White Paper

Cisco Bandwidth Control for Education Networks

Today's colleges and universities are using a new class of rich media applications to aid and enhance the educational process. However, as high-bandwidth voice, video, and e-learning applications proliferate, the demand for bandwidth resources grows as well. The popularity of recreational peer-to-peer (P2P) and gaming applications only exacerbates this situation and, in some cases, can even bring the entire campus network to a halt. Higher bandwidth networks also offer an attractive target for hackers, and emerging denial-of-service (DoS) and worm attacks can take advantage of bandwidth-rich environments to propagate faster than before, flooding network resources. To protect against these problems, many institutions are turning to new security strategies from Cisco Systems®. The Cisco® Bandwidth Control for Education Networks Solution provides a powerful set of tools to implement and enforce institutionwide bandwidth control policies and unlock the full potential of high-bandwidth academic networks.

INTRODUCTION

For students and faculty at today's academic institutions, high-speed network access and reliability have become baseline expectations. Never before has so much rich media content been used in educational material. At the same time, the recreation habits of students accustomed to broadband access at home from a young age present a challenge to educators needing to balance the recreational and educational needs of students.

In this environment, college and university IT departments are increasingly challenged to fairly distribute network resources and try to ensure that legitimate academic applications have the bandwidth they need. Over the past several years, many academic institutions have found that mission-critical and revenue-generating applications such as Blackboard/WebCT, GRID Research, and educational video systems were increasingly competing for bandwidth against recreational file-sharing applications. This file-sharing trend has evolved on two fronts: the use of P2P file-sharing applications for recreational (and often illegal) download of digital movies, music, games, and applications and the use of P2P tools for legitimate academic content downloads and research collaboration.

Recreational file sharing has become the biggest offender of excessive bandwidth consumption. According to a recent article in *The Chronicle of Higher Education*, several schools have encountered situations in which a small number of students using movie and music file-sharing applications were consuming more than 50 percent of campus bandwidth. Many institutions are also now facing the threat of their high-bandwidth networks serving as ideal launching pads for fast-propagating viruses and worms, which can generate enormous amounts of useless traffic to disrupt critical services.

Colleges and universities need new strategies to more fairly and effectively control bandwidth usage among students, faculty, and applications. By providing comprehensive bandwidth control, such strategies can help ensure that critical applications always have the resources they need, that recreational activities are permitted but appropriately prioritized, and that the potential effect of a virus or worm attack is limited. However, modern bandwidth control tools must be intelligent enough to allow educators to define rules with differing levels of granularity, depending on the type of traffic, the nature of the application, and the area of the network. (For example, the institution's primary high-speed connection to its ISP requires more granular control than a dormitory switch uplink.) To provide true, institutionwide bandwidth control, these tools must also be scalable and manageable.

All contents are Copyright © 1992–2006 Cisco Systems, Inc. All rights reserved. Important Notices and Privacy Statement.

Cisco Bandwidth Control for Education Networks

Faced with these issues, many colleges and universities are searching for new ways to control, protect, and optimize campus network resources. Cisco Bandwidth Control for Education Networks provides a suite of intelligent tools that bring granularity and flexibility to the deployment and enforcement of campus bandwidth control policies.

The Cisco Bandwidth Control for Education Networks Solution encompasses four primary components:

- The Cisco Service Control Engine (SCE), which provides granular control of high-value ISP connections
- Cisco IOS® Software Bandwidth Control Features, which implement rate-limiting mechanisms using Cisco Catalyst® switches deployed at the edge, distribution, and campus core networks
- CiscoWorks QoS Policy Manager (QPM) 3.2, which provides centralized, scalable implementation and management of quality-of-service (QoS) and bandwidth control features
- The Cisco Application and Content Networking System (ACNS), which optimizes the delivery of rich media content in bandwidth-constrained networks

Together, these tools give academic IT departments unprecedented control over critical academic resources.

CISCO SCE

The most heavily used link in an education network is the ISP connection. Controlling the use of the ISP link bandwidth requires a combination of hardware and software capable of packet-by-packet inspection at line rate. The Cisco SCE performs application-level and user-level processing of network traffic with dedicated application-specific integrated circuits (ASICs) and software that has proven itself in the most demanding high-speed networks.

Deployed at the institution's ISP connection, the system performs stateful and deep-packet inspection, classifying each network activity into application-level and user-level transactions and sessions. The solution then applies a variety of bandwidth control mechanisms to enforce institutional policies for each class of traffic and help ensure appropriate behavior on the network.

End-to-End QoS

A fundamental requirement of any bandwidth control solution is the ability to apply QoS mechanisms. These mechanisms control the bandwidth of specific users and prioritize traffic to help ensure appropriate handling of delay-sensitive applications. QoS capabilities are essential for carrying delay-sensitive IP voice and video traffic over an institution's ISP link, as well as for rate limiting recreational P2P traffic.

The Cisco SCE uses three levels of QoS:

- **Hierarchical bandwidth control:** The Cisco SCE supports granular bandwidth control by allocating part of a link's bandwidth for groups of specific application flows. Academic IT departments can define these groups according to categories such as "all P2P traffic," "browsing and streaming traffic," "all traffic flowing off net," and so on. In addition, colleges and universities can use the Cisco SCE to enforce minimum and maximum bandwidth limits and priorities for the total traffic that is produced by a given user, as well as for the specific applications (browsing, gaming, and so on) in which the user engages. These advanced mechanisms are used in a tiered fashion.
- **Differentiated Services (DiffServ) queuing:** Internet applications use DiffServ to help ensure that packets from delay-sensitive applications are prioritized over other packets. The Cisco SCE includes DiffServ-compliant transmit queues using "Best Effort Forwarding," four levels of "Assured Forwarding," and "Expedited Forwarding" for delay-sensitive applications.
- **DiffServ marking:** The Cisco SCE's advanced classification capabilities can also be used for marking the IP type of service (ToS)/DiffServ code-point (DSCP) byte of the associated traffic. Each flow or group of flows can be marked with a relevant DiffServ value based on the application or service. The next-hop Layer 3 device, such as a switch or router, then uses this marking to carry the delay-sensitive traffic appropriately. As a result, the Cisco SCE, crucial to the Cisco Bandwidth Control Solution, can serve as the ideal network element for classifying and marking application traffic for other DiffServ-enabled network elements.

© 2006 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

The Cisco SCE's purpose-built hardware and software was designed and optimized to operate in high-throughput links and high flow-capacity environments such as large universities and school districts. The solution can maintain reliable, hierarchical bandwidth control mechanisms and DiffServ marking and queuing processes while operating in multi-Gigabit Ethernet and OC-48 links carrying millions of IP flows.

Stateful Classification with the Cisco SCE

The core of the Cisco SCE is state-of-the-art hardware design specifically built to perform stateful application recognition and control of network traffic at wire speed. The programmable hardware core of the system yields the ideal balance of performance and flexibility needed for high-speed education networks.

In order to truly classify network traffic by application (HTTP, Session Initiation Protocol [SIP], Kazaa, and so on), service (local browsing, off-net-P2P, and so on), or content and destination, the Cisco SCE performs bidirectional, multiflow state analysis of application sessions. The solution maintains bidirectional flow state by tracking in both directions (client to server and service to client) of a network microflow as a single context. It also tracks the IP messaging that is associated with a given flow.

