TechBrief
Extreme Networks Migration Guide

Introduction

Mission-critical web applications – enterprise resource planning, e-commerce, voice-over-IP, server co-location and private Intranets – are redefining how we do business. Information is available from multiple sources, network traffic is unpredictable and performance demands are great. Change is rapid and constant. And the emergence of newer applications will keep the momentum of change moving.

Yesterday's legacy routers – with their limited packet-forwarding schemes -- cannot keep up with these new web applications. They are slow, expensive and difficult to manage, and often require learning and using proprietary routing protocols. What's more, they were designed to support yesterday's predictable traffic patterns instead of today's random, megabyte-intensive data flows.

Extreme Networks builds scalable, high-performance networks that adapt swiftly to change and empower the delivery of future applications for enterprise LANs, Internet service providers and web content providers. We combine Ethernet and the Internet protocol with wire-speed gigabit switching and Policy-Based Quality of Service so businesses can deliver more information faster and accommodate future growth.

Our Summit™ stackable and BlackDiamond™ chassis switches share a consistent hardware, software and management architecture for end-to-end simplicity, plus a non-blocking switch fabric to ensure ultra-fast throughput. Based on our internally developed Summit Chipset™ and ExtremeWare™ software, this architecture delivers Wire-Speed IP Routing at Layer 3 and full line-rate switching at Layer 2.

The Extreme Networks® architecture also implements standards-based routing using RIP, RIP v2, OSPF, IGMP v2, IGMP snooping, DVMRP and PIM Dense Mode. Consequently, our switches can be used with existing legacy routers to build large, fault-tolerant meshed networks that can scale to massive proportions. While our switches handle IP traffic at wire-speed, legacy routers can support non-IP traffic and WAN connectivity.

This guide demonstrates how businesses can cost-effectively migrate to a wire-speed Gigabit Ethernet solution that can scale well into the future and reduce the overall cost of network ownership.

Wiring Closet

1. Legacy router solution – Shared 10 Mbps or 100 Mbps Ethernet to the Desktop

In this scenario, Ethernet switches have been deployed in the wiring closet to segment Ethernet collision domains. However, desktop connections are still provided primarily by shared-media hubs, as shown in Figure 1. As the process of microsegmentation has progressed over time, the number of users per segment has been reduced to less than 10.

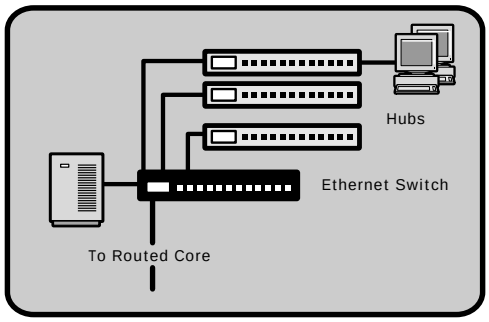

Figure 1

Now a new solution is required that provides more bandwidth per user, greater uplink bandwidth to support more centralized traffic flows, improved resiliency to ensure high availability of connectivity to the core, and more predictable bandwidth for mission-critical applications. Another drawback of shared-media networks is that they are not compatible with many emerging applications, such as enterprise resource planning (ERP), voice over IP (VoIP), videoconferencing and high-bandwidth streaming video.

2. Legacy router solution-- Switched 10 Mbps Ethernet to the Desktop

In the network shown in Figure 2, each desktop has been provided with a dedicated 10 Mbps switched connection with Fast Ethernet to the core. But because these are early second-generation switches, the desktop ports are restricted to half-duplex operation at 10 Mbps. Compared to the shared-media Ethernet in Scenario #1, bandwidth to the desktop and the network core has been increased significantly.

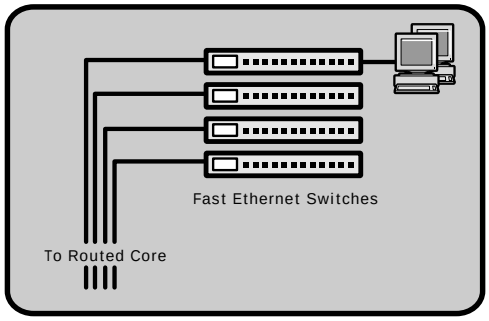

Figure 2

However, these switches do not offer a smooth migration path to higher-bandwidth connections or more predictable bandwidth. Because many switches of this generation have only one uplink port, dual redundant connections to the core are often not feasible. These older switches also lack features, such as intelligent IP multicast and Policy-Based Quality of Service, that are required to support today's emerging applications.

3. Extreme Networks Solution- Summit48 Clusters

A cluster of Summit48 switches provides a highly cost-effective solution for wiring closets serving 48 to 384 users. Each Summit48 switch provides 48 10/100 Mbps ports and two Gigabit Ethernet ports. In low-density configurations, the gigabit ports can be used for local server connections and connectivity to the core. As pairs of Summit48s switches are deployed in each closet, the switches can be interconnected with Gigabit Ethernet and dual-homed to the core of the network, as shown in Figure 3.

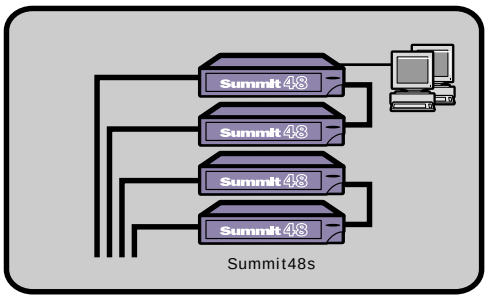

Figure 3

As a direct replacement for Layer 2 wiring closet switches, the Summit48 cluster may be configured for Layer 2 forwarding with the entire cluster serving as a single broadcast domain. With the uplinks forwarding at Layer 2, half of the cluster uplinks will carry traffic while the other half is in standby mode in accordance with Spanning Tree or the Extreme Standby Routing Protocol (ESRP).

In addition to the redundant uplinks, resiliency can be added to the cluster in the form of redundant interswitch connections. In the future, when more uplink bandwidth is required, Layer 3 forwarding can be enabled on the uplink ports. With Layer 3 enabled, all links to the core carry traffic, which can be load-balanced with OSPF or RIP with equal cost multipath (ECMP) routing.

Layer 3 uplinks can also improve resiliency when meshed to the core. Configuring the meshed topology is possible in any fashion and reconvergence time after a topology change is much faster at Layer 3 than at Layer 2. Another advantage of Layer 3 forwarding on the uplinks is that the Layer 2 broadcast domain serving the desktop ports can be segmented to localize broadcast traffic.

4. Extreme Networks Solution- Wiring Closet Stacks

The wiring closet stack of Summit1 and Summit48 switches provides a fully redundant solution for wiring closets serving 96 to 288 switched 10/100 Mbps desktops. In this configuration, the Summit1 switches provide wiring closet routing and the Summit48 switches are configured to forward traffic at Layer 2.

Each Summit48 switch is dual-homed to the pair of Summit1 switches with ESRP providing the protection from router or link failure. The Summit1 links to the core perform Layer 3 load sharing, while half of the redundant Layer 2 connections between each Summit48 and Summit1 switch are in standby mode.

