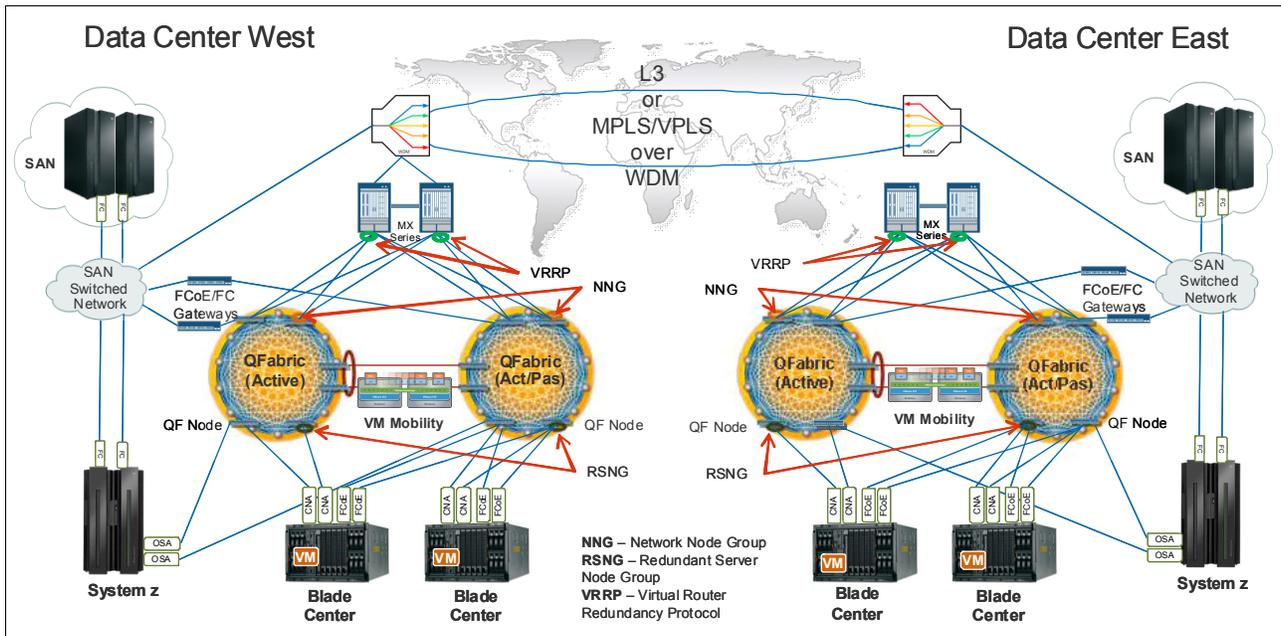


Figure 24 HA, resilient, flexible inter-data center scenario

Results

This use case, based on a two DC business deployment model, illustrates the functional requirements of business continuity and recovery. QFabric solutions support this business use case by applying different communication-based use cases at the DCN level as follows:

- ▶ Business-to-business use case to support communication between business-to-business resources
- ▶ Business-to-consumer use case to support communication between business-to-consumer resources
- ▶ Business-to-employee use case to support communication between business-to-employee resources
- ▶ Connectivity between data centers that involves enterprise networks and data center networks
- ▶ DCN access to support communications between various actors within the data center

We can discuss how QFabric deployment models successfully support this use case. For this use case scenario, resilience and availability non-functional requirements address those aspects of the supporting DCN infrastructure. While not directly affecting the system functionality seen by users, they can have a profound effect on how that system is accepted by both the users and those responsible for supporting the system. Greater flexibility, VM mobility, performance, and service requirements also provide additional support to concepts addressed in this use case.

Higher business resiliency and availability

QFabric supports business resilience and service availability in the following ways:

- ▶ Hardware resource resilience, for example, redundant fabric, Interconnect, and Directors.

- ▶ Data and control plane resilience, for example, redundant fabric links, load balanced routes, and redundant control ports.
- ▶ Separating data and control plane inside the QFabric data center network stretched outside the boundaries of a single box.
- ▶ Redundant management infrastructure, for example, redundant management connections and control networks.
- ▶ Leveraging QFabric redundant server node group (RSNG) configuration. See “Redundant server node group (RSNG)” on page 27 for details.
- ▶ Leveraging QFabric network node group (NNG) configuration to provide redundant external connectivity to routers, switches, and service appliance. See “Network node group (NNG)” on page 27 for details.

Greater flexibility and VM mobility

VM mobility allows flexibility in application deployment and recovery because you can move a running VM from one physical server to another with minimum impact to users. VM mobility operates only in an L2 environment intra-data-center and inter-data-center under specific requirements and configurations. Supporting this deployment requires dynamic resource optimization in the resource group and the network attributes in the following list. QFabric provides a flexible any-to-any network environment for VM mobility support that addresses the necessary attributes and characteristics, some of which include:

- ▶ Highest availability.
- ▶ Support for various performance requirements, such as 1 GbE and 10 GbE for server access.
- ▶ VLAN, anywhere in the DC, supporting large VM mobility domains.
- ▶ The QFabric key characteristic of low latency supports easier recovery, as in the example of VMWare fault tolerance (FT) that creates a duplicate, secondary copy of the virtual machine on a different host within the same DC, or across DCs.
- ▶ Support for large MAC tables to scale a number of VMs.
- ▶ Deploys network state mobility along with the VM mobility.

Enabling vMotion within a DC can improve business continuity; however, challenges can arise using current technology across DCs over distances.

Highest performance and service level requirements

MPLS and VPLS are proven technologies with widespread deployment that provide reliable L2 stretch and L3 segmentation for network services. VPLS provides VLAN extension over a shared IP/MPLS network. This setup requires a network with low latency and jitter, and high resiliency whenever you need to stretch L2. You need to consider time delay from long distance connection for VM mobility deployment as we mentioned previously.

QFabric satisfies these requirements with its key features and characteristics, such as ultra-low latency and non-dropping packets under congestion for FCoE flows.



Transform your DCN to be fabric-enabled

IBM understands that the first step to transform the network infrastructure is to develop a sound enterprise network architecture, one that takes the business and IT environments, security and privacy policies, service priorities, and growth plans into account.

In this chapter, we describe how to migrate to a smart data center with QFabric, and discuss the organizational considerations involved in the migration.

Transforming the DCN

In today's data center, the rapid growth of compute capacity and increased data volume demand a high performance and scalable data center network infrastructure.

Business are leveraging virtualization and increasing operational efficiency to gain business agility and reduce costs. Unfortunately, the complexity and inflexibility of many legacy DCN architectures are inhibiting scalability and business growth.

Juniper has developed QFabric architecture, a fundamental improvement in networking architecture, which collapses the DCN by flattening the legacy hierarchical model with any-to-any connectivity. QFabric's features enhance the user experience and improve economic factors while creating a next generation data center.

Keeping your business in a traditional DCN model does not allow you to respond to new business requirements and growth, but even so, moving to a new model might appear risky at first. As we show in this chapter, transforming your legacy DCN to a QFabric based one is a dramatic improvement in infrastructure, services, operations, and cost. We also highlight how you can accomplish the migration process in a structured and efficient way.

QFabric helps you smoothly transition to a smart DCN with the following features:

- ▶ QFabric seamlessly integrates with existing management systems and network devices such as routers, firewalls, proxies, WAN Optimization Devices, DWDM switches, and so forth.
- ▶ All QFabric components are centrally managed from the QFabric Director, which provides standards-based management interfaces for interoperating with leading third-party network management and orchestration systems.
- ▶ QFabric provides a familiar operational model and tools like ping, traceroute, and so forth.
- ▶ QFabric provides one logical device and one operating system abstraction.
- ▶ QFabric transforms the DCN so that it behaves and is managed as a single, logical device.
- ▶ QFabric builds the foundation for interconnecting a dynamic pool of resources across the entire data center, and seamlessly integrates with existing infrastructures, providing deployment flexibility without requiring wholesale replacement of the existing network or compute resources.

Network migration strategy

You can compare the migration process to that of manufacturing. In the initial stage, you design a clearly defined and validated product; it is approved; and then you define how it will be manufactured. So we determine our final product; create a strategy to move from the current product to the final one; and then we define the processes and tools to reach our goal.

The transformation of a DC is a complex global project that is critical for the enterprise. We describe the migration strategy from the network point of view with a focus on the design and implementation phases.

The process demonstrates how you can migrate your services and applications from your legacy network to a smart DCN based on QFabric in an smooth and efficient way. The key assumption in the process at this point is that you create a new network and then migrate the existing services and applications to this new network.

Considerations

Virtualizing the data center introduces more complexity into the network. For example, a new application can introduce unknown flows that might be only partially managed. Any complexity or flows in a network that are unknowns, unmanaged, or partially managed will affect reliability and availability. The key requirements for transitioning to a smart DCN are a clear and simple design²⁴, sufficient training, frequent testing, capacity planning, operation teams using an industrial process, and both active and passive monitoring. Consider all these requirements as mandatory in building and managing your new smart DCN.

Creating a service that will be offered to customers using data center resources requires that all DC teams—servers, storage, network, virtualization, and applications—work on this service development. With this approach, the DC becomes the factory of the IT and one of the most important tools of your business.

As you move through the migration process, keep in mind various concerns such as financial constraints (OpEx and CapEx), production impacts, technical issues like training and process changes, and management issues. Each of these concerns is an individual step in the process that must be carefully defined and accurately executed.

Technical considerations related to QFabric architecture

The purpose of building a single-tier architecture is to create an environment where the Spanning-Tree protocol is not required for Loop-Prevention. QFabric supports the following methods that address this environment:

- ▶ Link Aggregation Groups (LAG) through the LACP protocol.
- ▶ L3 Routing Protocols, such as OSPF, BGP, ISIS, and RIP to establish L3 routing sessions. These protocols run on the NNG QF/Nodes.
- ▶ L2 Protocols, such as the Spanning-Tree Protocol and the common variants. These protocols are run on the NNG, and are used to interconnect QFabric to other switches directly. However, in the majority of use cases, QFabric will not need to run STP because QFabric can use LAG to interconnect with other L2 devices, enabling all links, and eliminating STP.
- ▶ Internal Loop Prevention. The Internal Fabric Control Protocol handles all internal loop prevention automatically.

Consider the following points when migrating from an environment where STP is currently widely deployed:

- ▶ Verify what protocols the network end devices support. If 802.3ad/LACP is supported, we recommend that you migrate away from STP, and use LAG interfaces for multiple interconnections.
- ▶ QFabric can be installed, configured, and connected to interfaces on the Core Switches that are administratively down, until the time when QFabric should be brought into the active network.
- ▶ Because QFabric will not require STP in the majority of situations, QFabric can be connected to the Core Switch with a LAG interface and made active immediately. Because QFabric does not participate in STP in this scenario, it will not try to become the Root Bridge and change the L2 forwarding pattern. It will just be seen as another LAG interface on the Core Switch.

²⁴ “Make everything as simple as possible, but not simpler.”, Albert Einstein

QFabric supports the major L3 Routing Protocols. In DCs where the EIGRP protocol is used, you might need to convert from EIGRP to a standard-based routing protocol such as OSPF, or ISIS as IGP, and BGP. Follow these steps to move from EIGRP to OSPF:

1. Enable OSPF Routing Process on the current environment.
2. Use a route-policy or route-map to export the EIGRP routes to OSPF.
3. Bring up the new Juniper QFabric device and configure OSPF.
4. Stop EIGRP when all routes are migrated inside an OSPF instance.

See *Migrating EIGRP to OSPF Day One Booklet*, from Juniper Networks for more migration details at:

<http://www.juniper.net/us/en/community/junos/training-certification/day-one/networking-technologies-series/migrating-eigrp-to-ospf/>

The design of QFabric allows the network administrator to maintain the defined VLAN hierarchy. The VLANs can be extended from the legacy network to QFabric and have larger, flatter VLANs; this increases the scope of their Migration Domain in a virtualized environment. You should plan VLAN usage to limit the Broadcast Domain. QFabric allows you to create very large L2 Domains. Create these carefully because larger domains result in increased broadcast traffic. Whenever possible, only configure the VLANs on the required QF/Nodes and ports.

There are two main recommended ways to connect an appliance to QFabric. A simple question determines which methods to use: Will the device be running a Network Protocol (L2 or L3)?

- ▶ If yes, then connect the appliance to the NNG.
- ▶ If no, then connect the appliance to an RSNG.

Because QFabric has the ability to move massive amounts of data, the core access layer device should be capable of providing multiple 10 GbE ports to ensure you have the required capacity.

The Fibre Channel over Ethernet (FCoE) and Data Center Bridging Exchange (DCBX) supports enable QFabric to serve as a convergence point for both the SAN and the Ethernet network. QFabric, seen as an L2 FCoE Transit Switch, performs FIP-Snooping to ensure that the FCoE traffic is sent to the appropriate ports. The FCoE flows are sent to the FCoE-to-FC Gateway, which can be done by QFX3500 in stand-alone mode.

Assumptions

In our example, we describe a migration inside a data center. The legacy network and the new network, based on QFabric, are physically separated by a short distance, that is, the networks can be in the same building or in two separate buildings within a campus.

A migration over long distances implies some substantial constraints such as:

- ▶ FC migration and synchronous replication are not possible.
- ▶ Legal constraints might apply if migration is from one country to another.
- ▶ Latency can occur, depending on distance.

For clarity and simplicity, figures here do not show all network components but offer a logical view of the components and functions. For example, the router component in the Core layer is a representation of all the routers in this layer. The same is true with the other layers. Blue lines in the figures show the logical connections needed or built during each phase and also indicate changes between phases.

We focus on the network in our migration example. Keep a global view of the transformation mechanisms that you will use to migrate applications and servers. A transformation using a physical server to physical server (P2P) migration approach will not have the same constraints as an approach using a physical server to virtual server (P2V), or one using a virtual server to virtual server approach (V2V). You need to study these approaches carefully with the System team so everyone can clearly understand the impact on the network migration.

Migration phases

The migration process, or project, is based on four main phases with a transversal phase for management. These phases are:

- ▶ Phase 1: Building a smart DCN. Details how to implement the new DCN based on QFabric and to how to prepare the new network for deploying new production services.
- ▶ Phase 2: Connecting the new DCN with the legacy network. Details how the new DCN is interconnected with the legacy network to allow communications between the two.
- ▶ Phase 3: Migrating services and applications to the new DCN. This phase is the heart of the process; it describes each service and application that is migrated through an established project process based on design, test, deploy and support activities.
- ▶ Phase 4: Disassembling the legacy DCN. The last step of the process during which you remove the legacy network.
- ▶ Transversal Phase: Project Management. This phase is conducted over the entire duration of the project and conducts global governance of the project such as planning definition, resources management, milestones, and follow up. Efficient project management is crucial for a successful migration.

Phase 1: Building a smart DCN

This phase outlines transforming to QFabric following the project model of design, test, deploy, and support.

- ▶ Design the new DC by defining high level and low level designs.
- ▶ Qualify the designs with lab testing.
- ▶ Deploy a new network infrastructure without production flows for this stage.
 - Qualify the network with testing prior to production.
 - Interconnect this new infrastructure with the WAN network of the DC, permitting access to the new infrastructure from the outside.
 - Deploy pilot applications on this new infrastructure in pre-production mode, deploying the applications in a test or experimental area to validate the design.
 - Deploy pilot application on this new infrastructure in production, validating the infrastructure stability.
 - Conduct acceptance test plan, which opens the smart DC to migration and new deployment.
- ▶ Conduct operational testing, or regular verification in service ramp up. Add production traffic and supervise the smart DC to verify that the infrastructure can absorb loads.

At the end of this phase, you will have a new infrastructure coexisting with the legacy one. At this point, there is no interconnection between the two infrastructures because services were migrated to the new network. Figure 25 shows a view of this new network.

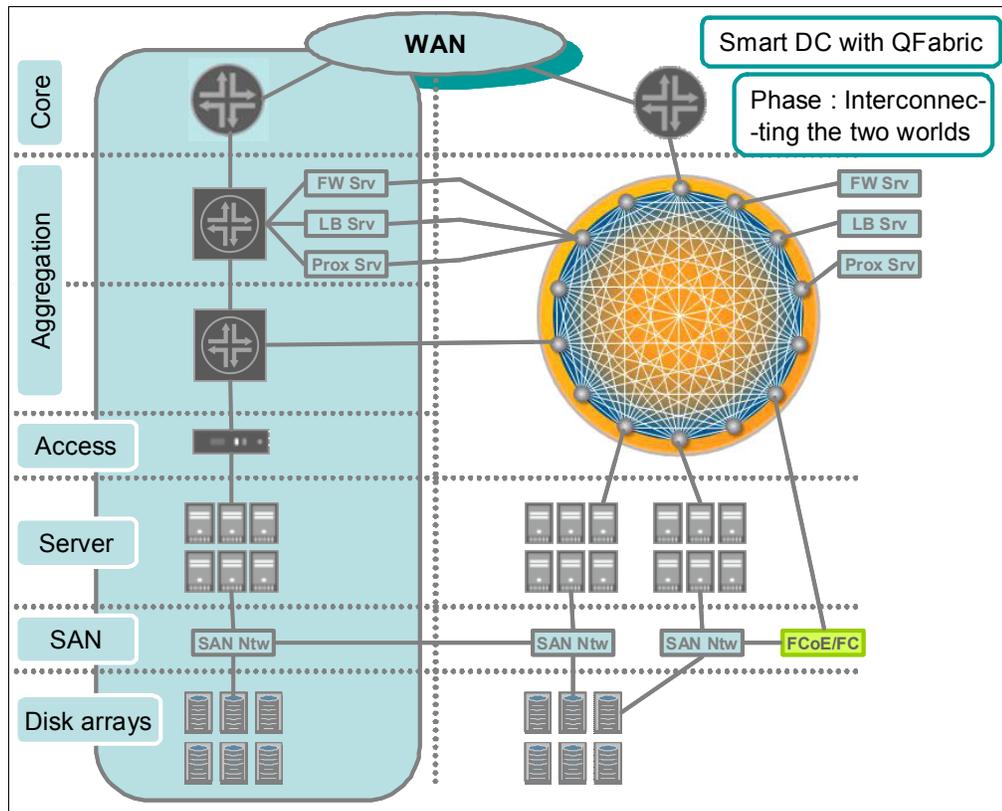


Figure 25 Network view of phase 1

Phase 2: Connecting the new DCN with the legacy network

Your next step is to connect your legacy network with your new smart DCN, one view of which is shown in Figure 26 on page 53. Connecting the two infrastructures allows open communication between them and prepares for the next phase.

Base this phase on the design, test, deploy and support project model, keeping in mind the following concerns:

- ▶ Define the interconnections into high level and low level designs.
- ▶ Test these interconnections in a lab environment.
- ▶ Build interconnects to access legacy resources.
 - L2 /L3 interconnections
 - SAN interconnections, if required.
 - Services interconnections, for example, security, load balancers, proxies, and so forth.
- ▶ Use network ready test plans to verify that all interconnections are working correctly.

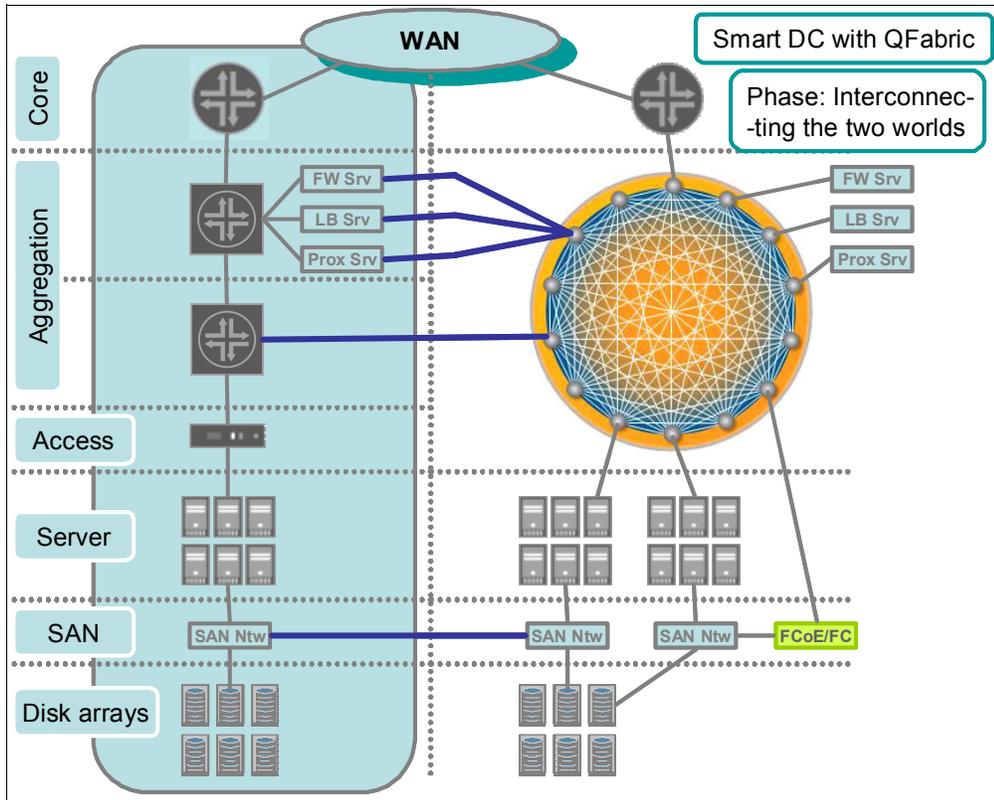


Figure 26 Interconnections between the legacy and new infrastructure

Phase 3: Migrating services and applications to the new DCN

This phase involves migrating services to the new smart DCN. The following points are critical in ensuring a successful migration:

- ▶ Prepare your migration design.

Study your existing environment to create an exhaustive view of all the components connected to the networks, including the flows between them for each service and application.

- Identify components, such as servers, storage, DB, security rules, and the flows between each component.
- Identify dependencies between components.
- Define non-functional requirements (NFR).

After thoroughly studying your existing environment, you should have a complete view of your applications and services, and fully understand constraints defined by the NFR. You need to define the target deployment of these services on the new infrastructure. After you define and document the target, you can determine the migration path.

- ▶ Define the migration scenario for each service, including steps and impacts. In some cases, you might need a temporary network for this phase.
- ▶ Each migration has unique constraints, complexity, and risks.
 - List constraints (for example, minimum or no impact allowed on a service).
 - Define a risk plan that includes risk assessment and action plans.
 - Define the scenario.
 - Define steps, methods, and scripts.

- ▶ Test the migration in a lab environment to validate the methods, process, and scripts. You might need to build the lab, recreating your existing environment and conducting test walk-throughs.
- ▶ Migrate services following the defined migration plan.
- ▶ Conduct operational testing, or regular verification, for the migrated service.

Figure 27 shows a network view where some services have been migrated to the new DCN.

Migration follow-up

Industrialize the migration of all the services from the legacy DC to the new DC using the same process as previously described, such as design, test, deploy, and follow-up.

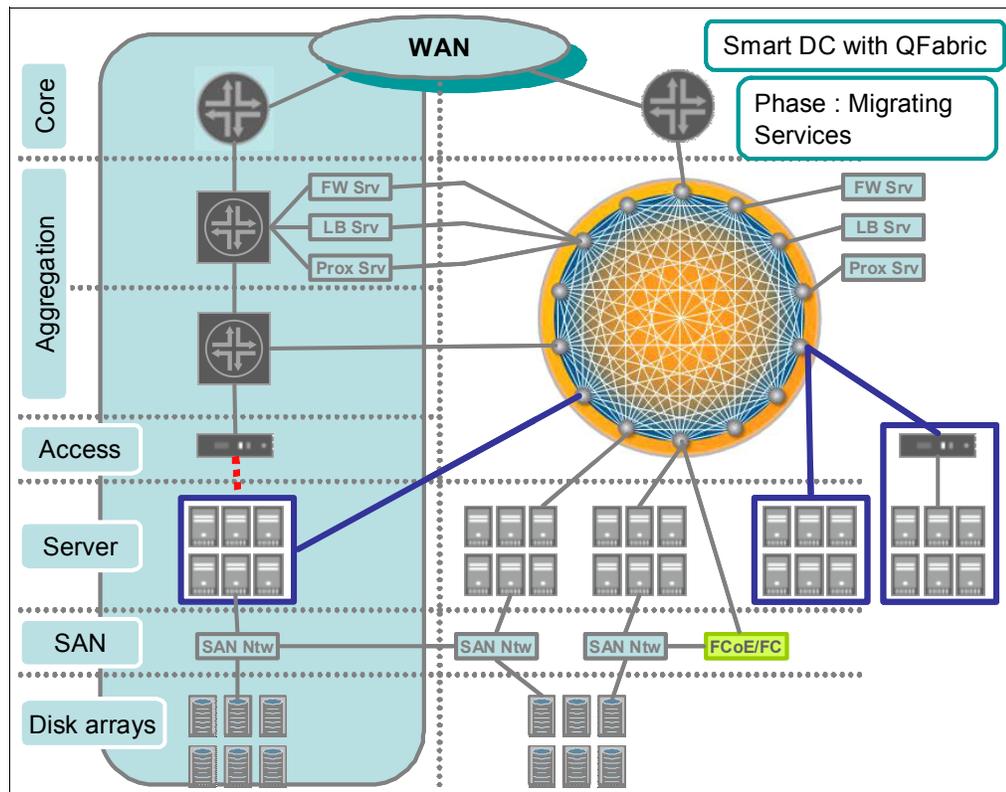


Figure 27 Migration view where some services are in the new infrastructure

Phase 4: Disassembling the legacy DCN

When you are finished migrating services, your legacy infrastructure should contain no more production flows. Your next step, after verifying stability on the new infrastructure, is to disassemble the legacy DCN, as shown in Figure 28 on page 55.

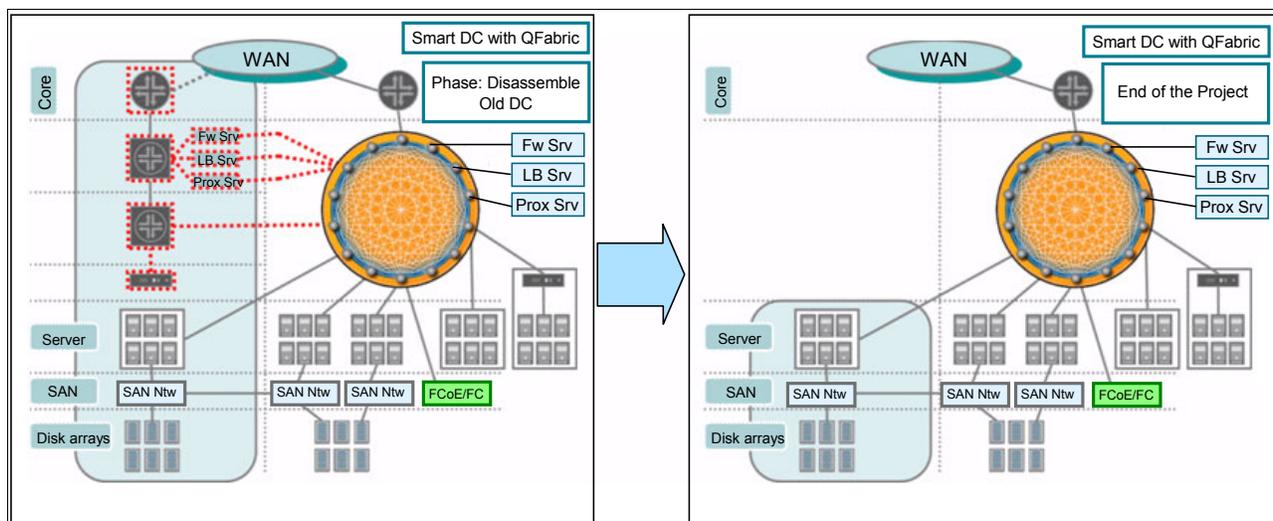


Figure 28 View showing disassembly and new DCN with QFabric

Transformation organization

As you transform your legacy DCN to a smart network with QFabric, you will learn about various technical concepts and trends, such as:

- ▶ Working out of the spanning tree
- ▶ Evolution to 10 GbE
- ▶ LAN/SAN convergence
- ▶ FCoE and related storage issues, (iSCSI / NAS / FC)
- ▶ Virtualization within the server
- ▶ VM mobility in the DCN
- ▶ DCN flows evolving from North/South to East/West
- ▶ Consolidating infrastructure networks
- ▶ Reducing power consumption and heat dissipation, thus reducing TCO.

Transforming your legacy network to a smart DCN also changes the role of IT in your organization. The cloud approach, virtualization, financial reduction, and legal constraints are the critical keys for the IT team as they manage servers, storage, network, db, security, and site facilities.

Moving to a QFabric environment requires careful planning and redesign at many levels, and must take into account many perspectives, including:

- ▶ Industry directions
- ▶ Technical issues
 - Horizontal integration - application delivery, security services, blade switches
 - Vertical integration - servers, storage, applications, cabling.
- ▶ Operational concerns
 - Increased importance of configuration and change management
 - Impact on operational processes
 - Organizational clarity on roles and responsibilities.
- ▶ Business demands
 - TCO benefits gained by reducing rack space, power, cooling, and Opex.

We discuss these topics further in the following chapter “Next steps: IBM and Juniper Networks can help”.



Next steps: IBM and Juniper Networks can help

High bandwidth applications, through the use of SOA and Web 2.0, have shifted the DCN paradigm. User and application demands will continue to force enterprises and service providers to take a more strategic approach in designing and deploying faster networks to link users, applications, and data. Companies will have to continue data center consolidation and use virtualization to efficiently move and replicate data and resources across networks. All these efforts will reduce capital and operational expenses, and give businesses competitive advantages in the marketplace.

A good strategic move is to leverage Juniper Networks' innovative fabric switching technology, which can provide a simple architecture with high speed connectivity to servers, storage, appliances, and network devices. QFabric offers unique market advantages and clear product roadmaps to meet your company's DCN requirements, while preserving current investments.

Over decades, IBM has built comprehensive technical expertise and a deep understanding of the evolving demands of network, server, storage, and desktop virtualization. IBM has extensive design and integration experience in complex data center networking infrastructures and cloud computing environments. Also, IBM has a global pool of skilled networking professionals with an in-depth knowledge of IT and networking infrastructure and world-class project management skills.

IBM and Juniper Networks have built a strong partnership that offers leading-edge network products and technologies. This partnership can help your company design and implement this unrivaled data center design and integration using one of the approaches described in this chapter, namely the comprehensive transition approach and the services lifecycle approach.

Comprehensive transition approach

Moving to a fabric environment requires careful planning and redesign on many levels. A comprehensive data center network design is critical when preparing for a smarter data center and cloud computing.

As a next step we suggest that you develop a more comprehensive approach to transition your data center network. You need to consider broad input from multiple teams that manage the IT infrastructure today. That input can be gathered across five key areas of architectural considerations:

- ▶ Design scope
- ▶ Security
- ▶ Management
- ▶ Organization
- ▶ Technology

To properly address these areas, greater cohesion will be required across the teams that manage the IT infrastructure. The key architectural considerations should be addressed in the context of the existing infrastructure, IT strategy, and overall business goals.

A transition to a cloud computing enabled data center network usually starts with virtualization scope, security, and management considerations to address some key questions:

- ▶ Design Scope
 - What are the existing and planned unique and overlapping virtualization capabilities of your heterogeneous system and storage platforms?
 - How should the network adjust and enhance the flow of traffic when virtual servers, in the form of virtual machines, are created, moved, or deleted at the click of a button?
- ▶ Security
 - Where will security devices, such as firewalls and intrusion protection and detection services, be placed and managed? Will they be physical, virtual, or both?
 - How do you support different types and levels of security for diverse application workloads on a shared, virtualized infrastructure?
 - How do you achieve fast provisioning and mobility of security services to match the speed of virtual machine creation, movement, and deletion?
- ▶ Management
 - How do you define a smooth migration path from a static network to one that can respond to a highly automated and dynamic infrastructure?
 - How do you manage the delivery of differentiated service levels to different application workloads on a shared infrastructure while meeting existing service level agreements with users?
 - What amount of monitoring is required, and at what level of granularity, for virtualized resources to enable required performance awareness, event correlation, and reporting?

After considering these issues, you can take the next step of considering organization capabilities and technology adoption that are also needed for your comprehensive transition approach. The following issues come into play at this point:

- ▶ Organization
 - Who will implement, configure, manage, maintain, and support network resources that reside in virtualized platforms or exist as virtual appliances?
 - How is the administration accountability and integrity maintained when mobility of the resources is incorporated into the architecture?
 - Does your team require enhanced interdisciplinary skills to support a virtualized environment?
- ▶ Technology
 - How do you decide where and when to adopt an emerging, disruptive technology such as Fibre Channel over Ethernet (FCoE), which enables the convergence of data and storage networks?
 - Which network protocols and standards do you adopt to address the requirements of your dynamic infrastructure, for example Multiprotocol Label Switching (MPLS), virtual routing and forwarding (VRF), Internet Protocol version 6 (IPv6), or pseudowires?
 - Which technologies are your peer teams (network, servers, storage, applications) adopting?

Both Juniper Networks and IBM provide services that can help you move to a fabric-enabled DCN. For more information, go to:

<http://www.juniper.net/au/en/products-services/consulting-services/>

http://www-935.ibm.com/services/us/en/it-services/integrated-communications-services.html?cm_re=masthead-_-itservices-_-communications

Services lifecycle approach

Based on a tested and refined lifecycle model for networking, IBM Network Strategy and Optimization Services and Network Integration Services can ensure your network can provide the level of availability and performance your business requires. This process is illustrated in Figure 29 on page 60.

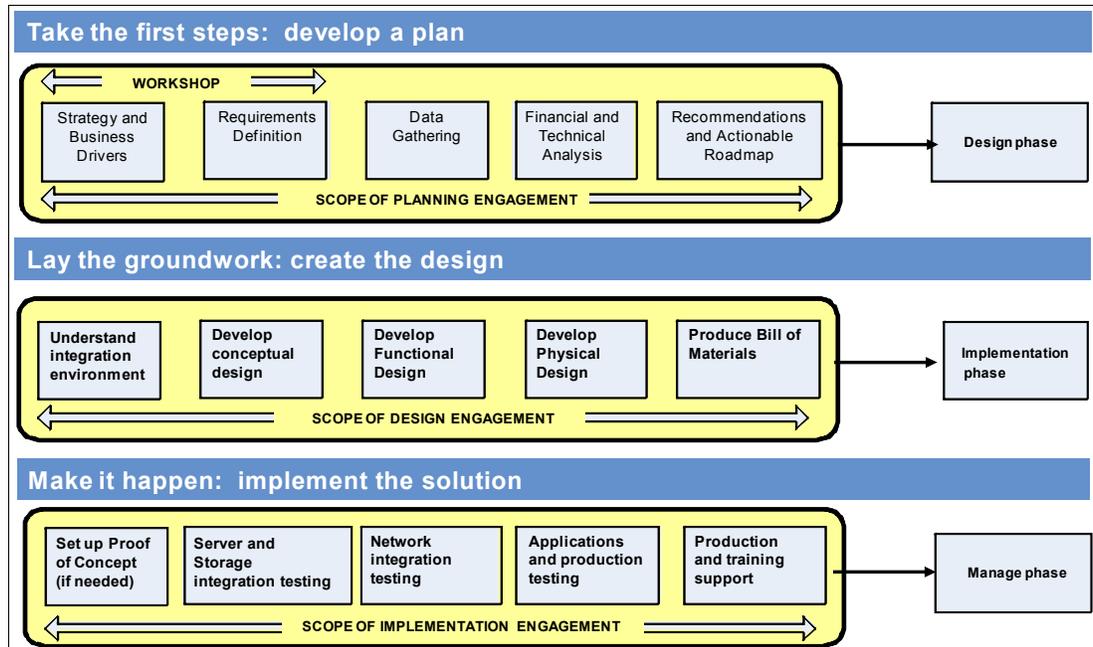


Figure 29 The IBM service lifecycle approach

IBM can walk you through the process of designing and building a smart DCN. Experienced IBM practitioners will help you to:

- ▶ Develop a plan that fits your unique needs:
 - Understand your current IT and networking environment.
 - Collect and document your requirements.
 - Identify performance and capacity issues.
 - Determine your options.
 - Compare your current environment to your plans.
 - Make recommendations for transition.
- ▶ Create a design to fulfill the requirements:
 - Develop a conceptual level design that meets the identified solution requirements.
 - Create a functional design with target components and operational features of the solution.
 - Create a physical design to document the intricate details of the solution, including vendor and physical specifications, so that the design can be implemented.
 - Produce a bill of materials and a plan for implementation.
- ▶ Provide a seamless integration:
 - Review the implementation plans.
 - Perform a site readiness survey.
 - Procure the equipment.
 - Develop installation procedures, solution testing, and certification plans.
 - Stage the solution.
 - Implement and test the solution.
 - Train your in-house support staff.

For details about this and other IBM Integrated communications services, go to:

<http://www-935.ibm.com/services/us/en/it-services/integrated-communications-services.html?lnk=mhse>

A partnership that makes a real difference

The value of working with IBM and Juniper Networks²⁵ lies in their complementary strengths. This 360-degree, broad-based strategic alliance builds greater business value and more robust solution offerings, as illustrated in Figure 30. The QFabric architecture has its origins in the ground breaking Stratus Project, a collaborative effort that drew from the best of both companies to deliver a quantum leap in networking technology. It is one example of what IBM and Juniper have been able to achieve by working together.

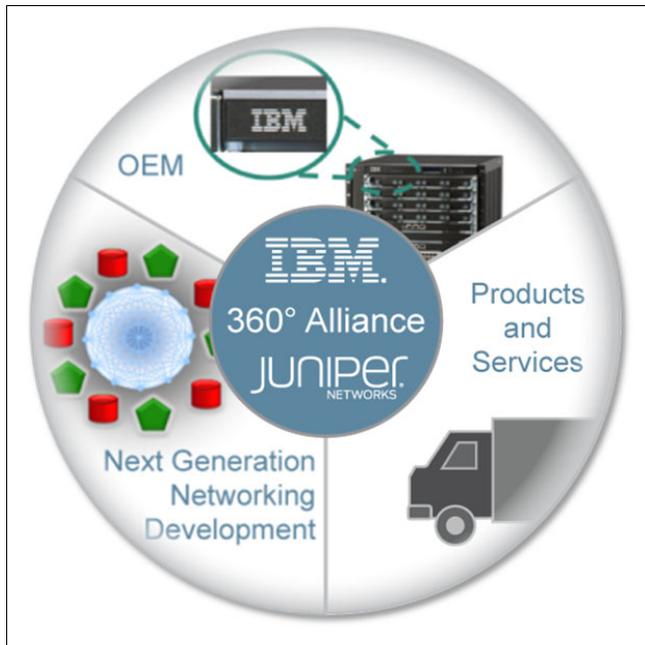


Figure 30 IBM and Juniper - a true 360-degree relationship

The relationship goes beyond product and solution development. It also allows easier management and more cost-effective engagements. A joint service relationship enables IBM to provide Juniper's Ethernet networking and security products and support within the IBM data center portfolio of products, simplifying the vendor engagement for network infrastructure customers. In addition, IBM Managed Maintenance Service offers remote and onsite IBM support for Juniper devices. This allows clients to add Juniper hardware to their existing IBM maintenance contract, simplifying management and accountability.

Together, IBM and Juniper Networks have the experience, resources, and solutions to deliver the highly responsive and efficient networking infrastructure that can become an important competitive differentiator for today's businesses. IBM also offers financing to accelerate your project's cash flow break-even point. The key is the combined strength of two industry leaders, working together to transform the way business is performed today and tomorrow.

Learn more about Juniper Networks' data center network solutions, at:

<http://www.juniper.net/us/en/dm/simplify/>

For a list and description of IBM and Juniper Networks solutions, refer to:

<http://www.ibm.com/solutions/alliance/us/en/index/junipernetworks.html>

²⁵ IBM and Juniper Network - Creating a trading edge for financial markets:
<http://public.dhe.ibm.com/common/ssi/ecm/en/jns03005usen/JNS03005USEN.PDF>

QFabric implementation at IBM

QFabric is operating today at the IBM site in Poughkeepsie, New York, in the IBM global test environment for System z (IBM zEnterprise™ System) and High Performance Computing (HPC) servers.

IBM and Juniper are working on joint technology solutions, products, and standards development, network management, cloud solutions, and managed security services.

The IBM Poughkeepsie Lab is an early field trial site for QFabric. QFabric is deployed to test IBM systems and software. As these tests are conducted, QFabric is used as the switching technology for connectivity. The goal is to see IBM systems and software in a quasi production environment, making use of the QFabric system as a networking infrastructure.

Figure 31 is an example of a test scenario currently set up in the IBM Poughkeepsie Lab.

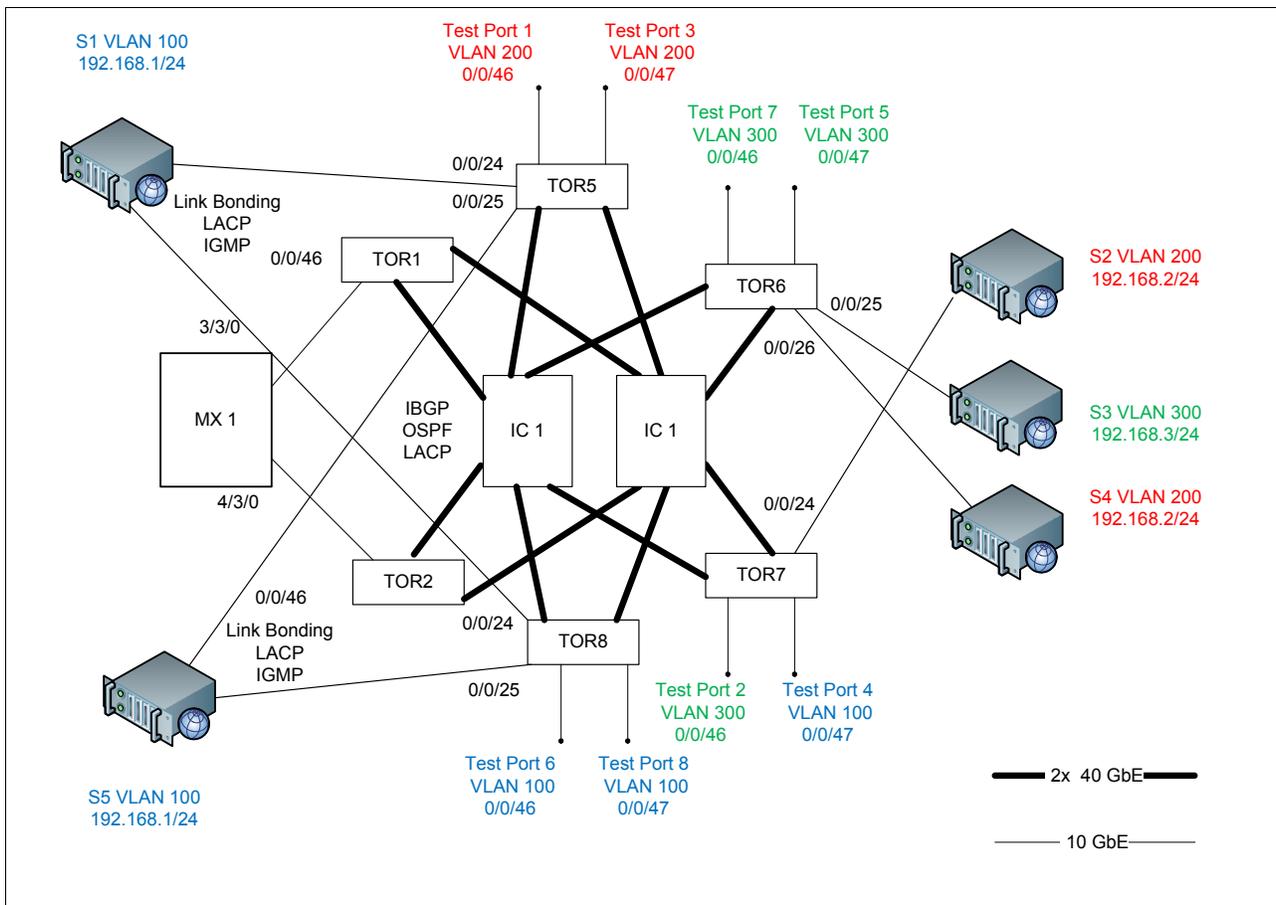


Figure 31 QFabric test environment at IBM Poughkeepsie, NY

In this test environment, the hardware consists of:

- ▶ QFabric components:
 - Two QFX3008 Interconnect Chassis (IC)
 - Two QFX3001 System Directors
 - Six QFX3500 Edge Nodes (TOR)

- ▶ MX-Series:
 - One MX960 Router/Switch (MX)
 - Two 10 GbE connection to the fabric
- ▶ Servers:
 - Three IBM 3550 M2 servers running RH 5.03 using dual 10 GbE Solarflare NICs
 - Two IBM 3550 M2 servers running RH 5.5 using dual 10 GbE Emulex CNA NICs
- ▶ Packet Generator:
 - Spirent Test Center using 8 CV ports

In this test environment, the software consists of:

- ▶ QFabric:
 - Junos 11.3B2
- ▶ MX-Series:
 - Junos any supported version
- ▶ Operating system on the servers:
 - Linux Red Hat 5.5
- ▶ The test software consists of:
 - Iperf is commonly used software to test TCP, UDP and IDP connections on a network, the software runs in many environments. Data sheets and downloads can be found at the sourceforge iperf site: <http://sourceforge.net/projects/iperf/?abmode=1>
 - Mgen is commonly used software for measuring network scalability and performance with multicast. Data sheets and downloads can be found at the US Navy site: <http://cs.itd.nrl.navy.mil/work/mgen/index.php>
 - Spirent Test Center - Eight ports are used on a model CV card to evaluate the performance of QFabric for both unicast and multicast traffic.

QFabric has become an essential part of the IBM Poughkeepsie Lab, which carries out end-to-end tests from the VM/servers to storage through all elements composing the data center. Recently, the IBM Poughkeepsie Lab completed STAC benchmark testing²⁶ of QFabric and compatibility matrix qualification for QFabric in the SAN environment.

The IBM/Juniper alliance team has also designed an environment to conduct client demos, perform Proof of Concept studies, and provide client tours. Contact your IBM representative for details about the QFabric Lab & Design Center, or request a briefing at:

<http://ibm-vbc.centers.ihost.com/schedule/request-briefing/>

²⁶ For a description and a link to more details, go to “STAC benchmark results” on page 20.



The team who wrote this guide

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

Bill White is a Project Leader and Senior Networking and Connectivity Specialist at the International Technical Support Organization, Poughkeepsie Center.

Hoi T Hon is an Infrastructure Architect in the IBM Global Services, US ICS/GTS division, and a member of the Data Center Networking, Center of Excellence. He specializes in network consulting, optical network, data center, and UC™ solutions using multi-vendor technologies and products. Hoi has 17 years of experience in the networking field and holds Juniper and many other industry certifications. He has a Bachelor's degree in Electrical Engineering from CCNY, CUNY. Hoi has written several IBM papers on network architecture and deploying optical network solutions.

Stephen Sauer is an Infrastructure Architect with IBM Global Services, France ICS/GTS division and a member of the Data Center Networking, Center of Excellence. He has over 20 years of experience in the field of network infrastructure and telecommunications. He has held various positions, including Support Engineer within data center operations and Technical Leader at a network integrator. Currently, Stephen works on the development of the network in cloud computing environments.

Raymund Schuenke is an IT Consultant with the Integrated Communications Services (ICS) practice in Global Technology Services®, IBM Germany. He has been with IBM for more than 10 years. His areas of expertise include data center networking, networking strategies and their alignment with overall IT strategies, managed services and processes, and service management. Raymund has led and contributed to many cross-brand client consulting engagements across industries.

Thanks to the following people for their contributions to this project:

Ella Buslovich, Alison Chandler, Irena Slywkanycz
International Technical Support Organization, Poughkeepsie Center

Faton Avdiu, Casimer DeCusatis, Peter Demharter, Michele Girola, John Havriluk,
Mark Lewis, and Wes Toman
IBM

Greg Bassett, Scott Burwell, Sean Capshaw, Jonathan Coleman, Vaishali Ghiya, Michael Goldgof, Ron Halbach, Mark Hinckley, George Hennessy, John Moore, Chris Rogers, Fraser Street, and Jeremy Wallace
Juniper Networks

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks® publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document, REDP-4830-00, was created or updated on March 7, 2012.



Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>



The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

BladeCenter®	Redbooks®	UC™
FICON®	Redguide™	WebSphere®
Global Technology Services®	Redbooks (logo)  ®	zEnterprise™
IBM®	System z®	
Netcool®	Tivoli®	

The following terms are trademarks of other companies:

Other company, product, or service names may be trademarks or service marks of others.