For example, the solution:

- Associates all packets in an individual TCP connection
- Associates a particular HTTP resource request ("GET") with a Web server's HTTP response ("200 OK")
- Identifies a P2P file-exchange communication by identifying the P2P servants' file-request message exchange
- Distinguishes a set of User Datagram Protocol (UDP) packets as part of a specific multimedia streaming session

In many IP sessions, the signaling protocol (used to initiate, tear down, and control the session) and the actual data (the file, video stream, or audio stream) is transmitted on a separate flow. Hence a single application session can often include one TCP control flow and multiple TCP, UDP, or Real-Time Transport Protocol (RTP) flows for the actual content or data. These content or data flows (and their port numbers) are negotiated in real time on the control flow. The Cisco SCE dynamically traces the message exchange on the control flow in order to bind the content and data flows to the control flow and then applies a coherent QoS and bandwidth control policy to these bundles of flows.

The Cisco SCE can also maintain state for different users using the network links on which it resides. This function is crucial for delivering per-user QoS policies such as rate limiting and prioritizing bandwidth by individual user and application (for example, limiting an individual user's P2P traffic to a certain bitrate, prioritizing certain users' streaming traffic over others, and so on). These capabilities are also essential for fairly allocating bandwidth to the various application flows, based on a stateful understanding of the application or service that each flow carries.

Controlling Link Bandwidth

With the Cisco SCE, academic IT departments can control the upstream and downstream bandwidth of groups of flows over the ISP links where the solution is deployed. Institutions typically deploy the solution inline between Ethernet or point-of-service (PoS) IP links, where it can limit the portions of capacity consumed by various types of IP flows. For example, institutions can enforce the following types of policies:

- Limit P2P upstream and downstream off-net traffic to 30 percent of the link
- Limit all P2P traffic on the link to 150 Mbps downstream and 50 Mbps upstream, except during weekends
- Limit all P2P traffic to 150 Mbps downstream and 50 Mbps upstream, except for traffic generated by users from a certain subnet
- Prioritize any other type of application traffic over P2P, without actually limiting the bandwidth of P2P flows

© 2006 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

Institutions can implement this bandwidth control mechanism through the use of definitions called global bandwidth controllers (GBWCs). Each GBWC is defined by three parameters:

- The type of applications/services of traffic that it controls
- The bandwidth limit in bits per second that the GBWC must enforce on the traffic
- The Cisco SCE interface applicable to the GBWC

After classifying traffic, the solution associates all of the flows of a given IP application or service with a specific GBWC. The solution enforces the bandwidth limit of this GBWC on all of these flows, as shown in Figure 1. Institutions can associate any combination of IP applications or services with the same GBWC and forward any traffic that is not classified by a specific GBWC to a default GBWC.

Figure 1. Global Bandwidth Controller

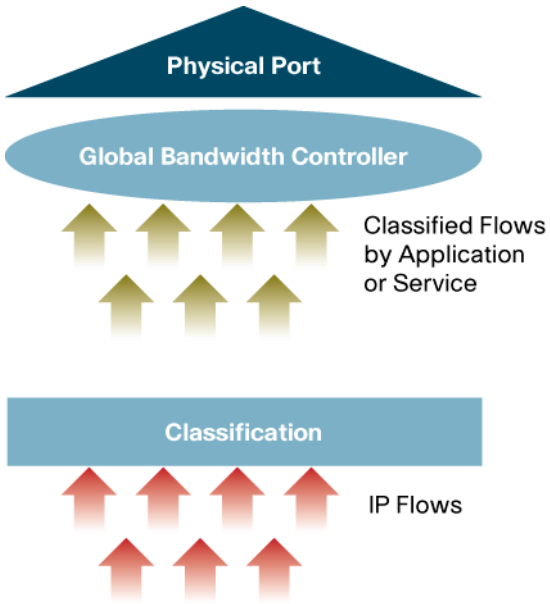

Controlling User Bandwidth

The Cisco SCE controls traffic not only at the application level, but at the user level as well. (Institutions can define a "user" as any managed entity on the user side of the device.) The solution can apply policy and accounting rules individually. A user could be a student, a faculty member, an administrator, a remote VPN user, or a cable-connected user.

The solution's per-user and per-service bandwidth control mechanisms provide ideal tools for allowing educators to enrich their service offering and generate a better network experience for all users. The per-user and per-service bandwidth control mechanisms use concepts employed by private-line Frame Relay networks, in which each user and service is granted a committed information rate (CIR) and peak information rate (PIR) and a level of priority related to other users and services.

Specifically, institutions can define the following parameters on a per-user and per-service basis:

- **PIR:** This is the maximum bandwidth rate enforced by the Cisco SCE for a particular service or user. PIR is separately defined for upstream and downstream traffic.
- **CIR:** This is the bandwidth rate that the Cisco SCE provides to the user or service when there is no congestion or oversubscription. When congestion occurs, the solution attempts to provide the various CIR rates according to assurance level or relative priority (described later). CIR is separately defined for upstream and downstream traffic.

© 2006 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

- **Relative priority for the user's PIR and CIR:** A metric used to divide bandwidth between different users in the event of bandwidth congestion on the link, in the upstream or downstream directions. This metric can have a value of 1 through 10, with the value representing the rate of bandwidth reduction from PIR to CIR during congestion.
- **Assurance level for the service's PIR and CIR:** A metric used to divide the user's bandwidth among the services that the user consumes in the event of congestion between these services. This metric can have a value from 1 to 10.

These parameters allow educational IT departments to define policies along the lines of the example in Table 1, for a hypothetical faculty user.

Table 1. Defined Service Parameters for a Faculty User

Service	Upstream		Downstream		Assurance Level
	CIR	PIR	CIR	PIR	
Browsing and Streaming	–	Unlimited	–	Unlimited	7
Gaming	200 Kbps	Unlimited	200 Kbps	Unlimited	10
P2P	–	300 Kbps	–	300 Kbps	3
Total user traffic	384 Kbps	512 Kbps	1 Mbps	1.5 Mbps	9

The Cisco SCE enforces per-service and per-user bandwidth configuration using an entity called a bandwidth controller (BWC). There are two types of BWCs: Normal BWCs enforce bandwidth for various IP services. User bandwidth controllers (SBWCs) enforce bandwidth for the total traffic that the user produces in each direction. The solution associates each flow or bundle of flows with a BWC according to the IP traffic type and with an SBWC according to the user originating the bundle of flows. Those flows that are not associated with a BWC are placed in a default BWC with administrator-defined PIR and CIR values. (See Figure 2.)

Figure 2. BWCs and SBWCs

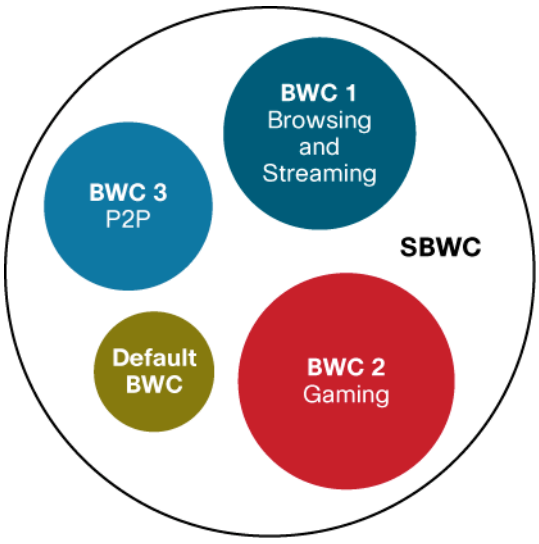

Academic IT departments can configure both BWCs and SBWCs for PIR oversubscription (that is, when the sum of the BWC's PIR is higher than the PIR of the SBWC). Each service's assurance level determines its bandwidth rate. As congestion increases, the Cisco SCE dynamically decreases the rate of flows in a way that helps ensure CIR and, if possible, PIR. The higher the congestion, the closer the enforced rate is throttled to CIR.

