BLADE NETWORK TECHNOLOGIES

10 Gigabit Ethernet Technology Overview and Applications for Enterprise Data Centers

Technology Brief

Contents

Introduction

As bandwidth-intensive applications such as Web 2.0, Virtualization, High-Performance Computing (HPC) and Network Attached Storage (NAS) continue to proliferate within the Enterprise Data Centers, 10 Gigabit Ethernet (10 GbE) provides a cost-effective way to increase throughput and seamlessly deliver customer service level agreements. Indeed, with prices falling below $500 per port, 10 GbE switches are fast becoming the technology of choice for the access, distribution and core layers of the data center. Where "point" technologies such as Myrinet for HPC and Fibre Channel for storage once dominated, 10 GbE is steadily displacing or "unifying" these technologies—providing the necessary performance, proven flexibility, and lower-cost advantage. Nonetheless, today's 10 GbE switches and routers must provide more than just speeds and feeds. Energy efficiency and cooling are critical factors in optimizing data-center economics—as companies begrudgingly balance performance with power consumption to meet energy budgets.

This Technology Brief is divided into two parts. The first part discusses the evolution of Ethernet, and then focuses on 10 Gigabit Ethernet and related transceiver and cabling technologies. The second part examines BLADE's RackSwitch product line, including 10 GbE applications within enterprise networks and data centers.

Part I: A Brief History of Ethernet

Although Ethernet has been around for over three decades, the technology has continually evolved keeping pace with just about every application trend and technology breakthrough on the Internet. The IEEE jumped on the Ethernet bandwagon in 1983 when it introduced 10 Mbps Ethernet. The IEEE's involvement set in place an official open-standard for the networking industry, and drove interoperability within the networking industry. 10 Mbps Ethernet evolved to Fast Ethernet (100 Mbps), then Gigabit Ethernet (1000 Mbps), and most recently 10 Gigabit Ethernet (10,000 Mbps)—the focus of this brief. Despite the evolution of Ethernet over the past three decades, the basic Ethernet frame format and principles of operation have remained virtually unchanged. As a result, networks of mixed speeds (10/100/1000 Mbps) operate uniformly without the need for packet fragmentation, packet reassembly or address translation. The advancements in Ethernet have focused mainly on speed and related cabling and transceiver technologies.

10 Gigabit Ethernet Overview

10 GbE is the latest speed phenomena in Ethernet's ongoing evolution. Ratified in 2002 as IEEE 802.3ae, 10 GbE supports 10 Gigabits per second transmission over distances up to 80 km. In almost every respect, 10 GbE is fully compatible with previous versions of Ethernet. It uses the same frame format, Media Access Control (MAC) protocol, and frame size. Thus, network managers can retain existing management tools and operational procedures as they introduce 10 GbE into their network. The differences between 10 GbE and previous versions of Ethernet are inconsequential as it pertains to compatibility. Firstly, unlike its predecessors, 10 GbE only operates in full-duplex mode. This attribute is a benefit rather than disadvantage since full-duplex operation provides faster, lower-latency responses to network transmissions. Secondly, the IEEE 802.3ae standard defines two different types of PHYs—LAN PHY and WAN PHY. The LAN PHY transmits Ethernet frames directly over a 10 GbE serial interface. LAN PHY is for standard 10 Gigabit per second transmission within an enterprise or local area-network (LAN) environment. The WAN PHY encapsulates Ethernet packets in SONET, OC-192c frames and provides a rate matching mechanism of 9.953 Gigabits per second. This SONET-compatible interface allows native 10 Gigabit Ethernet to be transported directly over an existing service provider SONET/SDH network. Table 1 provides a comparison of the key attributes between 1 GbE and 10 GbE.

Table 1: Comparison between 10-Gigabit Ethernet and 1-Gigabit Ethernet

1 Gigabit Ethernet (GbE)	10 Gigabit Ethernet (10GbE)
CSMA/CD and Full Duplex	Full Duplex only
Leveraged Fiber Channel Physical Medium Dependent (PMDs)	New Optical PMDs
8B/10B Coding	New 64B/66B Coding
Supports LAN Applications to 70km.	Supports LAN and WAN Applications up to 40km.

10 Gigabit Ethernet Cabling Technologies

10 GbE, like previous versions of Ethernet, supports both copper and fiber cabling. However, because of its higher bandwidth 10 GbE poses a number of implementation issues and added complexity especially where copper cabling is involved. For example, although 1 GbE operates seamlessly over Category 5 (Cat 5) cabling, this grade of cabling has difficulty meeting the bandwidth and crosstalk demands of 10G transmission rates, and is therefore not recommended for use with 10 GbE. The next two sections discuss copper and fiber cabling technologies for 10 GbE respectively.

