

Advanced QoS Services for the Intelligent Internet

Rapid growth in Internet and Intranet deployment and usage has resulted in a major shift in both corporate and consumer computing paradigms. This shift has resulted in massive increases in demand for network bandwidth, performance and flexibility to support both existing and emerging applications and services. However, this demand has often left Internet Service Providers (ISPs) with insufficient network capabilities to fully leverage the opportunity. In response to these demands, Cisco introduces new Quality of Service (QoS) extensions to Cisco IOS software network services. This set of services will enable a new Intelligent Internet business model wherein ISPs can generate profitable business growth by rapidly defining, deploying and charging for differentiated services targeted at specific customer requirements including efficient handling of both mission critical and bandwidth hungry web applications.

Critical Network Requirements

In order for Internet Service Providers to offer widespread value-added services which meet customer's demanding application requirements, the network infrastructure must handle a key set of technological and business requirements:

- **Services Scalability**—An increasing number of network capabilities and services will be activated to meet customer needs and to implement service provider network resource allocation policies. The network must maintain high packet throughput in this services rich environment.
- **Intelligent Congestion Control**—The network must actively seek to avoid congestion conditions which can lead to throughput degradation, to recover gracefully from

congestion situations and to distinguish between temporary traffic bursts and long term traffic overload conditions. The network must also be able to maintain preferential treatment for higher priority traffic without exacerbating congestion under overload conditions.

- **Investment Protection**—Service providers must be able to gain the benefits of new services and capabilities via leveraging the installed base of network elements and software and without requiring widespread, fundamental changes in the underlying network hardware and protocols.
- **Traffic Classification and Prioritization**—The network must be able to efficiently sort and classify packets into traffic classes or service levels for appropriate network handling to meet customer application requirements and willingness to pay for services. Both service providers as well as customer/applications must have the capability to classify traffic but the service provider must be able to override customer classifications under appropriate conditions.
- **Granular, Lightweight Metering**—The network must be able to gather highly detailed and accurate measurements with network layer visibility and make these available to external applications which support billing/accounting, network planning and network management and monitoring with minimal impact on network element performance.

- **Adaptive Application Support**—Customers must be able to request network resources (e.g. bandwidth allocation) which meet application requirements without modifying the installed base of applications. Applications must also be able to choose to actively leverage new network capabilities to reach the maximum base of users under a variety of network conditions.
- **Policy and Service Flexibility**—Service providers require sufficient technological building blocks and feature “knobs” to have the flexibility to define and offer highly differentiated services to target customers with minimal vendor interactions. Service providers must also have the flexibility to specify resource allocation policies at fine-grained levels including by physical port, by address and by application.
- **Revenue/Capacity-Unit Maximization**—Service providers need to be able to rapidly upgrade network capacity to meet growing bandwidth demand while simultaneously deriving maximum revenue from current installed capacity. Thus, service providers require the flexibility to implement network technologies and policies aimed at smoothing demand (e.g. lower evening rates) and maximizing revenue/carried bit (e.g. preferential treatment for premium class traffic).
- **Bandwidth Allocation**—Service providers must be able to provide and enforce bandwidth commitments to traffic sources and to flexibly determine packet handling policy when bandwidth allocations are exceeded.
- **Network Element Cooperation**—Network elements at the edge and the backbone of the network must actively cooperate to maximize throughput and performance, minimize the impact of congestion and provide end-to-end service levels and QoS policies to meet customer requirements. Service providers also require the means to provide end-to-end service quality from networks built with multiple technologies (e.g. routers, Frame Relay, ATM, tag switching) and to interwork with external networks while meeting QoS goals.
- **Network Layer Visibility and TCP/IP Tuning**—Network layer traffic visibility must be provided in order to apply QoS features to key granular traffic sources including application flows and IP addresses. The network must also be highly tuned for efficient handling of IP and TCP-based traffic and applications.

Cisco IOS Software QoS Services—Edge and Backbone

Cisco IOS software is a platform that delivers network services and enables networked applications that meet the critical requirements stated above. These network services include services that are the foundation of all networks including connectivity, security, scalability, reliability, and management, and services that enable leading-edge applications including IBM, multimedia, voice, and QoS services.