© 2006 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

When a BWC becomes congested, the solution fairly allocates bandwidth between the flows according to the widely employed “max-min” principle, which helps ensure that the flow with the lowest rate limit has the highest possible rate. This is a significant advantage of the Cisco SCE, because institutions can limit the bandwidth of various types of flows based on the intelligent classification of the flows instead of enforcing a strict limit on all flows. When an SBWC becomes congested, the solution allocates the bandwidth between the BWCs of the user in proportion to the defined PIR, CIR, and assurance-level values. First, the solution allocates the CIR, then PIR. If all services are reduced to their CIR and congestion remains, the assurance-level values determine the relative reduction of each BWC’s CIR.

The Cisco SCE also offers a more advanced option for managing the bandwidth of a given IP traffic type, called “extra pipes.” Institutions can use extra pipes to allocate a dedicated bandwidth portion to certain IP services, outside of users’ SBWC figures. Extra pipes have PIR and CIR values just like standard BWCs. The assurance level of an extra pipe is identical to the user’s overall relative priority. Extra pipes can be useful for reserving bandwidth for delay-sensitive applications such as voice and video over IP.

Putting It All Together: A Hierarchical Bandwidth Control Scheme

In most campus scenarios, educators are interested in controlling the amount of traffic that flows over the ISP link and belongs to a specific IP application. At the same time, they would like to control the bandwidth consumed by each user of the IP application. The Cisco SCE provides an ideal solution for this hierarchical bandwidth control model. The solution implements a hierarchical bandwidth control scheme using a three-stage process:

- **Stage one:** The solution classifies traffic flows by IP application and maps flows to specific users.
- **Stage two:** The solution uses the PIR, CIR, assurance-level, and relative priority parameters of BWCs and SBWCs to control the flows’ bandwidth by user traffic and the services the user consumes.
- **Stage three:** At the same time, the solution uses the classification flows by IP application to enforce per-link, per-application, and per-service bandwidth control limits, according to the defined policy.

Figure 3 illustrates the various levels of bandwidth control. The Cisco SCE uses stateful classification of flows into IP services or applications to define per-service BWCs and per-link GBWCs. The solution also maps each flow to a user.

© 2006 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

Figure 3. Hierarchical Bandwidth Control

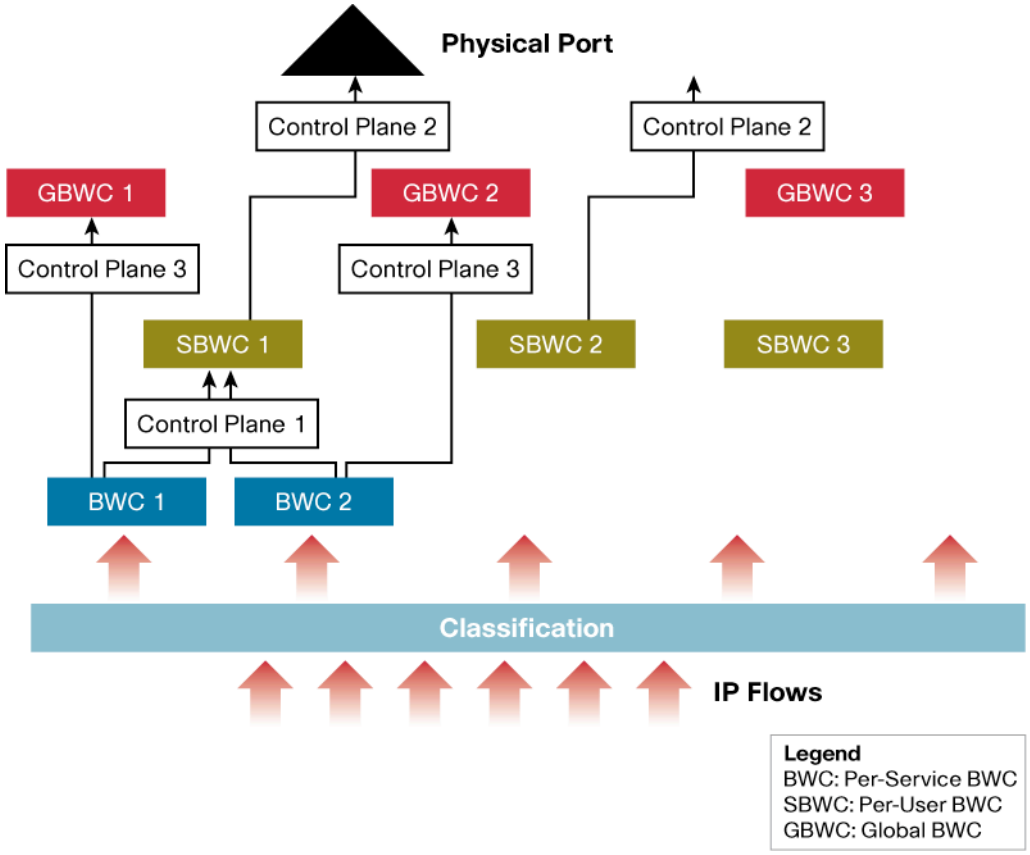

The three levels of bandwidth control define the system's behavior in case of congestion. Figure 3 identifies each behavior as control plane 1, control plane 2, or control plane 3. These control planes function as follows:

- Control plane 1 defines the link between per-IP-service BWCs and per-user SBWCs. If the SBWC becomes congested because a user tries to consume more bandwidth than the system allows, the BWCs reduce the bandwidth of that user according to CIR, PIR, and assurance level.
- Control plane 2 defines the actions taken if the solution's physical port becomes congested. In this scenario, the solution invokes a fair allocation of bandwidth scheme between the SBWCs of the users. Each SBWC controls the bandwidth of each user's IP services, according to the definitions in control plane 1.
- Control plane 3 defines the actions taken if a GBWC becomes congested. The BWCs of the underlying IP services controlled by the GBWC (for example, P2P) reduce bandwidth for that service according to the relative priority of the underlying SBWC.

Example: Hierarchical Bandwidth Control at a Large Institution

The following example illustrates how a large institution might implement hierarchical bandwidth control. In this scenario, the institution has deployed the Cisco SCE at Gigabit Ethernet or PoS links connecting to the ISP.

In an environment without bandwidth control, P2P traffic can occupy more than 65 percent of the capacity of the ISP links, with all of this traffic typically produced by just 5 to 20 percent of users. Academic IT departments might want to limit the percentage of P2P traffic on a Gigabit Ethernet ISP link to 30 percent. They might also want to limit browsing traffic on the link to 80 percent in order to reserve capacity for other education applications such as IP videoconferencing or live streaming of special events. In this example, the institution uses the Cisco SCE's GBWC feature to configure two GBWCs: one for P2P services with a 30 percent bandwidth limit and the other for Web browsing with an 80 percent bandwidth limit.

© 2006 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

What if the institution wants to prioritize bandwidth for faculty and research users over general students? Academic IT staff would use the Cisco SCE to build and configure the SBWC as shown in Table 2.

Table 2. SBWC Configuration to Prioritize Bandwidth for Faculty and Researchers

	SBWC Downstream		SBWC Upstream		Relative Priority
	CIR	PIR	CIR	PIR	
Faculty	10 Mbps	30 Mbps	5 Mbps	20 Mbps	9
Students	2 Mbps	4 Mbps	1 Kbps	2 Mbps	7
Guests	Lower CIR, PIR, and relative priority				

The institution can also use the solution to control Web browsing, online gaming, and media streaming services on a per-service level, helping ensure a baseline CIR for these services in case of network congestion. Table 3 illustrates the BWC configurations for Web browsing and online gaming services for the two classes of users. In this case, the solution uses the same classification to define the Web browsing service for the BWCs as it uses to define the Web browsing service for the GBWCs for the entire link.

