



# Policy Based Quality of Service

for enterprise LANs



## Executive Summary

Competitive pressures are driving corporations to adopt new network-based business models and to deploy e-business applications. Network delivery is the cornerstone of business-critical applications.

In this environment, network managers must be able to provide guaranteed performance – quality of service (QoS.) But increasing network traffic caused by new applications is making it more difficult to provide predictable service. Some have tried to address the QoS requirement by “throwing bandwidth at the problem.” This is not only expensive, but it does not provide performance guarantees.

What is needed is a two-pronged approach – not just more bandwidth, but bandwidth plus intelligent bandwidth management. Bandwidth management has typically involved sophisticated queuing and bandwidth reservation techniques, not to mention an alphabet soup of standards. But there is a simpler solution.

Policy-Based QoS management systems, such as Extreme Networks' policy management, allow network managers to implement QoS in terms of high-level performance and business objectives. Instead of getting involved in the details of queuing mechanisms and configuring routers and switches, a network manager can simply define a policy that says – give my SAP traffic guaranteed bandwidth and the highest priority. Extreme's policy management simplifies the details of policy implementation.

This paper examines current enterprise QoS requirements, the techniques that can be used to address these requirements, and how Extreme's policy management implements policy-based QoS management.

## **The Case for QoS**

The network is now the foundation for new business models and e-commerce initiatives. Mission-critical applications that were once confined to the corporate data center are now widely distributed, and the network is the glue that ties these applications together.


In this environment it is critical that networks be able to deliver predictable and consistent performance – in other words not just service, but quality of service. Poor network performance is no longer just an inconvenience, it directly impacts employee productivity and customer satisfaction. Because the network is an enabler of new business processes and increased productivity, network performance and QoS increasingly have top management visibility. Networks can no longer be best effort delivery systems. Quality of Service controls the expectations of users and application delivery.

## **QoS Challenges**

There is pressure on network managers to deliver high levels of performance – and do it consistently. But, this has become increasingly difficult for two key reasons. First, applications are becoming more bandwidth intensive and can consume bandwidth that is needed by mission-critical applications. When there is no more bandwidth, you must intelligently allocate the resources for your critical applications.

While legacy applications were primarily character-oriented and limited by the interaction speed of end users, new applications exchange data, graphics, voice, and video information between high performance clients and servers. Past raw bandwidth requirements, it is the interaction between the applications on the network that has the most impact on performance.

Mission-critical applications such as ERP and legacy SNA must share the enterprise network with bandwidth-intensive applications such as streaming media, remote database updates, and multicast conferencing. These new applications seem to have an insatiable appetite for bandwidth and network managers are finding that bandwidth is being consumed almost as fast as it can be deployed.



The second dimension of the QoS challenge is the fact that new applications are being deployed quickly and unpredictably. While legacy networks were designed to support applications that were largely under the control of the corporate IT department, application deployment has become a more decentralized function. Today, users are able, and even encouraged, to roll out new distributed applications without the knowledge of the corporate IT department and network operations staff.

It is increasingly difficult for network managers to stay ahead of this growth curve and deploy bandwidth when and where it is needed. As a result, QoS can fluctuate unpredictably. One day, users may be getting high levels of service while the next day they are bombarding the help desk with complaints about performance.

### **Diverse Application Requirements**

Applications differ in the way they use bandwidth and in their QoS requirements. The unpredictable mix of applications running on a network and the conflicts that occur when everyone tries to run their applications at once causes QoS problems. Dealing with these conflicts is the key challenge of QoS management.

There are four service characteristics that are commonly used as QoS metrics – bandwidth requirements, latency, jitter, and packet loss. In addition, resource availability must be considered with regards to tiered QoS, Server Load Balancing and web cache redirection.

### **Bandwidth**

Application bandwidth requirements can be categorized as sustained, bursty, and interactive. Streaming media applications, such as Microsoft's NetShow, require a sustained amount of bandwidth in order to provide users with high quality audio and video information. Some are designed to run over low-speed dial-up connections and require no more than 56 kbps while others, such as high quality MPEG-2 video, may require up to 10 Mbps. These applications can become virtually unusable if the required bandwidth is not available – even for very short periods of time.

