

JUNIPER NETWORKS EX SERIES ETHERNET SWITCHES: QOS-ENABLING THE ENTERPRISE

Assuring End-to-End Application Performance

Table of Contents

Executive Summary	1
Introduction	1
QoS Baseline Requirements	2
The Need for Queue Granularity	2
The QoS Tool Box: Classifying and Marking Traffic	3
Consistent Traffic Handling	4
Smoothing Flows with Traffic Shaping	4
Juniper Networks QoS Savvy Switches	5
Conclusion—Juniper Networks QoS Advantage	7
About Juniper Networks.	7

Table of Figures

Figure 1: All Juniper Networks EX Series switch platforms provide eight queues per port, giving enterprises the flexibility to accommodate numerous classes of traffic.....	5
Figure 2: JUNOS Software utilizes a single source code, follows a predictable release train, and employs a modular architecture to ensure the consistent implementation of control-plane features across product lines.	6

Executive Summary

The enterprise network is no longer a data-only transport. It has evolved into a multi-services infrastructure carrying business-critical applications, voice and video streams, and traffic from building automation and other control systems. The emergence of unified communications further spurs the trend toward converged networks.

IT's challenge is to ensure predictable performance for all applications and traffic types sharing the IP infrastructure. The right set of quality of service (QoS) tools is critical to meeting this challenge. QoS tools must be sufficiently granular to accommodate a rich mix of traffic types. Network switches and routers deployed at every level in the network must support a complementary set of QoS mechanisms to ensure consistent, predictable traffic handling end-to-end. And the QoS tools must be easy to use so that operations overhead doesn't balloon as IT begins implementing QoS policies.

To simplify the job of creating a converged infrastructure, Juniper Networks® designed its EX Series Ethernet Switches with a comprehensive set of QoS features providing granular controls that can be implemented consistently and easily across the entire network. With the QoS capabilities built into the EX Series switches, enterprises can be assured of predictable performance for current as well as future applications.

Introduction

One of the most significant IT trends is the push by enterprises to combine data, voice and video traffic on a single multiservice network. To date, the adoption of IP-based telephony (IPT) and video—including videoconferencing, video streaming and video-on-demand—has been the primary driver for converged networks. Many enterprises are also moving traffic from security cameras, badge readers, environmental controls and other building automation systems off of proprietary networks and onto their IP infrastructure. Looking ahead, unified communications (UC) will provide the next major impetus for building multiservice networks.

Analysts point to unified communications as the natural evolution of IPT, blending voice, instant messaging (IM), and other communications technologies with collaboration tools to enable new business productivity. According to Gartner, "By 2010 communications will be provided from within major enterprise IT applications (eCommerce, CRM, ERP, HR packages)."¹

To accommodate these mixed media applications, IT is tasked with building a network infrastructure capable of supporting converged real-time communications. Such an infrastructure must provide consistent end-to-end quality of service to ensure that each traffic type gets the bandwidth and handling it needs to meet user expectations and business objectives.

Understanding these requirements, Juniper Networks designed its new EX Series Ethernet Switches with a comprehensive set of QoS features, making it easy to deploy the same QoS controls across the entire network, from the wiring closet to the WAN edge. With the Juniper Networks EX3200 fixed-configuration switches, the EX4200 switches with Virtual Chassis technology, and the EX8200 switches, IT can deliver the predictable performance needed to support today's converged voice, video and data applications, as well as building automation traffic and emerging UC-enabled applications.