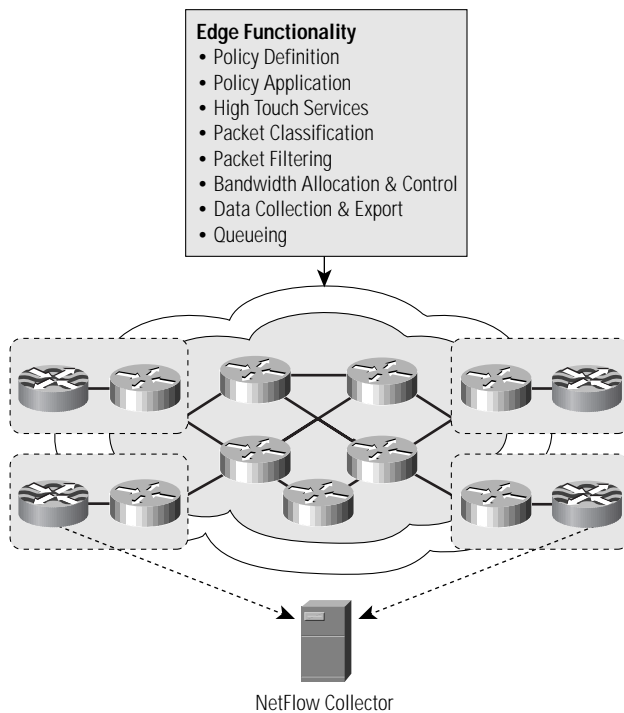
Cisco IOS QoS Services consists of three key components: a rapidly evolving set of service building blocks and capabilities, a highly flexible means of specifying policies which utilize the building blocks to control network resource allocation in support of customer and application requirements, and a distribution of functionality that optimizes Internet scalability for services deployment and bandwidth/capacity growth.

Cisco IOS QoS services provide service providers with the means of distributing network functionality and responsibility between edge functions and backbone functions. This distribution of functionality enables simultaneous performance and services scalability.

At the edge of the network, ISPs gain the capability to flexibly:

- specify policies that establish traffic classes and service levels
- specify policies that define how network resources are allocated and controlled to handle these traffic classes
- efficiently map packets into the traffic classes
- apply policies and “high touch” services to meet customer application and security requirements
- collect and export detailed measurements concerning network traffic and service resource utilization


Figure 1 Edge Functionality



In the backbone of the network, Cisco IOS QoS and supporting technologies provide the capabilities to:

- scale the network to provide extremely high capacity, performance and reliability
- provide policy administration and enforcement
- provide streamlined queuing and congestion management

Figure 2 Backbone Functionality



In the backbone, Cisco IOS QoS Services provide the capability to effectively control, manage and scale the high bandwidth network necessary to handle the demands of Internet traffic growth while meeting the QoS requirements of business and consumer applications. Cisco IOS QoS Services and supporting technology deliver backbone functionality focused on extremely high throughput and capacity scalability as well as policy administration and enforcement. The backbone is relieved of the responsibility of implementing high touch services on high speed interfaces, thus contributing to the reliability and stability of the network. Note that Cisco IOS QoS features at the edge of the network will interwork with either router or ATM technology in the backbone to provide end-to-end QoS capabilities.

The remainder of this section provides an overview of the new IOS QoS capabilities.

IP Precedence for Traffic Classification (edge)

IP precedence provides the capability to partition the traffic into multiple classes of service. The network operator may define up to 6 classes of service and then utilize Extended Access Control Lists (Extended ACLs) to define network policies in terms of congestion handling and bandwidth allocation for each class. The IP Precedence feature utilizes the 3 precedence bits in the Type-of-Service field in the IP header to specify class of service assignment for each packet. The IP Precedence feature provides considerable flexibility for precedence assignment including customer assignment (e.g. by application or access router) and network assignment based on IP or MAC address, physical port, or application.