Table 3. BWC Configurations for Web Browsing, Streaming Video, and Online Gaming

	BWCn Downstream		BWCn Upstream		Assurance Level
	CIR	PIR	CIR	PIR	
Faculty: Web Browsing	4 Mbps	20 Mbps	1.5 Mbps	10 Mbps	5
Faculty: IP Video	4 Mbps	20 Mbps	3.5 Mbps	10 Mbps	9
Students: Web Browsing	2 Kbps	3 Mbps	1 Mbps	2 Mbps	4
Students: Online Gaming	1 Kbps	2 Mbps	1 Mbps	2 Mbps	8

Differentiated Services and Delay-Sensitive Applications

The Cisco SCE offers two dedicated QoS mechanisms for processing packets of delay-sensitive applications, which differentiates the Cisco SCE from comparable offerings in the industry. These QoS mechanisms are based on the IETF RFC2475 standard, which defines the architecture for a DiffServ-enabled network. The DiffServ architecture is based on a simple model that classifies and assigns traffic to different categories. After classification, the network forwards packets according to the network characteristics required.

The first QoS mechanism allows the network to send packets from delay-sensitive applications to one of the DiffServ-compliant transmit queues in the Cisco SCE. The types of transmit queues include “Best Effort,” four levels of “Assured Forwarding,” and “Expedited Forwarding.” The Expedited Forwarding queue has strict priority over the other queues, and the solution always transmits packets in this queue before packets in the other queues. The solution transmits packets in the Assured Forwarding queues using a weighted “round-robin” algorithm and always transmits packets in the Best Effort queue last. (Institutions should forward packets from delay-sensitive applications to the Expedited Forwarding queue.)

As a second QoS mechanism for handling delay-sensitive applications, the Cisco SCE can mark ToS/DSCP bytes in outgoing packets, which next-hop network devices can use to activate QoS features. The IETF’s DiffServ Request for Comments (RFC) document defines the characteristics of per-hop behavior and queue structure for a DiffServ-enabled network, but the specification does not address the task of classifying traffic to the various queues. The Cisco SCE’s stateful classification capability provides an ideal vehicle for performing this classification. Institutions can configure the solution with a policy that maps applications and services to the relevant ToS/DSCP marking, with each marking corresponding to one of the DiffServ queues in the next-hop network device.

© 2006 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

The Cisco SCE offers an extremely thorough control capability for an institution's heavily used ISP connections, allowing administrators to achieve a truly fair allocation of network bandwidth. As a result, the solution provides an ideal starting point for colleges and universities that require institutionwide bandwidth control. And, by implementing the Cisco SCE to control bandwidth, institutions also gain a powerful tool to limit the potential impact on network resources of viruses and worms.

CISCO IOS SOFTWARE BANDWIDTH CONTROL CAPABILITIES

In addition to the granular bandwidth control mechanisms available through the Cisco SCE, colleges and universities can use several tools embedded within the Cisco IOS Software in Cisco switches and routers. Institutions can easily manage these capabilities through the CiscoWorks QPM solution, which provides tools for deploying an institutionwide network bandwidth policy.

To begin implementing campuswide control of recreational gaming and P2P traffic, institutions can establish a "Scavenger" policy class for such traffic. The Scavenger class, based on an Internet-II draft, can provide deferential services or "less-than-Best-Effort" services for certain applications that have little or no contribution to the organizational objectives of the education network. Such applications might include P2P media-sharing applications (such as Kazaa, Morpheus, Grokster, Napster, iMesh, and so on), gaming applications (such as Doom, Quake, Unreal Tournament, and so on), and any entertainment video applications. By assigning a minimal bandwidth queue to Scavenger traffic, institutions can limit such traffic to virtually nothing during periods of congestion, but make those services available when network bandwidth is not being used for education purposes. With this strategy, colleges and universities can implement flexible, nonstringent policy control that emphasizes educational activities, but also accommodates the recreational applications that students enjoy.

Enhancing Security with Scavenger-Class Services

The Scavenger class can also play a critical role in DoS attack and worm mitigation. Institutions can configure a Scavenger class of traffic as part of a broader QoS policy, as shown in Figure 4. (When provisioning for Scavenger-class traffic, institutions should mark Scavenger traffic to DSCP CS1. Scavenger traffic should also be assigned the lowest configurable queuing service.)

Figure 4. Using the Scavenger Class

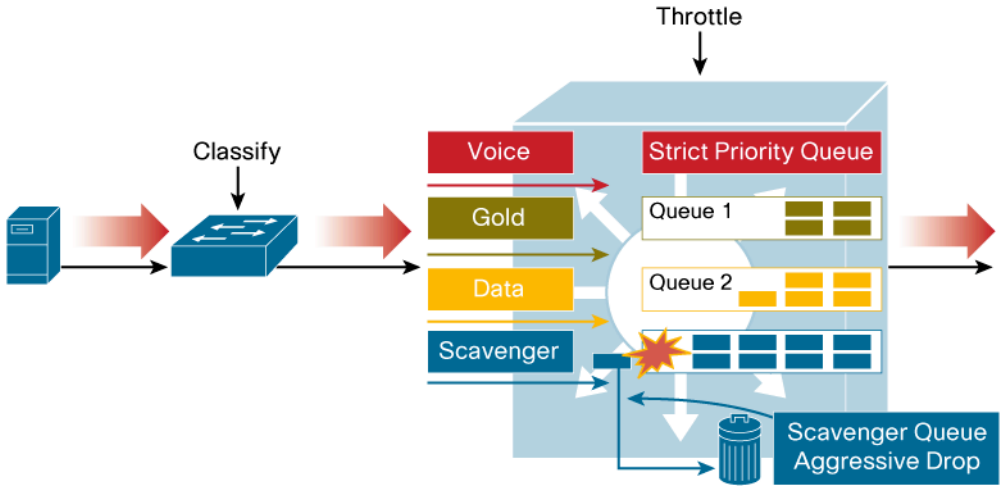

© 2006 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

Worms exploit security vulnerabilities in their targets to deliver harmful payloads, which typically include a self-propagating mechanism. Usually, worm attacks do not target the academic network infrastructure itself, but the network can become collateral damage as the worm propagates exponentially. The rapidly multiplying volume of traffic flows caused by the attack can eventually use up most of the CPU/hardware resources of the routers and switches in their paths, indirectly denying service to legitimate traffic flows, as shown in Figure 5.