Copper Cabling

10GBASE-CX4

The IEEE 802.3ak standard was approved in February 2004, and provides a means for 10GbE to operate over copper cabling. Also known as 10GBASE-CX4, the standard uses 8-pair, twin-axial, shielded copper cables—similar to Infiniband-style cabling. CX4 differs markedly from the more common RJ45 copper cabling. CX4 cabling is rigid, thick and rather cumbersome to deploy. The longer the length used, the thicker the cable. CX4 cables start at 30 American Wire Gauge (AWG) for short lengths to 24 AWG thickness for a full 15 meters. Whereas 1 GbE over RJ45 cabling can easily extend to 100 meters, CX4 is limited to 15 meters for 10 GbE.

10GBASE-T

Ratified in June 2006, 10GBASE-T (IEEE 802.3an) defines 10 Gigabit per second data transmission over unshielded twisted-pair (UTP) cabling, at distances up to 100 meters. Running 10 GbE over RJ45 cabling will enable higher density switching solutions, lower-cost 10 GbE NICs, and broader market acceptance. However, in legacy enterprise networks where Category 5 (Cat 5) cabling still dominates, an upgrade to Cat 5E or better is necessary for the proper operation of 10GBASE-T. This is because alien crosstalk, insertion loss, and echo to name a few, are more detrimental on signal quality for lower-grade cable at higher bandwidth. These inherent signal quality problems affecting 10 GbE operation over UTP cabling, can be mitigated and reduced through advanced cabling design, professional installation of cabling plants and use of only certified cabling and patch panels. From a cabling design standpoint, the IEEE has defined four link segment models, which vary in distance based on cable grade or quality.

Table 2: Cabling Types and Supported Distances

Cabling	Distance (meters)	References
Class E/Category 6	55 to 100	ISO/IEC TR-24750 TIA/EIA TSB-155
Class E/Category 6 (Unscreened)	55	ISO/IEC TR-24750 TIA/EIA TSB-155
Class E/Category 6 (Screened)	100	ISO/IEC TR-24750 TIA/EIA TSB-155
Class F/Category 7	100	ISO/IEC TR-24750
New Class E(A)/Augmented Category 6	100	ISO/IEC 11801 Ed. 2.1, TIA/EIA 568-B.2-10
http://grouper.ieee.org/groups/802/3/ae/public/index.html		

Category 6A and Class F/Category 7 cabling systems will reliably support 10Gbps data-rates on links up to 100 meters. Category 6 UTP cabling systems are capable of operating at 10Gbps, but limited to 55 meters. For Category 6 UTP and Category 6 FTP cabling systems, field testing and certification is generally recommended to ensure proper performance for these cabling systems at 10Gbps speeds. 10GBASE-T installations will follow traditional RJ45 copper deployments such as server farm aggregation in the data center and horizontal building cabling for the wiring closet.

Figure 1 10 Gigabit Ethernet Copper Cabling

(Source: Hubbell, http://www.hubbell-premise.com)

Fiber Cabling

The IEEE has defined a total of 7 optical layers for 10GbE as outlined in Table 3.

Table 3: IEEE 10 Gigabit Ethernet Optical Specifications

Device	Range	Optics	Fiber
10GBASE-LX4	300m MMF/10km SMF	1310nm WWDM	MM or SM
10GBASE-SR	300m	850nm	MM
10GBASE-LR	10km	1310nm	SM
10GBASE-ER	40km	1550nm	SM
10GBASE-SW	300m	850nm	MM
10GBASE-LW	10km	1310nm	SM
10GBASE-EW	40km	1550nm	SM
http://www.10gea.org/optical-fiber-10ge.htm			

10GBASE-LX4

10GBase-LX4, as defined by IEEE 802.3ae, uses wave division multiplexing technology to send signals over four wavelengths of light carried over a single pair of fiber optic cables. The 10GBASE-LX4 system is designed to operate at 1310 nm over multi-mode or single-mode dark fiber. The design goal for this media system is from 2 to 300m over multimode fiber or from 2 to 10km over single-mode fiber—with even longer distances possible depending on cable type and quality. LX4 is more expensive than both SR and LR because it requires four times the optical and electrical circuitry in addition to optical multiplexers. The number of components required to implement the LX4 limits its ability to fit into smaller form factors such SFP+. Today, LX4 is supported over the XENPAK and X2 transceiver form factors.