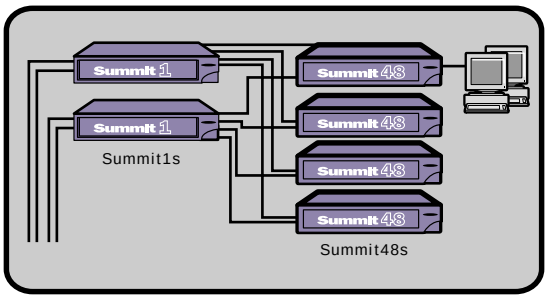

Figure 4

With this configuration, there is no single point of failure in the wiring closet. In fact, if desktop systems are equipped with dual-port network interface cards and each NIC port is connected to a separate Summit48, the wiring closet stack can provide fully redundant connectivity from the end system to the network core for up to 288 users.

5. Extreme Networks Solution- BlackDiamond 6800 in the Wiring Closet

The BlackDiamond 6800 switch delivers the functionality of the wiring closet stack in a 10-slot modular chassis switch that serves up to 256 switched 10/100 Mbps desktops or up to 224 10/100 Mbps switched desktops with four individual or aggregated Gigabit Ethernet uplinks to the core. BlackDiamond interface options include 4-port and 6-port GBIC-based Gigabit Ethernet, 32-port 10/100BASE-TX, and 32-port 100BASE-FX modules that can be installed in any of the eight user slots.

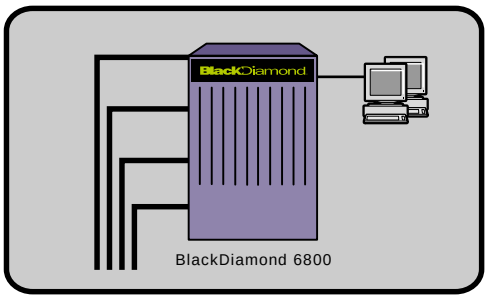

Figure 5

Each interface module has dual connections to the redundant BlackDiamond Management Switch/fabric Modules (MSMs), which eliminate single points of failure and, together with redundant power, provide very high fault tolerance in a single chassis. Compared to other Extreme Networks wiring closet solutions, the BlackDiamond 6800 switch can scale in larger increments, while reducing the complexity of interswitch cabling, multiswitch configuration and physical management.

6. Legacy Router Solution- Collapsed Backbone Router

As shown in Figure 6, multiprotocol bridge/routers have been deployed as the collapsed backbone for medium-sized enterprises or for site backbones in large campus LANs. Resiliency can be achieved by dual-homing the wiring closets to the core as shown.

The collapsed backbone can scale in size by interconnecting site or departmental routers with either point-to-point links or a Layer 2 LAN, such as Fast Ethernet or FDDI. A distributed backbone based on site routers and FDDI is shown in the next scenario.

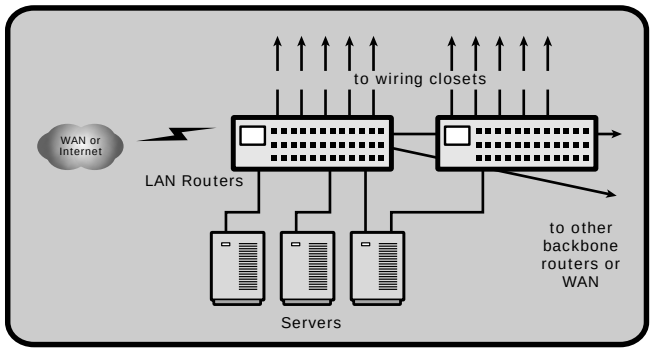

Figure 6

Traditional collapsed-backbone routers provide a simple, effective solution as long as the flow and volume of network traffic remain fairly predictable and within the capacity constraints of the router. When bandwidth demands exceed these boundaries, data bottlenecks and contention for limited forwarding resources can lead to excessive user response time and session time outs.

The performance capacity of traditional high-end legacy routers is limited in terms of the processing power of the CPU (100,000 to 200,000 pps) and the capacity of the backplane (~1 Gbps). These performance limitations prevent traditional routers from providing adequate support for Gigabit Ethernet interfaces or for high-density Fast Ethernet or FDDI interfaces.

7. Legacy Router Solution- Shared or switched FDDI

For large enterprises and campus infrastructures, site routers are commonly connected with a Layer 2 backbone, such as shared-media FDDI or switched FDDI, as shown in Figure 7. Resiliency in backbone router connections to FDDI requires dual-homing to pairs of wiring concentrators or to redundant core switches in the case of switched FDDI.

FDDI has been a popular backbone technology because it provided a 10-fold improvement in performance over shared 10 Mbps Ethernet, supported network topologies with large physical diameters and provided a high degree of fault tolerance due to the self-healing, dual-ring topology. However, FDDI does not support link speeds beyond 100 Mbps, which prevents it from meeting the need for greater backbone capacity.

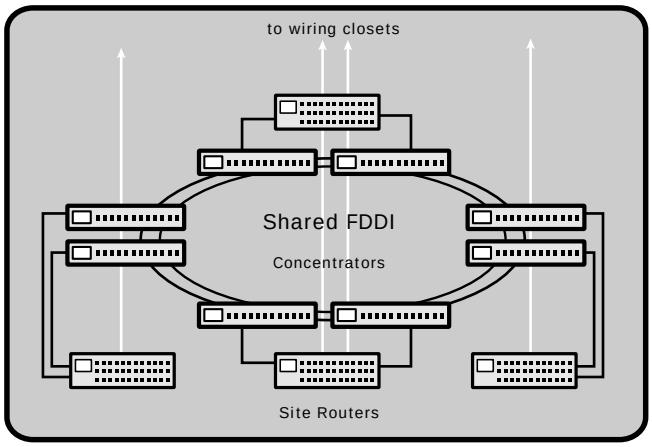

Figure 7

8. Legacy Router Solution- ATM LAN Emulation

In the mid-1990s, Asynchronous Transfer Mode (ATM) appeared to be a promising technology for LAN infrastructures. Hardware-based switching of fixed-length cells ensured that link speeds could scale indefinitely. In addition, ATM's hard quality of service (QoS) provided the service guarantees required to support data, voice and video traffic over a single LAN/WAN infrastructure.

However, the difficulty of converting frame-based network applications to native ATM's connection-oriented model prevented its adoption as an end-to-end desktop-to-desktop or desktop-to-server LAN solution. This saddled ATM LANs with the burden of handing frame segmentation and re-assembly (SAR) plus emulation of the connectionless behavior of frame-based LANs, as shown in Figure 8.

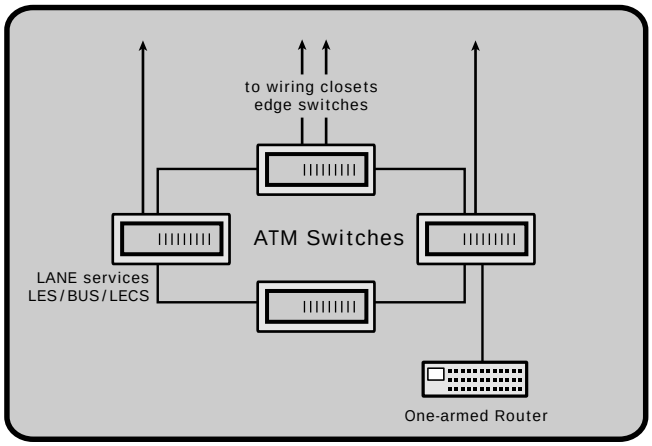

Figure 8

Many adopters of ATM LAN backbones have found that the hybrid network, based on a combination of connectionless and connection-oriented network technologies, has a number of drawbacks:

- ATM aggregate switch performance fails to keep pace with Ethernet switching
- ATM price per port and price per Gbps fails to keep pace with Ethernet switching
- LAN Emulation (LANE) has led to performance bottlenecks and anomalies
- End-to-end QoS on the ATM LAN is difficult to implement and manage
- IP multicast services are more difficult to implement than in a fully connectionless network
- The complexity of ATM and dual ATM/Ethernet networks make management more difficult and increase the total cost of network ownership.