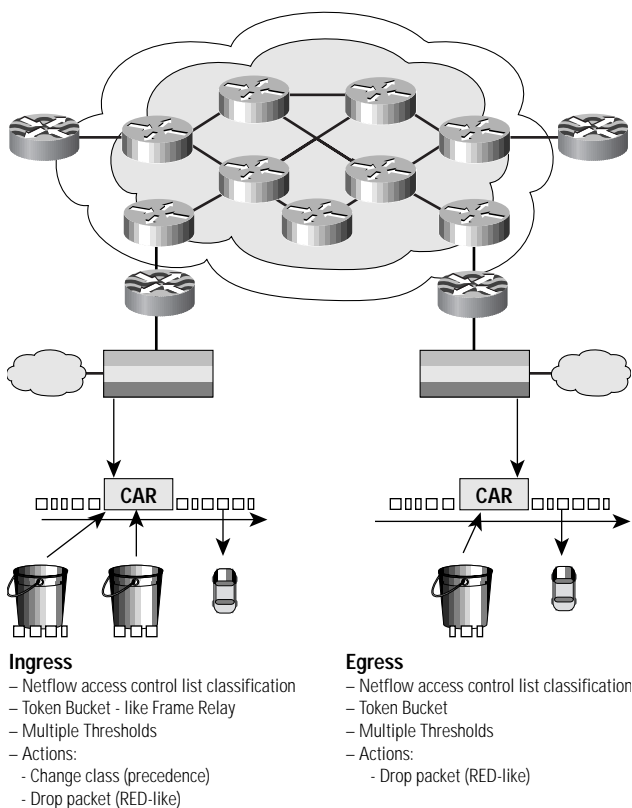
The IP Precedence feature enables the network to act either in passive mode (accepting precedence assigned by the customer) or in active mode utilizing defined policies to either set or override the precedence assignment. IP Precedence can be mapped into adjacent technologies (e.g. Tag Switching, Frame Relay or ATM) to deliver end-to-end QoS policies in a heterogeneous network environment. Thus, IP Precedence enables service classes to be established with no changes to existing applications and with no complicated network signaling requirements.

Committed Access Rate to Manage Access Bandwidth (edge)

Committed Access Rate (CAR) provides the network operator with the means to allocate bandwidth commitments and limitations to traffic sources and destinations while

specifying policies for handling traffic which exceeds the bandwidth allocation. CAR policies may be utilized at either the ingress or egress of the network. CAR thresholds may be applied by access port, by IP address or by application flow. The CAR feature uses token bucket filters to measure traffic load and limit sources to bandwidth allocations while accommodating the inherently bursty nature of IP traffic. For traffic which exceeds allocated bandwidth, CAR utilizes Extended ACLs to define policies including bandwidth utilization thresholds under which packet priority is modified or packets are dropped.

Figure 3 Committed Access Rate



Examples of CAR policy options include:

- Firm CAR - packets exceeding the allocated bandwidth are discarded
- CAR + Premium - packets exceeding allocated bandwidth are “recolored” with either higher or lower precedence levels
- CAR + Best Effort - packets exceeding the allocated bandwidth are recolored up to a burst threshold after which they are dropped

- Per Application CAR - different CAR policies are specified for different applications. For example, mission critical applications utilize a CAR + Premium policy, multimedia applications utilize a CAR + Best Effort policy

NetFlow Switching and Data Export (edge)

NetFlow Switching provides high performance for network layer services and lightweight, fine-grained data collection. NetFlow Switching enables “per-flow” application of Cisco IOS Network Services such as security and traffic accounting. NetFlow enables service scalability – once a flow has been established a single task simultaneously applies services, switches packets and collects data on a connection oriented basis. NetFlow switching performance is also invariant to access list size - thus, selected services may be widely deployed and activated in the network with minimal router performance impact.

NetFlow exports extensive flow-by-flow measurements for collection, post-processing and usage by accounting/billing, network planning and network monitoring processes. The data collected for each flow includes the following:

- source and destination IP address
- start-of-flow and end-of-flow timestamps
- packet and byte counts
- next hop router address
- input and output physical port interfaces
- source and destination TCP/UDP port numbers
- IP protocol type
- Type-of-Service field
- TCP flags
- source and destination autonomous system numbers
- source and destination subnet masks


Thus, the granularity of the NetFlow data provides a key metering mechanism for differential charging for services based on parameters such as time of day, class of service, application usage or traffic usage.

Random Early Detection for Congestion Management (backbone)

Random Early Detection (RED) provides network operators with the ability to flexibly specify traffic handling policies to maximize throughput under congestion conditions. RED works in conjunction with robust transport protocols (e.g. TCP) to intelligently avoid network congestion by implementing algorithms which:

- distinguish between temporary traffic bursts which can be accommodated by the network and excessive offered load likely to swamp network resources
- work cooperatively with traffic sources to avoid TCP slow start oscillation which can create periodic waves of network congestion
- provide fair bandwidth reduction to reduce traffic sources in proportion to the bandwidth being utilized. Thus, RED works with TCP to anticipate and manage congestion during periods of heavy traffic to maximize throughput via managed packet loss. RED also provides the network operator with considerable flexibility including parameters to set minimum and maximum queue depth thresholds as well as packet drop probability and MIBs for network management visibility into packet switching and packet dropping behavior.


Figure 4 Random Early Detection



Weighted RED for Class of Service Based Congestion Management (backbone)

WRED combines IP Precedence and RED capabilities to provide differentiated performance characteristics for different classes of service – thus providing preferential traffic handling for higher priority packets. WRED also provides preferential traffic handling under congestion conditions without exacerbating the congestion. Network operators have the flexibility to define minimum and maximum queue depth thresholds and drop probabilities for each class of service and MIBs are provided for each service level for network management visibility.


Figure 5 Weighted Random Early Detection



Weighted Fair Queuing Based on Class of Service (backbone)

Weighted Fair Queuing (WFQ) provides the capability to provide expeditious handling for high priority traffic requiring low delay while fairly sharing the remaining bandwidth between lower priority traffic sources. WFQ divides link traffic into high priority and low priority flows (based on metrics including IP Precedence and traffic volume). High priority flows receive immediate handling while low priority flows are interleaved and receive proportionate shares of the remaining bandwidth. Packets from low priority flows may be discarded during periods of congestion. The Cisco IOS QoS services provide distributed WFQ per Class of Service to enhance WFQ performance and scalability (note that WFQ is also supported in an ATM backbone environment).

Figure 6 Weighted Fair Queuing



QoS Enabling Building Blocks

Cisco provides several key technological building blocks to enable scaleable, flexible QoS capabilities:

Extended Access Control Lists (edge)

Extended Access Control Lists (Extended ACLs) provide the network operator with a highly granular filtering mechanism to define edge policies including:

- packet classification into classes of service
- network resource management including bandwidth allocation and congestion management associated with these classes.

Extended ACLs provide highly efficient list processing and matching capabilities to minimize router performance impact in service rich environments. Extended ACLs also provide the network operator with the flexibility to implement policy by physical access port (per customer), by user/host IP address or MAC address, by traffic class or by application flows.

Distributed Switching and Services (edge and backbone)

Distributed Switching and Services on the Cisco 7500 family of routers provides the network operator with a highly economical means of scaling both switching and services performance (in both the edge and the backbone of the network) via high density Versatile Interface Processor (VIP) cards. With Distributed Switching, each VIP card handles local switching decisions based on a local, pre-populated route cache with minimal dependency on centralized resources. With Distributed Services, processor intensive functions such as encryption, compression, queuing, IP multicasting, tunneling, fragmentation, congestion control and bandwidth allocation are offloaded to VIPs for local execution. Thus, router switching and services performance increases as additional VIPs are added to the platform enabling high packet throughput with widely deployed, value-added, Layer 3 services.

Tag Switching (backbone)

Tag switching is a multilayer switching technology for scaling router and switch backbones including the Internet and large corporate intranets. Tag switching combines the performance and traffic management capabilities of Layer 2 switching with the scalability and flexibility of network layer routing. Tag switching also allows ATM switches in the backbone of the network to take advantage of Layer 3 routing information, thus minimizing signaling overhead while increasing network scalability. The tag header used in tag switching incorporates the IP precedence information used to indicate class of service or priority. Thus, tag switching

allows advanced IP services based on the new Cisco IOS QoS capabilities to be optimally integrated into an ATM based backbone to provide end-to-end QoS support. Tag switching also enables traffic engineered paths that lead to both reduced congestion and improved bandwidth utilization.