Figure 5. Network Infrastructure Overload Resulting from a Worm Attack

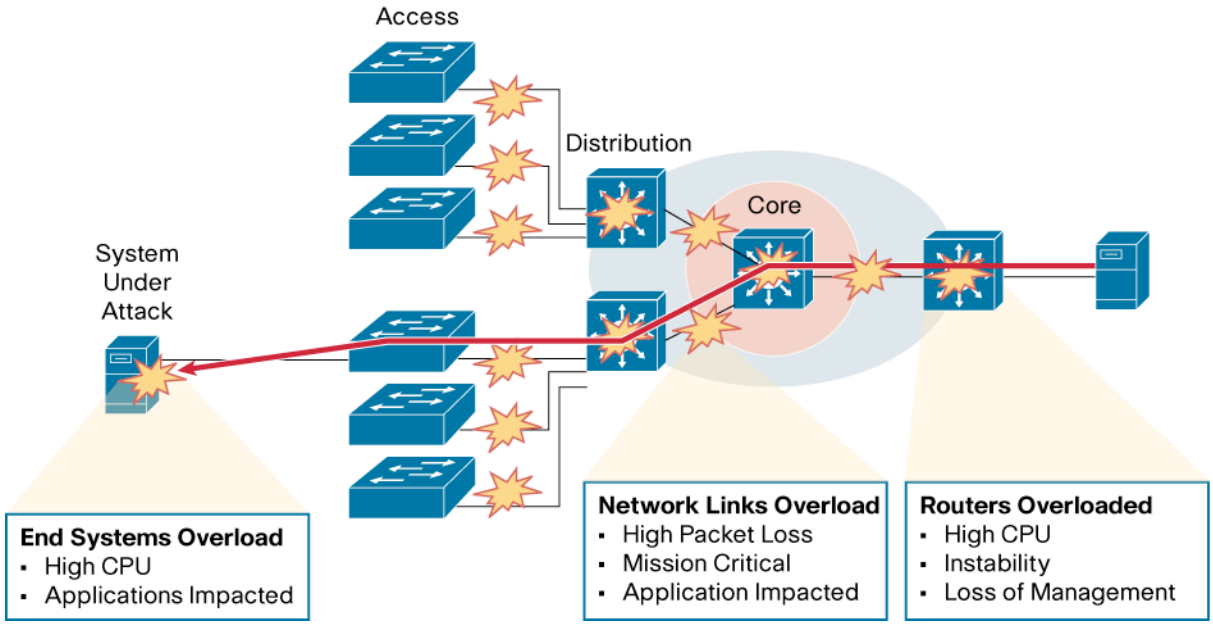

Institutions can proactively mitigate DoS/worm flooding attacks within the education network by immediately responding to out-of-profile network behavior indicative of a DoS or worm attack, using campus access-layer policers. Such policers meter traffic rates received from endpoint devices and mark down excess traffic spikes to the Scavenger class (DSCP CS1) when the traffic exceeds specified limits (at which point they are no longer considered normal flows).

In most education networks, it is quite abnormal (within a 95 percent confidence interval) for PCs to generate sustained traffic in excess of 5 percent of link capacity. In the case of a Fast Ethernet switch port, common in academic institutions, this means that it would be unusual for an end-user PC to generate more than 5 Mbps of uplink traffic on a sustained basis.

Of course, this does not mean that academic IT departments should police all traffic to 5 Mbps and automatically drop the excess. If that were the case, there would be no reason to deploy Fast Ethernet or Gigabit Ethernet switch ports to endpoint devices, because even a 10-BaseT Ethernet switch port would have more uplink capacity than the 5 Mbps policer-enforced limit. Furthermore, such an approach would severely penalize legitimate traffic that happened to exceed 5 Mbps on a Fast Ethernet switch port.

Institutions can take a less severe approach by coupling access-layer policers with hardware and software (campus/WAN/VPN) queuing polices, with both sets of policies provisioning a “less-than-Best-Effort” Scavenger class. Access-layer policers would mark down out-of-profile traffic to DSCP CS1 (Scavenger class) and then direct all congestion management policies (whether in Cisco Catalyst switches or in Cisco IOS Software) to provision a “less-than Best-Effort” queuing service for any traffic marked to DSCP CS1.

© 2006 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

Scavenger Class in Practice

Let us illustrate how this might work for both legitimate traffic exceeding the access layer's policer watermark and illegitimate excess traffic resulting from a DoS or worm attack. In the former case, assume that the PC generates more than 5 Mbps of traffic, perhaps because of a large file transfer or backup. Under normal operating conditions, the campus network rarely if ever experiences congestion, because abundant capacity is generally available to carry the traffic. (Most academic networks use Gigabit Ethernet uplinks to the distribution and core layers of the campus network and would require 1000 Mbps of traffic from the access-layer switch to congest.) If traffic is destined for the far side of a WAN/VPN link (generally at a rate of less than 5 Mbps), dropping occurs even without the access-layer policer, because of the bottleneck created by the campus/WAN speed mismatch. The network's TCP sliding windows mechanism would eventually find an optimal speed (less than 5 Mbps) for the file transfer. As a result, access-layer policers that mark down out-of-profile traffic to the Scavenger class would not affect legitimate traffic, aside from the obvious remarking. The policers would not cause any new reordering or dropping on the flows.

In the case of a DoS or worm attack, access-layer policers have a much different effect on traffic caused by the attack. As hosts become infected and traffic volumes multiply, even the campus network can become congested. If just 11 end-user PCs on a single switch begin spawning worm flows to the switch's maximum Fast Ethernet link capacity, the Gigabit Ethernet uplink from the access-layer switch to the distribution-layer switch will congest. In response, the network will begin queuing/reordering/dropping activities. At this point, the network grants voice-over-IP (VoIP), critical data applications, and even Best Effort applications priority over worm-generated traffic (because Scavenger traffic is dropped the most aggressively). As a residual benefit of this strategy, bandwidth-intensive P2P and gaming application traffic receives the same policing treatment, because the network does not explicitly classify these activities as high-priority traffic. At a basic level, the Scavenger-class strategy addresses low-priority gaming and P2P traffic and DoS/worm traffic with the same methodology.

Deploying a Scavenger-Class QoS DoS/Worm Mitigation Strategy

Institutions implementing the Scavenger-class QoS strategy to mitigate DoS attacks and worms should follow these steps:

- **Profile applications:** Education network administrators should begin by profiling applications to determine what constitutes normal as opposed to abnormal flows, within a 95 percent confidence interval. Thresholds demarking normal and abnormal flows vary from campus to campus and from application to application. Administrators should beware of overscrutinizing traffic behavior, because this can quickly exhaust time and resources, and this behavior can change daily. (Remember, this Scavenger-class QoS strategy does not penalize legitimate traffic flows that temporarily exceed thresholds. Only sustained, abnormal streams generated simultaneously by multiple hosts—which are highly indicative of a DoS or worm attack—are subject to aggressive dropping. And even then, such dropping occurs only after legitimate traffic has been serviced.)
- **Deploy access-edge policers:** To contain true abnormal flows, deploy campus access-edge network policers to remark abnormal traffic to the Scavenger class (DSCP CS1).
- **Deploy additional defenses at the distribution layer, if applicable:** When the campus network uses some types of Cisco Catalyst switches in the distribution layer, academic IT staff should deploy a second line of policing defense at the distribution layer using the rate-limiting mechanisms available in Cisco Catalyst switches, described later.
- **Use Scavenger-class queuing policies throughout the environment:** To support the Scavenger-class strategy's remarking policies, institutions must enforce end-to-end "less-than-Best-Effort" Scavenger-class queuing policies within the campus network, the WAN, and VPN connections.

Even when a college or university has deployed Scavenger-class QoS institutionwide, the strategy only mitigates DoS and worm attacks. It does not prevent them or remove them entirely. To provide comprehensive protection, institutions should overlay QoS-enabled infrastructures with additional firewall, intrusion detection, identity, and other security solutions. Institutions can also employ converged security solutions such as Cisco ASA 5500 Series Adaptive Security Appliances, which combine multiple security services in a single, integrated platform.

Bandwidth Control Features of Cisco Catalyst Switches

Many Cisco Catalyst switches can incorporate rate-limiting mechanisms to police and restrict traffic flows and allow institutions to implement more granular and flexible bandwidth control strategies. Cisco switching solutions with available Cisco IOS Software bandwidth control features include Cisco Catalyst 3750, 4500, and 6500 Series Switches.

© 2006 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

Cisco Catalyst 6500 and 4500 Series Rate Limiting

Colleges and universities can take advantage of User-Based Rate Limiting (UBRL) capabilities in Cisco Catalyst 6500 Series Switches with the Cisco Catalyst 6500 Series Supervisor Engine 720 and Cisco Catalyst 4500 Series Switches with the Cisco Catalyst 4500 Series Supervisor Engine V-10GE. These capabilities function similarly to the per-user control capabilities of the Cisco SCE. UBRL uses microflow policing capabilities to dynamically learn traffic flows and rate limit each unique flow to a specific rate. Academic IT administrators can apply UBRL to ingress traffic on routed interfaces with source or destination flow masks. A Cisco Catalyst 4500 with the Supervisor Engine V-10GE can support up to 100,000 individual flows and 512 different rates.