Extreme Networks Backbone and Core Solutions

9. Extreme Networks Solution- Summit1 or Summit7i Cluster

For medium-sized networks or campus sites with Gigabit Ethernet uplinks to the wiring closet, a partially meshed cluster of Summit1 or Summit7i switches provides a scalable, high-performance core solution. The Summit1 switch has eight Gigabit Ethernet ports and the Summit7i switch has 32 Gigabit Ethernet ports.

With the Summit1 switch, uplinks to the wiring closets can be dual-homed to the switch cluster as shown in Figure 9. This dual-homed configuration can use Layer 3 for load sharing and redundancy or Layer 2 with ESRP for router and link redundancy and protection. In the configuration, 16 Gigabit Ethernet ports are available for additional wiring closet connectivity, and 8 Gigabit Ethernet ports are available for site and departmental server connections or for connections to the remainder of the campus core network.

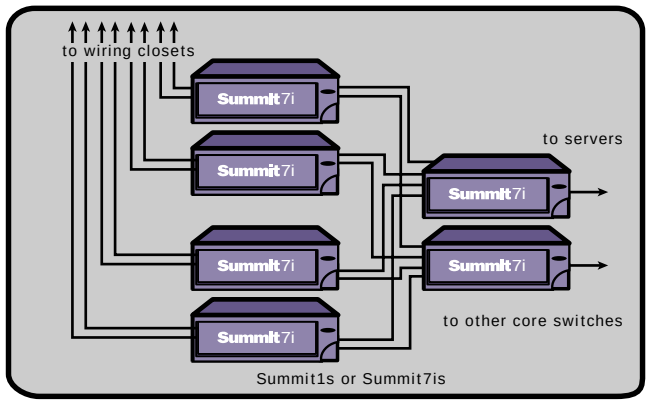

Figure 9

The resiliency features of the Summit1 cluster include optional redundant power supplies, partial meshing within the cluster and dual-homed uplinks and core connections. With the Summit 7i switch, the same configuration can scale using link aggregation for wiring closets and interswitch connections to the core or server switches. The Summit7i switch with 28 autosensing 100/1000BASE-X or 1000BASE-SX ports and four GBIC-based Gigabit Ethernet ports for fiber connections is ideal for directly attaching servers to the core, as shown in the upper right corner of Figure 10.

10. Extreme Networks Solution- BlackDiamond 6800 Switch Mesh

For large site networks or campus backbones, a partial mesh of BlackDiamond 6800 switches provides the highest level of scalability and redundancy for the most demanding Layer 3 core applications. For interswitch connections within the mesh, one set of aggregated links per BlackDiamond 6800 switch provides physical and logical redundancy and load sharing. Each BlackDiamond 6800 switch can accommodate up to 48 Gigabit Ethernet ports with a 3:1 switch fabric over-subscription ratio or up to 32 Gigabit Ethernet ports with zero over-subscription.

In the two-switch fully meshed configuration shown on the left of Figure 10, up to 92 Gigabit Ethernet ports are available for single-homed or dual-homed connections to the wiring closet, server and core. In the four-switch partial mesh shown on the right of Figure 11, up to 176 Gigabit Ethernet ports are available for additional connections.

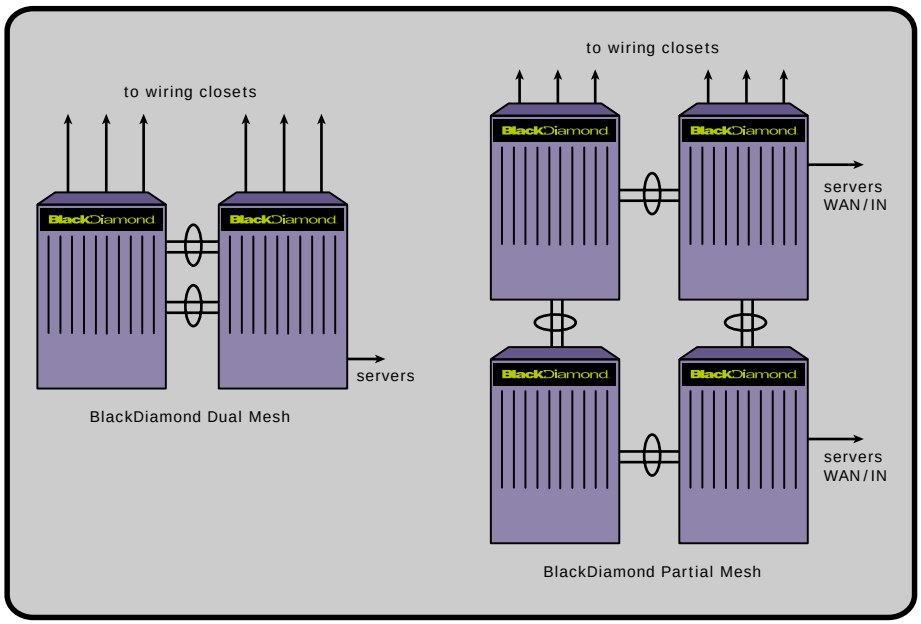

Figure 10

For even larger campus networks, the number of partially meshed BlackDiamond switches may be increased indefinitely. The resiliency features of the BlackDiamond chassis include redundant power supplies, redundant management switch/fabric modules and partial meshing within the cluster and dual-homed uplinks.

Data Centers and Server Farms

11. Legacy Router Solution- Backbone Router

With a collapsed router backbone, server farms usually connect servers directly to server site backbone routers or to routers using a "distribution" layer hub or Layer 2 switch. A non-redundant server farm can be based on a single router and switch as shown on the left in Figure 12. Redundancy can also be achieved using dual-homed server connections to redundant switches and routers, as shown in Figure 11.

High-speed router interfaces – 100 Mbps or higher – can be very expensive. This makes the distribution layer switch a viable alternative for connecting large server farms in fault tolerant configurations. But router-based server farms can suffer from the same performance limitations as traditional collapsed-backbone router. In this case, aggregated server farm traffic and congestion on router interfaces that support direct or indirect server connections can accentuate performance problems.

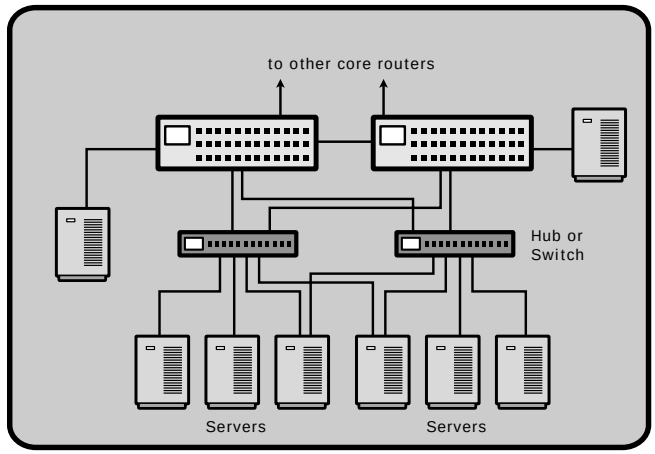

Figure 11

12. Legacy Router Solution- FDDI-Attached Servers

With FDDI backbones, servers are either directly attached to the ring with dual-homed connections to pairs of wiring concentrators or dual-homed to redundant core switches, as shown in Figure 12. Another alternative is to connect the servers to FDDI routers as shown in Figure 11.