Other applications, such as file transfers, are bursty. These applications attempt to grab as much bandwidth as they can to speed data delivery. This bursty traffic must be controlled because it is the most common cause of network congestion that adversely affects the performance of other applications. These bursty traffic characteristics are often the result of the end-to-end TCP protocols used by many applications.

## Latency

Some applications are sensitive to the latency, or delay, in transmitting data across a network. End-to-end latency is due to the latency of physical transmission media and delays introduced by intermediate routers and switches. Significant delays can be introduced when packets are queued for long periods of time. Some queuing mechanisms are designed to control these delays while others can magnify the problem. Real-time audio and video applications, including voice-over-IP, fall into this category. Latency also degrades the response times of interactive applications.

## Jitter

Jitter is related to latency because it refers to the variability of delay the data experiences in networks. Variable queuing delays in routers and switches can cause jitter, and some queuing techniques differ in the amount of jitter that they introduce. Excessive jitter can disrupt real-time video, audio, and voice over IP traffic flows.

## Packet Loss

Packet loss is typically caused by network congestion and it affects all applications because packet retransmission reduces the overall efficiency of networks and, therefore, the amount of bandwidth available to applications. The impact of packet loss differs from application to application. Some multimedia applications can become unusable when packet loss occurs while most business applications simply experience degraded performance.

Application Types	QoS Requirements			
	Bandwidth	Latency	Jitter	Packet Loss
ERP Applications	Moderate	Low	-	Low
Legacy SNA Applications	Low	Low	-	Low
Productivity Applications	Low-Moderate	Moderate	-	-
E-Mail	Low to Moderate	-	-	-
File Transfer	Bursty High	-	-	-
Thin Clients (Citrix, etc.)	Low to Moderate	Low	-	Low
Videoconferencing	Sustained High	Low	Low	Low
Voice over IP	Sustained Moderate	Low	Low	Low
Streaming Media	Sustained Moderate to High	Low	Low	Low
Server Load Balancing	QoS requirements are application and server dependent			

*There are four service characteristics that are commonly used as QoS metrics – bandwidth requirements, latency, jitter and packet loss. This QoS requirements table summarizes these characteristics for some common application types.*

## Techniques for Addressing QoS Requirements

### Throwing Bandwidth at the Problem

Until recently, network managers only had one way of dealing with the QoS problem – throw bandwidth at it. Ethernet has proven to be a scalable technology in terms of interface speed and bandwidth beginning with shared 10 Mbps, switched 10 Mbps, stepping up to 100 Mbps and now reaching 1,000 Mbps, or one gigabit. Gigabit Ethernet provides the higher bandwidth needed in today's enterprise networks but just adding more bandwidth is not enough.

While networks must always be adequately provisioned to deliver good performance under normal traffic loads, bandwidth alone is not sufficient to guarantee that specific applications will perform adequately under all traffic loads. Unless bandwidth is managed, mission-critical applications may experience degraded performance. Bandwidth must be intelligently managed based on application requirements and business priorities. Policy-based management translates business priorities into bandwidth management decisions.

### Extreme's Policy Management

Extreme's policy management service enables network managers to use sophisticated bandwidth management techniques without complexity and cost outweighing the benefits of QoS deployment.

Although this paper describes the techniques that Extreme uses to manage network bandwidth and QoS, network managers and the operations staff are not exposed to all of these details. Network managers only need to specify the QoS results that they want to achieve and Extreme's policy management handles the implementation details. That's what simple Policy-Based QoS management is all about.