UBRL provides an ideal solution for education environments that require a per-user rate-limiting mechanism at the distribution and core of the campus network. Administrators can define different outbound and inbound traffic rates per user. Administrators can also deploy UBRL on a per-group basis—for example, to rate limit an entire VLAN or subnet.

Cisco Catalyst 3750 Series Rate Limiting

All switches in the Cisco Catalyst 3750 Series support an extensive rate-limiting feature set that institutions can apply on both Gigabit Ethernet and Fast Ethernet interfaces. Education IT managers can set policies to allocate higher or lower bandwidth to certain users, groups of users, or applications. For example, an IT manager can decide to allocate faculty members connected on Gigabit Ethernet ports only 200 Mbps of bandwidth. Rate limiting, in conjunction with other features such as time-based access control lists (ACLs), can help academic IT departments save money by carefully managing how much bandwidth is available and when—delaying the purchase of additional WAN links.

With pilot rollouts of Gigabit-to-the-desktop LANs occurring at most universities today, rate limiting provides an important tool for IT managers to manage the bandwidth consumed by these Gigabit Ethernet ports until the distribution and core areas of the network are upgraded to support higher-bandwidth technology, such as 10 Gigabit Ethernet. With rate-limiting strategies in place, faculty members and students can send and receive only up to the bandwidth they have been individually allocated. Institutions enjoy maximum flexibility in choosing how to allocate this bandwidth. For example, a university could choose to give faculty and research students more bandwidth than undergraduate students or allocate greater bandwidth to application servers and workstations.

Rate Limiting in Campus Dormitories and Campus Buildings

Colleges and universities can employ the rate-limiting capabilities of Cisco Catalyst switches to more effectively control bandwidth in dormitories and campus buildings. For example, in dormitories, education IT staff can rate limit the bandwidth each dormitory student receives to that of a typical student. This rate limiting can be asymmetrical. (Since normal student usage consists of more downstream traffic than upstream traffic, the institution might provision more downstream traffic than upstream traffic to each student.)

Academic IT departments can keep track of users and computers with MAC addresses and can rate limit the amount of bandwidth the user or computer receives. IT staff can also make allowances for delay-sensitive applications. For example, a university might want to allow students to download videos of classroom lectures to view in their dorm rooms. However, video is one of the most bandwidth-intensive and jitter-sensitive applications the network can support. When IT managers rate limit the downstream and upstream bandwidth, they must provision enough downstream traffic for a student to watch a video (typically from 300 Kbps to 2 Mbps, depending on the encryption used). In this scenario, the IT manager should allocate more buffers to the video queue and assign the video traffic to the high-priority queues to minimize jitter.

In campus buildings, IT managers typically want to rate limit faculty traffic to a higher rate, and assign faculty traffic to a higher-priority queue, than student traffic. For example, an institution could rate limit faculty traffic on Gigabit Ethernet ports to 200 Mbps and assign such traffic to the second-highest priority queue. In student classrooms (where lecturing occurs), students can use the same Gigabit Ethernet ports, but have their traffic rate limited to just 500 Kbps. Academic IT departments can even employ more flexible rate limiting depending on the environment and type of user. For example, engineering graduate students often send large design files to servers to be processed, graphics students need to transfer large video

© 2006 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

files, and undergraduate students in workstation clusters can be working on their own assignments and projects. So, in graduate student research centers and student workstation clusters, IT managers can rate limit students to a higher rate, such as 4 Mbps.

In both dormitories and campus buildings, academic IT departments should shape and place IP phones on Strict Priority Queuing. Strict Priority Queuing helps ensure that the network always gives voice traffic the highest priority. Strict Priority Queuing provides the dedicated bandwidth that voice traffic needs to minimize latency and jitter. Although voice traffic is very latency- and jitter-sensitive, it does not require much bandwidth (less than 100 Kbps). Figure 6 illustrates a typical university network. Figure 7 illustrates an example in which voice traffic has been shaped to 25 Mbps, according to the anticipated volume of voice traffic.