10GBASE-(S/L/E)R

These optical specifications are designed for use in enterprise networks and inject a single 10Gbps signal through the optical laser. 10GBASE-SR uses the lowest cost optics—850nm—to support 10GbE transmission over standard multimode fiber (“MMF’), for distances up to 300m. The “SR” or short-reach specification is the cheapest optics of all the defined 10GbE optical types. 10GBASE-LR uses slightly higher cost optics—1310nm—than SR and requires a more complex alignment of the optical wavelength to support single-mode (“SMF”) up to 10km. “LR” is commonly referred to as “long reach” optics. 10GBASE-ER, commonly referred to as extended-reach optics, is the most expensive optics at 1550nm and supports distances up to 30 or 40km over SMF.

10GBASE-(S/L/E)W

The 10GBASE-W specifications are designed for use in wide-area networks (“WANs”). These standards operate at the same baud rate as OC-192/STM-64 SONET/SDH equipment. Similar to SR, LR, and ER standards, three lasers are supported for different distances.

10 Gigabit Ethernet Transceivers

So far, the previous section has discussed the various fiber and cabling technologies developed for 10 GbE. Network managers must also consider the myriad of interface options that connect the networking equipment to the cabling plant. Transceivers provide the interface between the switching equipment and the transmission cable. Transceivers are not specified by the IEEE 802.3 standard but by multi-source agreements (MSAs) among switch vendors and chip manufacturers. The MSA is an ongoing effort, driven by the need to meet demands for higher-port densities, lower-power consumption, smaller form-factors, higher performance, and lower price points.

Network architects designing a new 10 GbE network or managing an existing one, will encounter some if not all of the following transceiver types.

300-PIN

The first 10 Gigabit Ethernet transponders conformed to the 200- or 300-pin MSAs. These modules were primarily designed to support SONET/SDH applications by converting between a 10Gbps serial optical signal and 16 parallel 622Mbit/s electrical signals.

XENPAK

XENPAK provided a significant step forward from the 300-pin MSA form-factor. These hot-pluggable modules although smaller in size than the 300-pin MSA, are still large and bulky, but sufficed for the first-generation of low-density 10GbE switches. XENPAK supports the IEEE 802.3ak copper standard for CX4, as well as the LX4 standard. XENPAK is based on a parallel interface (XAUI).

XPAK and X2

Demands for lower power consumption and higher port density have spurred innovation for significantly smaller, hot-pluggable modules. The first of these are the XPAK and X2 MSA form-factors. These are 40 percent smaller than XENPAK, but still support the same XAUI interface. XPAK is used primarily in Network Adapter Cards (“NIC”) and Host Bus Adapter (“HBA”) markets for use in servers and storage devices.

XFP

The XFP MSA is a third-generation of 10GbE transceivers. Unlike XENPAK, X2 and XPAK, XFP has a 10Gbit/s serial electrical interface (XFI) that converts serial 9.95-10.7 Gbps electrical signals into external serial 9.95-10.7 Gbps optical or electrical signals. This eliminates Mux/Demux serial-to-parallel conversion logic chips inside the module and allows the serial 10 Gbps physical-layer IC (PHY) to be moved on to the PCB (away from optics-generated heat) and everything up to the XFI serial interface to be integrated into the CMOS media-access controller chip. Through its XFI interface, XFP modules can support 10GbE, Fibre Channel, and SONET/SDH.

SFP+

SFP+ is the next-generation transceiver module form-factor, specified by the ANSI T11 Group for 8.5- and 10-Gbps Fibre Channel, as well as 10 Gigabit Ethernet. The SFP+ module is even smaller (by 30%) than XFP, and consumes less power. It also houses fewer electronics within its smaller footprint, thereby reducing the cost per port for a 10GbE SFP+ module. SFP+ has been a boon to the 10GbE industry, allowing switch vendors to pack more ports into smaller switch form-factors, and lowering system costs through better integration of IC functions at the host card level.