### Where Overprovisioning Does Not Work

**Desktop Aggregation to Core Network**  
Over-subscription at the closet uplink

**Backbone to WAN edge link**  
Gigabit backbone into a Fast Ethernet WAN edge

**WAN Edge**  
Fast Ethernet to T-1, T-3, etc.

**Backbone to server farm link**  
100 Mbps desktops will cause congestion into the server farm

**The link into the server**  
Priority to different classes

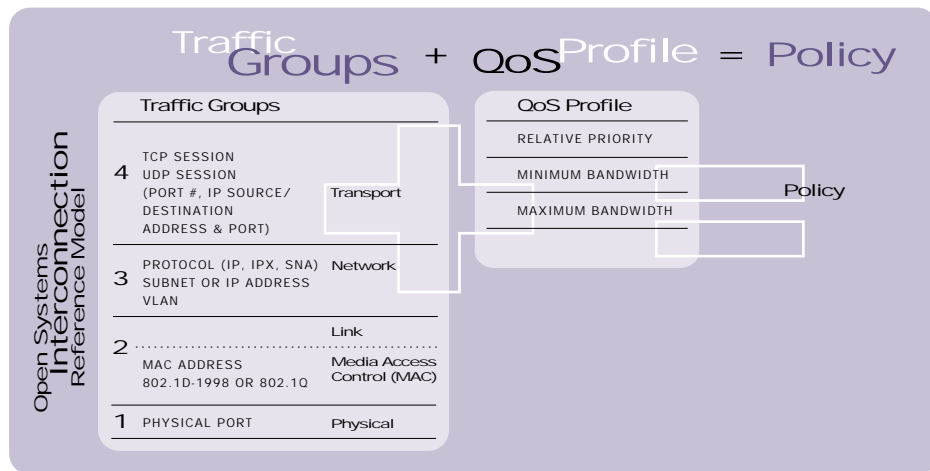
**Time-critical or 24 x 7 applications**  
Over-provisioning is not enough.  
There must be a guarantee.

**Logical topology changes (routing protocols)**  
Changes will affect traffic patterns and cause sub-optimal bandwidth situations

## Traffic Groups

Policies are sets of high-level rules that determine how network resources are allocated to applications. QoS policies identify traffic groups and define QoS profiles that specify the level of service that those traffic groups receive.

Traffic groups may include all of the manufacturing subnet, for example, within a server load balancing group, or all web servers or a power workgroup or even one particular server supporting a mission-critical application. Through a browser-based management tool the network manager decides ahead of time which traffic groups have priority and under what circumstances.



*For each traffic group, the network manager can define a QoS profile that allocates minimum and maximum bandwidth and relative priority.*

## QoS Profiles

QoS profiles based on these parameters enable network managers to specify the treatment that each traffic group will receive:

- **Policies that guarantee a minimum amount of bandwidth** For example, IPX traffic can be placed in a queue that is guaranteed no less than 25 percent of the available bandwidth.
- **Policies that guarantee a maximum amount of bandwidth** For example, non-routable traffic assigned to a port class can be placed in a queue that is allowed to consume up to 20 percent or less bandwidth. Maximum bandwidth is guaranteed as a maximum average bandwidth. Short bursts of data may actually exceed the maximum. However, the system will contain bursts to an average over time.
- **Policies that set relative priority to establish a pecking order for traffic** For example, a low priority might be given to routine data backups. Although data backup traffic is allowed to consume its reserved minimum bandwidth, when additional bandwidth is available it will be allocated, up to the assigned maximum, to all traffic groups based on priority.

## A Framework for Implementing Policy-Based QoS

Extreme Networks' Summit stackable and BlackDiamond chassis switches offer bandwidth with intelligence.