With switched FDDI, there is more server farm bandwidth than with the router-connected servers shown in Figure 11. But there is still no Gigabit speed migration path for server or backbone connections. Another drawback of large FDDI server farms is that many servers are often in the same Layer 2 backbone broadcast domain, which wastes server CPU cycles during extraneous broadcasts and reduces the effectiveness of router-based security.

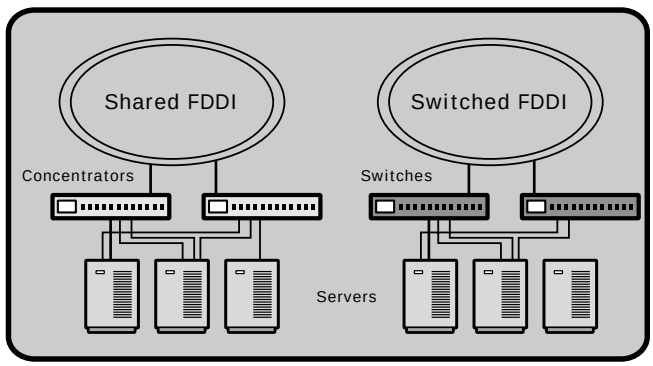

Figure 12

13. Legacy Router Solution- ATM-Attached Servers

In ATM environments, large server farms generally have Fast Ethernet-attached servers that are dual-homed to distribution-layer edge switches. These edge switches are in turn dual-homed to ATM switches, as shown in Figure 13. The distribution layer offers economical server connections and helps avoid LAN emulation, which would be required if the servers were directly ATM-attached. Layer 3 switching occurs in the distribution layer, while the ATM backbone performs Layer 2 switching.

When high-performance Layer 3 edge switches with ATM OC-48c and Gigabit Ethernet become available, server farm performance will scale immensely in this ATM/Fast Ethernet network. However, some early adopters of ATM in the LAN are reconsidering this approach, given the higher complexity and higher cost of ownership of a cell-to-frame network.

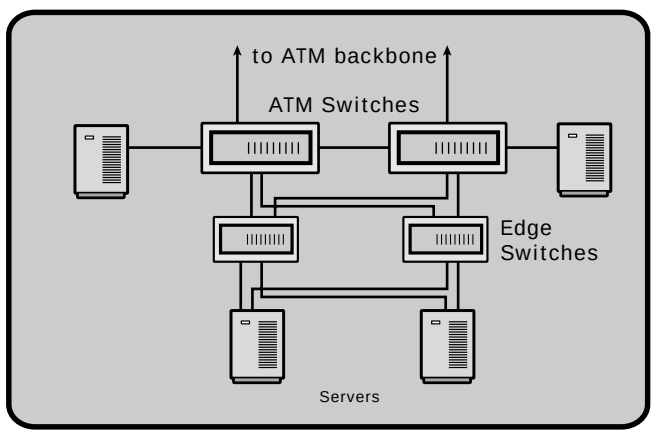

Figure 13

14. Extreme Networks Solution- Dual Summit Switches

In small server farms, servers can be connected directly to a collapsed-backbone Summit switch or to the Summit cluster core shown in Figure 9. In more robust server farm designs, availability requirements dictate that servers be dual-homed to pairs of server farm switches that are connected to separate core switches as shown in Figure 14.

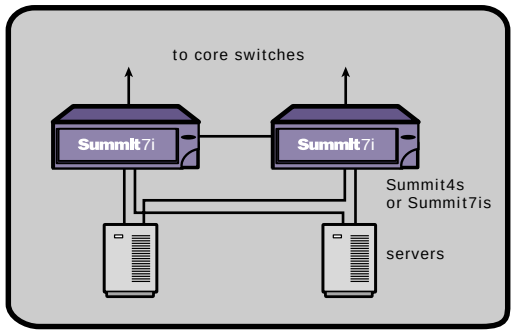

Figure 14

Pairs of Summit4 switches are a good choice for smaller scale high-availability server farms because they allow up to eight servers to be dual-homed using Fast Ethernet to the core and up to four more servers to be dual-homed using Gigabit Ethernet. Optional redundant power supplies provide additional server farm resiliency for the Summit4 switches.

The Summit7i switch can be used instead of the Summit4 switch when more server connections are required and when 100 Mbps to 1000 Mbps scalability of server-to-switch bandwidth is desirable. With a pair of Summit7i switches, up to 28 servers can be dual-homed to the core using 100/1000BASE-X switch ports.

15. Extreme Networks Solution- Direct Core Attachment

In medium-sized networks where core switches and the data center are co-located, a server farm can perform wire-speed Layer 3 switching by directly connecting servers to dual meshed BlackDiamond core switches, as shown in Figure 15. The BlackDiamond switches have sufficient port density and redundancy features to support the core network, as well as provide reliable connections for large server farms without requiring additional dedicated server-distribution switches. For this type of direct-attachment solution, the server farm inherits the performance and resiliency characteristics of the core network design.

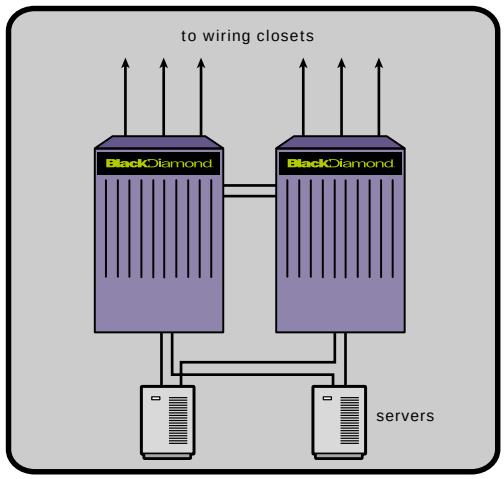

Figure 15

16. Extreme Networks Solution- Dual BlackDiamond and Summit48 Mesh

For large campus server farms, a meshed pair of BlackDiamond 6800 switches can be dedicated to the server farm site, while a distribution layer of Summit48 switches provide high-density Fast Ethernet server connections, as shown in Figure 16.

In this design, each Summit48 switch performs Layer 2 switching and is dual-homed using the Extreme Standby Router Protocol (ESRP) to the BlackDiamond switches. Servers may also be dual-homed to the Summit48 switches using redundant server adapters. With 6 Gigabit Ethernet ports allocated to core and interswitch connections, over 40 Summit48 switches can be dual-homed to the BlackDiamonds. This means the basic redundant configuration can support over 960 dual-homed Fast Ethernet servers. For resiliency and high availability, BlackDiamond supports redundant power supplies and redundant management switch/fabric modules.

