

Using Content Networking to Provide Quality of Service

Introduction

Content Networking

Cisco Content Networking delivers the network agility required by the enterprise to deploy new Internet business applications critical to securing competitive advantage by increasing revenue while reducing operating costs. By creating the end-to-end intelligent network services required for Internet business applications such as e-commerce, supply chain management, and workforce optimization, Cisco Content Networking integrates the enterprise with customers, suppliers, and business partners.

Cisco Content Networking is an intelligent network architecture that dynamically recognizes Internet business applications and engages network services to achieve end-to-end security, performance, and availability. This architecture has three components:

- Intelligent network classification and network services delivered through Cisco IOS[®] software
- Intelligent network devices that integrate Internet business applications with network services
- Intelligent policy management framework for configuration, monitoring, and accounting

With these three components, the dynamic architecture of Cisco Content Networking delivers the intelligent network services required to drive the next-generation Internet business model.

New Internet Business Applications

Agile enterprises are deploying a dizzying array of new applications that support a wide variety of business processes, including workforce optimization, supply chain management, e-commerce, and customer care. The network must be able to dynamically recognize each application and provide the associated set of services. For example, when implementing a workforce optimization solution, remote sites need to access Web pages hosted at corporate headquarters. The Web pages need to be protected through security controls and the bandwidth over the remote link must be properly managed. The network needs to provide a variety of services, including security and quality of service (QoS).

Cisco intelligent network devices provide a wide range of services necessary to deploy Internet business applications. Many of these applications utilize advanced client/server technologies and leverage Web technology. This makes it difficult for the network to properly classify them. Fortunately, Cisco intelligent network classification features, including context-based access control (CBAC) and network-based application recognition (NBAR), enable the network to recognize and provide the proper set of services for these applications.

NBAR, a new Cisco IOS software feature, provides intelligent network classification for use with the broad array of QoS features available.

Public

Copyright © 1999 Cisco Systems, Inc. All Rights Reserved.

Page 1 of 8

Types of Quality of Service

Overview

QoS is a fundamental service required for next-generation business traffic. As enterprises shift mission-critical applications to the network and create a unified multiservice architecture for data, voice, and video, the ability to manage delivery terms becomes increasingly critical. Enterprises need QoS to deploy on-demand bandwidth-intensive applications such as videoconferencing, as well as time-sensitive information delivery applications such as stock transactions. QoS also reduces wide-area network (WAN) costs through more efficient use of network links. With QoS in place, network administrators worry less about “healthy” levels of oversubscription and concern themselves instead with managing the traffic mix to ensure that time-sensitive and mission-critical traffic is not delayed. QoS allows organizations to converge their data, voice, and video traffic.

The quality of a network connection is described in terms of availability, latency, jitter, and capacity. Availability is the assurance that traffic will reach its destination successfully, and forms the basis of most service-level agreements (SLAs) today. Latency is the delay that traffic experiences as it travels across the network; jitter is the change in this latency over time. Capacity is the total amount of bandwidth available on a link.

Establishing a particular QoS level for a connection is a complex process, in part because of the stateless, best-effort paradigm upon which the Internet is based and the fact that one must balance all of the QoS parameters discussed above. There are two main approaches to QoS: the integrated services model and the differentiated services model.

The integrated services model, or intserv, negotiates a particular QoS at the time it is requested. Before exchanging traffic, the sender and receiver request a particular QoS level from the network. Upon acceptance, the intermediate network devices associate the resulting traffic flow with a specific level of jitter, latency, and capacity. Resource Reservation Protocol (RSVP), a protocol for signaling QoS requirements for a particular traffic flow, is a key component in this model.

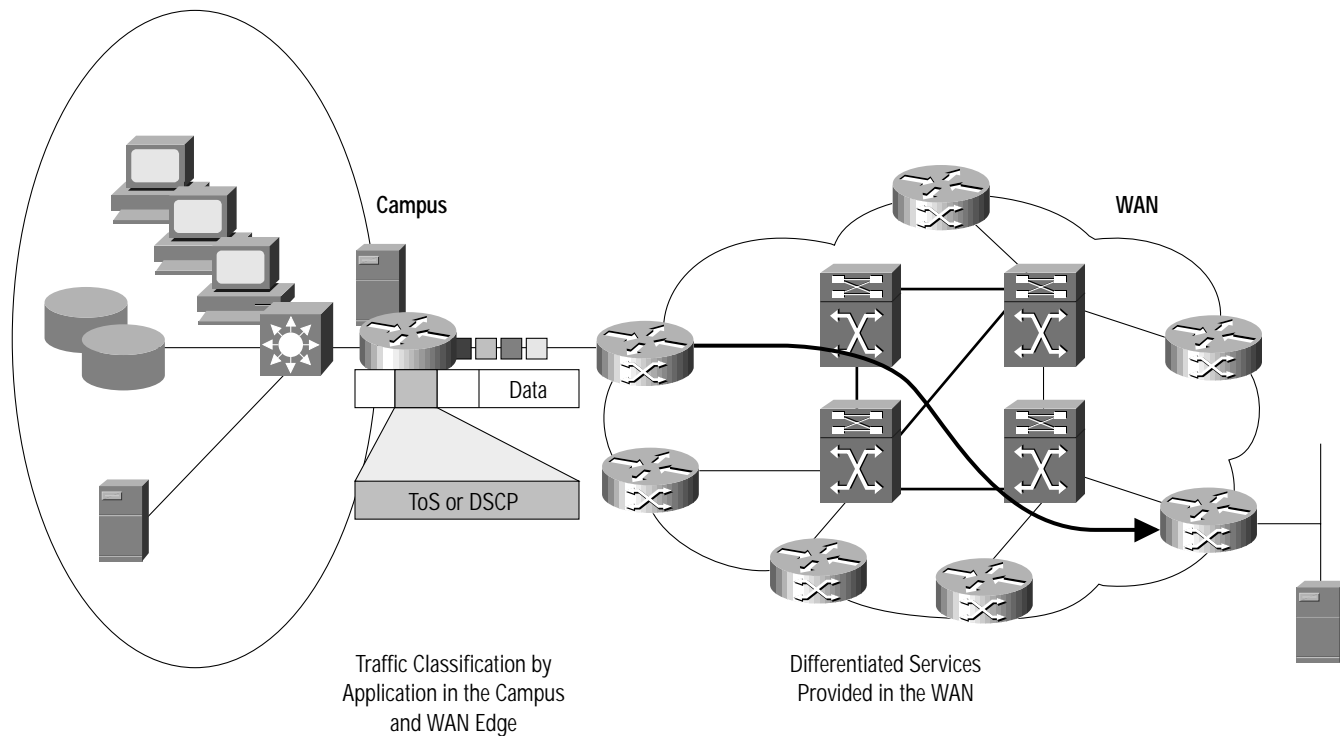
The differentiated services model, or Diff Serv, takes a different approach. A few, coarse classes of traffic handling—similar to gold, silver, or bronze levels of frequent flier cards—are established by the network administrator. When the sender needs a particular kind of handling, it marks the individual packets accordingly.

Differentiated Services

The differentiated services architecture is based on a simple model where traffic entering a network is classified and possibly conditioned at the boundaries of the network. The class of traffic is then identified with a DS code point or bit marking in the IP header. Within the core of the network, packets are forwarded according to the per-hop behavior associated with the DS codepoint.

One of the primary principles of the differentiated service model is that one should mark packets as close to the edge of the network as possible. It is often a difficult and time-consuming task to understand to which class of traffic a given data packet belongs, so you want to classify the data as few times as possible. By marking the traffic at the network edge, core network devices and other devices along the forwarding path will be able to quickly determine the proper class of service (CoS) to apply to a given traffic flow. Figure 1 illustrates the application of differentiated services in a network. In this figure, marking occurs at the WAN-edge router so that the WAN core can provide three different classes of service.

Figure 1 Differentiated Services



Another option is for the originating application or device to mark the packet. This provides even greater scalability and flexibility as each device will only handle its own data stream and will be able to include more contextual information in deciding how to mark the packet. However, this method also gives control to the end user. There is nothing to prevent them from marking all of their traffic as “Gold Service.” Network-based marking provides greater control over how data is treated on the network, as it is not possible for end users to erroneously mark packets with the wrong priority. It is also easier to manage a smaller number of network devices than the orders of magnitude larger number of individual end systems. New and emerging technologies may make it possible to push classification and marking out to the end systems, but for now most enterprises will probably choose to mark packets within the network infrastructure.

Another important aspect of differentiated services is that the markings must be consistently interpreted from end to end as the marking does no good unless all the devices in the network path understand the per-hop behavior that should be applied to a given class of traffic. The connection quality will only be as good as the weakest link, so if one of the routers in the path does not act appropriately, the overall service for the given packet may not be as desired.

Network-Based Application Recognition

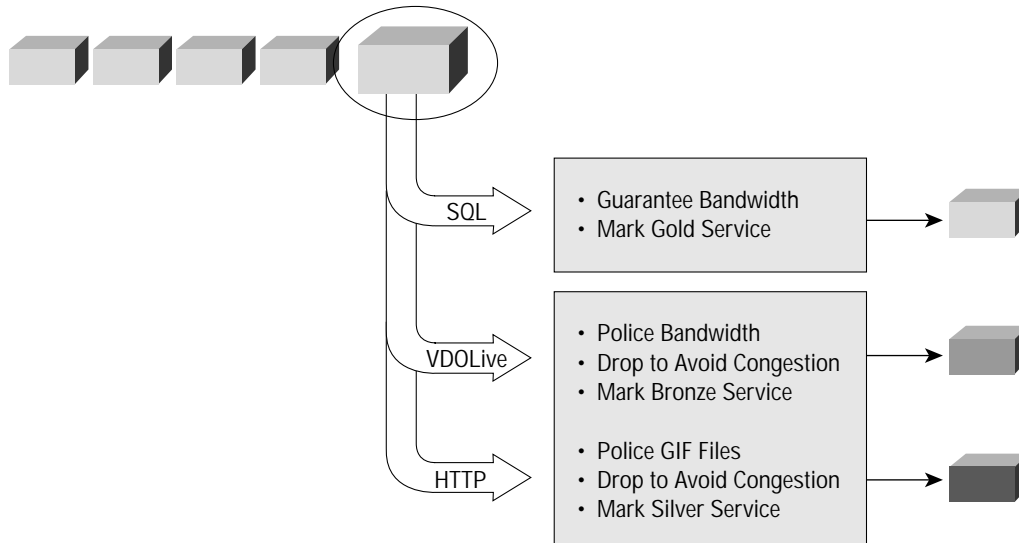
Overview

Classifying applications sounds like a simple task, but it is often a hard undertaking. The difficulty is that today’s Internet-based and client-server applications make it difficult for the network to identify and provide the proper level of control you need.

NBAR, a key component of the Cisco Content Networking architecture, solves this problem by adding intelligent network classification to your infrastructure. NBAR is a new classification engine that can recognize a wide variety of applications, including Web-based and client/server applications that dynamically assign Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) port numbers. Once the application is recognized, the network can invoke specific services for that particular application. NBAR currently works with QoS features to ensure that the network bandwidth is best used to fulfill your objectives. These features include the ability to guarantee bandwidth to critical applications, limit bandwidth to other applications, drop selective packets to avoid congestion, and mark packets appropriately so that your network and the service provider’s network can provide the proper QoS from end to end.

For example, your customer service representatives require a fast response when they query for an order status in the corporate data warehouse hosted on an Oracle database server. Unfortunately, if others on the network are using high-bandwidth applications, such as VDOLive or viewing large GIF files, then the SQL*NET transaction to the Oracle database may be delayed. NBAR addresses this problem by properly classifying the applications and then providing guaranteed bandwidth to the SQL*NET queries while simultaneously policing the other applications. Figure 2 illustrates this solution.

Figure 2 NBAR Provides Intelligent Network Classification



Details about NBAR Classification

NBAR supports a wide range of network protocols, including these stateful protocols that were once difficult to classify, including:

- HTTP classification by URL, Host and MIME type
- Citrix published application
- Oracle SQL*NET
- Sun RPC
- Microsoft Exchange
- UNIX r commands
- VDOLive
- RealAudio
- Microsoft Netshow
- File Transfer Protocol (FTP)
- StreamWorks
- Trivial File Transfer Protocol (TFTP)

NBAR can also classify static-port protocols such as those currently classifiable with access control lists (ACLs). New protocols support can be quickly and easily added via packet description language modules (PDLM) from Cisco. PDLMs contain the rules used by NBAR to recognize an application and in most cases can be loaded without the need for a new Cisco IOS software image or even a reboot.

In addition to classification, NBAR can assist you in developing your QoS policy. This is the set of rules that you want applied to your applications. For example, you might want to provide a certain amount of guaranteed bandwidth for mission-critical applications while policing other applications to a maximum bandwidth. NBAR’s protocol-discovery feature shows you the mix of applications currently running on the network so that you can have some knowledge of network traffic conditions before crafting the policy.

QoS Services Supported by NBAR

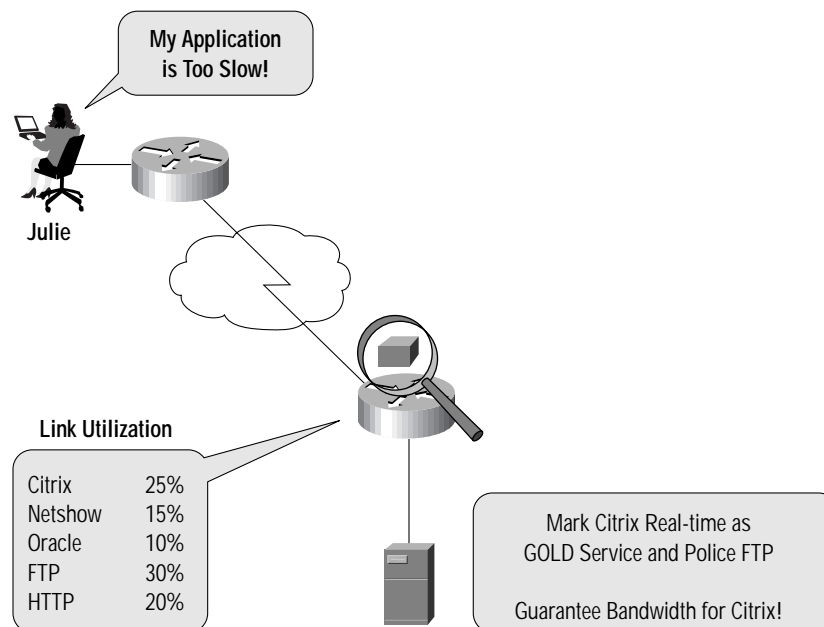
NBAR is just one component of the solution. Once the traffic is classified, QoS mechanisms must be deployed to provide the proper level of service. The following features may be applied to traffic that's been classified by NBAR:

- Guaranteeing bandwidth with class-based weighted fair queuing (CBWFQ). CBWFQ is a mechanism to provide guaranteed bandwidth to particular traffic classes while still fairly serving all other traffic in the network. To characterize a class, you assign it bandwidth, weight, and maximum packet limit. The bandwidth assigned to a class is the minimum bandwidth delivered to the class during congestion.
- Marking for differentiated service downstream using the type of service (ToS) bits or Diff Serv code points (DSCP). As discussed above, packet marking is usually done at the edges of the network so that the proper policies may be applied later in the routing path. The older ToS byte markings provide three bits in the IP header, of which six combinations are usable. The newer DSCP mechanism utilizes six bits in the IP header, providing 64 separate markings. The bits used by both mechanisms overlap, so one would generally utilize either ToS or DSCP values in the network, but not both.
- Enforce bandwidth limits via policing. The police feature allows one to limit the bandwidth available to a particular application. For example, one can limit FTP traffic to at most 1 mbps on a DS3 WAN link. Any traffic over the policed rate will be dropped or the ToS or DSCP values may be changed to a lower class of service.
- Congestion avoidance policy (random early detection [RED]). Congestion -avoidance techniques monitor network traffic loads in an effort to anticipate and avoid congestion at common network bottlenecks. Without RED, when the output queue is full, packets are dropped until the congestion is eliminated. Unfortunately, global traffic synchronization may occur as waves of congestion crest only to be followed by troughs during which the transmission link is not fully utilized. This occurs when multiple TCP hosts simultaneously reduce their transmission rates in response to packet dropping, and then increase their transmission rates once again when the congestion is reduced. RED takes advantage of TCP's congestion-control mechanism by randomly dropping packets before the link is completely congested. This avoids the problem of dropping packets from all the hosts at one time while still signaling that they should temporarily slow down their packet transmission.

Example Usage

Before describing more of the benefits and example applications for NBAR, it may be helpful to give a quick example.

Figure 3 Using NBAR to Provide QoS



In Figure 3, Julie is trying to complete her white paper in Microsoft Word. Her firm doesn't load Word locally on her desktop, but instead uses Citrix to centrally host the application. Unfortunately, the application is too slow and so Julie calls the help desk to complain.

Using NBAR, the network manager is able to quickly determine which protocols are running on the link to the Citrix server. As seen in the diagram, the link is 100 percent full and FTP is using 30 percent of the link. Using NBAR to classify Citrix real-time data, such as screen updates, the network manager can mark Citrix packets for Gold Service and use CBWFQ to ensure those packets get priority in the network. At the same time, he can police FTP so that it does not use up so much bandwidth. The end result is that Julie can complete her white paper without further delay.

Application Details

In this section we'll explore how NBAR and the associated QoS services can be used to enable Internet business applications.

Mission-critical applications, such as Oracle, Citrix, Microsoft Exchange, or the new breed of Web-based applications, must perform well to ensure your success in today's fast paced e-business environment. Bottlenecks may occur in many places, including the network. These bottlenecks often occur even though you've budgeted what you thought was more than adequate bandwidth for each application.

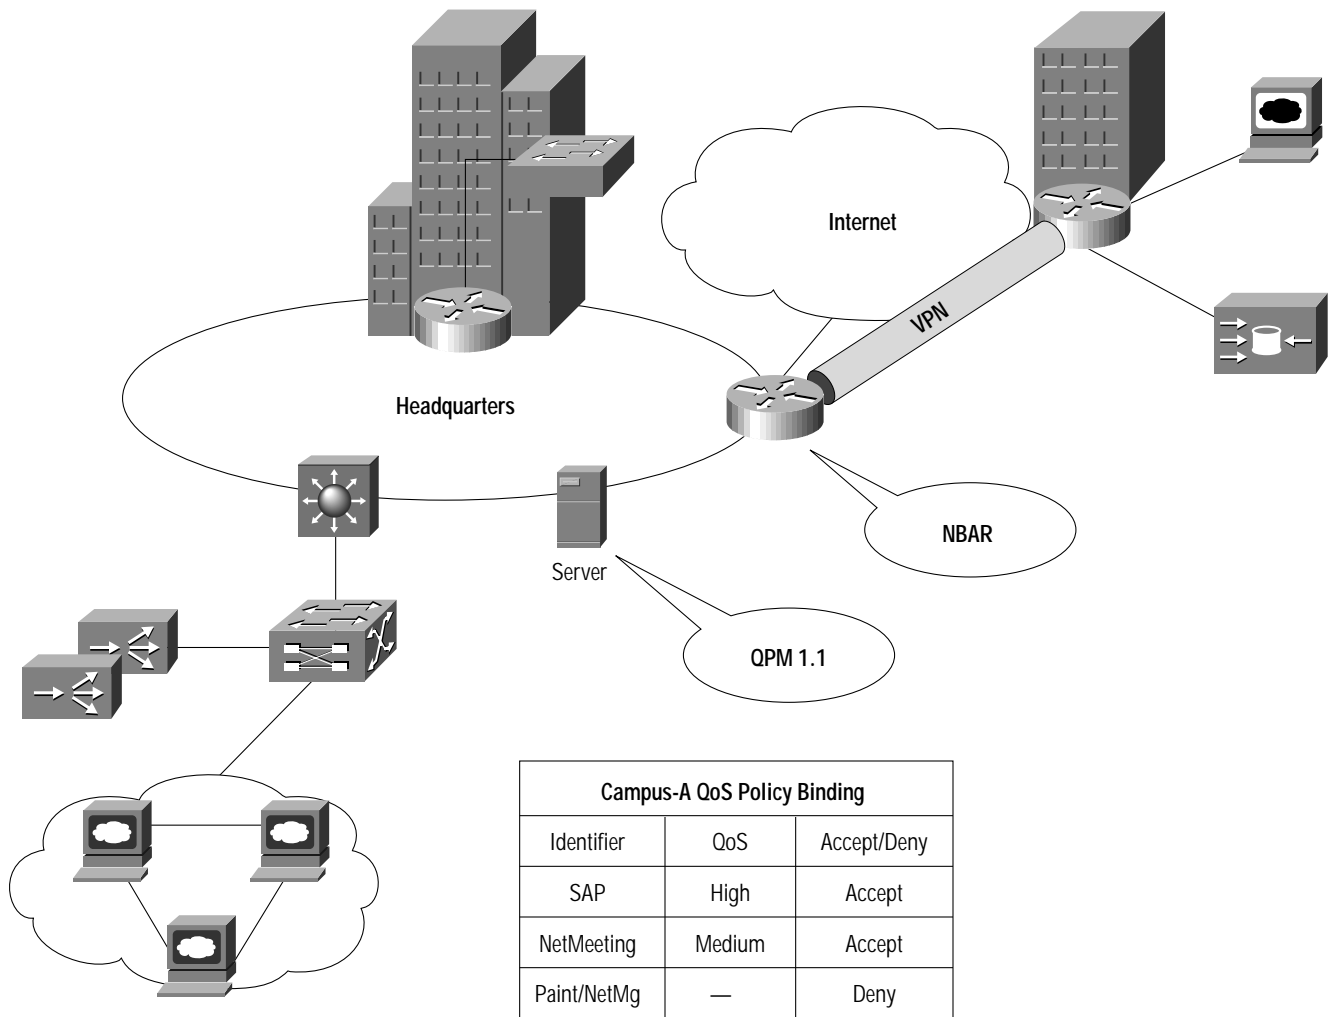
What often happens is that employees take advantage of new Internet applications, such as streaming audio and video or downloading of new programs. All of these applications can quickly consume your WAN bandwidth. Unfortunately, these are not typically the mission-critical applications you want to give priority to on the network.

NBAR, by intelligently classifying applications, allows the network to provide differentiated services to each application. You can provide absolute priority and a guaranteed amount of bandwidth to your mission-critical applications, such as Oracle or an application that runs on a particular Web page. At the same time you can limit the bandwidth consumed by the "bandwidth hogs." The end result is that users can access their mission-critical applications with minimal delay without the need to upgrade costly WAN links or cutting off access to commonly used, but not mission-critical, applications.

Figure 4 illustrates an intranet used for workforce optimization. The remote office is connected via a virtual private network (VPN). VPNs reduce networking costs while providing increased flexibility. Unfortunately, the service quality in a VPN is often difficult to guarantee. Running NBAR and VPN concurrently in the same router solves this problem by identifying mission-critical traffic before it is encrypted, allowing the network to apply the appropriate QoS controls. By running both VPN and NBAR concurrently, it is ensured that the packets are processed in the correct order to achieve both maximum security and the appropriate QoS level. In this example, QoS Policy Manager Version 1.1 is used to configure the QoS policy for the network.

If a remote employee accesses a critical enterprise resource planning (ERP) application, NBAR will identify the packet, mark it as a Gold Service packet, and, using CBWFQ, place it into a priority queue. The VPN processes will then tunnel and encrypt the packet while maintaining the Gold Service marking on the new packet. When used with service providers that offer differentiated services on their networks, the ERP application will receive priority treatment as it travels through the VPN. Other less critical applications accessed by the same employee may be given lower priority service or be restricting with policing.

Figure 4 Workforce Automation Network



NBAR also allows you to provide priority service to Web pages and types of Web content that you deem critical. For example, sales tools can be given absolute priority and guaranteed bandwidth, ensuring that your sales force never has to wait for a price quote because another employee is watching the firm's new television commercial on streaming video. In addition you can police the amount of bandwidth consumed by certain classes of content, such as JPEG pictures to further ensure that the mission-critical Web pages are easily and quickly viewed.

Combining data, voice, and video on your network is another way to optimize your workforce. Unfortunately, each of these services requires different network characteristics. NBAR is able to intelligently identify the type of each packet, and in combination with the QoS features, ensure that each application is treated appropriately.

For example, if you deploy a training system that utilizes streaming video, such as Cisco IP/TV[®] technology, you'll want to ensure that employees see a clean picture, not one that is choppy and hard to see. With NBAR, the network can easily recognize the streaming video traffic and assign it to a higher priority class of traffic that receives a minimum guaranteed bandwidth. Other traffic, such as e-mail, can be assigned to a lower priority class, as e-mail must be delivered, but doesn't have the latency and bandwidth constraints of the streaming video. The end result is that the trainee receives their video training on demand with high quality while the network concurrently serves other applications.

Citrix WinFrame, another application you may be running on your network, allows employees to use their familiar PC-based applications while providing the manageability and flexibility of centralized server-based solutions. NBAR can classify Citrix Independent Computing Architecture (ICA) traffic and perform subport classification of Citrix traffic based on Citrix published applications. The appropriate level of service can then be applied to the various applications deployed in your Citrix environment.

Finally, NBAR helps you control WAN costs. In many parts of the world, and especially between countries, telecommunications links can still be prohibitively expensive. This leads to a dilemma for the network manager: on the one hand you need to provide access to new client-server and Internet-enabled applications, while on the other hand you need to control WAN service costs. NBAR provides a solution to this problem by enabling you to intelligently utilize WAN bandwidth so that you can provide acceptable service levels with the minimum possible bandwidth. As discussed before, the higher priority traffic can be given guaranteed bandwidth using CBWFQ while lower priority traffic can be policed or selectively dropped using weighted RED (WRED) to avoid congestion.

Conclusion

NBAR, a key component of the Cisco Content Networking architecture, adds intelligent network classification to your infrastructure to ensure that your mission-critical applications receive the bandwidth they need. The combination of intelligent network classification as provided by NBAR with intelligent network services provides the network agility your organization needs to be successful in the fast-paced Internet business environment.



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