

DESIGNING SERVICE PROVIDER CORE NETWORKS TO DELIVER REAL-TIME

SERVICES

More companies and individuals are relying on the Internet for their business and personal communication. At the same time, service providers are migrating their circuit-based networks to single multiservice IP infrastructures to lower total cost of ownership and to deliver profitable Internet services.

To support all forms of communication (business applications, voice, video, and e-mail), the Internet infrastructure must allow flexible and scalable traffic discrimination based on a user or application performance requirement.

The Cisco 12000 Series router offers guaranteed priority packet delivery, a set of features that allows service providers to provision, across a single infrastructure, multiple discriminated Internet communications. The Cisco 12000 Series uniquely delivers priority-based congestion control, dedicated low latency queuing, and packet sequence integrity under all conditions required by premium services such as voice over IP (VoIP). More importantly, these features are supported seamlessly on Cisco platforms across access, edge, and core multiservice IP networks.

MULTISERVICE IP: MARKET DYNAMICS

With the growing acceptance of the Internet as the one service platform to facilitate interpersonal and interbusiness communication, Internet Protocol (IP) is now firmly established as the de-facto, global standard internetworking protocol—and the best way to deliver high-bandwidth communication services.

Internet access has already become mission critical to many consumers and companies. The communication industry is now moving from the traditional time-division multiplexing (TDM) model where bandwidth or transit services are the products, to a new model where different services and the associated prices are discriminated based on user or application performance level, traffic priority, and time of day. In the new model, bandwidth represents only one type of service (known as virtual leased line [VLL]). Most carriers and today's competitive network and application service providers (NASP) realize that their most profitable business opportunity is to offer integrated and managed Internet services (data, voice, and video) over a single but scalable IP multiservice infrastructure.

Instead of charging users based on a fixed or flat-rate basis, service provider customers need the ability and flexibility to provision Internet connections to address different market segments and to support different types of service-level agreements (SLAs). A premium service is suitable for mission-critical and voice traffic, a gold service is suitable for certain Web traffic and file exchange, and a silver service suitable for e-mail and news traffic.

Carriers and NASPs that offer these services will have a competitive advantage over those providers who still believe their networks should utilize circuit-based technology optimized for traditional telephony—the old model.

In response to this transition from the old bandwidth model to the new differentiated services model, the strategic and competitive direction of NASPs is to build single multiservice infrastructures with IP as the foundation. Such infrastructures must be capable of simultaneously supporting mission-critical business applications, real-time voice and video communications, better “best-effort” services as well as legacy “best-effort” type of services. The same infrastructure must be also capable of carrying traffic from legacy connection-oriented services (such as high-level data link control [HDLC], dedicated private line, ATM, public switched telephone network [PSTN]) for migration purposes.

The benefits obtained by implementing a multiservice IP infrastructure include:

- Increased revenue and profit opportunity through an infrastructure that supports new value-added differentiated services over data, voice, and video. The desire is to charge accordingly for these new services based on actual usage and level of performance delivered.
- Increase of bandwidth efficiency resulting from the statistical multiplexing and traffic engineering performed by IP routers. Bandwidth overprovisioning is becoming too expensive.
- Lower total cost of ownership (TCO) by eliminating complex, multiple (connection oriented) overlay networks, therefore saving on operation costs. Service providers are particularly interested in achieving more efficient management and control of network resources to keep operational costs in check.