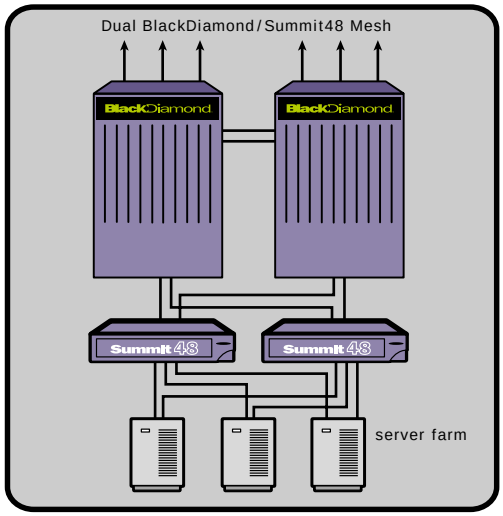

Figure 16

17. Extreme Networks Solution - Large Campus Server Farm

For large campus server farms, a meshed pair of BlackDiamond 6800 switches can be dedicated to the server farm site while a distribution layer of Summit48 switches provide high-density Fast Ethernet server connections, as shown in Figure 17.

In this design, each Summit48 switch performs Layer 2 switching and is dual-homed using the Extreme Standby Router Protocol (ESRP) to the BlackDiamond switches. Servers may also be dual-homed to the Summit48 switches using redundant server adapters. With 6 Gigabit Ethernet ports allocated to core and interswitch connections, over 40 Summit48 switches can be dual-homed to the BlackDiamond switches. This means the basic redundant configuration can support over 1000 dual-homed Fast Ethernet servers. For resiliency and high availability, BlackDiamond 6800 switch supports redundant power supplies and redundant management switch/fabric modules.

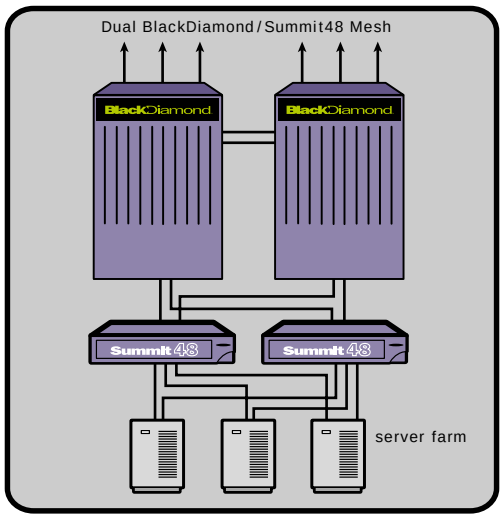

Figure 17

18. Extreme Networks Solution- BlackDiamond 6800 Web Hosting Site

Far different from enterprise server farms, web-hosting sites have vast arrays of servers that require the ultimate in reliability and scalability to keep pace with the explosive growth of e-commerce and other traffic over the Internet.

To meet these requirements, a mesh of BlackDiamond 6800 switches can provide fully redundant server distribution to support user access to web content. It can also provide fully redundant back-end server connectivity for content replication, backup and restore functions.

As web-hosting sites and bandwidth demands grow, additional BlackDiamond switches can be added to both the distribution and back-end meshes. For larger web-hosting sites, new tiers of meshed BlackDiamond switches can be added to increase server and router connections. By doing so, the web-hosting site can be expanded to support a virtually unlimited number of users and web sites.

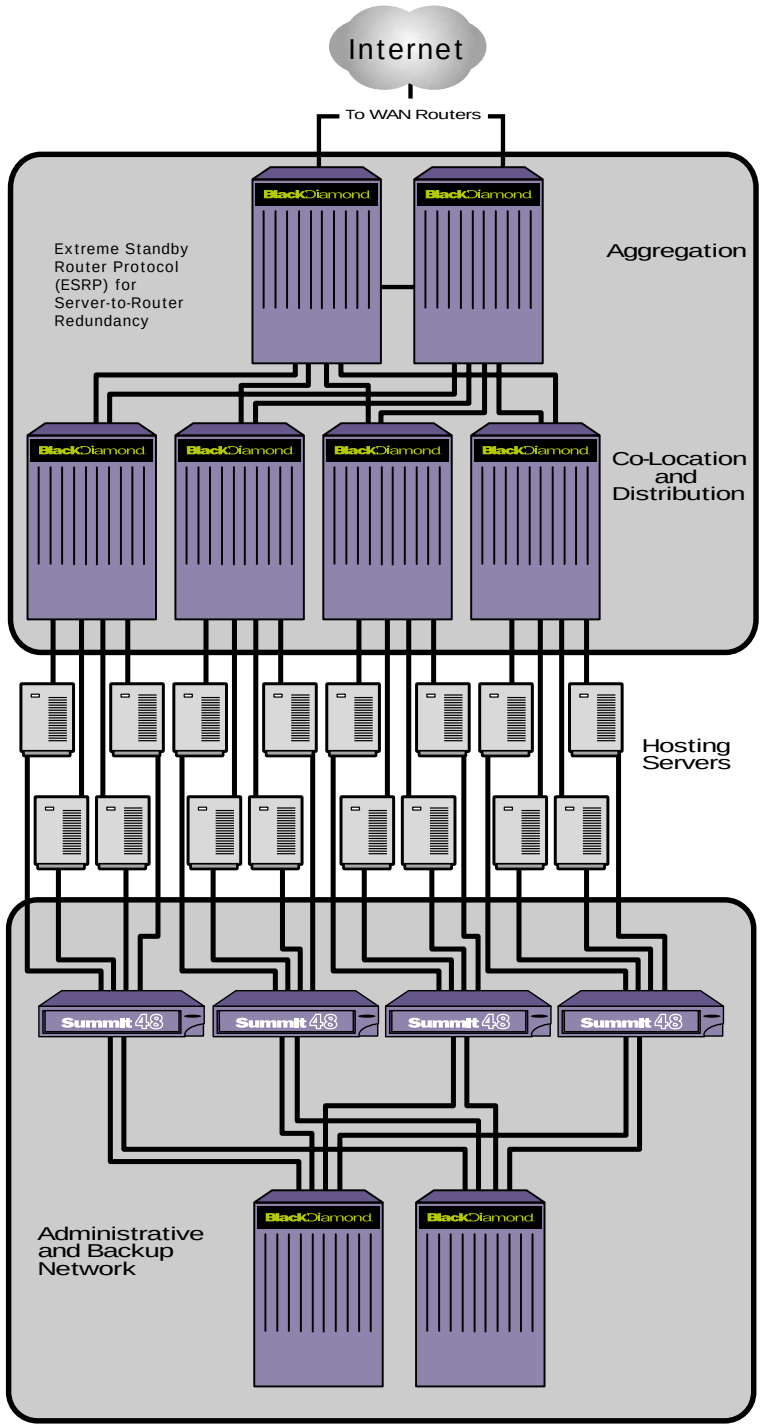

Figure 18

Summary

Transitioning the wiring closet to fully switched 10/100/1000 Mbps Ethernet is relatively simple when starting from shared or second-generation switched Ethernet. Close attention should be paid to ensuring that end-station network interface cards can handle full-duplex operation and 10/100 Mbps autosensing. For an interim period, Fast Ethernet can be used for wiring closet connections to traditional multiprotocol routers that constitute the core.

When faced with sweeping changes – such as going from Token Ring to Ethernet in the wiring closet or transitioning backbone and server farm access methods – a more carefully planned, staged migration away from the legacy technology is in order.