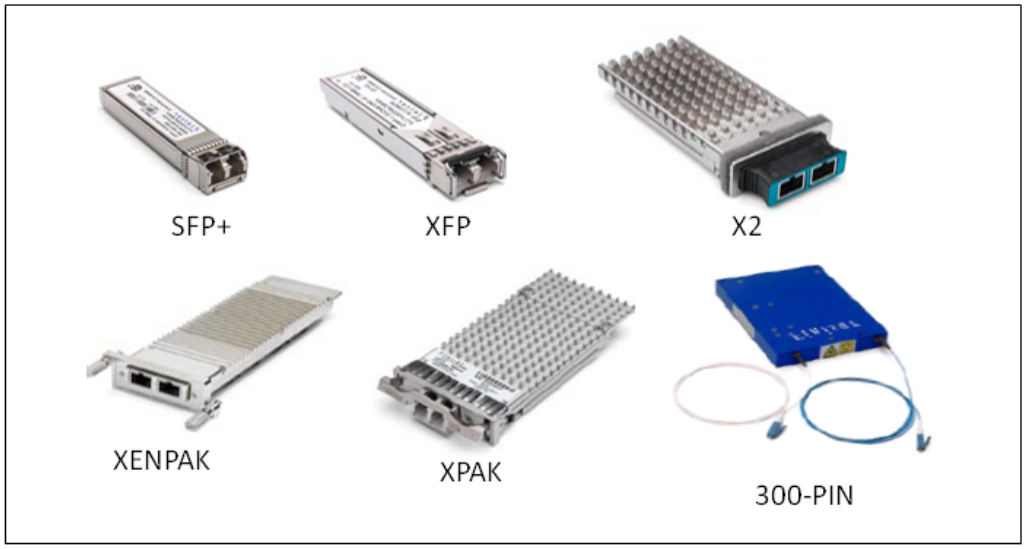

Figure 2 10 Gigabit Ethernet Transceivers (Not drawn to scale) (Source: www.finisar.com)

Part II: A Closer Look at BLADE's RackSwitch Product Family

Overview

BLADE's RackSwitch product line consists of line-rate, low latency and energy efficient 1-Gigabit Ethernet and 10-Gigabit Ethernet switches for data center and high-performance computing applications. Designed as a result of BLADE's years of experience successfully building embedded blade switches for the blade server market, the RackSwitch product line offers tight integration and elements in common with the servers it supports. Rear-to-front or front-to-rear cooling is consistent with server designs, and facilitates data center energy efficiency schemes such as hot- and cold-aisle layouts. Switch ports can be front or rear-facing to allow closer proximity to server ports and thus, keep cable lengths short and performance optimized. From a centralized management platform, network managers can point and click to configure, monitor and manage both BLADE's embedded switches and RackSwitches in a single rack enclosure or across a cluster of racks and servers.

BLADE's RackSwitch G8000 48-Port GbE Aggregation Switch

The RackSwitch G8000 is a 1-RU, fixed-configuration switch, with 48-ports of 10/100/1000 Mbps Ethernet, plus four optional 10 GbE uplink ports. Four of the 48 1GbE ports can be figured with SFP optics to support fiber connectivity.

Figure 3 BLADE 48-port 1-GbE RackSwitch G8000—for Top-of-Rack Applications

The switch is designed with a non-blocking 176 Gbps switch fabric and an average port-to-port latency of **5.1 microseconds**[i].

Two expansion or modular slots are available to each support a 2-port 10 GbE module. The 10 GbE module comes in two flavors—2-port CX4 or 2-port SFP+. The purpose of the 10 GbE module is to enable stacking and or uplink connectivity to another switch. The G8000 is powered by 1+1 redundant AC power supplies, and consumes only 124W of power. The switch has a mini-USB console port and an out-of-band 10/100/1000 Mbps management port.

BLADE's RackSwitch G8100: 24-Port 10 GbE Low Latency Switch

The RackSwitch G8100 is a 1-RU, fixed-configuration switch, with 24-port 10 GbE. It has 20 CX4 ports and four SFP+ ports. The CX4 ports are recessed for simplified cable management and better bend radius.

Figure 4 BLADE RackSwitch G8100—for High-Performance, Low-Latency Cluster Computing

The G8100 is ideal for high-performance cluster computing applications that require extremely low latency and high bandwidth. A 480 Gbps switch fabric delivers non-blocking, line-rate performance on all 24 ports, as well as an average port-to-port latency in of **360 nanoseconds**.[ii] Like the G8000, the G8100 is designed for low power consumption, consuming only **120W** of power. It is powered by 1+1 redundant AC power supplies. The switch has a DB9 console port and an out-of-band 10/100/1000 Mbps management port.

BLADE's RackSwitch G8124: 24 Port 10GbE SFP+ Low Latency Switch

The RackSwitch G8124 is a 1-RU fixed-configuration switch, with 24-port 10 GbE SFP+ ports. A 480 Gbps switch fabric delivers non-blocking, line-rate performance on all 24 ports, as well as an average port-to-port latency in of 680 nanoseconds.[iii]

Figure 5 BLADE RackSwitch G8124—for High-Performance, Low-Latency Cluster Computing with cost-saving SFP+

The G8124 is ideal for datacenters requiring the flexibility of supporting either 1GbE or 10GbE through its SFP+ interface. Like the G8100, the G8124 is low power, consuming between 115W to 168W depending on the speed of the port (1G/10G), type of transceivers (SR or DAC) and number of active ports. Power supply and console port are the same as the G8100.

BLADE Transceiver and Cabling Options

As noted earlier, 10 GbE supports a host of transceiver and cable options. The RackSwitch product family supports SFP transceivers for 1GbE, and SFP+ pluggable optics for 1GbE or 10 GbE interfaces. The RackSwitch G8000 48-port GbE switch supports standard RJ45 interfaces for the 10/100/1000 Mbps connections, as well as CX4 or SFP+ 10 GbE uplinks. The G8100 can be either CX4 based or SFP+ for node connections.

