

OPPORTUNITIES AND CHALLENGES WITH THE CONVERGENCE OF DATA CENTER NETWORKS

10GbE, Standards-Based DCB, Low Latency
Ethernet, and FCoE Are the Starting Points
for True Network Convergence

Table of Contents

Executive Summary	3
Introduction	3
Requirements for Converged Data Center Networks	4
Three Different Technologies for Three Different Networks	5
Ethernet	5
Enhancements to Ethernet for Converged Data Center Networks: DCB	5
DCB: A Set of IEEE standards	5
Infiniband	7
Low Latency Ethernet, iWARP, and RoCEE	7
Fibre Channel	7
Enhancements to Fibre Channel for Converged Data Center Networks: FCoE	7
Future Direction for FCoE	8
Meeting the Scalability Requirements	8
A Clear Return on Investment	9
Prerequisite: Physical Infrastructure	9
Phase 1: Ethernet I/O Consolidation Within a Rack	9
Phase 2: Consolidation of Infiniband and Fibre Channel to an Ethernet-Based Infrastructure	10
Cautionary Notes	11
Phase 3: Fabric Convergence Throughout the Entire Data Center	12
Conclusion	12
About Juniper Networks	13

Table of Figures

Figure 1: The legacy approach of deploying and managing different network infrastructures has created complexity, compromised performance, and prohibited scalability in the data center	4
Figure 2: PFC ETS and QCN	6
Figure 3: Ethernet I/O convergence and network simplification	10
Figure 4: Server access layer convergence options	11
Figure 5: The convergence of LAN, SAN, and HPC clusters represents the third and final phase in the adoption of a unifying next-generation network fabric in the data center	12

Executive Summary

Ethernet and Fibre Channel are mature standards that have evolved over time thanks to substantial long-term investments, and both—along with Infiniband—have significant market share in the data center. However, due to the drive for larger, more efficient data centers, each of these technologies has faced significant scaling challenges. The development of extensions to these standards—in particular Data Center Bridging (DCB) and Fibre Channel over Ethernet (FCoE)—have for some time been promoted as a way of providing the necessary capabilities for true network convergence.

In 2011, customers will finally be able to invest in convergence enabling equipment and begin reaping the benefits of convergence in their data centers. With the first wave of standards now complete—both the IEEE DCB enhancements to Ethernet and the INCITS T11 FC-BB-5 standard for FCoE—enterprises can benefit from server- and access-layer I/O convergence while continuing to leverage their investments in existing aggregation, core LAN, and Fibre Channel backbones.

Also in 2011, products will emerge that not only provide 10-gigabit Ethernet (10GbE) but also 40GbE at the edge. These will be accompanied by fabric technologies to provide the necessary bandwidth for connecting the edges into a network that is optimized not just at rack scale but for larger regions of the data center. The price premium for 10GbE is slowly dropping, with lower cost converged network adapters and lower cost optics that allow data centers to reap the benefits of lower capital and operational expenses, thanks to convergence.

However, there is much more to come. Ten gigabit LAN on Motherboard (LoM) is largely a roadmap item. 10GBASE-T is still not power efficient enough to allow widespread deployment. There is no such thing as a standards-based fabric to connect it all together, and as networks converge, scalability issues get even more difficult. Other challenges also arise that need to be addressed before full network convergence can become a reality.

That said, enterprises can begin converging multiple network types at their own pace and within their particular cost and complexity constraints. As Ethernet continues to evolve, there will be many more opportunities—as well as many challenges—for network convergence. One such opportunity will be the ability to choose what to converge and at what cost. At a minimum, products can be purchased today with the confidence that they will continue to meet convergence requirements for years to come. This gives organizations the opportunity to evolve their networks over time, steadily gaining the benefits of convergence and dealing with the complexities of one step before moving to the next.

Juniper, which continues to be actively involved in both the DCB and FCoE standards setting efforts, has released its first generation of convergence products, and is fully committed to continuing the development of standards and products that will eventually make a fully converged data center network a reality.

Introduction

The data center network is the critical enabler of all services delivered from the data center. A simple, streamlined, and scalable data center fabric can result in greater efficiency and productivity, as well as lower operating costs. Such a network also allows the data center to support much higher levels of business agility and not be a bottleneck that hinders a company from releasing new products or services.

For both technical and business reasons, most organizations have multiple networks based upon multiple network technologies. These organizations have separate LAN segments for the campus and within the data center for management, backup, and inter- and intra-application communications, including clustering and virtual server mobility. In addition, many organizations have multiple Fibre Channel storage area networks (SANs) for disk and tape, different operating systems, or for different business applications or functions. Some organizations may also have one or more Infiniband networks to support their high-performance compute clusters.

Each of these networks is large enough to be built from multiple oversubscribed switching layers, each using different technologies from different companies, different hardware platforms with different architectures, and even different operating systems—oftentimes from the same vendor.

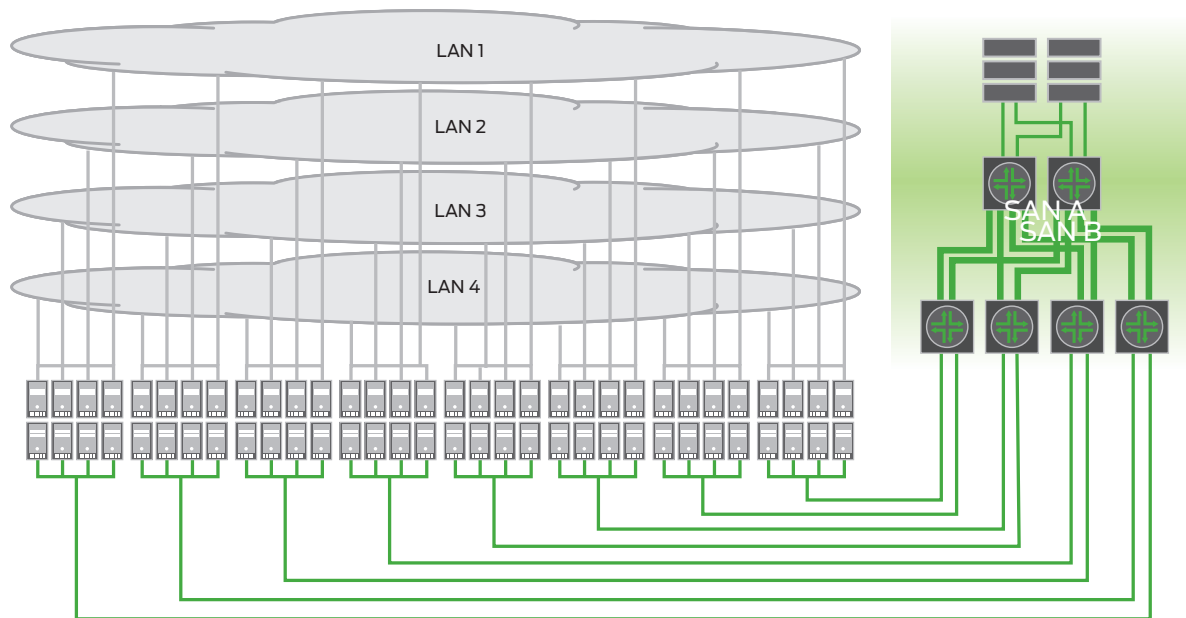


Figure 1: The legacy approach of deploying and managing different network infrastructures has created complexity, compromised performance, and prohibited scalability in the data center.

Businesses have tried to curb this complexity by deploying applications in dedicated infrastructure silos. However, while this allows optimization within individual silos, it results in poor overall resource utilization and higher operating costs across the data center. On the Fibre Channel SAN side, even though customers have high levels of storage pooling, those levels are often not as high as they would like. Attempting to converge existing SANs into the data center LAN, whether using network attached storage (NAS), Internet Small Computer System Interface (iSCSI), or FCoE (without changes to the existing approach), would split and isolate pools of storage into the much smaller silos than those that already exist on the Ethernet side, actually reducing efficiency and making compliance and data protection harder.

Put simply, merging multiple networks—each of which is struggling to scale and operate efficiently—to deliver convergence without first rethinking and simplifying the network is a recipe for disaster. To achieve the lowest total cost of ownership (TCO) in the data center, businesses must simplify the network, reduce the number of devices, collapse multiple switching tiers, use fewer operating systems across all network devices, and unfurl one unified, converged network fabric across all data center resources. To stand any chance of full network convergence, this simplification is absolutely essential.

Requirements for Converged Data Center Networks

IT organizations should consider a number of key factors before embarking upon any effort to achieve I/O or network convergence, or move towards any form of data center fabric, let alone try to do both at the same time. Instead, organizations need to ask themselves the following questions:

1. Is the network ready for cabling solutions designed for 10/40/100GbE?
2. Has the effort required to deploy a large flattened data center network been considered?
3. Have all of the requirements of network convergence been fully considered?
4. Are the implications of FCoE, iSCSI, and NAS fully understood?
5. Have the operational and organizational implications been considered?
6. Has the future trajectory of the standards been taken into account?

Three Different Technologies for Three Different Networks

The three primary networking technologies currently in use in today's data centers—Ethernet, Infiniband, and Fibre Channel—were created to address distinct sets of problems. As each has continued to evolve independently, it has constructed its own ways of performing the same networking tasks such as learning, addressing, discovery, monitoring and management, and quality of service (QoS). Each also has an entire ecosystem of vendors and tools.

Moving forward, each of these are evolving to support the needs of the new generation of data centers, including public and private clouds that include, but are not limited to, the needs of a fully converged network. In fact, they are all evolving and are only part way down the path required to deliver a fully converged network.

Ethernet

Of these technologies, Ethernet is the most prominent in global deployments and has evolved to address the most diverse networking requirements in local, metro, and WANs. Ethernet is deployed in data centers in much the same way it has been deployed in campus networks for years. Today, Ethernet's primary challenge in data centers is to provide greater scale, improved performance, and lossless transport. Many of Ethernet's shortcomings are avoided in the data center by deploying servers with multiple network interface cards (NICs), providing separation of campus traffic, backup traffic, management traffic, cluster traffic, server virtualization traffic, and storage traffic on separate logical networks—each configured to meet the needs of the given service. Other shortcomings of Ethernet are avoided by simply treating the data center not as a single large resource pool but as many small resource pools within which individual business systems must be implemented. Many Ethernet companies have been developing various standards and technologies, both independently and through IEEE and IETF, to address Ethernet's challenges within the data center. These new approaches, along with the move towards 10, 40 and 100GbE will lead to the creation of a new type of Ethernet data center network.

Enhancements to Ethernet for Converged Data Center Networks: DCB

Ethernet, originally developed to handle traffic using a best-effort approach, has mechanisms to support lossless traffic through 802.3X Pause, but these are rarely deployed. When used in a converged network, Pause frames can lead to cross traffic blocking and congestion. Ethernet also has mechanisms to support fine-grained queuing (user priorities), but again, these have rarely been deployed within the data center. The next logical step for Ethernet will be to leverage these capabilities and enhance existing Ethernet standards to meet the needs of convergence and virtualization, propelling Ethernet into the forefront as the preeminent infrastructure for LANs, SANs, and high-performance computing (HPC) clusters.

These enhancements benefit general Ethernet I/O convergence (remembering that most servers have multiple 1GbE NICs not for bandwidth but to support multiple network services), and existing Ethernet- and IP-based storage protocols such as NAS and iSCSI. These enhancements also provide the appropriate platform to support FCoE. In the early days of the development of these standards prior to them moving under the ownership of IEEE the term Converged Enhanced Ethernet (CEE) was used to identify them.

DCB: A Set of IEEE standards

Ethernet needed a variety of enhancements in order to support I/O and network convergence and server virtualization. The server virtualization aspects are covered in other white papers, even though they are part of the DCB protocol set. With respect to I/O and network convergence, the development of new standards started with the following existing standards:

1. User Priority for Class of Service - 802.1p, which already allows identification of eight separate lanes of traffic (used as-is)
2. Ethernet Flow Control (pause, symmetric and/or asymmetric flow control) - 802.3X, which is leveraged for PFC
3. MAC Control Frame for Priority-based Flow Control - 802.3bd, to allow 802.3X to apply to individual user priorities (modified)

A number of new standards that leverage these components have been developed. These standards have either been formally approved or are in the final stages of the approval process:

1. Priority Flow Control (PFC)—IEEE 802.1Qbb, which applies traditional 802.3X Pause to individual priorities instead of the port
2. Enhanced Transmission Selection (ETS)—IEEE 802.1Qaz, a grouping of priorities and bandwidth allocation to those groups
3. Ethernet Congestion Management (QCN)—IEEE 802.1Qau, a cross network as opposed to point-to-point backpressure mechanism
4. Data Center Bridging Exchange Protocol (DCBx)—which is part of the ETS standard for DCB auto-negotiation

The final versions of these standards specify minimum requirements for compliance, detail the maximum in terms of external requirements, and also discuss in some detail the options for implementing internal behavior and the downside of some lower cost but standards compliant ways of implementing DCB. It is important to note that these standards are separate from the efforts to solve the L2 multipathing issues that are not technically necessary to make convergence work. Also, neither these standards nor those around L2 multipathing address a number of other challenges that arise when networks are converged and flattened.

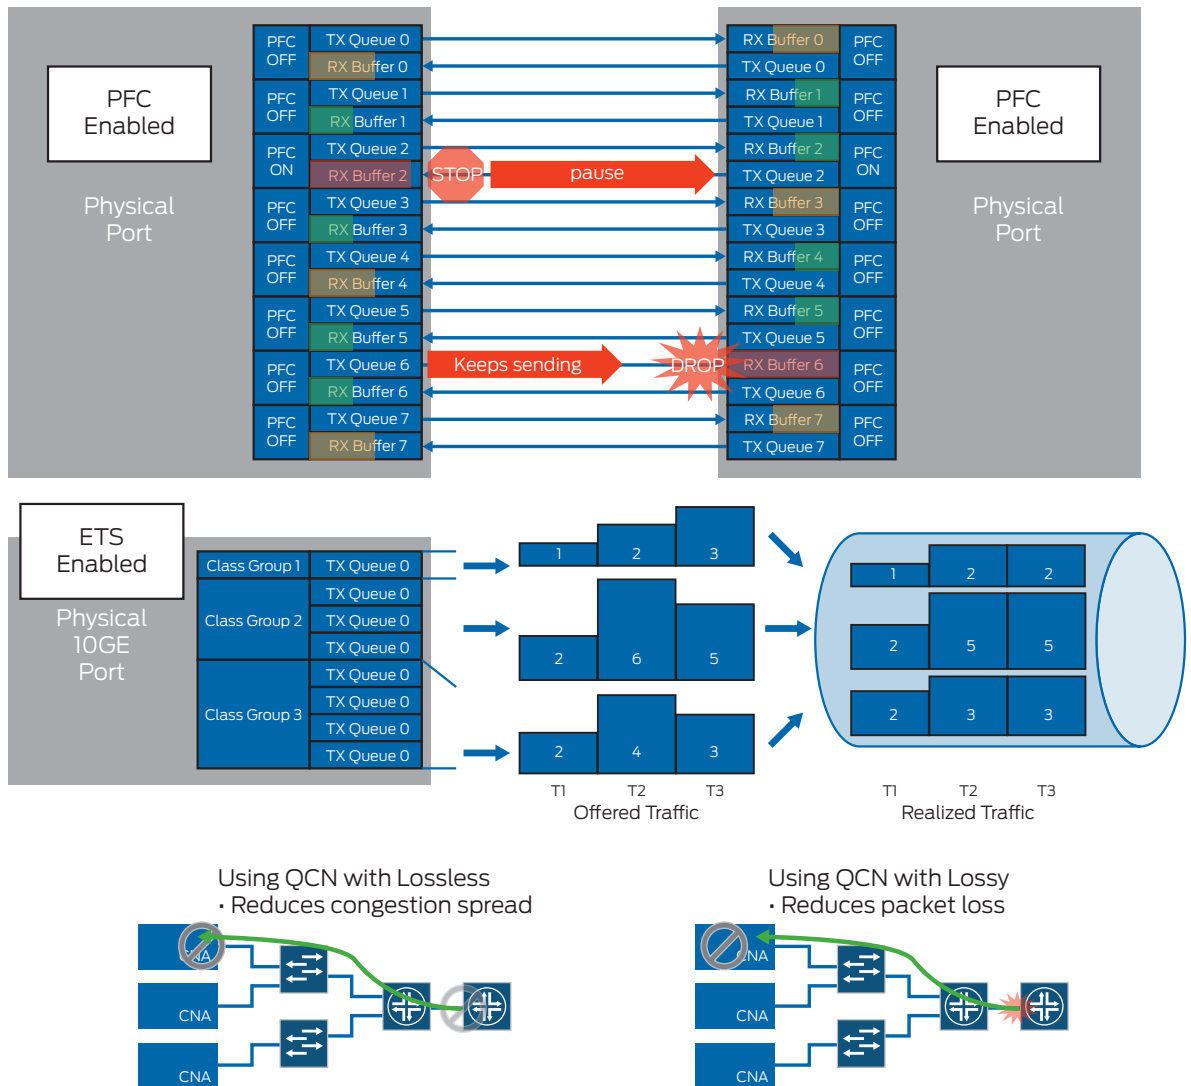


Figure 2: PFC ETS and QCN

Infiniband

Originally designed to replace the Peripheral Component Interconnect (PCI) that links I/O devices to servers, Infiniband has emerged as the de facto standard for low latency server-to-server interconnections in high-performance compute HPC grids and other cluster topologies. However, scalability and security issues continue to plague Infiniband in data centers today, and advances in Ethernet technology are quickly closing the gap in terms of throughput, latency, and cost. For a considerable time, Infiniband had clear advantages over GbE with respect to latency and performance, although this is no longer the case.

Low Latency Ethernet, iWARP, and RoCEE

Although Infiniband has successfully replaced a number of proprietary cluster network technologies, it is incorrect to characterize it as being the dominant, let alone the only switching technology used for HPC. Indeed, many clusters today do run on Ethernet switches. Clearly, the release of higher speed Ethernet switches (particularly 10GbE switches), along with switches built from a single ASIC and supporting cut-through switching, makes Ethernet a viable technology even for ultra low latency HPC.

The Ethernet switch, however, is only one part of the solution. Just as important is the implementation of Ethernet-based low latency HPC NICs that provide the operating system, middle layer, or application layer interfaces required to allow easy migration from Infiniband to Ethernet. The leading HPC I/O card vendors have led the development of such cards based around two technologies—Internet Wide Area RDMA Protocol (iWARP) and RDMA over CEE (RoCEE). From an Ethernet switch perspective (apart from providing the appropriate L2/L3 functionality needed by the customer), all that is necessary is to provide full Ethernet interoperability support with these interfaces, including at a DCB level where HPC cluster traffic desires a lossless Protocol Field Protection (PFC) priority for guaranteed low latency delivery.

Clearly, a robust implementation of DCB not only allows the deployment of HPC on an Ethernet platform but allows that network to be converged along with the traffic that previously would have required separate NICs to avoid cross traffic interference impacting HPC performance.

Juniper has worked closely with the leading HPC NIC vendors to ensure comprehensive support of both iWARP- and RoCEE-based solutions.

Fibre Channel

Fibre Channel was initially designed as a transport mechanism for the Small Computer System Interface (SCSI) protocol to connect and share a set of external disks with a group of servers. In doing so, the protocol is gaining the benefits of pooling by bringing together the benefits of the Ethernet network with the deterministic nature of the storage channel. The technology has been progressively enhanced over time, and today it includes sophisticated services and management tools that make it the preferred network technology for pooling storage within medium and large data centers. Scale, performance, and cost pose the greatest challenges to the future viability of Fibre Channel in the modern data center, particularly as server virtualization drives a rapidly increasing need to attach servers to shared storage.

While standards-based solutions for iSCSI and NAS are readily available and widely deployed across Ethernet networks today, Fibre Channel remains the dominant protocol in larger data centers. Technologies such as iSCSI and NAS are largely used within data centers in the small and medium business market. NAS is used largely for external access of pooled storage services by campus client PCs. Even where these other Ethernet-based technologies would make sense in a greenfield deployment for many organizations, the cost of change from an operational perspective is high and any change would lead to a long-term coexistence of multiple technologies. Fibre Channel over Ethernet was designed to allow such organizations to move to Ethernet-based storage while in theory minimizing the cost of change.

Enhancements to Fibre Channel for Converged Data Center Networks: FCoE

The proposed FCoE protocol has been developed by the T11 Technical Committee—a subgroup of the InterNational Committee for Information Technology Standards (INCITS)—as part of the Fiber Channel Backbone 5 (FC-BB-5) project. The standard was passed over to INCITS for public comment and final ratification in 2009, and it has since been formally ratified. In 2009, T11 started development work on Fibre Backbone 6 (FC-BB-6), which is intended to address a number of issues not covered in the first standard, and develop a number of new deployment scenarios.