In these more challenging transitions, it is often desirable to install a new solution in parallel with an existing one. For example, when the network core is transitioned from multiprotocol legacy routers to meshed wire-speed IP switches, the two backbones can operate in parallel during a transition period.

But over time, key IP services and application server resources – which make up a large majority of network traffic – should be shifted to the wire-speed IP switches in the core. This leaves multiprotocol legacy routers to handle of non-IP traffic, such as AppleTalk and DECnet.

As shown in Figure 19, businesses can truly maximize the performance of their network infrastructures by deploying wire-speed Layer 3 switching, 10/100/1000 Mbps Ethernet and IP throughout the entire network – in wiring closets, the backbone and server farms.

Extreme Networks' Summit stackable and BlackDiamond chassis switches share a consistent hardware, software and management architecture for end-to-end simplicity. This consistent product functionality and manageability across the entire network minimizes administrative complexity and can significantly reduce the total cost of network ownership.

With managed bandwidth, consistent end-to-end functionality and streamlined management, the stage is set for business-optimization of the network by protecting mission-critical applications with Policy-Based Quality of Service. Once the delivery of mission-critical applications can be guaranteed, the introduction of newer network services and applications, such as videoconferencing and voice over IP, can be considered in earnest.

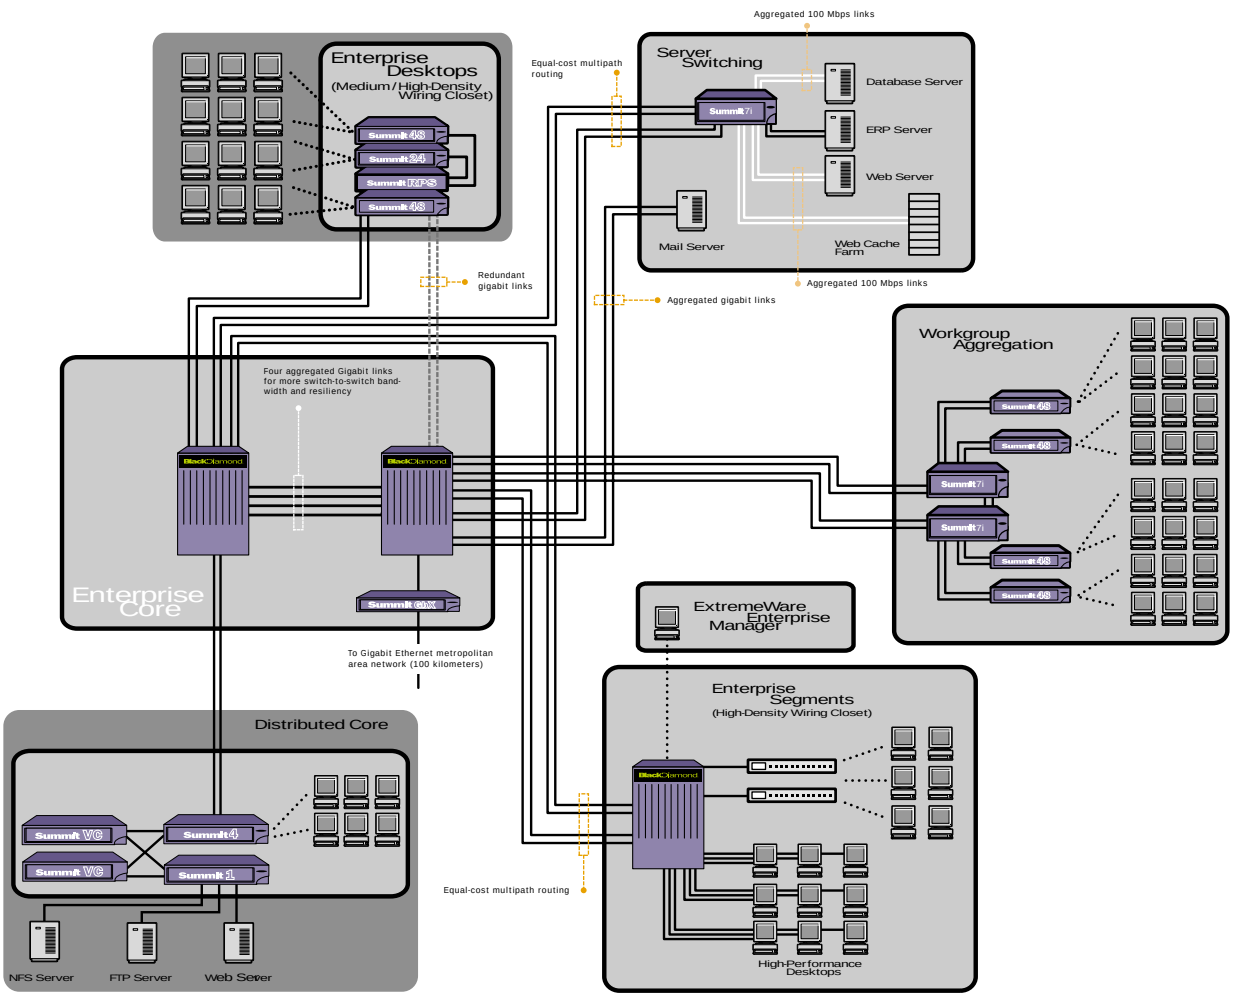

Figure 19

BlackDiamond™ 6800 Switch

Wire-speed Gigabit Ethernet performance, high availability, high port-density, fault tolerance and redundancy for the world's largest mission-critical networks.

- Ideal for core, wiring closet and server switching
- 64 Gbps non-blocking switch fabric and 48 million packets per second throughput
- Wire-Speed IP Routing at Layer 3 and wire-speed switching at Layer 2
- Spanning tree, multipath routing and redundant router protocol increase system reliability
- Up to 256 10/100 Mbps Ethernet ports or 48 Gigabit Ethernet ports in a single chassis
- Fault-tolerant: hot-swappable module and components, load-sharing power supplies and management modules, link aggregation, dual software images and system configurations

Summit48™ and Summit24™ Enterprise Desktop Switches

These wire-speed Enterprise Desktop Switches provide high-density 10/100 Mbps Ethernet connections to the desktop and Gigabit Ethernet links to the core.

- 48 10/100 Mbps Ethernet ports and 2 Gigabit Ethernet ports or 24 10/100 Mbps Ethernet ports and 1 Gigabit Ethernet port
- Summit24 - 8.5 Gbps non-blocking switch fabric and 5.1 million packets per second forwarding rate
- Summit48 - 17.5 Gbps non-blocking switch fabric and 10.1 million packets per second forwarding rate

Summit1™ Switch

This one's all gig. Pure wire-speed performance and non-blocking switch capacity for network backbones.

- 8 Gigabit Ethernet ports
- 17.5 Gbps non-blocking switch fabric
- 11.9 million packets per second Wire-Speed IP Routing

Summit4™ and Summit4/FX™ Switches

Ideal as a server and segment switch. It has just the right mix of 10/100 Mbps and Gigabit Ethernet ports. Plus the usual wire-speed performance and non-blocking capacity.

- 16 10/100 Mbps Ethernet ports
- 6 Gigabit Ethernet ports
- 17.5 Gbps non-blocking switch fabric
- 11.3 million packets per second Wire-Speed IP Routing

Summit7i™ Internet Data Center Switch

32 gig ports and many other chassis-like features packed into a 7-inch high stackable. Plus gigabit over copper, access policies, server load balancing and more.