In order to overcome the 15 meter distance limitation of CX4, the G8100 supports Active CX4 cables, extending the distance to 100 meters. The active CX4 cable snaps directly into the switches' CX4 port which then supplies power to the cable in order to drive the 10G signal over 100 meters.

For in-rack or adjacent rack cabling deployments, the RackSwitch family supports SFP+ direct attached Twinaxial copper cabling. The copper cable-plus-SFP+ connector combo provides a low-cost, low-latency and energy-efficient solution, well-suited for in-rack connections of less than 10 meters between server and top-of-rack switches. As far as optical transceivers options, the G8000 and G8100 support the SR and LR optical specifications.

BLADE is closely monitoring **10GBaseT**—given its potential to radically reduce the cost per port of 10 GbE and enable widespread deployment—and will implement when the power, latency and cost is in line with customer expectations. Many of the same benefits of twisted pair are available with Active CX-4 and/or SFP+ twinax cables.

Table 4: Supported Transceivers and Cables for RackSwitch

Model	1GbE Transceiver	10GbE Transceiver	10GbE Cable	Serial Cable
G8000	BLADE SFP 1000Base-T Copper Transceiver (p/n BN-CKM-S-T) BLADE SFP 1000Base-SX Short Range Fiber Transceiver (p/n BN-CKM-S-SX) BLADE SFP 1000Base-LX Long Range Fiber Transceiver (p/n BN-CKM-S-LX)	SFP+ 10GBase-SR Short Range Transceiver (p/n BN-CKM-SP-SR)	Passive or Active 10GbE CX4 cables*	BLADE DB9 to mini-USB cable (p/n BN-SB-SRL-CBL)
G8100	N/A	BLADE 10GBase-SR SFP+ (p/n) BN-CKM-SP-SR	Passive or Active 10GbE CX4 cables*	Standard DB9 Null-modem cable
G8124	BLADE SFP 1000Base-T Copper Transceiver (p/n BN-CKM-S-T) BLADE SFP 1000Base-SX Short Range Fiber Transceiver (p/n BN-CKM-S-SX)	BLADE 10GBase-SR SFP+ (p/n) BN-CKM-SP-SR	SFP+ copper direct attach cables (DAC), 1 to 10 meters: BN-SP-CBL-1M BN-SP-CBL-3M BN-SP-CBL-7M BN-SP-CBL-10M	BLADE DB9 to mini-USB cable (p/n BN-SB-SRL-CBL)
***See "Recommended Transceivers and Cables for 10G Ethernet Switches" available at http://www.bladenetwork.net/media/PDFs/WP_Gen_Rec_10GbTrans+Cables.pdf for details on passive and active CX4 cables supported by BLADE.**				

Network Management

Network management tools are necessary to efficiently monitor, configure and maintain large switch installations with consistently applied firmware upgrades and configuration changes. BLADE's BLADEHarmony™ Manager is a tool that centralizes and simplifies management of its blade server switches and rack switches through a web interface. BLADEHarmony Manager can provide a quick, at-a-glance views and reports on all discovered RackSwitches, and embedded blade switches on the network, and allow an administrator to perform configuration changes, upgrades and other scheduled maintenance tasks on individual switches or groups of switches.

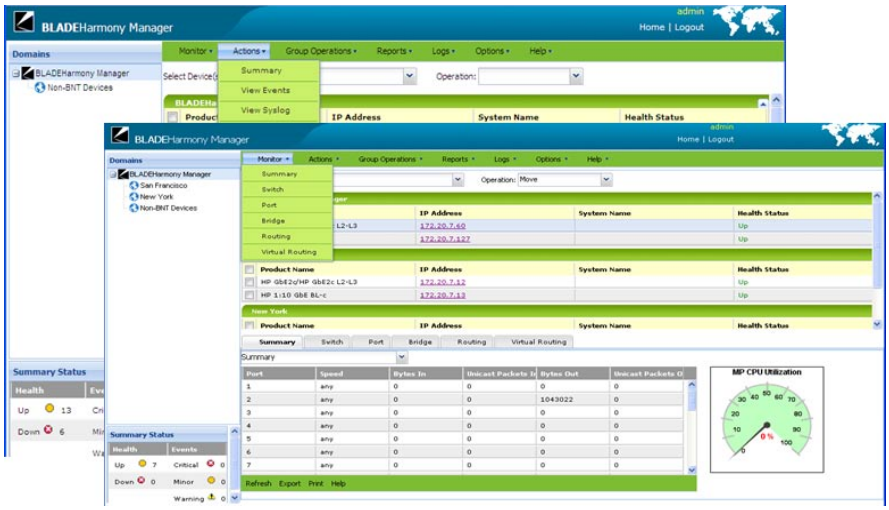

Figure 6 BLADE Harmony Manager, GbE/10GbE "Interconnect" Management & Provisioning

Interoperability

BLADE's G8000 G8100 and G8124 top-of-rack switches offer standards based IEEE and IETF protocols for interoperability with existing networks. In addition, the BLADE switches support an industry-standard command line interface (CLI), allowing for immediate configuration "out-the-box," and retention of existing network management tools, scripts, and operational procedures.