¹Gartner: The Big Picture: IP Telephony Leads Communications-Enabled Business Processes; Publication Date: 7 November 2007 ID Number: G00150905 - by Geoff Johnson

QoS Baseline Requirements

QoS technologies can be quite esoteric. However, evaluating the QoS capabilities of network devices needn't be complex. When building a converged network, IT should look for the following QoS capabilities:

- **Sufficient queue granularity to accommodate current and emerging applications:** Different types of traffic must be handled differently in order to perform optimally. Therefore, network switches should support enough queues per port to accommodate all of the various classes of traffic, and at least two algorithms for servicing queues to ensure differentiated handling for each traffic class. For example, some vendors support priority queuing for latency-sensitive applications such as IPT and videoconferencing where allocated bandwidth can be quantified and enforced, and some form of round-robin queuing that provides qualitative handling for data applications.
- **Consistency across network devices:** The goal of QoS technology is to deliver predictable application performance throughout the network. However, achieving this goal is only possible if every switch and router in the network supports a common set of QoS capabilities. For example, if an access switch supports four queues per port while an aggregation switch supports two, the granularity of traffic handling will be lost as multiple classes of traffic from the access layer are squeezed into two "generic" classes at the aggregation layer.

Likewise, if access switches operate at Layer 2 and use Ethernet-based QoS markings, routers and other devices operating at Layer 3 may ignore those markings and repeat the process of classifying and marking each packet using IP-based QoS markings. Some IT organizations may choose to re-mark packets at administrative boundaries. However, inconsistent QoS functionality across network devices shouldn't force such re-marking, which can result in unpredictable behavior as devices re-evaluate the type of handling a packet needs at each hop and how that handling maps to the device's particular QoS capabilities. For the most predictable performance, each class of traffic should be marked as close to the source as possible—preferably by the access switch—and handled in exactly the same way at each hop along the path.

- **Tools that simplify QoS configuration and management:** In some cases, IT may wish to dive down and configure individual QoS parameters. However, most enterprises will benefit from predefined templates that simplify configuring QoS across their network. Juniper, for example, offers templates for key traffic classes so that IT need only identify which switch ports are being used for IP phones, PCs or uplinks; the switch will automatically set the appropriate QoS parameters.

Building a converged network with equipment from a single vendor can reduce administrative overhead if the switches share a common set of QoS capabilities and configuration tools. If each platform has a different set of QoS features and means of configuring them, IT will find its job unnecessarily complicated.

The Need for Queue Granularity

Applications must perform predictably to be usable. As enterprises prepare their networks for voice and video services, the need for QoS to support these types of traffic is clear. Analysts at the Burton Group, for example, note that "it cannot be stressed strongly enough that the ability to effectively manage IPT performance is the key to meeting user expectations."

According to the Burton Group, for an IPT implementation to be successful, the LAN must provide sufficient baseline bandwidth to support a call, as well as uniformity of bandwidth so that bandwidth is always available. "An organization that considers high-quality voice performance essential to its business will likely make whatever investments are necessary to implement sufficient QoS guarantees to ensure acceptable performance levels," notes Burton Group analyst Irwin Lazar.²

But voice isn't the only class of traffic that needs QoS guarantees. Many types of traffic suffer when large file or bulk data transfers or other bandwidth-hogging applications dominate LAN links, causing congestion at oversubscribed aggregation points. Network control protocols (for example, Spanning Tree BPDUs, OSPF route updates), network management information, transaction-based enterprise applications and video can also be squeezed, resulting in potential data loss and network instability.

In building a converged infrastructure, enterprises need network devices with sufficiently granular QoS to ensure predictable performance for all traffic classes. Many organizations will find they have six or more classes of traffic that require differentiated handling based on their importance to the enterprise and their bandwidth, loss and delay requirements.

²"Preparing Your Network for IP Telephony" – August 2006

For example, voice and videoconferencing have distinct and clear-cut bandwidth and latency sensitivities that will become more nuanced as unified communications applications are rolled out, indicating the need for at least two QoS classes to accommodate these traffic types. Network management information, which is crucial for network health, must be considered a third class of traffic requiring expedited handling. Organizations that have folded traffic from building automation systems onto their IP infrastructure may want to place transmissions from physical plant security systems, environmental control systems and the like into a fourth QoS class.

Many enterprises will also want to divide data traffic into two or three broad classes to ensure that business-critical applications, such as order entry and reservation systems, aren't compromised by less time-sensitive traffic such as bulk file transfers or email. By establishing a multi-tiered traffic model, for example, enterprises can mark those applications for which there's a response-time expectation with a high-priority "gold" designation while "best effort" email traffic is assigned a "silver" classification. "Bronze" could be reserved for non-business-related Internet traffic (for example, YouTube or Facebook) that may be allowed on the network but is considered the lowest priority.