Figure 6. Typical University Network

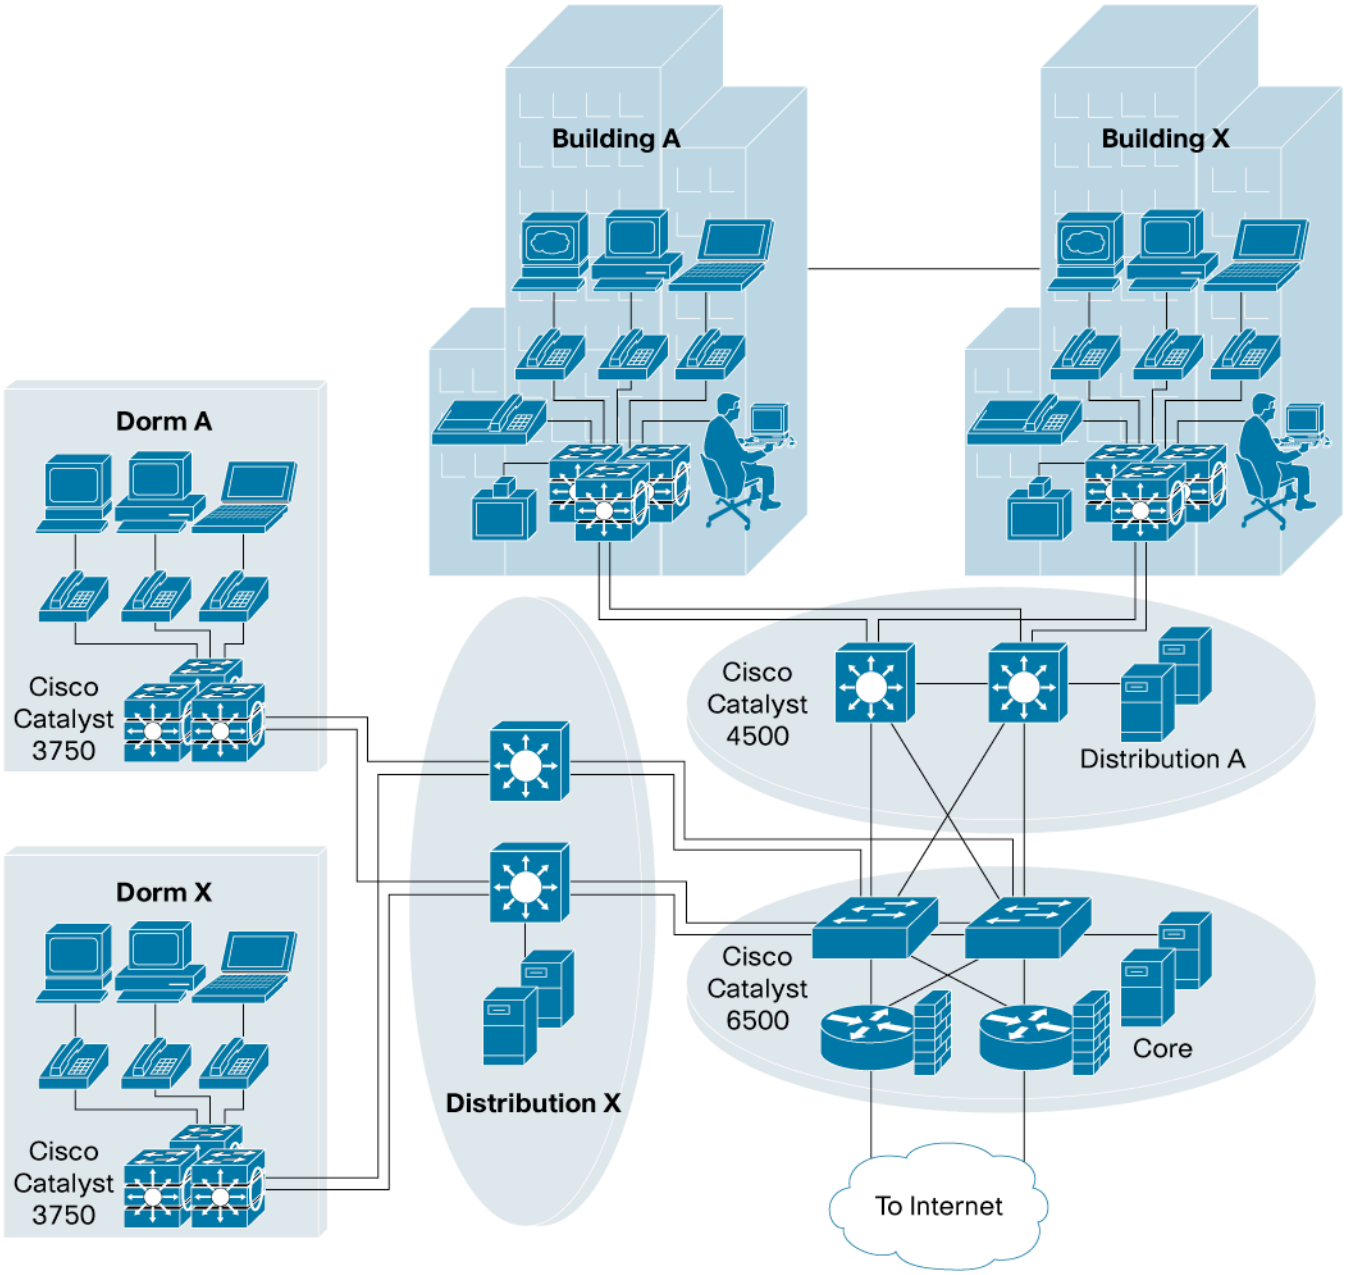

© 2006 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

Figure 7. Prioritizing Voice Traffic

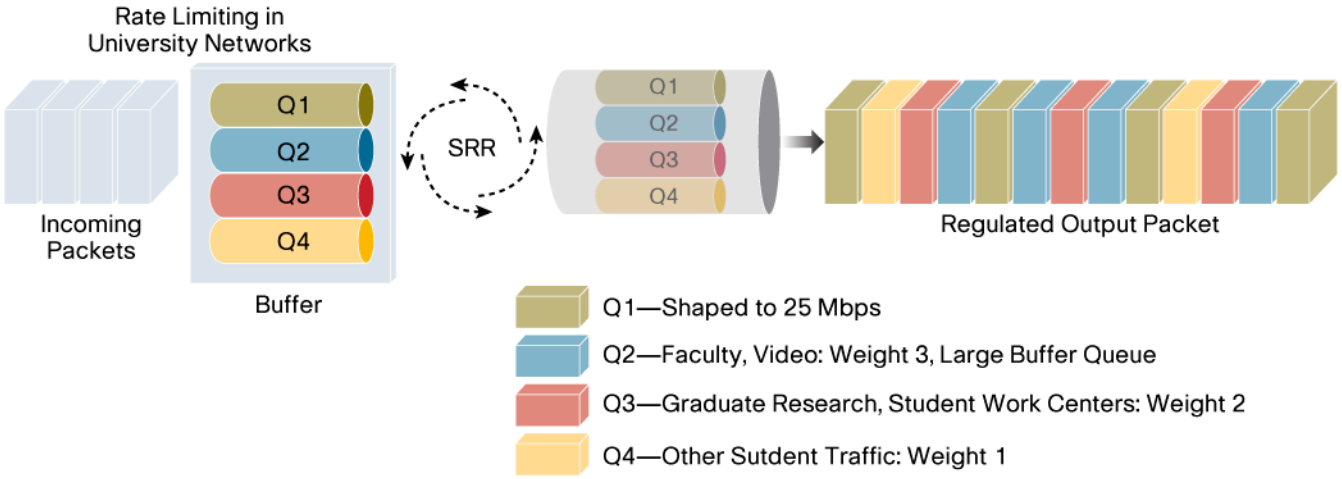

Education IT managers can preferentially give higher bandwidth to select buildings or distribution centers, depending on bandwidth requirements. (They can also use other criteria, such as granting higher rate limits to departments that pay more for the additional bandwidth.) In general, IT managers allocate more bandwidth from the campus building distribution network to the core network than from the dormitory distribution network to the core. Work-intensive faculty offices, graduate student research centers, and student work centers that are housed in campus buildings typically require higher network bandwidth and have a greater need for the WAN link to the Internet.

CISCOWORKS QPM 3.2

As powerful as Cisco bandwidth control strategies might be, they are of no practical use if academic IT departments cannot easily implement them and manage them in a scalable way. CiscoWorks QPM 3.2 offers a secure, Web-based tool for providing end-to-end QoS for converged data, voice, and video networks. As part of the CiscoWorks family of network management solutions, CiscoWorks QPM 3.2 combines traffic monitoring with configuration of differentiated services across the IP infrastructure by taking advantage of the Cisco IOS Software and Cisco Catalyst operating system QoS mechanisms built into LAN and WAN switching and routing equipment from Cisco Systems. These capabilities allow institutions to more easily scale bandwidth control services across the network and enforce institutionwide bandwidth control policies.

With CiscoWorks QPM, academic IT departments can:

- Perform baseline monitoring of critical traffic flows to define policies
- Classify applications into service classes
- Provision QoS with networkwide enforcement
- Validate QoS settings and results

Education IT managers can use CiscoWorks QPM to gain greater visibility into network operations through traffic flow monitoring. This insight can be used to configure the appropriate policies to help ensure application performance and to automate multiple service levels across any network topology. (See Figure 8.) The solution also provides centralized QoS analysis and policy control for voice, video, and data networks and supports networkwide, content-based DiffServ and campus-to-WAN automated QoS configuration and deployment.

© 2006 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

After QoS has been deployed, QPM monitoring can help academic IT administrators determine whether policies are having the desired effect by providing packet or bitrate measurements at WAN interfaces for inbound and outbound traffic. Network administrators can also view QoS graphs, including line and bar charts, next to policy descriptions. Administrators can troubleshoot performance problems by examining traffic patterns relative to QoS enforcement mechanisms, including policing, queuing, shaping, and dropping (Figure 8). Administrators can even use a date and time "zoom" function to scan QoS data over different time periods and use file export functions to perform additional analysis with other tools.

Figure 8. Configuring Bandwidth Policies with CiscoWorks QPM

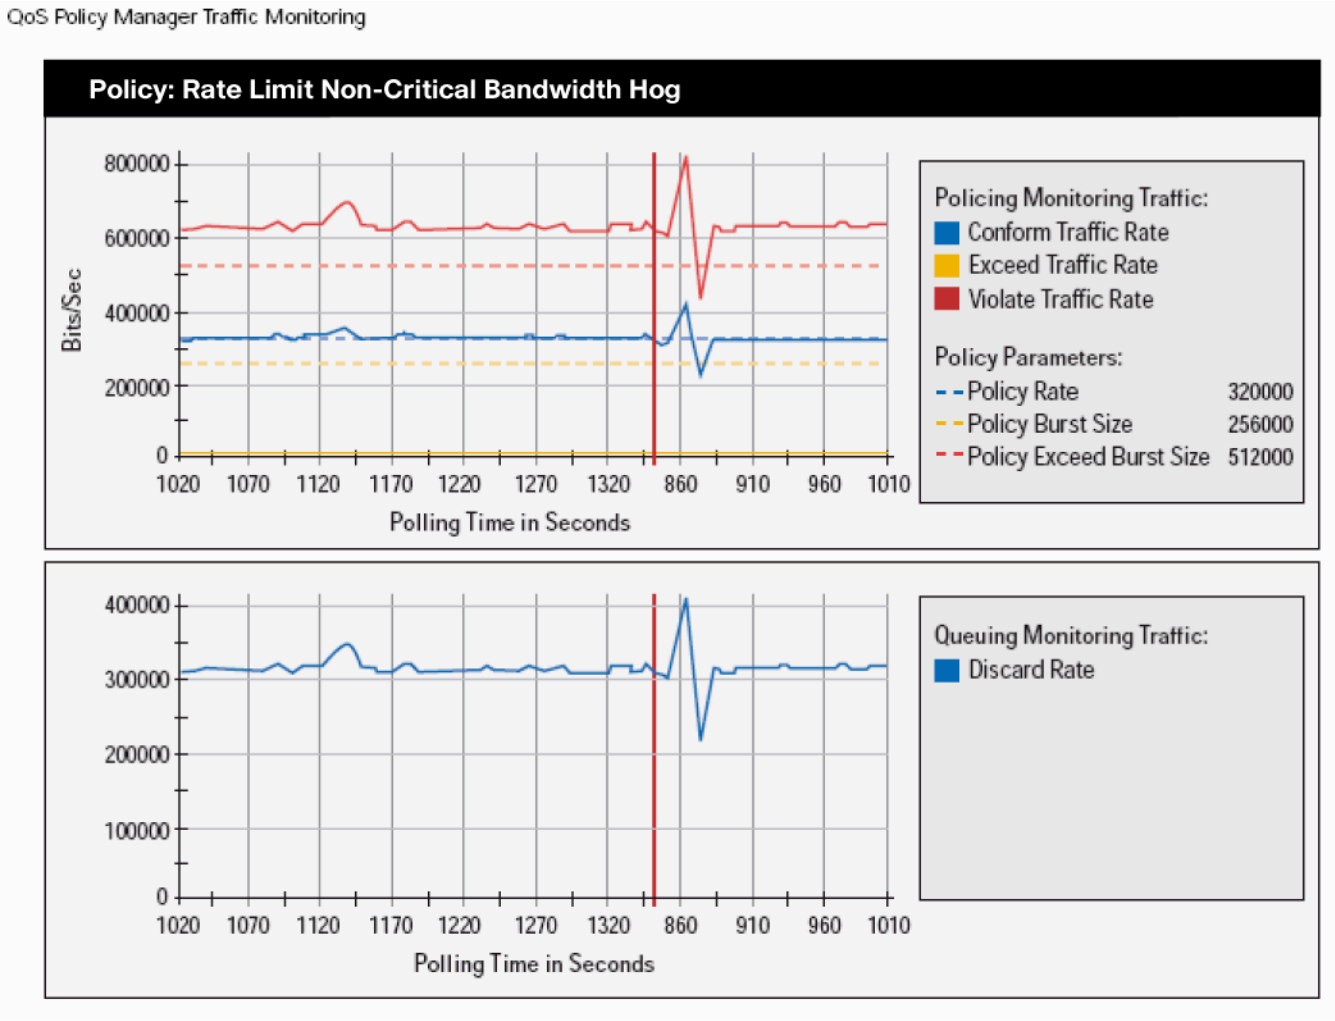

© 2006 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

CISCO ACNS

In some regions of the world, the high cost of high-speed WAN links forces colleges and universities to rely on T1, E1, or even smaller circuits to support the academic network. In these environments, institutions can have a difficult time delivering rich media content and innovative educational applications. The Cisco ACNS solution optimizes content delivery in bandwidth-constrained networks, reducing network congestion by storing and delivering content at the network edge.

Cisco ACNS software combines the technologies of demand-pull caching and prepositioning of Web applications, objects, files, and streaming media to accelerate content delivery. The solution runs on Cisco Wide Area Application Engines (WAEs), Cisco Content Distribution Manager (CDM), and Cisco Content Router platforms.

Together, these intelligent hardware and software components provide:

- Content-edge delivery: Institutions can use the Cisco WAE appliance or Cisco WAE network module to avoid WAN congestion by storing and delivering content at the network edge.
- Central content management capabilities: Academic IT departments can use the Cisco CDM appliance, as well as the network and device management capabilities of the CiscoWorks software suite, to centrally manage and control content distribution.
- Content routing capabilities: Institutions can use the Cisco Content Router appliance for HTTP routing, as well as the Web Cache Control Protocol (WCCP) embedded in Cisco routers and Cisco Catalyst switches with Cisco IOS Software to optimize content routing and caching.

For more information about the Cisco ACNS solution, visit http://www.cisco.com/go/acns.

ENHANCING THE ACADEMIC NETWORK

Combining a variety of intelligent strategies and scalable tools, Cisco Bandwidth Control for Education Networks provides a comprehensive solution for implementing and enforcing institutionwide bandwidth control policies. With these strategies, academic IT departments can more fairly and effectively allocate network resources than ever before.

Cisco Bandwidth Control for Education Networks provides:

- **Effective bandwidth control in each functional area of the network:** The Cisco SCE can manage heavily used ISP links, while Cisco IOS Software rate-limiting capabilities such as Scavenger-class QoS, UBRL, and edge switch-based rate-limiting mechanisms control bandwidth in the network core, distribution, and edge layers. CiscoWorks QPM 3.2 allows institutions to easily scale uniform bandwidth control policies across the campus network.
- **Reduced costs:** Colleges and universities can more intelligently transition to higher-bandwidth services and eliminate the need to upgrade network infrastructure prematurely, simply to accommodate growing P2P and gaming applications. Cisco Bandwidth Control for Education Networks also allows institutions to reduce P2P application usage on more expensive network links and peering points.
- **Reduced network congestion:** Cisco Bandwidth Control for Education Networks offers the flexibility, granularity, and intelligence to tightly control recreational P2P applications and throttle down bandwidth to prevent worms from flooding network resources with useless traffic. The Cisco ACNS solution can preposition rich media content at the edge of bandwidth-constrained networks to enhance application performance and the user experience.
- **Enhanced protection against DoS and worm attacks:** Cisco bandwidth control strategies provide a critical tool for defending against DoS and worm attacks by allowing academic IT administrators to throttle down bandwidth of any abnormally behaving endpoint or traffic flow.
- **Improved network and application performance:** By restricting nonessential P2P and gaming traffic and helping ensure that critical applications and users always have the network resources they need, academic IT departments can provide a better-performing, more productive campus network environment.
- **Maximum scalability:** Using tools such as the CiscoWorks QPM, academic IT departments can easily scale QoS and bandwidth control policies across the entire institutional network and more effectively preserve and control campus bandwidth.

© 2006 Cisco Systems, Inc. All rights reserved.
Important notices, privacy statements, and trademarks of Cisco Systems, Inc. can be found on cisco.com.

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 Website 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

Copyright © 2006 Cisco Systems, Inc. All rights reserved. CCSP, CCVP, the Cisco Square Bridge logo, Follow Me Browsing, 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 Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, FormShare, GigaDrive, GigaStack, HomeLink, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MeetingPlace, MGX, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, ProConnect, RateMUX, ScriptShare, SlideCast, SMARTnet, The Fastest Way to Increase Your Internet Quotient, and TransPath 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 Website 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. (0601R)

Printed in the USA C11-348270-00 05/06