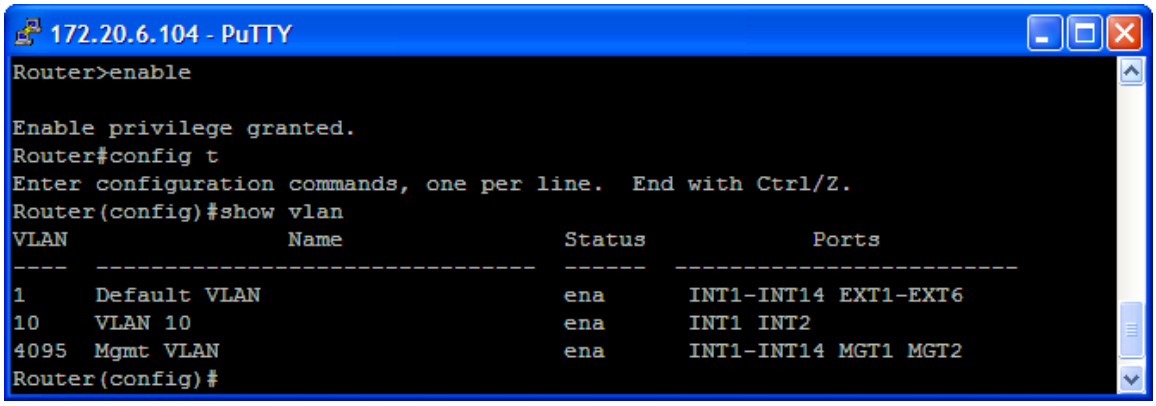

Figure 7 BLADE RackSwitch—Industry Standard CLI

RackSwitch Deployment Scenarios

The BLADE RackSwitch product family of 1 GbE and 10 GbE switches support a variety of deployment scenarios to include top-of-rack switch collocated with servers within a data center rack, and standalone campus switch for access and distribution.

Top-of-Rack Access Switch

The RackSwitches can be deployed as top-of-rack access switches, providing the necessary densities to aggregate large numbers of servers within a rack (see Figure 8). Multiple switches can be trunked or stacked to increase the switch ports within a given rack. In cluster configurations, besides server connectivity, there may be other I/O needs such as storage and management which increase the required switch port count per rack. Multiple switches may also be necessary for redundant server connections and to support high availability protocols such as NIC Teaming.

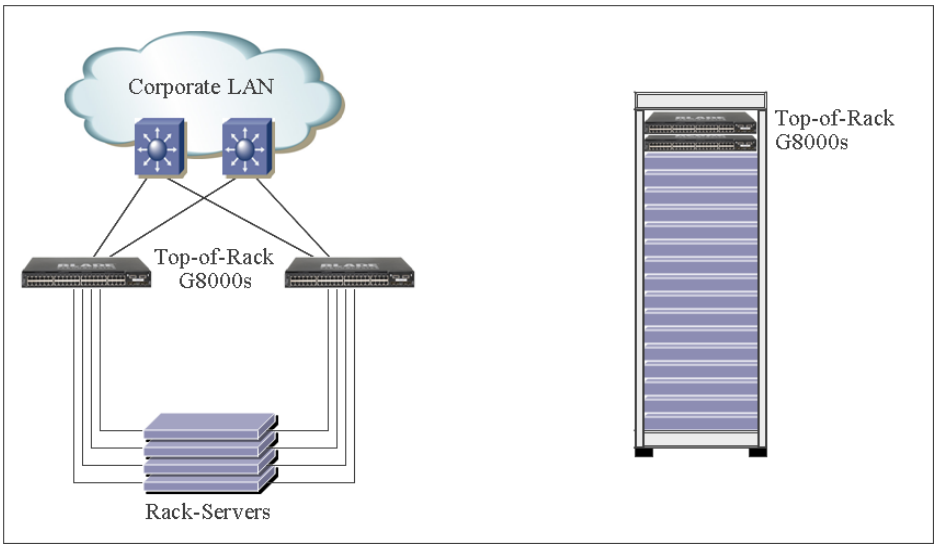

Figure 8 BLADE Top-of-Rack G8000—Logical Configuration & Physical Layout

Top-of-Rack Aggregator

When clusters of servers span multiple racks, the top-of-rack switches must be uplinked to an aggregation switch. The RackSwitch G8100 low latency10 GbE switch can serve as an aggregator for multiple top-of-rack G8000 switches (see Figure 9). Each RackSwitch G8000 can support up to four 10 GbE uplink ports, which would connect to the RackSwitch G8100.

Campus Aggregation

The G8100 can function as a low-cost, Layer 2, 10 GbE aggregation switch within the distribution layer of the enterprise network. The layer 3 functionality can be provided by an existing router, 10 GbE attached to the G8100. This is an unobtrusive way to increase bandwidth and performance in oversubscribed areas of the network.

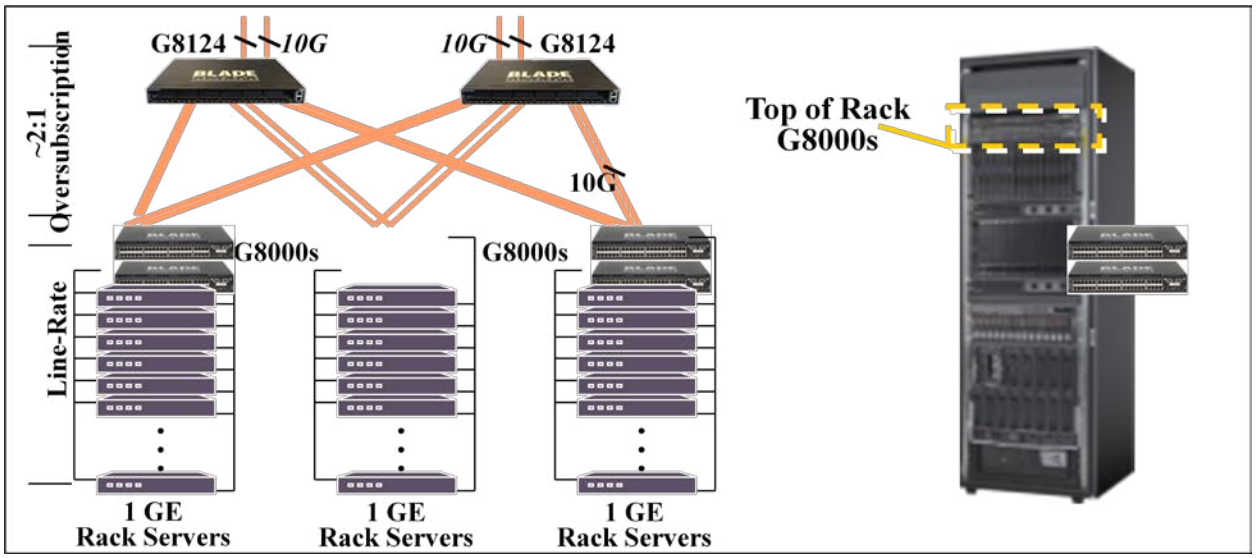

Figure 9: 10 Gigabit Ethernet Aggregation

10 Gigabit Ethernet Applications

High Performance Computing (HPC)

High performance computing involves the use of supercomputers or clusters of powerful processors to solve computationally-intensive problems. In order for a large number of processors to work together, clusters require interconnects that support high-bandwidth and very low latency communication. BLADE's RackSwitch G8100 delivers all the known benefits of Ethernet including ultra-low latency in the sub-300 nanosecond range and 10 Gigabit per second data transmission. The RackSwitch G8100 is a lower-cost alternative to expensive proprietary interconnects such as Infiniband and Myrinet. See Figure 10 for an example of how RackSwitch G8100 or G8124 can be deployed as the interconnects in an HPC cluster.