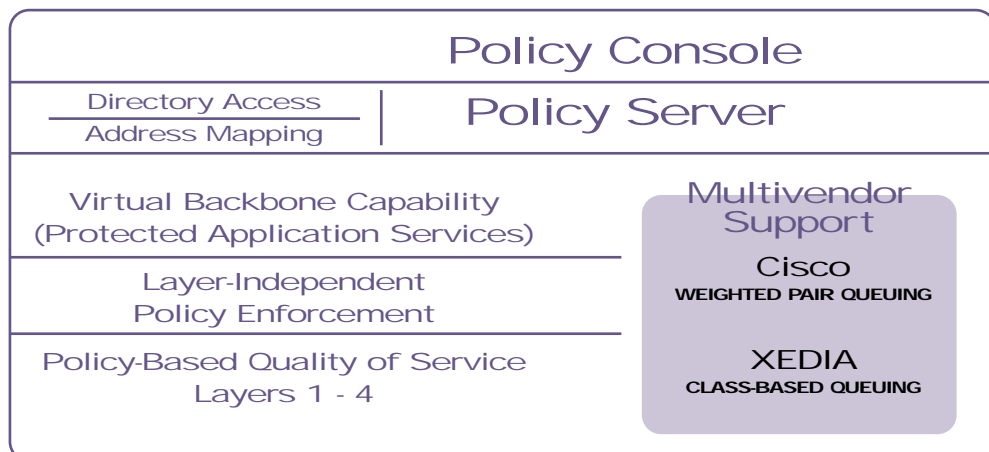
As part of its ExtremeWare Enterprise Manager software suite, Extreme Networks offers policy configuration tools that meet the application-specific needs of today's networks. Policies are easily configured and enforced to protect the performance of mission-critical applications.

These policy management tools, provide the following benefits:

- Simple web-based policy console tool for deploying protected user and application services
- An intelligent policy server for configuring end-to-end network quality of service (QoS) policies
- Policy editing and maintenance
- Policy conflict-resolution for setting policy precedence
- Policy scoping to define where policies are installed and enforced
- Dynamic user policies with Dynamic Link Context System (DLCS), which maps user and desktop device names to lower-layer network information
- Multivendor policy configuration
- Open standards-based implementation with the Common Open Policy Service (COPS) protocol

## Components of Extreme's Policy Management

Extreme's policy management tools consist of the policy console, a policy server, layer-independent policy enforcement, Extreme Networks switches, and multivendor support. Together, these elements deliver protected user and application services for your network. All of these elements are managed utilizing Extreme's policy management.



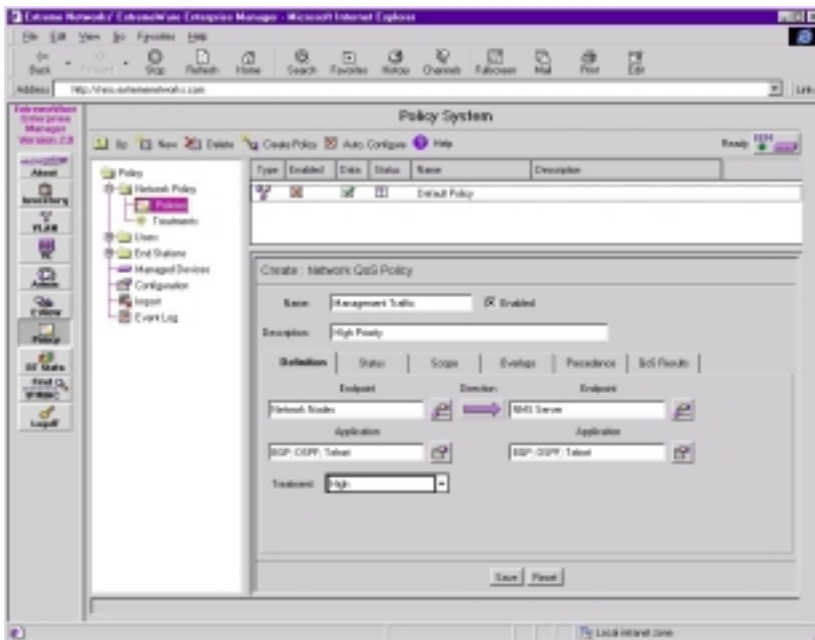
*Extreme's policy management tools consist of the policy console, a policy server, layer-independent policy enforcement, Extreme Networks switches, and multivendor support.*

## Policy Console

The policy console is a simple web-based user interface for configuring quality of service policies. This tool supports key aspects of policy creation such as standard policy editing, application and service definition, and policy scoping. The policy scoping function determines where policies will be installed.

The policy console leverages a wizard-like interface to intuitively guide you through the process of creating QoS policies for:

- Application server traffic – i.e. all traffic going to and from an e-commerce site
- Client/server traffic – i.e. ERP system activities
- Control traffic – i.e. limiting broadcast and multicast on the network
- User group policies – i.e. the engineering group has medium priority and a 20% minimum bandwidth
- Network-wide policies – i.e. all web traffic will not exceed 10% of overall capacity



*The policy console is a simple, web-based user interface that intuitively guides you through the process of creating QoS profiles.*


## Policy Server

The policy server is responsible for the installation and tracking of QoS policies on the network. This includes policy interpretation from high-level definition into low-level device configuration, policy conflict-resolution, device communication, policy installation and policy storage.

## Multivendor Support/Policy Conflict-Resolution

Today's networks are a heterogeneous mix of switches and routers from multiple vendors. Extreme's multivendor support ensures that intelligent policy creation and enforcement works across the entire network, and not just on switches from Extreme Networks. Extreme's policy management offers multivendor policy configuration for all Extreme Networks' switches, Lucent Internet Access and VPN routers, and select Cisco switches and routers.





On multivendor networks, policy conflict-resolution is critical. When two policies collide, there must be a mechanism to set precedence between the two. For example, if there is a conflict between a policy based on a destination subnet, and a policy based on a specific host on that subnet, which policy gets implemented?

Extreme's policy management provides tools for setting policy precedence. As a result, if two policies are in conflict, Extreme's policy management intelligently assigns one policy priority over another.

### **Dynamic Link Context System**

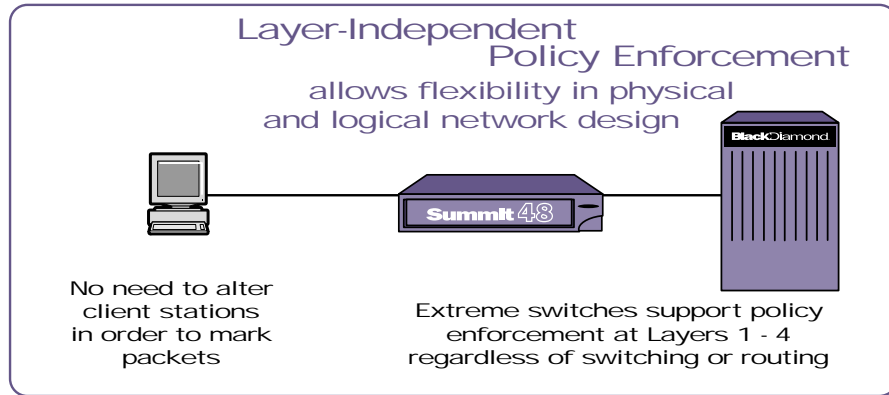
The Dynamic Link Context System (DLCS) is an integral component of Extreme's policy management, allowing Extreme Networks switches to resolve how a physical link is being used. By determining what users and workstations are attached to individual physical ports on the device, a "link context" is developed. This link context provides the information necessary to specify policies using high-level identifiers such as user and device names.

Unlike many policy implementations, the Dynamic Link Context System allows Extreme Networks switches to automatically map user and machine name, to dynamic low-layer information, such as IP address or physical port. This allows the Dynamic Link Context System to support the dynamic tracking of policies by user name. As a user physically moves within the network or has an IP address changed, policies will change to reflect the new physical port or IP address.

The first implementation of the Dynamic Link Context System supports Windows Internet Naming Service (WINS) snooping. By snooping the Windows Internet Naming Service, the Dynamic Link Context System eliminates the need to install special policy services, like enhanced DHCP servers, in the network environment. The Dynamic Link Context System stores user and device names in a resolution database and automatically maps them to an IP address, MAC address, physical switch port and switch. The resolution database is the key to defining roving user policies.