- 32 Gigabit Ethernet ports
- 64 Gbps non-blocking switch fabric
- 48 million packets per second Wire-Speed IP/IPX Routing

Glossary of Networking Terms

10BASE-T
IEEE networking standard for 10 Mbps twisted-pair Ethernet.

100BASE-TX
IEEE networking standard for 100 Mbps twisted-pair Ethernet; also called Fast Ethernet.

100BASE-FX
IEEE networking standard for 100 Mbps Ethernet over multimode fiber optic cabling. This is one version of Fast Ethernet.

1000BASE-SX
IEEE networking standard for a variant of Gigabit Ethernet over multimode fiber optic cable at an 850 nm wavelength.

1000BASE-LX
IEEE networking standard for a variant of Gigabit Ethernet over multimode and single-mode fiber optic cable at a 1330 nm wavelength.

1000BASE-T
IEEE networking standard for a variant of Gigabit Ethernet over unshielded twisted pair cable.

Address
A set of characters that identifies an individual network node.

Address Table
The database maintained by a switch of all addresses it has learned and the switch ports through which these addresses can be reached. It is used by the switch to make packet forwarding and filtering decisions.

ASIC
Application-Specific Integrated Circuit. A chip designed for a particular application. ASICs are commonly used in networking devices to maximize performance with minimum cost.

ATM
Asynchronous Transfer Mode. Cell-based network technology capable of transmitting data, voice, video, and frame-relay traffic.

Auto-Negotiating
Two-part process by which a network device automatically senses the speed and duplex capability of another device.

Autosensing
Process during which a network device automatically senses the speed of another device.

BGP4
An extension of the border gateway protocol, exchanges routing information between gateway hosts, each with its own router, in a network of autonomous local systems. Known as an exterior gateway protocol (EGP), it has many dynamic characteristics that make it the best choice as an interface between interior networks using RIP and OSPF and the Internet backbone. In addition to enabling network administrators to configure cost metrics based on policy management information, BGP4 supports classless inter-domain routing (CIDR) so that multiple IP addresses can be mapped to a single route.

Backbone
Interconnection in a LAN or WAN between subnetworks or workgroups. The high-speed connection to lower speed subnets. For example, a Gigabit Ethernet backbone connected to Fast Ethernet subnets.

Backplane
Bus or switching matrix that resides within a switch or hub chassis; all traffic crosses the backplane at least once.

Bandwidth
The maximum amount of data that can be transmitted in a fixed amount of time; usually expressed in bits-per-second or bytes-per-second.

Broadcast
Message forwarded to all devices within a network. Broadcasts exist at Layer 2.

Category 5
Networking standard certifying that a copper wire cable can carry data at up to 100 Mbps. See also UTP.

CLI
Command Line Interface. An interface that allows the user to interact with the operating system by entering commands and optional arguments.

Client / Server
Distributed computing model where desktop "clients" can access and share information resources from multiple "servers."

Collapsed Backbone
LAN architecture in which the subnetwork interconnection is concentrated within a Layer 3 switch or router.

Collision
Concurrent Ethernet transmissions from two or more devices on the same segment.

Concentrator
Device used in a LAN to combine transmissions from a cluster of clients and/or servers; often called a hub.

COPS
The Common Open Policy Service protocol is used with RSVP and policy-based networking for the communication between a network device and an authoritative policy management entity. Typically, this is a policy server or call admission control server process. It defines the transport and formatting of data used in this communication.

Data-Link Layer
See Layer 2.

Destination Address
The IP or MAC address of the node that is to receive the packet.

DHCP
Dynamic Host Control Protocol. An effective way to dynamically assign and reuse a fixed number of IP addresses when there are more devices on the network than addresses available. A DHCP server dynamically assigns IP addresses to devices requesting them. These address assignments expire after a time specified by the network manager. The DHCP server then reassigns these addresses to other devices as needed. DHCP is an extension to BOOTP in which the address assignments are static.

DiffServ
An IETF standard developed to help solve IP quality problems. DiffServ operates at Layer 3 and allows out-of-band negotiation. Diffserv relies on traffic conditioners sitting at the edge of the network to indicate each packet's requirements.

Duplex
A communication mode in which a device can send and receive data over the same link. The device can operate in full duplex and half duplex.

Dynamic Link Context System
The Dynamic Link Context System allows policy setting based on user names or desktop devices, and automatically maps them to lower-layer addresses.

DVMRP
The Distance Vector Multicast Routing Protocol is used for the communication and distribution of multicast routing table information. It is based on the RIP protocol used in unicast routing. (see IETF document draft-ietf-dvmrp-v3-07)

Enterprise Desktop Switch™
An Enterprise Desktop Switch combines the low cost and simplicity of a stackable edge device with the enterprise-class features found in more expensive chassis-based switches.

ERP
Enterprise Resource Planning. A business management system that integrates all facets of the business, including planning, manufacturing, sales, and marketing. As the ERP methodology has become more popular, software applications have emerged to help business managers implement ERP.

ESRP™
The Extreme Standby Router Protocol enables host devices to continue communicating even if a physical router fails.

Ethernet
An IEEE networking standard, originally developed by Xerox, for transmitting data at 10 Mbps.

Fast Ethernet
An IEEE networking standard for transmitting data at 100 Mbps. See 100BASE-TX.

Fault-Tolerance
The ability of a device to prevent or recover from network and internal failures. Key elements of fault tolerance include hot-swappable modules, redundant load-sharing power supplies, passive backplanes, and redundant cooling systems.

FDDI
Fiber Distributed Data Interface. An ANSI networking standard for 100 Mbps fiber-optic LANs; widely used as a backbone technology to interconnect several Ethernet or Token Ring networks.

Filter
An action of the switch to discard certain types of data packets.

Firmware
Software routines that are permanently written onto read-only memory.

Full-Duplex
The communication mode in which a device simultaneously sends and receives over the same link, doubling the bandwidth. A full duplex 100 Mbps connection has 200 Mbps. of bandwidth. A full duplex 1000 Mps connection has 2000 Mbps of bandwidth.

GBIC
Gigabit Interface Connector. The physical connection to Gigabit Ethernet media.

Gbps
Billions of (or giga) bits per second.

Gigabit Ethernet
Networking standard for transmitting data at 1000 Mbps.

Half-Duplex
The communication mode in which a device is capable of either sending or receiving, but not simultaneously.

Hardware Address
A device's physical or media access control (MAC) address.

Header
Special information contained in the beginning of a frame.

Host
Any entity on the network that can initiate a transmission. A router, a server or a workstation.

HTTP
Hypertext Transfer Protocol. Defines how requests for HTML and graphics files which make up a WWW page are handled between the web server and the client browser.

Hub
An unintelligent network device that sends one signal to all of the stations connected to it.

ICMP
Internet Control Message Protocol. The part of the IP protocol that handles error and control messages. The switch issues ICMP messages to report IP datagram problems back to the their source.

IGMP
Internet Group Management Protocol. A protocol that hosts use to keep local routers informed of their membership in multicast groups. When all hosts leave a group, the router no longer forwards datagrams that arrive for the group.

IGMP Snooping
IGMP Snooping provides a method for intelligent forwarding of multicast packets within a Layer 2 broadcast domain. By snooping IGMP registration information, a distribution list of workstations is formed that determines which end-stations will receive packets with a specific multicast address.