How an organization classifies its traffic is a business decision. IT's job is to provide an infrastructure capable of distinguishing the different traffic classes and giving each the network resources required for optimum performance or—in the case of worms, denial of service attacks and other unwanted traffic—the ability to prevent undesirable traffic from consuming an inordinate amount of bandwidth. In addition to providing "positive" traffic handling, enterprises can use QoS to throttle anomalous flows that may indicate an attack.

QoS tools allow IT to manage bandwidth, delay, jitter and loss parameters. The right tools are crucial for building a multi-services network capable of handling an organization's evolving traffic mix.

The QoS Tool Box: Classifying and Marking Traffic

QoS can be broken into three broad functions: identifying traffic, marking traffic and applying special handling to a flow based on its marking. Marking can be done at Layer 2, Layer 3 or both, depending on whether a network device supports Layer 2 and/or Layer 3 functions. Some vendors, such as Juniper Networks, also allow IT to specify QoS handling based on other criteria, such as the port a packet is traversing, a packet's MAC or IP address, a TCP/UDP port number, or a combination of any of these parameters. How granular IT can get in defining QoS policies will depend on the number of access control list entries (ACEs) a vendor supports. The higher the number of ACEs, the more control IT will have in distinguishing how the network should handle, say, Oracle and SAP application traffic relative to IP telephony and email marked as high priority.

At Layer 2, the IEEE 802.1p specification defines a way to use Ethernet virtual LAN (VLAN) tags to distinguish eight classes of service. At Layer 3, QoS markings use bits within the type of service (TOS) byte in the IP header. Originally, the three IP Precedence bits within the TOS byte were set aside for marking QoS. This method has been superseded by the Differentiated Services (DiffServ) Code Point (DSCP) field, a six-bit area within the TOS byte which allows for more granular QoS while also being backwards-compatible with IP Precedence.

With these mechanisms, traffic coming into an edge switch or router can be identified, matched against a QoS policy list, and marked for handling by subsequent network devices. Industry best practices indicate that traffic marking should be done as close to the traffic source as possible, preferably in the access switch. This practice ensures IT controls packet marking rather than users, applications or devices such as IP phones. While it may seem convenient to allow end nodes to mark packets directly, this approach can lead to abuses, such as savvy users marking all of their traffic as "high priority."

Various QoS best practices guidelines recommend that enterprises use Layer 3 DSCP markings end-to-end to ensure consistent traffic handling across the network. Routers, for example, will ignore Layer 2 markings and remark the packet with a DiffServ code point.

In a mixed environment of Layer 2 and combined Layer 2/Layer 3 switches, combination switches may be able to map between 802.1p and DSCP markings. However, re-markings and mappings between markings is highly inefficient and can make it difficult to achieve predictable application performance. Having Layer 3-aware devices in the access, aggregation and core layers of the network makes it possible to consistently mark, and therefore consistently handle, all classes of traffic as they traverse the network.

Consistent Traffic Handling

There are a variety of QoS mechanisms that can be applied to network traffic as it traverses a network device. These mechanisms can be broadly grouped into three categories: admission control, queuing and congestion management.

Admission control determines whether traffic is allowed onto the network. Most switch vendors support the IEEE 802.1X standard for port-based admission control on the LAN. Users must be authenticated—by presenting valid credentials such as a password, one-time token or digital certificate—and their devices must be in compliance with corporate policy on endpoint security posture before they gain access to the network. If authentication or a host posture check fails, an 802.1X-compliant switch will block network access. 802.1X can be seen as both a security and QoS mechanism. Not only does it prevent unauthorized network access, but 802.1X can help improve network uptime—and therefore application availability—by helping to keep bandwidth-consuming malware off the network and limiting what resources users can access once admitted to the network.

