

Quality of Service for Virtual Private Networks


This paper aims to address the topic of quality of service (QoS) and the tools available from Cisco Systems for designing a Virtual Private Network (VPN) with appropriate service levels for mission critical applications. This paper is not meant to be a design guide but more of a glimpse into some of the QoS technologies available from Cisco to help implement a successful VPN. The audience is expected to be familiar with VPN related issues such as security, firewalls, and routing/switching offerings from Cisco.

This paper also provides an introduction to the end-to-end QoS architecture for VPNs from Cisco, including rich classification, policing, shaping, queuing and congestion avoidance. The paper also touches upon the future of QoS policy deployment using Common Open Policy Service (COPS) and Cisco QoS policy servers.

The Need for VPNs

VPNs aim to give the remote corporate user the same level of access to corporate computing and data resources as the user would have if she were physically present at the corporate headquarters. By reducing the costs of transporting data traffic and by enabling network connections in locations where they would not be affordable, VPNs reduce the total cost of ownership of a corporate network.

Figure 1 VPN



VPNs can be categorized into three types:

- Remote-access VPNs—to connect telecommuters and mobile users to the enterprise WAN.
- Intranet VPNs—to connect branch offices and home offices within an enterprise WAN.
- Extranet VPNs—to give business partners limited access to the corporate WAN.

VPNs offer a cost-effective, scalable, and manageable way to create a private network over a public infrastructure such as the Internet or over a service provider's Frame Relay, ATM, or IP network. However VPNs will not be a viable alternative unless they can guarantee a predictable bandwidth, reliability, and security to users. With different traffic types on a WAN link vying for scarce bandwidth the irony is that at times the WAN links may be under-utilized while at other times users might experience the effects of extreme congestion, especially during peak hours. What is called for is a set of standards-based QoS tools which can provide appropriate QoS treatment for the customer's traffic.

The Need for QoS

Users of a widely scattered VPN do not usually care about the network topology or the high level of security/encryption or firewalls that handle their traffic. They don't care if the network implementers have incorporated IPSec tunnels or GRE tunnels. What they care about is something more fundamental, such as:

Do I get acceptable response times when I access my mission critical applications from a remote office?

Acceptance levels for delays vary. While a user would be willing to put up with a few additional seconds for a file transfer to complete, the same user would have less tolerance for similar delays when accessing a database or when running voice over an IP data network.

QoS aims to ensure that your mission critical traffic has acceptable performance. In the real world where bandwidth is finite and diverse applications from videoconferencing to ERP database lookups must all vie for scarce resources, QoS becomes a vital tool to ensure that all applications can coexist and function at acceptable levels of performance.

Cisco QoS for VPNs

Cisco Systems offers a comprehensive tool chest of QoS features bundled in Cisco IOS® that are applicable to VPNs. The primary QoS building blocks of VPNs are:

- Packet classification (using Committed Access Rate [CAR])
- Bandwidth management (policing with CAR, shaping with GTS/FRTS, bandwidth allocation with WFQ)
- Congestion avoidance (with WRED)
- Continuity of packet priority over Layer 2 and Layer 3 VPNs (with tag switching/Multi Protocol Label Switching [MPLS])


Packet Classification

The aim of packet classification is to group packets based on predefined criteria so that the resulting groups of packets can then be subjected to specific packet treatments. The treatments might include faster forwarding by intermediate routers and switches or lesser probability of the packets being dropped due to lack of buffering resources.

It is necessary that traffic be classified before tunneling and encryption since otherwise the tunnel header that is appended to the IP packet would make the QoS markings in the IP header invisible to intermediate routers/switches, which need to read this information and act upon it. Classification brings into question the right match criteria. There are a number of criteria based upon which we may classify traffic before it enters the VPN:


- IP addresses
- TCP/UDP port numbers
- IP precedence (3 bits in the type of service (ToS) field of the IP packet header)
- URL and sub-URL
- MAC addresses
- Time of day

Figure 2 Classification at network ingress



Once we classify packets based on the above criteria the next step is to “mark” or “color” packets with a unique identification to ensure that this classification is respected end to end. The simplest way of doing this is via the IP ToS field in the header of an IP datagram. In the near future the Internet Engineering Task Force (IETF)- sponsored Differentiated Service Code Points (DSCP) could become the classification criterion of choice. The purpose behind this type of marking of packets is to ensure that downstream QoS features such as scheduling and queuing may accord the right treatment for packets thus marked. In some cases the service provider whose backbone is being used for the VPN might provide differentiated services, classification allows you to leverage these services.