Express Forwarding (edge and backbone)

Express Forwarding technology for IP is a scalable, distributed, layer 3 switching solution designed to meet the future performance requirements of the Internet and Enterprise networks. Express Forwarding evolved to best accommodate the changing network dynamics and traffic characteristics resulting from increasing numbers of short duration flows typically associated with Web-based applications and interactive type sessions. Express Forwarding offers significant benefits in terms of performance, scalability, and resilience particularly in large, complex networks having transient traffic characteristics:

1. **Performance**—Express Forwarding implements a Cisco patent-pending expedited IP look-up and forwarding algorithm to deliver maximum layer 3 switching performance. Additionally Express Forwarding is less CPU intensive than route-caching, therefore it allows more CPU horsepower to be dedicated to packet forwarding
2. **Scalability**—Express Forwarding technology is optimized for information distribution allowing it to take advantage of the distributed architecture of the Cisco 7500 router family. Distributed Express Forwarding delivers scalable switching capacity by providing each of the VIPs with an on-card copy of the forwarding information database enabling them to autonomously perform Express Forwarding between local Port Adapters
3. **Resilience**—Express Forwarding offers an unprecedented level of switching consistency and stability in large dynamic networks. This results from the fact that the Express Forwarding look-up table contains all known routes, therefore eliminating the potential for “cache-misses” which occur with demand caching designs.

Supporting Capabilities

Network Monitoring

Cisco IOS QoS services and NetFlow data collection and export capabilities provide a new paradigm for ISP's to efficiently monitor network operations, drilldown to diagnose and resolve network troubles and provide both internal personnel and customers with data concerning service level agreement metrics. ISPs may now utilize a flow-based paradigm to visualize traffic flows in the network including aggregate traffic flows, flows by application and flows by class of service. Enhanced network monitoring based on layer 3 services MIBs will also provide the capability to observe bandwidth allocation and excess usage determined by CAR functionality and packet loss from WRED policy application.

Scalable Resource Reservation

The Resource Reservation Protocol (RSVP) allows applications to request a specific quality of service. Hosts and routers use RSVP to deliver these requests to the routers along the paths of the data flow and to maintain router and host state to provide the requested service. RSVP specifies that each router between network endpoints participate in the RSVP signaling sessions to reserve, teardown and manage appropriate resources. This model provides end-to-end, fine-grained resource control but leads to serious scalability difficulties in large, high speed networks (such as the ISP backbone). The Cisco IOS QoS services provide the functionality to build a highly scaleable resource reservation service that utilizes RSVP signaling and delivers RSVP benefits - edge routers interpret the RSVP messages and map them into high priority packets which are tunneled to the egress edge of the network for full RSVP negotiation with the far end CPE router. WRED provides preferential treatment for appropriate RSVP flows in the network backbone during periods of congestion and WFQ is applied to manage timely delivery of egress flows at the edge of the network.


This methodology avoids the overhead of RSVP in the backbone of the network while delivering the RSVP benefits of dynamic, customer (or application) controlled signaling to meet special QoS requirements in IP networks. RSVP-enabled traffic flows are then mapped into appropriate classes of service for transport across the backbone of the network. The Cisco IOS QoS services also provide high performance distributed queuing and multicasting

throughout the network to provide scaleable support for RSVP traffic flow transport and delivery. Metering for RSVP-enabled applications is provided by NetFlow data collection and export.

Figure 7 RSVP Options

1. Explicit RSVP Transport

Process RSVP through network
on each network node



2. RSVP "Pass-through" & CoS Transport

Map RSVP to CoS at network edge
Pass-through RSVP request to egress

QoS Applications in the Internet

The Cisco IOS QoS services will enable service providers to offer and support a variety of new and enhanced network services and applications.