FCoE was designed to allow organizations to move to Ethernet-based storage while, at least in theory, minimizing the cost of change. To the storage world, FCoE is, in many ways, just Fibre Channel with a new physical media type; many of the tools and services remain the same. To the Ethernet world, FCoE is just another upper level protocol riding over Ethernet.

While the development of FCoE as an industry standard will bring the deployment of unified data center infrastructures closer to reality, FCoE by itself is not enough to complete the necessary convergence. Many additional enhancements to Ethernet and changes to the way networking products are designed and deployed are required to make it a viable, useful, and pragmatic implementation. Many, though not all, of the additional enhancements are provided by the standards developed through the IEEE DCB committee. In theory, the combination of the DCB and FCoE standards allows for full network convergence. In reality, they only solve the problem for relatively small scale data centers. The challenges of applying these techniques to larger deployments involve the use of these protocols purely for server- and access-layer I/O convergence through the use of FCoE Transit Switches (DCB switches with FIP snooping) and FCoE-FC Gateways (using N_Port ID Virtualization to eliminate SAN scaling and heterogeneous support issues).

Juniper Networks® EX4500 Ethernet Switch and QFX3500 Switch (both 10GbE switches) support an FCoE Transit Switch mode. The QFX3500 also supports FCoE-FC Gateway mode. These products are industry firsts in many ways:

1. The EX4500 and QFX3500 are fully standards-based with rich implementations from both a DCB and FC-BB-5 perspective.
2. The EX4500 and QFX3500 are purpose-built FCoE Transit Switches.
3. The QFX3500 is a purpose built FCoE-FC Gateway which includes fungible combined Ethernet/FC ports.
4. The QFX3500 features a single Packet Forwarding Engine (PFE) design.
5. The EX4500 and QFX3500 both include feature rich L3 capabilities.
6. The QFX3500 supports low latency with cut-through switching.

Future Direction for FCoE

There are two key initiatives underway within FC-BB-6 which will prove critical to the adoption of FCoE for small and large businesses alike.

For smaller businesses, a new FCoE mode has been developed allowing for a fully functional FCoE deployment without the need for either the traditional FC services stack or FC Layer 3 forwarding. Instead, the FCoE end devices directly discover each other through a pure L2 Ethernet infrastructure. This can be as simple as a DCB-enabled Ethernet switch, with the addition of FC Initiation Protocol (FIP) snooping for security. This makes FCoE simpler than either iSCSI or NAS, since it no longer needs a complex Fibre Channel (or FCoE) switch, and because the FCoE endpoints have proper discovery mechanisms.

For larger customers, a set of approaches is being investigated to remove the practical Fibre Channel scaling restrictions that currently limit deployment sizes. As this work continues, it is hoped that the standards will evolve not only to solve some of these scaling limitations, but also to more fully address many of the other challenges that arise as a result of blending L2 switching, L3 Fibre Channel forwarding, and Fibre Channel services.

Juniper fully understands these challenges, which are similar to the challenges of blending L2 Ethernet, L3 IP forwarding, and higher level network services for routing. As part of its 3-2-1 data center architecture, Juniper has already demonstrated many of these approaches with the Juniper Networks EX Series Ethernet Switches, MX Series 3D Universal Edge Routers, SRX Series Services Gateways, and Juniper Networks Junos® Space.

Meeting the Scalability Requirements

Building “a single converged network fabric” involves connecting the full diversity of virtualized end nodes to one unified data center network. Unfortunately, today’s data center networks are incapable of accommodating this model by design. Existing equipment and practices continue to divide large network domains into smaller subnetworks in order to deal with the limitations of legacy Ethernet architectures. In fact, today’s Fibre Channel storage networks have much greater levels of pooling than possible in today’s Ethernet networks due to their high cross-sectional bandwidth and low latency.

This approach runs counter to the notion of a single large virtualized pool of resources—long considered the *Holy Grail* of data center productivity and efficiency. As a result, it is imperative for next-generation data center networking technologies and architectures to support such unprecedented inclusiveness and scale.

Next-generation data centers must be able to accommodate large numbers of servers without having to divide the network into many smaller subnetworks. The use of Spanning Tree Protocol (STP) should also be minimized whenever possible. In addition, next-generation data center architectures must decouple server addresses from their locations in data centers. And there should be no compromise between the secure isolation of services and scalable performance.

For additional examples of drivers toward such flat, high-performance designs, we need only look at modern software applications and programming models that pose additional challenges to data center network scale. Distributed applications, such as map-reduce and memory cache, change the profile of communication patterns within application logic. They significantly increase server-to-server communication, require no assumptions of network locality to keep software development simple and efficient, and thrive on predictable and uniformly low latencies.

Another driver of architectural improvements in the data center network is server virtualization. Server virtualization software uses IP subnets to define the network range for virtual machine (VM) live migration. Today, limitations in legacy Ethernet's architecture restrict IP subnets to just a few hundred servers, so there is a conflict between creating large, scalable networks for VM migration and managing large, scalable networks hampered by legacy Ethernet architectures. This conflict must be resolved in order to reap the full benefits of a large-scale virtualized server environment that runs on a converged Ethernet fabric, thus making any-server/any-service computing a reality.

As indicated earlier, there are many challenges in building the next generation of data centers and the networks to support them. Using current approaches, there are many conflicting requirements, and only by taking a different approach can these all be addressed concurrently. The next-generation data center network must:

1. Provide high cross-sectional bandwidth and low latency—optimization and flexibility at scale is not possible without this.
2. Move higher level network services out of the primary data path—otherwise, a single point of failure is introduced that prevents complete resource optimization.
3. Provide large L2 domains, but also provide routing between these domains and between domains and the outside world.

A Clear Return on Investment

As legacy servers get phased out, or as new racks are deployed in new or existing facilities, data center architects have the opportunity to thoroughly evaluate their options and consider the benefits of upgrading to a converged network fabric.

Being mindful of the potential rewards and risks, customers who migrate to a converged network fabric in the data center will likely do so in a number of phases.

Prerequisite: Physical Infrastructure

By understanding the nature of 10GbE as well as 40GbE and 100GbE in terms of connectivity options in the copper and optical domains, it is possible for organizations to immediately deploy future-proofed structured cabling within the data center that not only meets the needs of current 1GbE edge deployments, but also provides a nondisruptive migration to the next generation of 10GbE at the edge data centers with very high cross-sectional bandwidth. Juniper has worked closely with many of its customers and partners in this area.

Phase 1: Ethernet I/O Consolidation Within a Rack

Converged Network Adapters (CNAs) (which are I/O cards supporting DCB in the server) along with the corresponding top-of-rack switch consolidate multiple traffic types from the server onto a single wire while maintaining separation. They do the same for service-level agreements that existed when each of these was running on its own NIC and on its own L2 domain, switch, or network. Consolidating multiple 1GbE connections from server to network to a smaller number of 10GbE DCB connections provides many benefits to the customer. In addition, it allows network teams to work through many of the challenges involved before converging the traditionally non-Ethernet traffic types onto the new network. By implementing I/O convergence within the Ethernet domain in parallel with simplifying the network infrastructure—reducing the number of tiers and pooling other network services—organizations are preparing themselves for the next phase while gaining immediate OpEx and CapEx benefits.

For those organizations focused on Ethernet- or IP-based storage such as iSCSI or NAS, this phase also provides many benefits to better scaling and performance. Juniper has produced a wealth of material in terms of white papers and best practices guides that have already helped customers move along this path. Additional materials are now available that discuss the implications of DCB and fabric technology.




Figure 3: Ethernet I/O convergence and network simplification

Phase 2: Consolidation of Infiniband and Fibre Channel to an Ethernet-Based Infrastructure

Deployment of DCB technology at the server access layer provides the ability to also converge HPC using iWARP and RoCEE for cluster traffic, and to converge the first-hop Fibre Channel-to-Ethernet through FCoE before separating this traffic back out to an existing SAN backbone. Based upon the final full FC-BB-5 standard, there are actually two key deployment options for server access layer convergence:

Option 1—FCoE Transit Switch. In this model, the SAN team FCoE enables its backbone SAN fabric, while the network team deploys a top-of-rack DCB switch with FIP snooping. Servers are deployed with CNAs, and blade servers are deployed either with pass-through modules or embedded DCB switches. These are connected to the top-of-rack switch, which then has Ethernet connectivity to the LAN aggregation layer and Ethernet connectivity to the FCoE ports of the SAN backbone.

A rich implementation of an FCoE Transit Switch provides strong management and monitoring of traffic separation, allowing the SAN team to monitor FCoE traffic. FIP snooping provides perimeter protection, ensuring that the presence of an Ethernet layer in no way impacts existing SAN security. The SAN backbone can be simply FCoE-enabled with either FCoE blades within chassis-based systems or FCoE-FC Gateways connected to the edge of the SAN backbone, administered by the SAN team, and shared across the lower cost DCB-enabled edge. This architecture has several advantages:

- Low cost top-of-rack DCB switch
- Rich monitoring of FCoE traffic at top of rack (QFX3500 switch)
- FCoE enablement of the SAN backbone (FCoE blades or FCoE-FC Gateway) managed by SAN team
- Load balancing carried out between CNAs and FCoE ports of SAN fabric
- No heterogeneous support issues as top of rack is L2 connectivity only

Option 2—FCoE-FC Gateway. In this model, the SAN and Ethernet teams agree jointly to deploy an FCoE-FC Gateway at top of rack. From a cabling perspective, the deployment is identical to Option 1, with the most visible difference being that the cable between the top-of-rack FCoE-FC Gateway and the SAN backbone is now carrying native Fibre Channel traffic rather than being Ethernet carrying FCoE traffic. The FCoE-FC Gateway uses NPIV technology so that it presents to the servers as an FCoE-enabled Fibre Channel switch, and presents to the SAN backbone as a group of Fibre Channel servers. Advantages of this approach include:

- Clean separation of management through role-based access control (QFX3500 switch)
- No need for FCoE enablement of SAN backbone
- Fine-grained FCoE session-based load balancing (at the VM level for NPIV-enabled hypervisors) (QFX3500 switch)
- No heterogeneous support issues as an FCoE-FC Gateway presents to SAN fabric as a Fibre Channel-enabled server
- Available post deployment as license upgrade and fungible port reconfiguration with no additional hardware (QFX3500 switch)
- Support for upstream DCB switch such as an embedded switch in blade server shelf (QFX3500 switch)

Option 3—FCoE-FC FCF. For any deployment of any size, there is no value to local switching, as any rack is either pure server or pure storage. In addition, although the SAN standards limit any deployment to 239 switches, the supported limits are typically in the 16 to 32 switch range (in reality, most deployments are kept well below these limits). As such, this option has limited value in production data centers.




Figure 4: Server access layer convergence options

Cautionary Notes

Avoid pre-standards or partial standards products. Although in 2011, there are some products (CNAs, DCB switches, FCoE-FC Gateways, and FCoE-based Fibre Channel switches) that are fully compliant with the respective standards, there are still many products shipping that are pre-standard, are only partial implementations, or, if compliant, have very limited capabilities. Even if these products meet initial deployment requirements, they severely limit flexibility in terms of future capabilities and are therefore a poor investment choice.

Investment protection. Many of today's products were developed before work on FC-BB-5 was finalized, let alone before the FC-BB-6 project was started. This limits the customer use cases and inhibits many of these products from providing customers with the investment protection they want from FCoE-capable products. Any decision should be made with a clear understanding of future direction.

Phase 3: Fabric Convergence Throughout the Entire Data Center

Meaningful productivity and efficiency can only be achieved when data center optimization and network convergence occurs throughout the entire data center, not just within a single rack. This is the final phase in the adoption of a converged data center network fabric.




Figure 5: The convergence of LAN, SAN, and HPC clusters represents the third and final phase in the adoption of a unifying next-generation network fabric in the data center.

Conclusion

Converged data center networks will require a robust and complete implementation of FCoE and DCB standards to viably support the critical application and data integrity requirements of data center applications. Given the timing of the FCoE and DCB standard ratification process, and the apparent evolution of those standards over the next few years combined with the incremental progress in their cost effectiveness, the ability to implement a converged data center will likely occur in phases. In phases one and two, convergence within the rack will enable many gains while supporting separate LAN and SAN infrastructures through the use of FCoE-FC Gateways between the two. Subsequent phases will enable networks to be fully converged and unified through a combination of standards development and the development of new classes of products.

By maintaining active participation in the related standardization efforts and rethinking the technology and economics of the data center network from the ground up, Juniper Networks provides customers with a strategy to develop a single, converged data center fabric that delivers flexibility and performance, while at the same time continuing to drive down cost and complexity.

About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at www.juniper.net.

Corporate and Sales Headquarters

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters

Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters

Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

Copyright 2012 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Junos, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

2000315-005-EN Jan 2012

 Printed on recycled paper