Figure 3 ToS field in the IP Packet header



Differentiated services allow certain network traffic to receive premium treatment at the expense of other less-critical traffic on the same wide area network (WAN) link. This idea is similar to what we find in airlines where a first-class passenger may receive better treatment or service than an economy-class passenger while they both physically reside on the same airplane.

Bandwidth Management


Once traffic has been classified the next step is to ensure that it receives special treatment in the routers. This brings into focus scheduling and queuing.

Before we get into the subject of queuing it might be good to step back and consider what we mean by a flow. For this discussion a flow would be a group of packets which share a common criteria whether that criteria is a source/destination IP address or a TCP/UDP port number or a protocol or a type of service (TOS) field. Cisco provides two implementations of weighted fair queuing (WFQ): Flow-based WFQ and Class-based WFQ.

In flow-based WFQ, packets are classified by flow. Each flow corresponds to a separate output queue. When a packet is assigned to a flow, it is placed in the queue for that flow. During periods of congestion, WFQ allocates a portion of the available bandwidth to each active queue.

Class-based WFQ aims for providing weighted fair queuing functionality among traffic classes defined by the user. A user could create traffic classes using mechanisms like Access Control Lists (ACLs) and then assign a fraction of the output interface bandwidth to each of these traffic classes. The primary difference between flow-based WFQ and class-based WFQ is the fact that in flow-based WFQ bandwidth allocation is relative to other flows. But in class-based WFQ bandwidth allocation is absolute. Class-based WFQ allows the user to assign bandwidth to a class based upon a percentage of the available bandwidth or a fixed kbps value.

Figure 4 Weighted Fair Queuing



When to Use Class-Based WFQ Versus Flow-Based WFQ

Flow-based WFQ as it existed in Cisco IOS did not differentiate between traffic classes. As far as flow-based WFQ was concerned a packet was part of a flow. The flow could be based on source/destination address, TCP/UDP port number or some other criteria. There was no real bandwidth guarantee since the weights were assigned based on IP Precedence. There was no way to ensure that Hyper Text Transport Protocol (HTTP) based web traffic would have a higher guarantee of bandwidth over traffic conforming to FTP (File Transfer Protocol). Class-based WFQ gives users the following benefits which were not possible with flow-based WFQ:

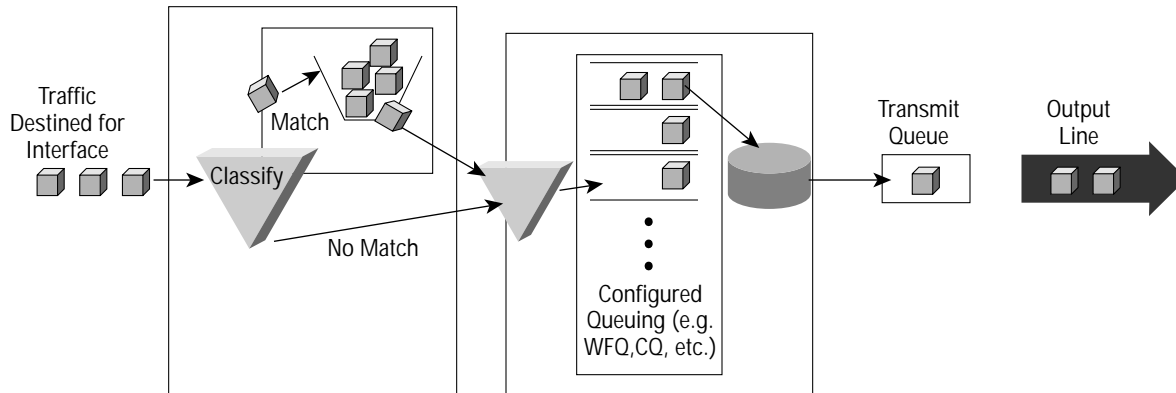
- Bandwidth guarantees for an application
- User defined traffic classes

In conclusion flow-based WFQ provides QoS guarantees that are relative to other flows whereas class-based WFQ provides for absolute QoS guarantees.