The Dynamic Link Context System reduces the complexity of policy development. The associations between high-level elements and low-layer network information are automatically made. This lowers administration costs by allowing policies to react to the changing network environment without user intervention.

## Layer-Independent Policy Enforcement



*Layer-independent policy enforcement empowers the network designer to build a network without having policy networking requirements dictate whether an Extreme switch is used at Layer 2 or Layer 3.*

Extreme Networks delivers Policy-Based Quality of Service for Ethernet networks through the combination of its non-blocking switch hardware architecture and its ExtremeWare software suite. A key feature of this software is layer-independent policy enforcement.


Layer-independent policy enforcement means that Extreme Networks' Summit and BlackDiamond switches can enforce policies at Layers 1-4 regardless of whether they are configured as Layer 2 switches or Layer 3 routers. This unique capability is the result of the Extreme Networks Policy-Based QoS architecture that allows any switch to enforce policies at Layers 1-4. This means an Extreme switch operating in Layer 2-only mode can enforce policies at Layers 1-4.

## Classifying Traffic

In order to enforce policies traffic must first be classified. Packets are inspected and sorted into traffic groups based on identifying information in packet headers.

Defining Traffic Groups with ExtremeWare's Policy-Based QoS			
Layer	When policies are based on:	Define traffic groups by:	OSI Layer
4	Applications	TCP Session (source port/destination address)  UDP Session (source port/destination address)	Transport Layer
3	Topology or Groups of Users	RSVP Flow  Protocol (IP, IPX, SNA) Subnet or IP Address VLAN	Network Layer
2	Individual Station Applications	MAC Address 802.1p or 802.1Q	Link Layer
1	Topology	Physical Port	Physical Layer

*Extreme switches can classify traffic based on Layer 1-4 information. This gives network managers the flexibility to identify traffic groups and implement policies based on a variety of criteria.*



Extreme switches can classify traffic at Layers 1-4 regardless of whether they are configured as Layer-2 switches or Layer-3 routers.

Once classified, packets can optionally be marked with prioritization codes appropriate for the class. All packets are placed into the output queues that correspond to each class. Queuing mechanisms can then provide the QoS required for each traffic class.

### **Marking Traffic**

After traffic is classified it can optionally be marked with codes that indicate the handling that they are to receive from routers and switches. Industry-standard packet markings include:

- Layer 2 – 802.1p
- Layer 3 – IP Type of Service, Diffserv

Once packets are marked, legacy routers and switches with no access to policy information can give them the appropriate levels of priority. Marking can also speed up processing by policy-aware devices because they can bypass the classification procedure and simply act based on the markings. It should be noted that while traffic classification slows down most routers and switches, Extreme switches are capable of classifying traffic and marking traffic at wire speed.

Extreme’s “i” series switches perform prioritization and bandwidth management based on Layer 1-4 implicit packet information or by using explicit 802.1p packet markings. They also support the observance and wire-speed manipulation of DiffServ, also known as IP ToS bits, and code points as a mechanism to determine end-to-end LAN/WAN class of service.

Extreme’s “i” series switches support eight individual hardware queues per physical port, which may be independently configured using prioritization and bandwidth management parameters. As with all Extreme switches, you can use the implicit Layer 1-4 classification or explicit 802.1p and DiffServ classification with the “i” series switches. These classifications can be assigned QoS profiles that perform bandwidth management, including minimum bandwidth for rate guarantees and maximum bandwidth for rate shaping and policing.

The ability to combine Policy-Based QoS with other capabilities like web cache redirection, wire-speed server load balancing and Wire-Speed IP Routing makes the Extreme Networks solution truly unique and far simpler to manage.

## Using Prioritization to Enforce Policies

Associated with each Extreme switch port is a set of queues. A queuing mechanism called weighted fair queuing provides priority and bandwidth controls, such as minimum and maximum bandwidth.

The queue length/size is dynamically allocated out of the shared central memory switch architecture. Each queue is assigned a priority relative to the other queues on the port. Using these queues, QoS profiles can be defined for each traffic group on every Extreme switch.

The types of queuing mechanisms used by routers and switches directly affects the quality of service that they are able to deliver. For example, simple round-robin queuing gives all traffic classes fair access to the network, but does not allow any prioritization. Strict priority queuing, in contrast, does enforce priorities, but during congestion all traffic classes except the highest priority can be completely shut down. This can result in long delays and jitter, and lower priority packets may be dropped as their queues fill up.

The weighted fair queuing technique used in Extreme switches is capable of enforcing priorities while allowing lower priority applications enough bandwidth to provide users with acceptable performance. Weighted fair queuing can also be used to allocate minimum or maximum amounts of bandwidth.



## How Queuing Affects Bandwidth Management

A variety of queuing techniques are currently implemented in networking products to manage network congestion. They differ in the ways they are implemented and in the levels of service they can provide to network users. Some of these queuing schemes give all traffic equal service, while others provide relative prioritization, and still others can allocate bandwidth to each traffic class.

The following is a brief summary of key queuing techniques and their capabilities.

### First-In, First-Out Queuing

All packets are sent in the exact sequence that they were received. First-in, first-out queuing is simple to implement, but it is not able to prioritize traffic or allocate bandwidth to specific flows. High priority packets could be stuck behind lower priority data.

### Round Robin Queuing

Round robin queuing classifies and queues traffic based on packet content. Each queue is serviced in turn to give each class of traffic an equal level of service. Round robin queuing provides users with fair access to network bandwidth, but it does not prioritize traffic or allocate bandwidth to specific flows.

### Priority Queuing

Priority queuing also classifies and queues traffic based on packet content. Queues are assigned priorities ranging from high to low priority. Higher priority packets always take precedence over lower priority traffic. During periods of heavy congestion all but the highest priority traffic may be completely blocked by high priority flows.

### Weighted Fair Queuing

Weighted fair queuing prioritizes traffic flows while ensuring that lower priority flows receive an acceptable level of service. Traffic is classified and put into queues that are assigned relative weights (priorities). The queues are serviced on a round-robin basis and the amount of data removed from each queue is in proportion to the relative weight of the queue. Higher priority queues get a larger portion of the bandwidth, but even the lowest priority queue always receives some bandwidth. This reduces the possibility of an application time-out. If a queue is empty, data will be sent from the next queue that has data to send.

Weighted fair queuing also improves the performance of low volume traffic flows, such as interactive applications and some streaming media formats, by queuing them ahead of high volume flows.

### Class-Based Queuing

Class-based queuing goes beyond prioritization by providing assured bandwidth to traffic flows. Traffic is first classified and put into queues. Each queue is assigned a number of bytes that will be forwarded each time the queue is serviced. The queues are serviced on a round-robin basis and the assigned amount of data is forwarded from each queue. Each traffic class is therefore allocated a percentage of bandwidth on the outbound link and may burst above their allocated bandwidth if the bandwidth is not being used by other traffic classes.

## Web Cache Redirection

The performance of a networked application is determined not only by the infrastructure, but also by speed at which servers can satisfy user requests for information. Many enterprises and service providers are using caching servers to improve the performance of applications and optimize WAN usage. Caching devices offload WAN links and servers by storing content that is accessed repeatedly. Requests for such content can then be redirected by the switch to high-performance caching devices, which spares servers and any associated WAN links from repetitious queries.

For content providers, caching significantly lowers repetitive hits on servers and allows content-rich pages to be served more quickly to customers. For service providers, this lowers the cost of WAN bandwidth consumption outside the point of presence (POP) because information can be cached within the POP facility. And for enterprise networks, frequently accessed content stays local, thus conserving WAN bandwidth and its associated costs.

Extreme Networks complements its bandwidth management capabilities by building web cache redirection into its “i” series switches. Extreme delivers web cache redirection with a twist – we do it transparently and at wire speed. Traffic is redirected at wire speed using Layer-4 criteria, such as http Port 80, to one or more load-shared ports across several web cache servers. All this occurs transparently, which means users do not need to reconfigure browser applications.

User session persistence can also be maintained in IP source and destination addresses. By employing persistence on IP destination addresses, service providers and web content providers can optimize repetitive web hits to specific content.

### Web Cache Redirection “Firsts”

Extreme’s wire-speed redirection capability is the first to be integrated with wire-speed Layer 2 switching and Layer 3 routing on a single hardware platform. This integration allows for effective web caching integration without radical changes to the network design.

Additionally, the point in the network where transparent redirection creates higher efficiency is often an aggregation point where traffic loads are excessive. Therefore, the device must have awesome data handling capabilities, such as the Summit7i’s 64 Gbps non-blocking switch fabric.

Extreme is also the first to combine web cache redirection and offer eight hardware queues per port, thus combining true bandwidth management to and from cached services. Bandwidth management allows the setting of eight independent classes of service consisting of reserved bandwidth minimums, specified bandwidth maximums and eight classes of prioritization on every port.

## QoS Management Does Not Have to be Complicated

Network managers are sometimes intimidated by the apparent complexity of QoS management. This is often due to the confusing array of technologies, buzzwords, and products that surround QoS and bandwidth management. But it doesn't have to be complicated if you adhere to a few basic rules.

First, the foundation of your QoS management initiative should be policy-based management, which allows you to focus on application and business requirements rather than the inner workings of QoS mechanisms and configurations of routers and switches. Extreme's policy management allows you to define priorities and bandwidth allocations for specific applications, stations, or groups of users via an easy-to-use web-based graphical interface.

Next, use a top-down approach to define policies. It may be tempting to develop a large set of very granular policies that address every combination of users and applications running on your network, but this is complicated and usually unnecessary.

In fact, QoS management is like most endeavors – you can achieve 80% of the desired results by expending 20% of the effort. Start with broad policies applied to your most important applications and users. A simple policy like giving high priority to SAP servers is almost trivial to implement, but can yield significant results. Policies can then be fine tuned after evaluating the initial results.



*A very effective QoS strategy is to first allocate bandwidth to broad classes of traffic and then prioritize traffic within each traffic class. For example, 10% of bandwidth could be allocated to ERP applications, 2% to Voice over IP, 15% to remote conferencing and streaming media, and the remainder shared by all other traffic.*

This first step effectively isolates these traffic classes from one another. For example, Voice over IP is protected from congestion due to bursty file transfers. Prioritization could then be used within a virtual backbone, for example, to give some critical ERP applications priority over less important ERP applications.

The bottom line is that with good policy-based tools, and a step-by-step, top-down approach, QoS management can be simple and very effective.



3585 Monroe Street Santa Clara, CA 95051-1450 Phone 408.579.2800 Fax 408.579.3000  
Email [info@extremenetworks.com](mailto:info@extremenetworks.com) Web [www.extremenetworks.com](http://www.extremenetworks.com)

©2000 Extreme Networks, Inc. All rights reserved. Extreme Networks and BlackDiamond are registered trademarks of Extreme Networks, Inc. in certain jurisdictions. ExtremeWare, ExtremeWare Vista, ExtremeWorks, ExtremeAssist, ExtremeAssist1, ExtremeAssist2, PartnerAssist, Extreme Standby Router Protocol, ESRP, SmartTraps, Summit, Summit1, Summit4, Summit4/FX, Summit7i, Summit24, Summit48, Summit Virtual Chassis, SummitLink, SummitGbX, SummitRPS, Velocity, the BlackDiamond logo and the Extreme Networks logo are trademarks of Extreme Networks, Inc., which may be registered or pending registration in certain jurisdictions. The Extreme Turbodrives logo is a service mark of Extreme Networks, Inc., which may be registered or pending registration in certain jurisdictions. All other registered trademarks, trademarks and service marks are property of their respective owners. Specifications are subject to change without notice.