The current ISP business model typically includes flat-rate network access and connectivity services combined with a number of value-added services such as web hosting, security, Intranet outsourcing or web commerce applications. The new Cisco IOS QoS services will enable ISPs to generate profitable revenue growth by defining, customizing, delivering and charging for differentiated, value-added network services.

The Cisco IOS QoS services will also enable ISPs to offer new services which support mission critical and bandwidth hungry web applications for enterprise customers.

Class of Service

Internet Service Providers may now offer multiple service tiers with differential pricing policies - e.g. premium priced Platinum service, mid-priced Bronze service and low-priced Standard service. During periods of congestion, higher priority traffic is given preferential treatment according to policies defined by the ISP. Detailed network measurements provided via NetFlow data export allow ISP's to introduce pricing policies based on usage, time-of-day and traffic class.

Classes may be assigned by customer, by IP address or by application. Key Cisco IOS QoS technologies utilized include IP Precedence for traffic classification and Weighted Random Early Discard (WRED) for congestion control policy.


Business Class Virtual Private Network Intranets and Extranets ISPs may now offer comprehensive, differentiated Virtual Private Network (VPN) services (including both intranets and extranets) to businesses via the Internet. VPNs may be utilized to connect branch offices, mobile users and telecommuters to the main campus as well as providing flexible, efficient connectivity to suppliers and customers.

Cisco IOS software provides a comprehensive set of features required to implement the VPN Extranet services including:

- prioritized traffic classes for VPN traffic via IP Precedence traffic classification
- preferential treatment for VPN traffic (including mission critical applications) via WRED policies
- ingress and egress bandwidth limitations on VPN traffic via Committed Access Rate [CAR] policies
- dedicated or dial access to VPN services including traffic prioritization by physical port, address or application and dial priority interworking
- tunneling - enabling new IP protocols to be encapsulated and delivered unmodified across an Internet backbone
- network-layer encryption and firewalls - enabling secure transport and controlled access

NetFlow captures and exports detailed measurements including intra-company, supplier and customer traffic volume and application usage patterns for billing/accounting and network planning purposes.

Figure 8 Business Case VPN




Multimedia

Multimedia applications may now leverage network capabilities to maximize customer reachability under a variety of access bandwidths and network congestion conditions. Key examples include:

- streaming MPEG video delivery to multiple sites with different connectivity bandwidth and CAR limits - e.g. one site with T3 connectivity, one site with T1 connectivity and one site with ISDN connectivity. Different MPEG frame types (i.e. I, P and B) are assigned different priorities (via IP Precedence) by the video server. Only high priority frames are delivered to the ISDN site, whereas all frames are delivered to the site with T3 connectivity. During network congestion, higher priority frames are given preferential treatment and delivered to all sites although the other frames may be discarded.

Figure 9 Multimedia



- web servers delivering “high priority” pages - e.g. HTML pages with key corporate messages, high bandwidth content or high advertising content with premium traffic classification.
- Webcast traffic is classified as low priority and multicast to multiple destinations. During periods of congestion, mission-critical traffic receives preferential treatment over webcast traffic

Access Bandwidth Management and Fractional Provisioning

The network operator provisions a T1 for the business customer but provides a CAR of 768kb/sec and usage-based pricing. The customer benefits by paying for only the required bandwidth being utilized. The customer (and the service provider) also benefits by receiving detailed usage measurements via NetFlow data export and by easy reconfiguration to higher service levels when required. In a related example, the network operator provisions a T1 for a business customer as above with a 1.2Mb aggregate CAR. The customer requests further subdivision of the bandwidth by subscribing to 256kb CAR for Web traffic with medium priority, 256kb CAR for mission critical database transactions with premium priority and 512kb aggregate CAR for several multimedia applications. Each traffic class is allowed to burst above its allocated CAR when bandwidth is available. Thus, mission critical traffic is assured of sufficient bandwidth when needed as well as preferential treatment within the network operators service domain. Applications

exceeding their CAR may either be upgraded in precedence (for premium charging), downgraded in precedence (for less preferential network treatment), or have packets discarded.

Summary

QoS Business Implications for Service Providers

The Cisco IOS QoS services provide the service provider with the means to maximize revenue generated per unit of bandwidth capacity by implementing services and policies which motivate customer behavior. For example - time and usage based billing via NetFlow measurements provide the service provider with a means of encouraging (or shifting) demand during periods of light network loading by offering off peak discount pricing. Traffic classes and prioritization allow the network operator to encourage the customers to classify their traffic and then to transport the highest value bits during peak usage periods and heavy congestion conditions. Bandwidth allocations via the CAR feature enable the network operator to carefully engineer network capacity and to meet bandwidth commitments during periods of congestion. Finally, ISPs are enabled to market “Express” services whereby premium customers receive superior network performance.

Thus, Cisco IOS QoS services enable ISP’s to pursue a new Intelligent Internet business model for profitable revenue growth via:

- offering and charging for targeted, differentiated services
- maximizing network utilization
- maximizing revenue/carried bit
- traffic growth from new applications

The new Cisco IOS QoS services also enable enterprise intranets and extranets to handle the growing demands of mission critical and bandwidth hungry web applications over the wide area network. Enterprise customers may also leverage Virtual Private Networks and other Advanced Services provided by the ISPs to optimize communications with customers, suppliers, branch offices and mobile/telecommuting employees

Cisco IOS QoS services provide Internet Service Providers with service building blocks and policies to flexibly allocate network resources to meet customer’s application requirements. The Cisco IOS QoS services model distributes functionality between the edge and the backbone of the network, thus allowing for wide-scale deployment of network services while concurrently providing backbone

scalability to meet extremely high packet throughput requirements. Finally, the following table provides a summary view of Cisco IOS QoS capabilities meeting the critical requirements for Internet Service Providers:

Requirement	IOS QoS Functionality
TCP/IP Tuning	RED and WRED and CAR excess bandwidth policies

Requirement	IOS QoS Functionality
Services Scalability	NetFlow Switching, Distributed Switching and Services, efficient matching and access list processing
Intelligent Congestion Control	IP Precedence, RED & WRED, CAR excess bandwidth usage policy
Investment Protection	QoS features based on software upgrades to existing routers, solution based on familiar IP Protocols
Traffic Classification	Extended Access Lists, IP Precedence network and customer precedence assignment, active and passive precedence, CAR excess bandwidth precedence changes
Granular, Lightweight Metering	NetFlow Data Collection & Export
Policy and Service Flexibility	CAR excess bandwidth policies, RED/WRED thresholds and probability parameters, extended ACL policy specifications by address, physical port and application granularity
Revenue/Capacity Unit Maximization	NetFlow metering, IP precedence assignment, RED & WRED preferential detection, CAR rate limits
Bandwidth Allocation	CAR policies via extended ACLs, RED/WRED congestion policy
Network Element Cooperation	Edge and backbone functionality distribution, interworking IP Precedence with Frame Relay Discard Eligibility, interworking IP Precedence with ATM Cell Loss Priority, interworking IP Precedence with Tag switching networks and technology



Corporate Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
World Wide Web URL:
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems Europe s.a.r.l.
Parc Evolic-Batiment L1/L2
16, Avenue du Quebec
BP 706-Villebon
91961 Courtaboeuf Cedex
France
Tel: 33 1 6918 61 00
Fax: 33 1 6928 83 26

Americas

Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
Tel: 408 526-7660
Fax: 408 526-4646

Asia Headquarters

Nihon Cisco Systems K.K.
Fuji Building
3-2-3 Marunouchi
Chiyoda-ku, Tokyo 100
Japan
Tel: 81 3 5219 6000
Fax: 81 3 5219 6010

Cisco Systems has more than 190 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the

Cisco Connection Online Web site at <http://www.cisco.com>.

Argentina • Australia • Austria • Belgium • Brazil • Canada • Chile • China (PRC) • Colombia • Costa Rica • Czech Republic • Denmark
Finland • France • Germany • Hong Kong • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Malaysia • Mexico
The Netherlands • New Zealand • Norway • Philippines • Poland • Portugal • Russia • Singapore • South Africa • Spain • Sweden
Switzerland • Taiwan, ROC • Thailand • United Arab Emirates • United Kingdom • Venezuela