Traffic Shaping


Traffic shaping becomes necessary when Layer 3 traffic must be shaped to a desired set of rate parameters to enforce a maximum traffic rate. The result will be a smooth traffic stream¹. Traffic shaping queues and forwards data streams (as opposed to dropping excess traffic) so as to conform to agreed upon Service Level Agreements (SLAs).

Figure 5 Generic Traffic Shaping



- The idea behind traffic shaping is that if bursty traffic (characterized by fits and starts) is queued then the TCP senders will realize this and in turn will back off and ensure that subsequent transmissions conform to a desired rate. This type of traffic is commonly referred to as adaptive traffic. The end result of traffic shaping is a smoothed packet stream.

Figure 6 Policing traffic



When to Use a Traffic Policer Versus a Traffic Shaper

Policing literally means to drop excess traffic, shaping on the other hand allows the excess traffic to be queued. To an application shaping is usually a better choice since shaped traffic does not require re-transmission while dropped traffic would require a re-transmission. In such cases Cisco Generic Traffic Shaping (GTS) is the tool of choice.

1. Layer 3 in the OSI model conforms to the network layer. In this instance the IP layer.

However if you shape in every instance then you might end up with very deep queues in a router which might result in re-transmission by the sender due to perceived delay. Policing/dropping of excess traffic is better suited to IP multicasts or to TCP-based traffic related to non-mission critical applications.

Congestion Avoidance

Congestion avoidance could be defined as the ability to recognize and act upon congestion on the output direction of an interface so as to reduce or minimize the effects of that congestion.

Congestion produces adverse effects in a VPN and should be avoided. With this in mind Cisco Systems provides IOS based tools like Weighted Random Early Detection (WRED) which is a Cisco implementation of the Random Early Detection (RED) algorithm. WRED provides for differential treatment of traffic by adding per-class queue thresholds which determine when packet drops will occur. The thresholds are user-configurable and set using the command line interface (CLI) in Cisco IOS.

Figure 7 Weighted Random Early Detection






Figure 8 WRED Drop Thresholds



Packet dropping is based upon the premise that adaptive flows such as TCP will back off and retransmit if they detect congestion. By monitoring the average output queue depth in the router and by dropping packets from selected flows WRED aims to prevent the ramp up of too many TCP sources at once. Unchecked this ramping up could result in problems such as TCP synchronization.