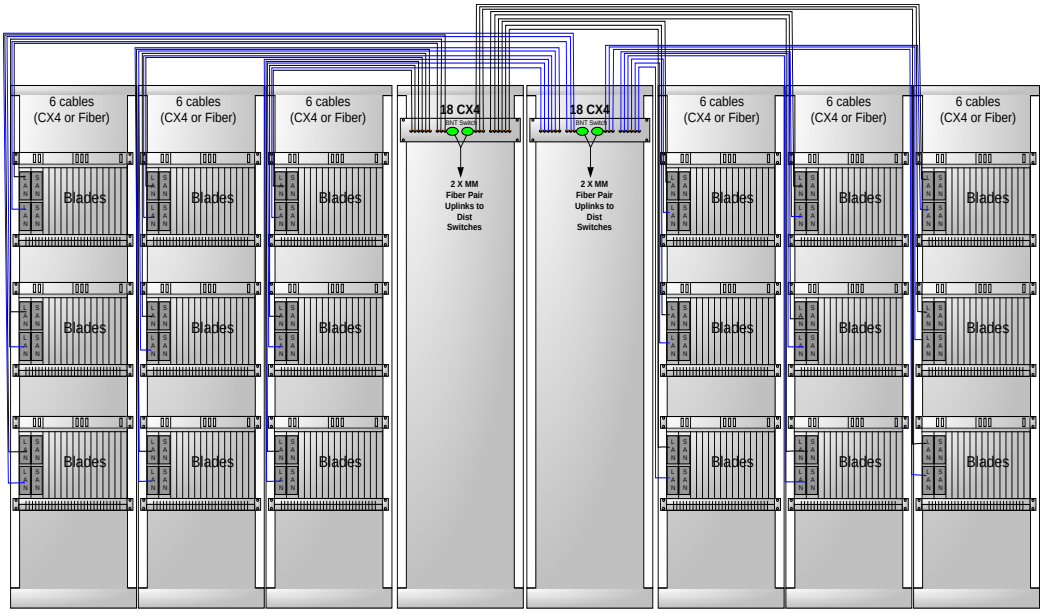

Figure 10 Full 10GbE High Performance Computing Cluster using RackSwitch G8100 or G8124

IP Storage Networking

As enterprise data centers continue to deploy more powerful processors and clustering for compute capacity, the need for terabits of storage capacity is becoming standard. IP storage solutions such as NAS, iSCSI, and FCIP, can all benefit from the high-bandwidth, low-latency, and scalability of 10 GbE. BLADE's G8100 provides a practical and cost-effective 10 GbE solution for storage networking as follows:

- Consolidation of native SANs over a common high-speed, low-latency 10 GbE backbone. IP networks provide improved manageability, interoperability and cost-effectiveness. Consolidation over the IP and 10 GbE infrastructure achieves increased capacity, higher availability and better pooling of storage resources using virtualization. Fibre Channel SANs can be extended and accessed remotely. Enterprise servers that were previously not a part of the storage network can be iSCSI-enabled and consolidated into an enterprise-wide IP storage strategy.
- Elimination of congestion and I/O bottlenecks. For example, high-performance NAS filers using clustered, parallel file systems require 10GbE I/O throughput for improved performance. Similarly, 10GbE-based iSCSI is essential for improved application response time and throughput as iSCSI commands are transported across metro and wide-area networks.
- Disaster recovery and remote site backup operations rely on the optimal performance of the intervening IP and Ethernet network. The mere existence of a network is not sufficient for real-time storage applications or mirroring and replication of critical data at remote sites. High-throughput, high-availability and low-latency are crucial requirements of the IP storage network, which can readily be provided by the BLADE G8100.

10 Gigabit Ethernet Campus Networking

The past five years have seen a steady rise in the number of 1 GbE connections in corporate networks—both in the user network and data center. This widespread deployment of 1-GbE has driven the need for higher-capacity uplinks in wiring closets and bigger pipes in the aggregation and core layers of the network. While many networks have implemented link aggregation (grouping of multiple 1-GbE links into a larger logical pipe) to increase bandwidth, 10 GbE provides not only a 10x increase in bandwidth, but also simplicity of design and better scalability than link aggregation.

Conclusion

In summary, 10 Gigabit Ethernet empowers companies to expand application capabilities, reduce time to solve complex financial and scientific applications, and quickly respond to changing customer needs and market conditions. High bandwidth, low-latency, and energy efficiency are the key differentiators of BLADE Network Technologies'10 Gigabit Ethernet switch offerings. BLADE's RackSwitch product family is designed for data centers—implementing high-performance compute clusters; and enterprise networks—consolidating disparate SAN solutions into a common, low-cost, and high-speed 10 Gigabit Ethernet, IP storage network.

About BLADE Network Technologies

BLADE Network Technologies is the leader for 1 and 10GbE solutions for virtualized, cloud-ready data center networks. BLADE's "virtual, cooler and easier" RackSwitch family demonstrates "Rackonomics," a low cost, low power and efficient way for scaling out data center networks. BLADE's customers include over 300 of the Fortune 500.

MKT090731

---

[i] Tolly Group., 29 October, 2008. "BLADE Network Technologies, Inc. RackSwitch G8000 Competitive Performance Evaluation versus Cisco Catalyst 4948-10GE Switch." Available from http://tolly.com/DocDetail.aspx?DocNumber=208348

[ii] Tolly Group #209115, May 2009. "BLADE Network Technologies, Inc. RackSwitch G8100 & G8124: Competitive Performance Evaluation versus Cisco Catalyst 4900M Switch." Available from http://tolly.com/DocDetail.aspx?DocNumber=209115

[iii] ibid