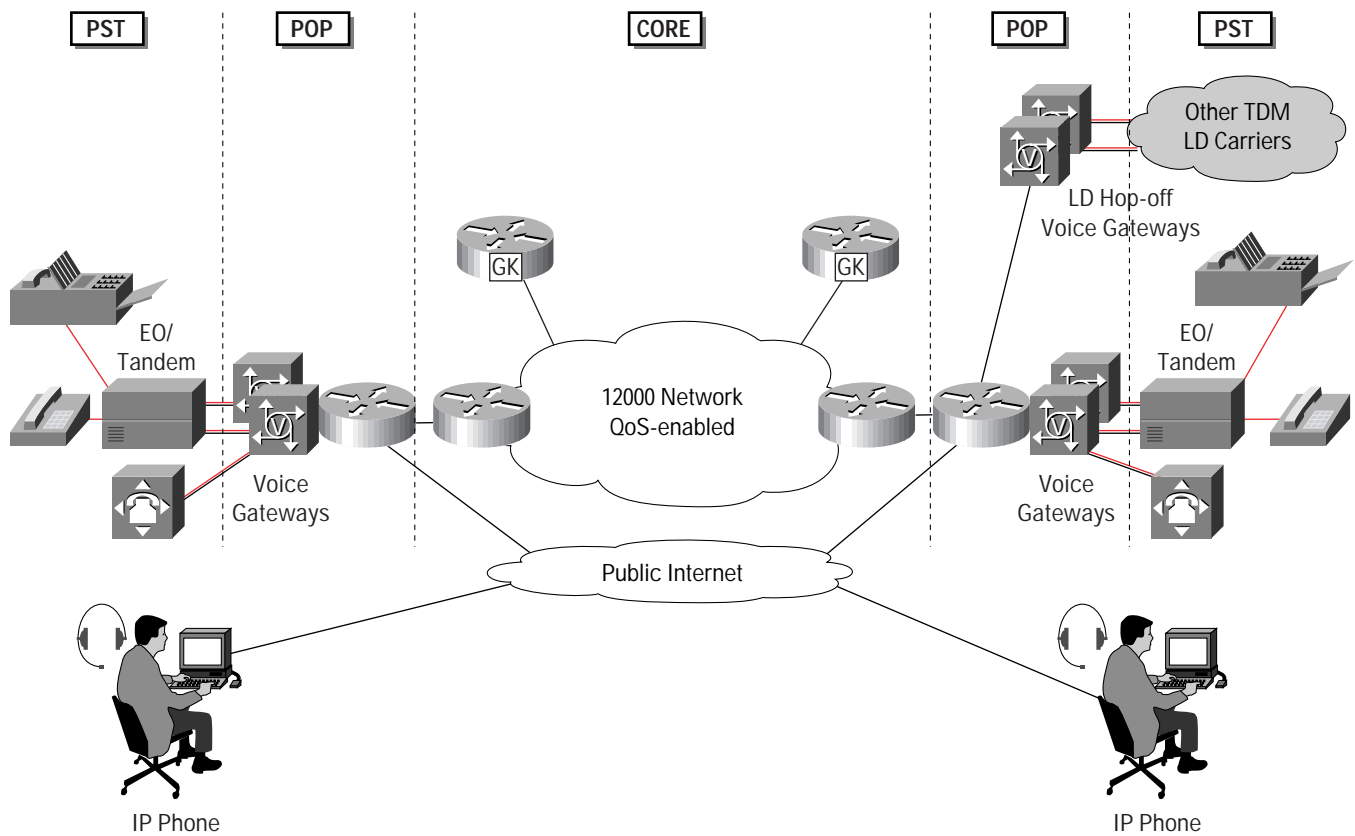
IP router vendors must address several economical and technical challenges to help service providers to build multiservice IP infrastructures. A shared IP network must provide high availability and reliability of service performance for real-time services (virtual leased line, IP telephony, and digital music and video services) and nonreal-time services. Along with voice and video streaming, multiple discriminated IP services provide customers with a range of service levels at a range of prices (pay more, get more). Although most real-time applications do not need guaranteed resource commitments from a core network such as a static amount of bandwidth, their requirements do conflict with nonreal-time services in terms of network resource usage.

Multiservice IP infrastructure must be enabled with IP traffic-measurement features to feed several applications such as traffic engineering, billing, routes analysis, performance monitoring, or service attack analysis. Control and efficiency of network resources are becoming an objective of Service Providers (SPs) to address increased competition and decreasing profit margins. These economical challenges translate into service requirements and challenges that an IP router vendor must address. These service requirements translate into network design requirements and IP platform requirements.

This document focuses on platform requirements that help the network to predict the service performance of a real-time signal. The terms IP traffic management and measurement are used to refer to the collective intelligence (specifically, hardware and software functions) within the Cisco 12000 Series router that allows provisioning and monitoring of services that require predictability such as VoIP. Traffic measurement capabilities are covered in a separate white paper dedicated to Cisco NetFlow.

Figure 1 shows a typical service provider topology using the Cisco 12000 Series Router to carry long haul or international VoIP traffic.

Figure 1
Cisco 12000 Series Network supporting VoIP



REAL-TIME SERVICE REQUIREMENTS

By their nature, voice and video applications do not need the network to deliver a fixed amount of bandwidth. What these applications do need is for the network to minimize transit delay and to keep delay within a reasonably narrow range. Within a VoIP connection, one endpoint takes a voice stream, in digitized format, packetizes it, and then transmits it over the IP network. The network introduces some variation in the delay (jitter) with each packet delivered. The receiving endpoint de-packetizes the voice stream, buffers it, and plays back the original signal. Buffering cancels network-induced jitter and the voice signal can be played back at a steady rate as long as transit delay is contained within a narrow range. Packets arriving before playback time can be used to reconstruct a source signal. Packets arriving after playback time are useless in reconstructing a real-time signal.

Real-time applications require predictable service from the network; that is, they require a bound (known a priori) on the delivery delay of each packet. In general, lower delay is preferred. Typically, a one-way transit delay of up to 150 ms does not cause a perceivable degradation in voice quality for most telephony applications. To set playback time (essentially defining total application delay), a telephony application needs to have some information about the maximum delay (the statistical bound) that each packet will experience. Since all data is buffered until playback time, an application is indifferent as to when data is delivered as long as it arrives before playback time. Certain telephony gateways (such as Cisco Voice Gateways) use an adaptive algorithm to size jitter buffers. The adaptive algorithm uses the current value of jitter experienced at the terminating gateway to size the jitter buffer.

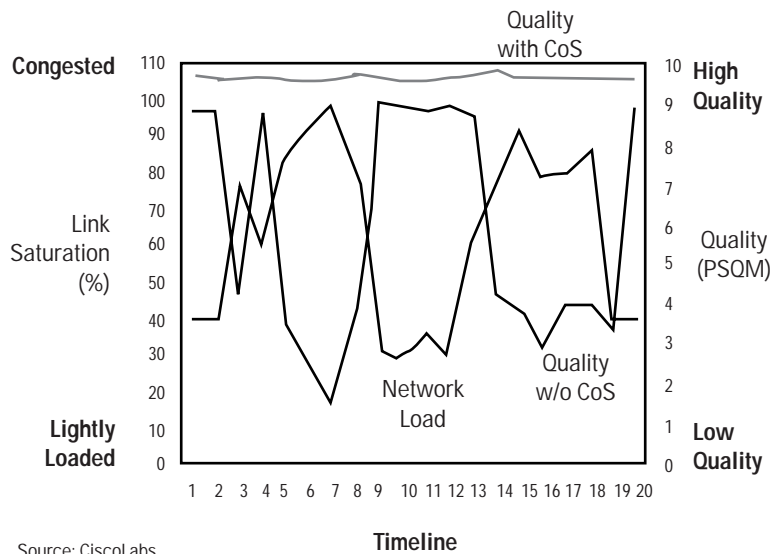
A large contributor of network-induced jitter is the queuing delay that each packet accumulates in routers. There are two components of this jitter. One is caused by contention for resource between packets from multiple real-time applications, such as voice calls. The other component is contention for resource between a real-time packet and non real-time packet. This jitter must be bounded and minimized by the network in order to support real-time communication.

Network resources (bandwidth and buffers) are shared between voice and video streams, as well as transmission control protocol (TCP) traffic that is bursty and bandwidth-consuming. To provide predictable and reliable service for voice and other real-time applications, IP routers must implement IP traffic management features including intelligent scheduling algorithms (traffic isolation), congestion avoidance (traffic protection), and network outage recovery mechanisms to protect from node or path failures (traffic engineering and fast reroute).

Figure 2 shows the relationship between voice quality with and without QoS when congestion occurs in a network. The center shows how voice quality degrades significantly as a direct consequence of increases in traffic load without QoS. The line at the top shows a network with QoS capabilities maintaining voice quality even during heavy network load conditions.

Figure 2

VoIP Quality with and without IP Traffic-Management Features



Source: CiscoLabs

⚠ Under load, a router without traffic management capabilities does not maintain Voice quality

TECHNICAL CHALLENGES

One major challenge in providing end-to-end QoS in a shared infrastructure is supporting scalability. Traditional connection-oriented approaches such as TDM and ATM that provide hard end-to-end QoS guarantees, which involve per-flow signaling, buffering, and scheduling, and are difficult to implement in high-speed core routers. IP traffic management methods based on traffic aggregation are required to control end-to-end service performance. A key element for planning and managing these enhanced services is having a scalable IP traffic measurement capability. It is important that these methods be consistent across different nodes within a network, and consistent across different media interfaces along a given path.

CISCO SOLUTIONS FOR SUPPORTING IP-BASED REAL-TIME SERVICES

Continuous innovations from Cisco in IP technology enable service providers to build powerful multiservice IP infrastructures using carrier-class Internet routing platforms that meet the economic and technical challenges of delivering competitive IP service offerings. The Cisco 12000 series Internet router offers scalable and consistent IP traffic management and measurement functions across access, edge, and core multiservice IP networks.

The following sections examine different components of IP traffic management features built in the Cisco 12000 Series routers and the methods (architecture) for deployment within a network to achieve predictability of service performance within a scalable and efficient infrastructure. These features consist in a control plane imbedded in Cisco IOS® Software and in a data plane fully implemented in hardware. The following features will be examined:

- Classification
- Policing and marking
- Shaping
- Scheduling
- Congestion manager

END-TO-END IP QoS ARCHITECTURE

Contention for network resources between applications with conflicting requirements is addressed by separating traffic into classes for different treatment. Classes can be defined to meet specific requirements such as delay/jitter limit, packet loss limit, and so on.

Admission control to a class is needed to make sure that no more traffic gets admitted than the resources allocated to that class can satisfy. Contention for resources between traffic with the same class (same requirement) is addressed by either using call admission control (voice, video) or congestion avoidance mechanisms (TCP flows).

The Integrated Services standard (IntServ) defines fine-grained (flow-based) methods of performing IP traffic admission control that uses Resource Reservation Protocol (RSVP). The Differentiated Services standard (DiffServ) defines methods of classifying IP traffic into coarse-grained service classes and defines forwarding treatment based on these classifications. The end-to-end Cisco QoS architecture is based on a combination of DiffServ architecture and the integration of resource admission control from RSVP/IntServ.

INTSERV/RSVP

The Integrated Services model inherits the connection-oriented approach from telephony network design. Every individual communication must explicitly specify to the network its traffic descriptor as well as requested resources. The edge router performs an admission control to ensure available resources are sufficient in the network. The IntServ standard assumes that routers along a path set and maintain state for each individual communication.

The IntServ model has not been widely adopted in the Internet because the connection-oriented approach assumes a “flat” model of the Internet; that is, one that is administratively homogeneous. The number of connections required to handle all traffic of the Internet leads to a state explosion in the core routers. However, the resource admission control concept defined within IntServ is a useful tool to manage application level traffic with strict QoS requirements.

The role of RSVP in the Cisco QoS Architecture is to provide resource admission control for VoIP networks. If resources are available, RSVP accepts a reservation and installs a traffic classifier in the QoS forwarding path. The traffic classifier tells the QoS forwarding path how to classify packets from a particular flow and what forwarding treatment to provide. The installation of a traffic classifier and flow treatment is the interface between RSVP and DiffServ. RSVP is a control plane feature that limits accepted VoIP load to what the network can support. Integration of resource-based admission control with DiffServ network (RSVP aggregation) aims at achieving scalable strict QoS guarantees for VoIP Calls.

DIFFSERV

The DiffServ approach divides QoS into a number of functional elements to be implemented in network interfaces. Each element provides a traffic control function or forwarding treatment (called per-hop behavior) on a packet by packet basis.

The standard defines the three types of functional elements:

- A small set of per-hop forwarding behaviors (PHB). Typically, a PHB represents the scheduling and discard priorities a packet should receive on a router interface. Each PHB is identified by a DS code point (DSCP), which occupies 6 bits of the type of service (TOS) byte in the IP header.
- Packet classification.
- Traffic conditioning functions including metering, marking, shaping, and policing.

This model, as shown in Figure 3, achieves scalability by implementing complex classification and conditioning functions at network boundary interfaces, and by applying a PHB to aggregates of traffic (behavior aggregate) in the core. A simple traffic reclassification may be required at the core. An important component of the Cisco QoS Architecture is Multiprotocol Label Switching (MPLS) QoS. MPLS may be deployed within an IP network, for example, to provide a scalable VPN solution, traffic engineering/tunneling, and other applications. Because MPLS is not an end-to-end host protocol, it is required that an MPLS infrastructure support IP QoS features rather than extending them to a new QoS. It is important to note that QoS is an end-to-end characteristic of an application or host that should not be altered by the transport technology (in this case, MPLS).

Figure 3
Differentiated Services PHB Functionality

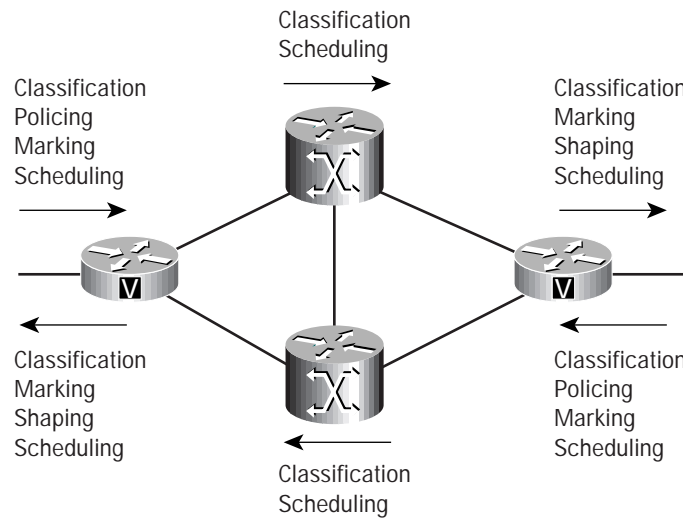
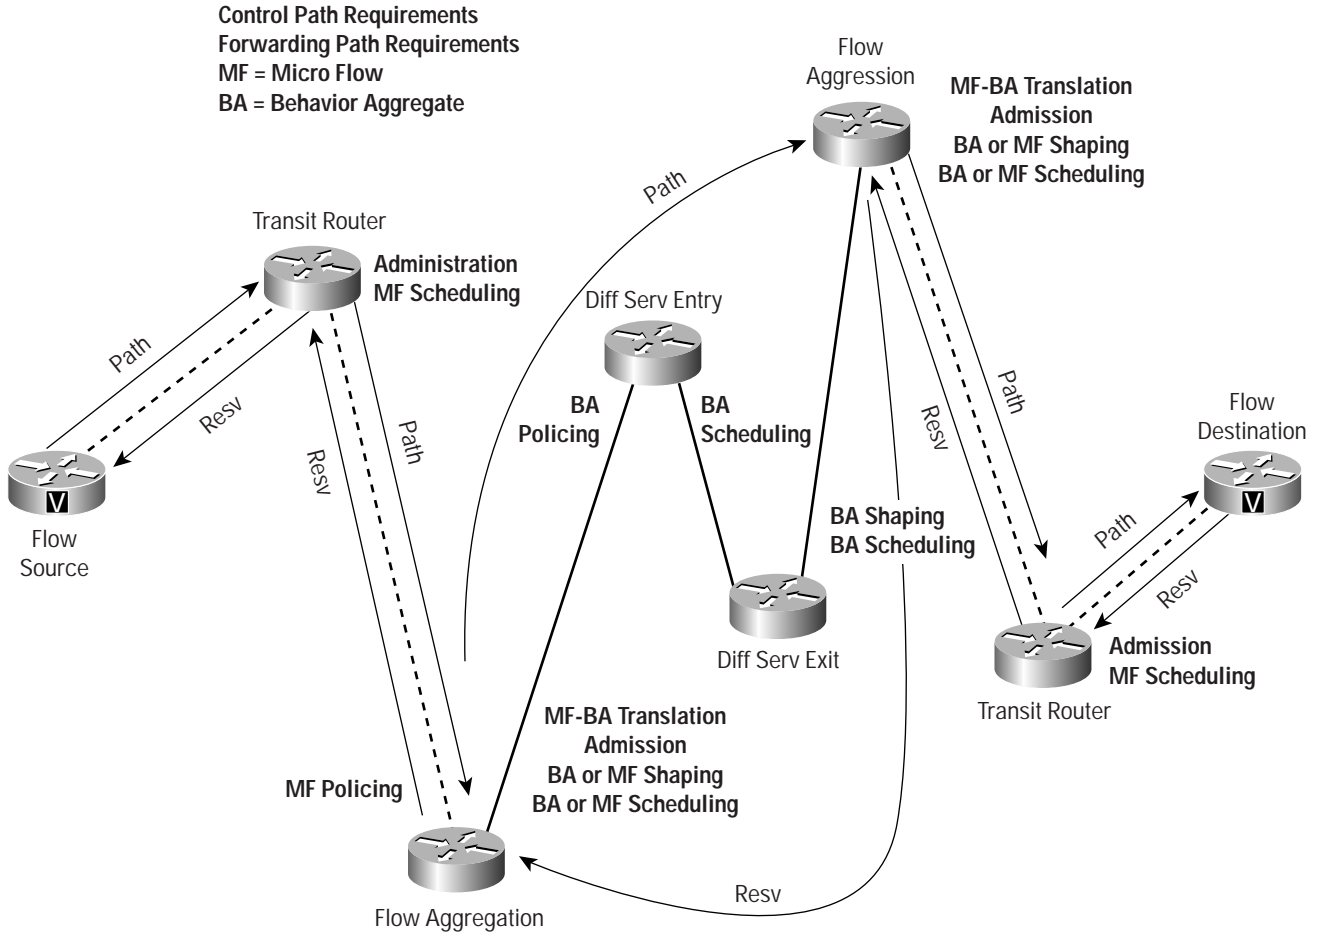


Figure 4

An Example of DiffServ and RSVP Integration



CLASSIFICATION

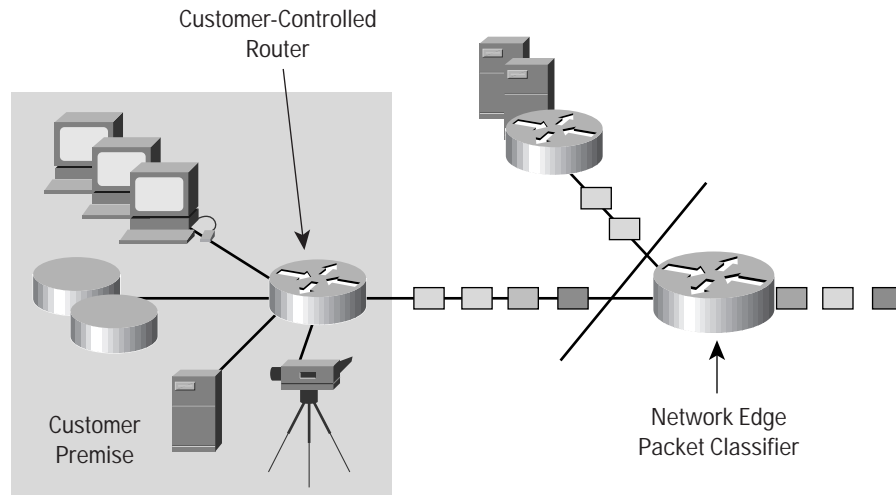
Classification is the process of defining a class of traffic and identifying packets that belong to the same class of traffic, or more simply, packets from applications having the same performance requirements.

TCP/IP traffic classification can be done in two modes: Behavior aggregate (BA) or multifield (MF). BA classification is based only on the DiffServ code point (DSCP). For best scaling, this mode is more suitable in the core.

MF classification selects packets based on a combination of one or more fields based on the TCP/IP header. All fields are user-programmable and include input interface, MAC address, DS field, source and destination IP address/prefix, source and destination port, and protocol type. MF mode is typically specified by access control lists (ACLs), to which the classification criteria will match. Its implementation is recommended as close to the edge as possible.

Cisco IOS Software offers flexible and multipolicy classification solutions that address current requirements and scale to future needs.

Figure 5
Classification at the Edge



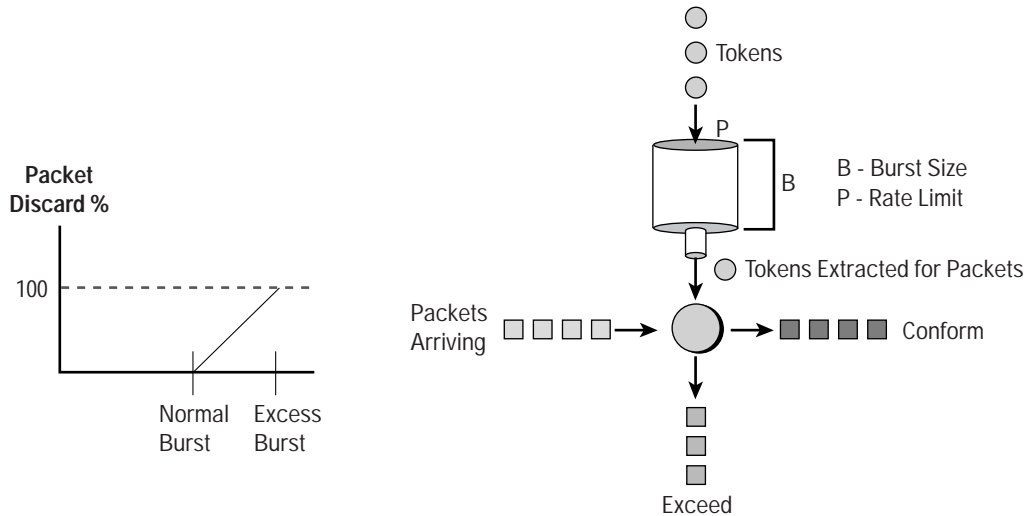
METERING

Traffic metering provides traffic controls that accommodate temporary bursts while limiting traffic sources to a long-term average rates. Cisco IOS Software uses a token bucket technique to measure traffic rate. Token bucket rate control does not intend to shape traffic streams. Token buckets are specified by defining three traffic parameters:

- *Committed rate*—Measured in bits per second, this is the long-term average rate permitted for the traffic source. Tokens are inserted into the bucket at the committed rate.
- *Normal burst*—Measured in bytes, allows for temporary bursts of packets that are deemed to conform to the token bucket limit. The normal burst represents the bucket's depth.
- *Excess burst*—Provides a “bonus round” where excess packets are gradually dropped to warn the violating traffic source to slow down before groups of packets are discarded due to exceeding the rate limit.

When traffic arrives, if sufficient tokens are available, then the traffic is said to conform; if not, the traffic is said to *exceed*. The appropriate action policy is then executed.

Figure 6
IP Traffic Metering



POLICING AND MARKING

Policing is the process of ensuring that incoming traffic belonging to a given class is conforming to the traffic profile defined (either signaled or provisioned) for that class. Policing happens at the ingress of a service provider network (domain). Typically, the profile is specified in terms of traffic type, incoming interface, average rate, and instantaneous rate limits. A flexible policing implementation can also be used to protect against attack. If the traffic does not conform—because the source is exceeding its allocated resource, then corrective/proactive actions are needed to protect conforming traffic. Options include dropping packets in excess of the contract, marking them as lower priority in the scheduler, or just passing them out.

In Cisco IOS Software, the policer and marker components are implemented via committed access rate (CAR). Typically, the user specifies metering parameters along with action parameters. Policers/marker can operate in four modes:

- *Policing off (transmit)*—Every packet is transmitted
- *Policing only*—Non-conforming packets are dropped
- *Policing and marking (set precedence and transmit)*—Nonconforming packets are marked
- *Marking*—All incoming packets are marked

From the operation perspective, CAR allows the user to specify an action to be taken for packets that either conform to or exceed the specified rate limit. The set of conform-actions and exceed-actions from which drop policies are selected are:

- *Transmit*—Switch the packet
- *Set precedence and transmit*—mark the precedence or DCSP bits and then switch
- *Continue*—Evaluate the next rate limit in a chain of rate limit statements
- *Set precedence and continue*—Mark the precedence DCSP bits and then evaluate the next rate limit in the chain
- *Drop*—Discard the packet

The packet classification functionality of CAR is an outcome of traffic matching and rate measurement functionality; it is not a standalone capability.

Multiple policing actions can be specified for an interface as long as they are associated with token bucket parameters. If a packet matches the traffic matching specification associated with a rate limit, then the associated token bucket is examined and the conform or exceed action is executed. Once the transmit action is selected and executed, subsequent policing statements are not evaluated. CAR policing statements can be nested by utilizing the *continue* key word. The default action at the end of a list of one or more rate limit statements is to transmit the packet.

TRAFFIC SHAPING

Traffic shaper delays some or all outgoing IP packets in order to bring a stream into compliance with the traffic profile associated with the output link it will go out on. Typically, traffic shaping happens at the egress of a service provider network (domain). A shaper uses a buffer to store the incoming packets in excess (burst) and to delay their transmission to the output interface. Typically, the shaper uses the result of the traffic meter to decide whether a packet should be expedited or delayed.

Distributed traffic shaping (DTS) is the implementation of the Cisco IOS traffic shaping scheme on a per-line-card basis. DTS on the Cisco 12000 series Internet router provides traffic shaping regardless of the encapsulation configured on the interface. DTS on the Cisco 12000 series Internet router can shape the output traffic to a specified bit rate and buffer excessive packets in the shape queue to transmit out later. Each shape queue uses a first in/first out (FIFO) as the scheduling scheme.

Traffic shaping parameters (traffic descriptors) are user-configurable. Traffic descriptors include three components:

- Committed information rate (CIR) specified in bits per second to sustain
- Committed burst size (Bc) specified in bits per burst
- Excess burst size (Be) specified in bits of queuing maintained in the pipeline.

Traffic shaper smoothes incoming traffic into the average bit rate defined by the CIR before it puts it out on the outgoing link. The Bc is used to derive the time interval over which the shaper monitors arriving traffic (Time interval = Bc / CIR). The user defines the average bit rate that is expected and burst size that is acceptable on that shape entity. Excess burst size defines the acceptable number of bits permitted to go over the burst size.

POLICY-BASED ROUTING


Policy-based routing (PBR) is a Cisco IOS QoS feature that allows users to override routing results obtained from a standard IP routing engine. Customers can define policies that selectively cause packets to take a different path than the one computed by the routing algorithm. The selected traffic can then benefit from a high bandwidth or low delay link. PBR works in conjunction with a Cisco IOS Software classifier that filters traffic to which the policy should be applied. PBR may also work with a marker to set IP precedence bits before forwarding a packet to a defined path. PBR is typically used at network peering interfaces: Incoming traffic that matches a defined user policy is directed to a private Internet connection (private peering link) and the rest of the traffic is directed to a public connection (public peering link).

TRAFFIC SCHEDULING

The Cisco 12000 series uses two-stage input scheduling to the switching fabric and output scheduling from the switching fabric. For each stage, a combination of virtual output queuing (VOQ), round robin, and an enhanced version of modified deficit round robin (MDRR) are used to control and isolate voice and video traffic. RED/WRED is used for traffic protection within a class.

Each traffic class is assigned a queue that isolates its traffic from traffic of different classes. The role of the scheduler is to provide a traffic class with a level of service that is independent from other traffic behavior. The scheduler uses multiple queue service disciplines to assign an amount of link bandwidth to a class. The scheduler must support head-of-line blocking (HOL) avoidance. As shown in Figure 7, HOL blocking occurs when a packet at the head of a queue is blocking traffic to all destinations because its destination line card is unavailable. Typically, the scheduler is provisioned so that voice and video classes have controlled latency and jitter.

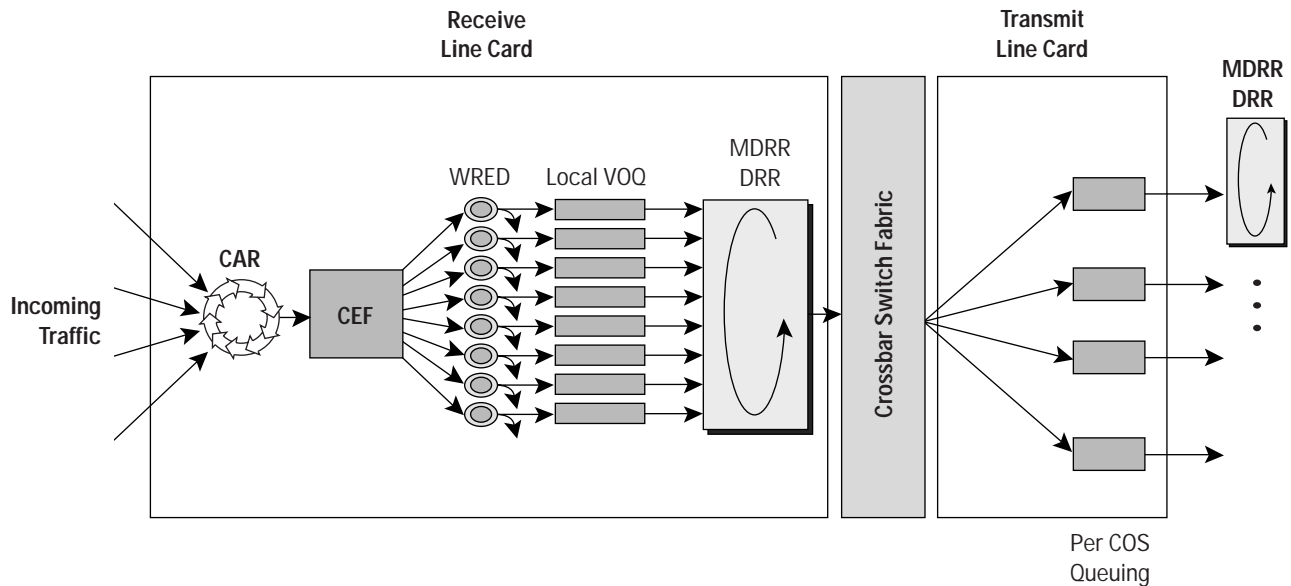
Figure 7
Head Of Line Blocking



Without virtual output queuing, the second packet in the queue 4 is blocked although the line card 2 (LC 2) is ready to receive traffic.

Virtual output queuing (VOQ) is a queuing technique in which packets from the routing engine or switching fabric are sorted according to output interface and IP precedence (Figure 8). At the receive side, an output port is the destination slot of a packet and the queue is called local output queue (LOQ). At the transmit side, the port is a physical port connected to a network and the queue is called output queue (OQ). Along with these queues, there are eight multicast queues per line card. The VOQ must be implemented to avoid the head of line blocking.

Figure 8
Virtual Output/COS Queuing in the Cisco 12000 Series Router

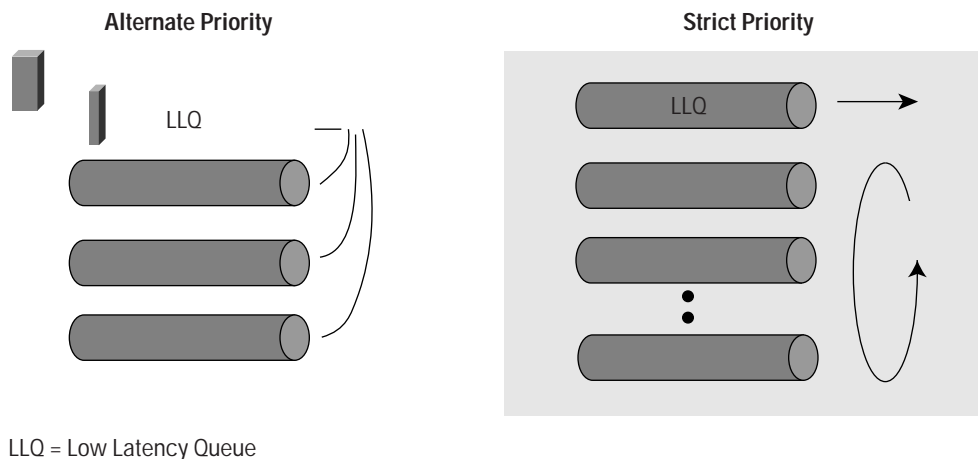


LOQs are serviced by hardware using a combination of RR and MDRR. The RR algorithm cycles through the queues one after the other, transmitting one packet before moving to the next queue. Each LOQ, along with multicast, is serviced round robin. Within a group, the queues are serviced MDRR. MDRR is an enhancement to DRR by adding “high priority-HP.” DRR servicing discipline tracks the byte deficit (specifically, the difference between the number of bytes ought to have been sent and the number of bytes that have been sent) for each queue and uses it to regulate long-term bandwidth assigned to the queue.

The MDRR service algorithm can work in two modes: Strict-priority mode and alternate-priority mode.

- *Strict-priority mode*—The high priority queue is serviced first. Only when all high-priority traffic is clear will other queues be considered. Remaining queues are serviced in DRR. This mode has the advantage of guaranteeing a minimum latency for the high-priority queue.
- *Alternate-priority mode*—A quantum of data is taken from the high-priority queue, then one quantum from one of the other queues (chosen via DRR), then back to a quantum of the high-priority queue again.

Figure 9
MDRR Operation



CONGESTION MANAGER

The previous section describes the way that the scheduler provides isolation for multiple traffic classes. The scheduler has no way to differentiate between traffic in the same queue. The congestion manager, known as drop preference, provides protection of one traffic class (an application or a user) from the misbehavior of other traffic classes within the same queue. The Cisco 12000 Series Routers use WRED as a mechanism for adjusting the congestion notification to TCP sources without causing TCP global synchronization. WRED is an active queue-management technique that uses a weighted value of a queue's average occupancy, discard thresholds, traffic priority (such as DCSP or MPLS Exp bits), and a random function to decide whether to drop a packet during network congestion. When multiple traffic classes are combined into the same queue, each class uses a different set of RED parameters. Typically, the drop threshold for higher-priority traffic is set above those used for lower-priority traffic. This means that as the average depth of the queue increases, the drop probability increases for the lower-priority traffic first, while the drop probability for the upper priority remains the same.

Unlike other network equipment vendor implementations of WRED, the Cisco 12000 Series Router computes the drop probability of an incoming packet based on a weighted average queue length instead of the current queue length. By using an average queue length, it allows RED to not react to a short (and transient) burst and react only to persistent congestion. At the other end of a connection, drop probability based on current queue length adversely affects QoS and network stability.

TRAFFIC ENGINEERING

Network traffic volumes can vary over time and by circumstance. Real-world variables such as time of day, day of the week, or natural disasters can create extra traffic loads that cannot be easily handled by the network paths computed by the IP layer. Traffic engineering is the ability to control an entire path that traffic will use to move from an ingress point to an egress point within a network. This means that each individual flow can be forwarded independently from the rest of traffic and independently from the paths chosen by Layer 3. Based on the offered load, a portion of traffic can be dynamically shifted over less-utilized paths. An added benefit is higher utilization of available bandwidth.

The connection characteristics of MPLS provide a practical way to exercise traffic engineering within an IP network. MPLS tunnels are created explicitly between an ingress router and an egress router of a network to carry a portion or all IP traffic between the ingress and egress routers. The hierarchical and connection characteristics of MPLS provide the ability to not only set up a traffic path, but also to define a protected path around a link to cover a link or node failure. This concept is referred as fast reroute.

TRAFFIC MEASUREMENT

The most common uses of IP traffic measurement are:

- *Traffic peering*—NSAPs selects the peering partners (public and private peering) based on the history of traffic exchanged with a given autonomous system
- *Traffic engineering*—Within an autonomous system, network designers may use traffic history to understand the traffic trend to load-balance traffic across alternate paths
- *Resource tracking*—P resource characterization such as IP address distribution per continent/per country, traffic breakdown per protocol/application (for example VoIP, multimedia)
- *Flexible usage*—based billing-Including destination autonomous system (AS)-based, class-of-service-based billing
- *(SLA) measurement and usage-based billing*—Service providers need can provide reports to their customers on the service level that users are getting (packet loss per class of traffic, peak and mean round-trip time, and so on)

The Cisco 12000 Series Router implements a set of traffic-measurement features ranging from performance monitoring to application/flow level accounting and billing. The distributed Cisco NetFlow accounting feature, for example, captures all, or a portion of IP packets (of the same flow) flowing across a router to generate high-quality performance traffic statistics. These traffic statistics include MPLS labels, IP addresses, protocol type, protocol ports, interfaces, type of service (ToS), and other TCP/IP header information. Statistics can be exported either to a Cisco NetFlow Collector Workstation or to a third-party application for further processing. Because the export interface is open, several software applications are available to customize exported statistics for specific customer requirements such as security, billing and performance. Source MAC and border gateway protocol (BGP) accounting are other measurement features typically suitable at the peering points connected through an Ethernet or FDDI interface. In this case, the MAC address identifies the partner.

Cisco continues to build on its technical leadership in high-end Internet routing with the Cisco 12000 Series routers that provide extensive hardware and software support for real-time services. The Cisco 12410 and 12416 are the two highest-capacity Internet routers available at 200 Gbps and 320 Gbps, respectively, and feature guaranteed priority packet delivery (GPPD). An industry exclusive, GPPD provides priority-based congestion control, dedicated low-latency queuing, and packet sequence integrity under all conditions. Service providers can seamlessly scale IP network core infrastructure to the highest traffic volumes provided by optical technology using the industry's most extensive support for real-time services to deliver high-revenue premium services demanded by an increasingly sophisticated Internet customer base.

**Corporate Headquarters**

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

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

All contents are Copyright © 1992–2004 Cisco Systems, Inc. All rights reserved. CCIP, CCSP, the Cisco Arrow logo, the Cisco Powered Network mark, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MGX, MICA, the Networkers logo, Networking Academy, Network Registrar, *Packer*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0402R) MC/N2/LW5799 0304