Figure 9 Global TCP synchronization



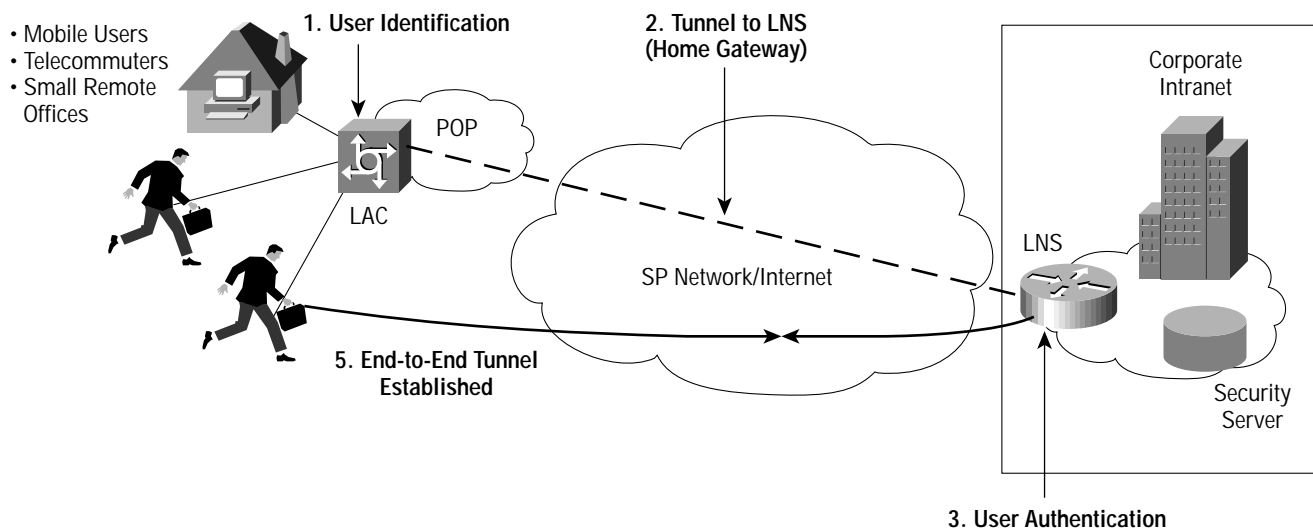
WRED provides for differential treatment by dropping packets from low priority traffic before it begins to drop packets from high priority traffic. WRED allows the user the option to select up to six such traffic classes (standard and premium being just two used for illustrating the point).

QoS for VPN tunnels

The QoS issue here is that the QoS parameter normally found in the header of the IP packet should be reflected in the tunnel packet header regardless of the type of tunnel in use. Consider the four primary tunneling protocols relevant to VPNs:

- Layer 2 Tunneling Protocol (L2TP) Tunnel
- IP Security (IPSEC) Tunnel
- Layer 2 Forwarding (L2F) Tunnel
- Generic Route Encapsulation (GRE) Tunnel