Queuing is a key traffic handling mechanism. According to best practices guides, the only way to provide service guarantees is to enable queuing at any node that has the potential for congestion, regardless of how rarely this may occur. And for predictable performance, network devices should support the same level of queuing at each level of the network.

As we noted earlier, QoS granularity is directly tied to the number of queues per port supported on a network device and the types of algorithms used to move packets out of queues. LAN-based network devices typically support two, four or eight queues per port, while WAN gear often supports up to 64 queues per port. Given the current mix of traffic on converged networks and the emergence of UC, enterprises would be wise to deploy eight queues per port at the access, aggregation and core layers of the network.

Vendors have defined various queuing algorithms to accomplish different goals. Priority or “strict” priority queuing, for example, ensures the lowest latency and is typically used for IPT, audio conferencing and other voice-enabled applications as well as certain types of video, such as videoconferencing. Priority queuing works by servicing queues sequentially, from highest priority to lowest priority. Once the highest priority queue is empty, packets from the next lowest queue are serviced, and so on.

Priority queuing is fairly widely supported by vendors, as is some form of class-based queuing, which encompasses schemes known as round robin, weighted round-robin and deficit round-robin. Whereas priority queuing is generally used by latency-sensitive applications, class-based queuing schemes are used to prioritize the forwarding of all other traffic on a port.

With simple round robin queuing, forwarding priority is allocated to subchannels on a rotating basis. With weighted round-robin, packets are placed in queues according to their class and serviced based on their priority. For example, four packets of “gold” traffic might be serviced for every two packets of “silver” and every one packet of “bronze” traffic.

Smoothing Flows with Traffic Shaping

Queuing is a way of managing traffic flows within a device; vendors also support mechanisms for managing traffic flows into and out of their equipment. These mechanisms allow IT to better manage bandwidth and alleviate congestion. Rate limiting, for example, can be used to throttle traffic as it enters or exits a device.

While its primary purpose is to prevent an application from over-consuming bandwidth, rate limiting can also help prevent malicious traffic from compromising network performance. For example, if a switch port has been configured to expect traffic from an IP phone at about 85 kilobits/second (kbps), and it receives a flow above 200 kbps, it can be assumed that something other than IP telephony traffic is being sent. The switch will drop traffic above the set rate limit, thus preventing potentially malicious traffic from blasting across the network to overwhelm or attack server resources.

Congestion management mechanisms are another means to control traffic flows. With one such mechanism, known as tail drop, the switch drops arriving packets as queue buffers become filled or begin to overflow. Policing mechanisms may also come into play if congestion occurs.

Policers determine whether a flow is adhering to its assigned traffic rate and either drop packets outright if they’re in violation or, if there’s excess network capacity available, mark them for handling further in the network. Should congestion occur further along, network devices will discard packets based on their drop precedence.

QoS best practices include policing traffic close to its source. This helps to conserve network resources by reducing the transmission of traffic that will be dropped subsequently. Policing also helps prevent the spread of worms, denial of service (DoS) attacks and other malware by throttling or shutting down excessive flows.

Juniper Networks QoS Savvy Switches

Juniper Networks designed its new EX Series Ethernet Switches with a granular, consistent set of QoS capabilities that ensure enterprises predictable application performance across any combination of traffic types. Running the same Juniper Networks JUNOS® Software as Juniper Networks routers, the fixed-configuration EX3200, EX4200 with Virtual Chassis technology, and the Terabit-chassis EX8200 switches enable enterprises to build a converged network knowing that a consistent set of QoS mechanisms is guaranteeing service levels from the desktop to the data center and beyond.

Juniper switches address all key QoS requirements, including:

- **Granularity:** All Juniper Networks EX Series switches provide eight queues per port, giving enterprises the flexibility to accommodate numerous classes of traffic. In addition, Juniper supports 7,000 access control list entries (ACEs) per EX3200 and EX4200 switch, and up to 64,000 per EX8200 switch, giving IT the flexibility to define very granular QoS policies. With a Juniper infrastructure, enterprises can deploy IP telephony, video and emerging unified communications applications, while ensuring that business-critical enterprise applications, network management information and building automation traffic aren't compromised.

Juniper routers and switches support latency-sensitive video and voice applications with strict priority queuing. Whereas many vendors support only one priority queue per port, Juniper lets IT define multiple priority queues per port to ensure that voice- and video-enabled applications get the proper handling. For data and other traffic, Juniper provides shaped deficit-weighted round robin queuing, granting priorities to certain traffic/queues over others via the assignment of different weights to the various queues, while preventing bandwidth starvation of lower priority queues by the higher priority ones.

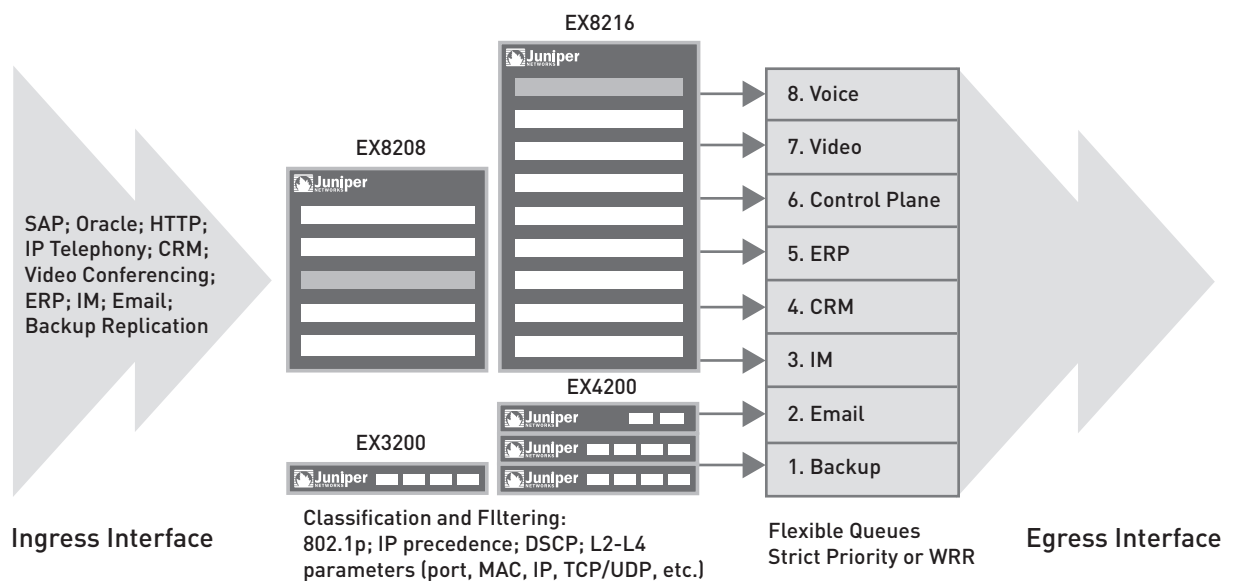


Figure 1: All Juniper Networks EX Series switch platforms provide eight queues per port, giving enterprises the flexibility to accommodate numerous classes of traffic.

In addition, JUNOS delivers additional traffic controls through rate limiting; tri-color traffic marking and policing for congestion control; and tail drop for congestion avoidance.

- **Consistency:** Juniper supports the same set of granular QoS controls on all of its EX Series switches, ensuring consistent traffic handling end-to-end. Each EX Series switch includes a full suite of Layer 2 and Layer 3 switching capabilities as part of the base license; the EX3200, EX4200, and EX8200 switches come standard with both an application-specific integrated circuit (ASIC)-based packet forwarding engine and a Routing Engine for complete control plane functionality.

As a result, enterprises have the option of using either Layer 2 or Layer 3 QoS mechanisms. At Layer 2, Juniper supports 802.1p marking and queuing and honors these markings as packets pass from switch to switch. At Layer 3, Juniper can mark packets using either the DSCP or IP Precedence bits and can classify, rewrite and queue packets based on these markings. Beyond these standard marking techniques, Juniper gives enterprises the flexibility to classify and mark traffic based on its ingress port, IP or MAC address, VLAN tag, TCP/UDP port number, or any combination.

In addition, each Juniper switch platform runs the same JUNOS Software, which ensures a consistent implementation and operation of each control plane feature across the entire Juniper product line.

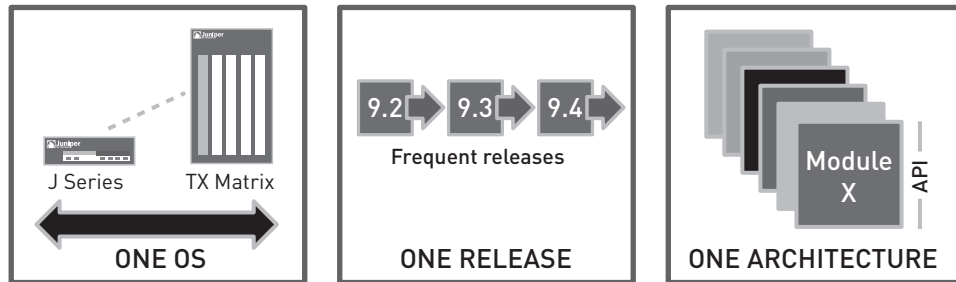


Figure 2: JUNOS Software utilizes a single source code, follows a predictable release train, and employs a modular architecture to ensure the consistent implementation of control-plane features across product lines.

- Ease of management:** Juniper’s support for a consistent set of QoS capabilities, a consistent operating system—JUNOS—and control plane, and a single management interface across all of its switch platforms greatly simplifies IT’s job. Rather than having to learn how to configure and manage QoS for each individual class of platform, IT can use the same configuration and management tools across all EX Series switches and achieve the same, predictable behavior from platform to platform.

To further simplify QoS implementation, the EX Series switches include templates that automatically configure QoS, security and other parameters based on the type of device connected to a port. Five preconfigured profiles are available: desktop port, desktop plus IP phone port, wireless access point port, routed uplink port, and Layer 2 uplink port. IT also has the option to create custom profiles and apply them through the JUNOS command-line interface (CLI), the J-Web device management system or the Juniper Networks Network and Security Manager (NSM). In addition, Juniper supports the use of access control lists (ACLs) to set QoS policies, with thousands of ACL entries available per switch.

The new EX Series switches also fully integrate into the Juniper Networks Unified Access Control (UAC) to provide standards-based 802.1X port-level access control as well as Layer 2-4 policy enforcement based on user identity or location. As part of the authentication process, the switches can configure VLAN assignments, enforce QoS and security ACL rules based on individual user identity and credentials, keep users from accessing specific applications through dynamic filters, and rate-limit user traffic to protect network resources from over consumption.

Conclusion—Juniper Networks QoS Advantage

Building a multiservice network is a strategic initiative that shouldn't be stymied by QoS pitfalls. Juniper Networks designed the EX3200, EX4200 and EX8200 switches to simplify creating a converged infrastructure. The EX Series switches enable enterprises to deliver predictable performance for any combination of traffic with the least possible administrative overhead.

By supporting a common set of QoS features from the access layer to the core, Juniper allows enterprises to follow best practices for ensuring consistent end-to-end QoS. For example, robust QoS features in access devices (whether the EX3200 or EX4200) means that network traffic can be classified and policed at the source. Likewise, standard Layer 3 support on all Juniper switch and router platforms means enterprises have the option to use DSCP marking end-to-end.

And by supporting the modular JUNOS control plane and a common management interface, Juniper switches greatly reduce operational overhead, shortening the learning curve for IT and streamlining network architecture and design.

Juniper Networks EX Series switches provide enterprises with the granular, consistent QoS capabilities they need to ensure predictable performance for current as well as future applications.

About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at www.juniper.net.

Corporate And Sales Headquarters

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

APAC Headquarters

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

EMEA Headquarters

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

Copyright 2009 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, JUNOS, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. "Engineered for the network ahead" and JUNOSE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

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