IEEE 802
Set of Institute of Electrical and Electronic Engineers standards for defining methods of access and control on LANs.

IP
A Layer 3 (network layer) protocol that contains addressing information and control information that allows packets to be routed.

IPX
Internetwork Packet Exchange. A networking protocol used by Novell® NetWare® operating systems. Like UDP/IP, IPX is a datagram protocol used for connectionless communications.

ISO
International Standards Organization.

Latency
Any delay introduced into the network that prevents packet forwarding at wire speed.

LAN
Local Area Network. A network where computers are connected in close proximity, such as in the same building or office park; a system of LANs connected at a distance is called a wide-area network (WAN).

Layer 1
The first, or physical, layer of the open systems interconnection (OSI) model. Delivers data across a network link. This layer must regulate signaling and keep the signal strong. Hubs, repeaters and concentrators operate at Layer 1. All packets received are repeated on the wire.

Layer 2
The second, or data-link layer, of the open systems interconnection (OSI) model. The media access control (MAC) layer. Transmits packets across a Layer 1 physical link by reading the hardware or MAC source and destination addresses in each packet. Switches operate at Layer Two. Switches have a forwarding table of the hardware addresses of the devices connected to them. When packets arrive, the switch reads the Layer 2 address and if it matches one in the table, forwards it to that port. Otherwise, it forwards or "floods" the packet to all ports.

Layer 3
The third, or routing, layer of the open systems interconnection (OSI) model. The network layer routes data to different LANs and WANs based on network address.

Link Aggregation
The grouping of multiple network links into one logical high bandwidth link. By grouping four 100 Mbps Ethernet NICs into one logical link, you can create up to 800 Mbps of throughput between the server and the switch.

MAC
Media Access Control. Layer 2 of the open systems interconnection OSI model. The data-link layer responsible for scheduling, transmitting and receiving data on a local area network.

MAC Address
Media Access Control address. The unique physical address of each device's network interface card.

MAN
Metropolitan Area Network. A network spanning a geographical area greater than a LAN, but less than a WAN.

Mbps
Millions of (or mega) bits per second.

Meshed Topology
A network built with a mixture of different network topologies. For example a high bandwidth backbone network that connects to a collection of slower segments.

MIB
Management Information Base. A database of information that the switch makes available to network management systems. For example, traffic statistics and port settings.

Multicast
A packet, or transmission, destined for many clients.

NIC
Network Interface Card. An expansion board which goes into a workstation or server and provides the connection to a network.

Non-Blocking
A switch's ability to transmit and receive packets on all ports simultaneously, at wire-speed.

Over-Subscription
Over-subscription or over-subscription ratios deal specifically with points in a network where bottlenecks occur. The impact of improper over-subscription ratios is congestion, which causes packet loss. Over-subscription ratios are calculated by adding the potential bandwidth requirements of a particular path and dividing the total by the actual bandwidth of the path. Although a ratio larger than 1:0 is considered over-subscribed, it does not necessarily mean congestion will occur.

OSPF
Open Shortest Path First. A routing protocol that maintains a map of all other routers and the networks to which they connect. Sends short messages asking whether a neighbor is alive and reachable. More efficient and scalable than vector-distance routing protocols that maintain tables of all known destinations and number of hops to reach them.

PIM Dense Mode
Protocol Independent Multicast Dense Mode. A multicast protocol similar to DVMRP in that it uses Reverse Path Forwarding but does not require any particular unicast protocol.

PIM Sparse Mode
A multicast protocol that works by defining a rendezvous point that is common to both sender and receiver. Sender and receiver initiate communication at the rendezvous point, and when flow begins it occurs over an optimized path.

Policy-Based Quality of Service
A network service that provides the ability to prioritize different types of traffic and manage bandwidth over a network.

Port Density
Number of ports, either physical or logical, per network device.

Port Mirroring
A switching feature that allows one port's media access control (MAC) layer data to be replicated to another port for monitoring by a network analyzer.

RIP
Routing Information Protocol. A protocol defined by RFC 1058 that specifies how routers exchange routing table information. With RIP, routers periodically exchange entire tables.

RMON
Remote Monitoring. A network management protocol that allows network information to be gathered at a single workstation.

Routing
The process of delivering a message across a network or networks.

Router
A network device that forwards packets to destinations based on Layer 3, IP addresses. A router implements various protocols to maintain information on the location of other routers. A router reads the Layer 3 network address information in every packet that it receives and determines whether it should be dropped or forwarded. If it is to be forwarded, the router looks in its routing table to find the best route between a sender and receiver.

RSVP
Resource Reservation Protocol. An IETF standard used to provide quality of service by reserving bandwidth before packet transfers to insure its availability.

Segment
Section of a network that is bounded by bridges, or switches; dividing an Ethernet into multiple segments is a common way to increase bandwidth on a LAN.

SNMP
Simple Network Management Protocol. A standard for gathering statistical data about network traffic and the behavior of network components. SNMP uses management information bases (MIBs), which define what information is available from any manageable network device.

Snooping
Looking into the packet to obtain information.

Source Address
The IP or media access control (MAC) address of the node issuing the packet.

Spanning-Tree
A process used to eliminate redundant data routes and increase network efficiency.

Store-and-Forward
Switching feature where the receiving port receives the entire incoming frame and stores it in the buffers before forwarding it to the destination port.

Switching Fabric
A term used to specify the maximum bandwidth of a switch at the backplane.

Subnet Addressing
A method that a manager can use to span multiple physical networks using a single IP network address. Local routers and intelligent switches use extensions of the IP network address to identify and route traffic to local, physical segments.

Subnet Mask
A number that a manager enters to tell the switch how to filter incoming packets. For example, a subnet mask of 255.255.0.0 for the address 192.3.1.254 tells the switch to only accept traffic destined for IP addresses that begin with 192.3. All other packets are dropped.

Switch
A network device that filters and forwards packets between LAN segments and/or desktops.

TCP / IP
Transmission Control Protocol/Internet Protocol. The suite of communications protocols used to connect hosts on the Internet.

Telnet
The protocol within the TCP/IP protocol suite that provides a terminal emulation function.

Topology
The physical or logical layout, or configuration of a network.

Uplink
A connection from a lower to higher device. A hub to a switch, a switch to a router, a router to a server.

UDP
User Datagram Protocol. A connectionless protocol that, like TCP, runs on top of IP networks.

Unicast
A packet destined for only one address.

UTP
Unshielded Twisted Pair. Cabling with wires that are twisted around each other; the individual wires are uninsulated. See also Category 5.

VID
VLAN Identifier. A number identifying a specific VLAN.

VLAN
Virtual LAN. A logical, not physical, group of devices, defined by software. VLANs allow network administrators to resegment their networks without physically rearranging the devices or network connections.

WAN
Wide Area Network. A network that uses telecommunications technology to connect computers or networks over long distances.

WINS
Windows Internet Naming Service. A system that determines the IP address associated with a particular network computer.

Wire Speed
The theoretical maximum rate at which packets can be transmitted and received on a network interface.

Workgroup
Collection of computers that are grouped for sharing resources such as data and peripherals.

3585 Monroe Street Santa Clara, CA 95051-1450 Phone 408.579.2800 Fax 408.579.3000
Email info@extremenetworks.com Web www.extremenetworks.com

12.99 EN-249