Figure 10 L2F/L2TP Operation

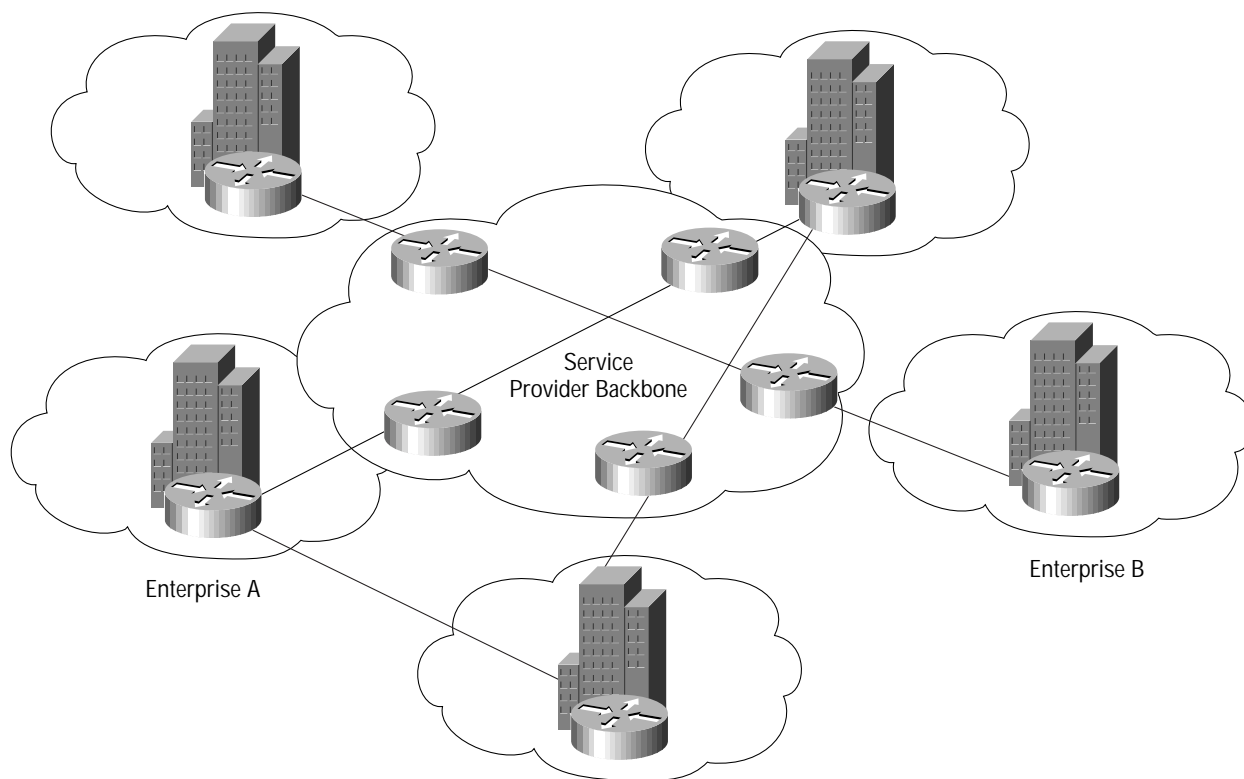


Layer 2 Tunneling Protocol (L2TP) is commonly used for node-to-node applications where the tunnel terminates at the edge of the user's network. L2TP is an IETF based standard which merges Cisco's Layer 2 Forwarding (L2F) Tunnel protocol with Microsoft's Point-to-Point Tunneling (PPTP) protocol. L2TP relies on 3rd party security schemes like IPSEC to secure packet level information. L2TP was designed primarily for Point to Point Protocol (PPP) traffic.

Generic Route Encapsulation (GRE) tunnels based on RFC 1702 allows any protocol to be tunneled in an IP packet. Today Cisco offers support for encapsulation of data using either IPSEC or Generic Route Encapsulation (GRE). In either of these cases Cisco IOS offers the ability to copy the IP ToS values from the packet header into the tunnel header. This feature which appears in IOS ver 11.3T allows the Type of Service (ToS) bits to be copied to the tunnel header when the router encapsulates the packets using GRE.

It allows routers between GRE-based tunnel endpoints to adhere to precedence bits thereby improving the routing of premium service packets. Now Cisco IOS QoS technology such as policy routing, WFQ, and WRED can operate on intermediate routers between GRE tunnel endpoints.

Figure 11 GRE Tunnel architecture



IETF Differentiated Services

The Internet Engineering Task Force (IETF) is currently working on a QoS model called "Differentiated Services" or more commonly DiffServ. DiffServ redefines the IP Type of Service (ToS) byte into the DiffServ Byte ("DS Byte"). This is used to signal the required QoS level for a packet. It is also used to identify packets as belonging to one class or another. DiffServ defines Per-Hop Behaviors (PHBs) which will foster common QoS behaviors in the network. The aim is to provide the basis for standards-based QoS in a VPN from end-to-end.

Conclusion

Many tools are available today from Cisco to help you manage your WAN links. These range from QoS tools for Traffic Classification, Policing / Shaping, Bandwidth Allocation, and Congestion Avoidance. The IETF DiffServ efforts provide the basis to extend Enterprise QoS policy into the SP network and Cisco intends to participate and monitor these efforts to ensure that our QoS products meet this new model's criteria for end-to-end QoS.



Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<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
Villebon, BP 706
91961 Courtaboeuf Cedex
France
<http://www-europe.cisco.com>
Tel: 33 1 69 18 61 00
Fax: 33 1 69 28 83 26

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

Asia Headquarters
Nihon Cisco Systems K.K.
Fuji Building, 9th Floor
3-2-3 Marunouchi
Chiyoda-ku, Tokyo 100
Japan
<http://www.cisco.com>
Tel: 81 3 5219 6250
Fax: 81 3 5219 6001

Cisco Systems has more than 200 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/offices>.

Argentina • Australia • Austria • Belgium • Brazil • Canada • Chile • China • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE
Finland • France • Germany • Greece • Hong Kong • 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 • Singapore
Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela