




Medianet Reference Guide

Last Updated: October 26, 2010




Solution Authors



John Johnston

John Johnston, Technical Marketing Engineer, CMO Enterprise Solutions Engineering (ESE), Cisco Systems

John has been with Cisco for 10 years, with previous experience as a network consulting engineer in Cisco's advanced services group. Prior to joining Cisco, he was a consulting engineer with MCI's Professional Managed Services group. John has been designing or troubleshooting enterprise networks for the past 15 years. In his spare time, he enjoys working with microprocessor-based electronic projects including wireless environmental sensors. John holds CCIE certification 5232. He holds a bachelor of science degree in electrical engineering from the University of North Carolina's Charlotte campus.



Sherelle Farrington

Sherelle Farrington, Technical Leader, CMO Enterprise Solutions Engineering (ESE), Cisco Systems

Sherelle is a technical leader at Cisco Systems with over fifteen years experience in the networking industry, encompassing service provider and enterprise environments in the US and Europe. During her more than ten years at Cisco, she has worked on a variety of service provider and enterprise solutions, and started her current focus on network security integration over four years ago. She has presented and published on a number of topics, most recently as the author of the SAFE WebEx Node Integration whitepaper and as one of the authors of the SAFE Reference Guide, the Wireless and Network Security Integration Solution Design Guide, and the Network Security Baseline document.



Roland Saville

Roland Saville, Technical Leader, CMO Enterprise Solutions Engineering (ESE), Cisco Systems

Roland is a Technical Leader for the Enterprise Systems Engineering team within Cisco, focused on developing best-practice design guides for enterprise network deployments. He has 14+ years of Cisco experience as a Systems Engineer, Consulting Systems Engineer, Technical Marketing Engineer, and Technical Leader. During that time, he has focused on a wide range of technology areas including the integration of voice and video onto network infrastructures, network security, and wireless LAN networking. Roland has a BS degree in Electrical Engineering from the University of Idaho and an MBA from Santa Clara University. He has co-authored the Cisco TelePresence Fundamentals book and has six U.S. Patents.



Tim Szigeti

Tim Szigeti, Technical Leader, CMO Enterprise Solutions Engineering (ESE), Cisco Systems

Tim is a technical leader at Cisco, where he has spent the last 10 years focused on quality-of-service (QoS) technologies. His current role is to design network architectures for the next wave of media applications, including Cisco TelePresence, IP video surveillance, digital media systems, and desktop video. He has authored many technical papers, including the QoS Design Guide and the TelePresence Design Guide, as well as Cisco Press books on End-to-End QoS Network Design and Cisco TelePresence Fundamentals. Szigeti holds CCIE certification 9794 and holds a bachelor of commerce degree with a specialization in management information systems from the University of British Columbia.



CONTENTS

CHAPTER 1

Medianet Architecture Overview 1-1

Executive Summary 1-1

Business Drivers for Media Applications 1-2

Global Workforce and the Need for Real-Time Collaboration 1-2

Pressures to be “Green” 1-2

New Opportunities for IP Convergence 1-3

Transition to High-Definition Media 1-3

Media Explosion 1-4

Social Networking—Not Just For Consumers Anymore 1-4

Bottom-Up versus Top-Down Media Application Deployments 1-5

Multimedia Integration with Communications Applications 1-5

Demand for Universal Media Access 1-5

Challenges of Medianets 1-6

Understanding Different Media Application Models 1-6

Delivery of Media Applications 1-8

Prioritizing the Right Media Applications, Managing the Rest 1-8

Media Application Integration 1-9

Securing Media Applications 1-10

Solution 1-10

The Need for a Comprehensive Media Network Strategy 1-10

Architecture of a Medianet 1-11

Common Requirements and Recommendations 1-12

Network Design for High Availability 1-12

Bandwidth and Burst 1-14

Latency and Jitter 1-15

Application Intelligence and Quality of Service 1-17

Admission Control 1-21

Broadcast Optimization 1-23

Securing Media Communications 1-23

Visibility and Monitoring Service Levels 1-24

Campus Medianet Architecture 1-24

Design for Non-Stop Communications in the Campus 1-25

Bandwidth, Burst, and Power 1-26

Application Intelligence and QoS	1-26
Broadcast Optimization with IP Multicast	1-27
Leveraging Network Virtualization for Restricted Video Applications	1-27
Securing Media in the Campus	1-28
WAN and Branch Office Medianet Architecture	1-29
Design for Non-Stop Communications over the WAN	1-30
Bandwidth Optimization over the WAN	1-31
Application Intelligence and QoS	1-31
Broadcast Optimization for Branch Offices	1-32
Data Center Medianet Architecture	1-33
Design for Non-Stop Communications in the Data Center	1-34
High-Speed Media Server Access	1-34
Media Storage Considerations	1-34
Conclusions	1-34
Terms and Acronyms	1-35
Related Documents	1-36
White Papers	1-36
System Reference Network Designs	1-37
Websites	1-37

CHAPTER 2

Medianet Bandwidth and Scalability 2-1

Bandwidth Requirements	2-1
Measuring Bandwidth	2-2
Video Transports	2-3
Packet Flow Malleability	2-3
Microbursts	2-5
Shapers	2-6
Shapers versus Policers	2-8
TxRing	2-11
Converged Video	2-12
Bandwidth Over Subscription	2-13
Capacity Planning	2-15
Load Balancing	2-17
EtherChannel	2-20
Bandwidth Conservation	2-21
Multicast	2-21
Cisco Wide Area Application Services	2-21
Cisco Application and Content Network Systems	2-22

Cisco Performance Routing 2-23

Multiprotocol Environments 2-23

Summary 2-24

CHAPTER 3

Medianet Availability Design Considerations 3-1

Network Availability 3-1

Device Availability Technologies 3-5

Cisco StackWise and Cisco StackWise Plus 3-5

Non-Stop Forwarding with Stateful Switch Over 3-7

Network Availability Technologies 3-10

L2 Network Availability Technologies 3-10

UniDirectional Link Detection 3-11

IEEE 802.1D Spanning Tree Protocol 3-11

Cisco Spanning Tree Enhancements 3-13

IEEE 802.1w-Rapid Spanning Tree Protocol 3-15

Trunks, Cisco Inter-Switch Link, and IEEE 802.1Q 3-15

EtherChannels, Cisco Port Aggregation Protocol, and IEEE 802.3ad 3-17

Cisco Virtual Switching System 3-18

L3 Network Availability Technologies 3-22

Hot Standby Router Protocol 3-23

Virtual Router Redundancy Protocol 3-25

Gateway Load Balancing Protocol 3-26

IP Event Dampening 3-28

Operational Availability Technologies 3-29

Cisco Generic Online Diagnostics 3-30

Cisco IOS Embedded Event Manager 3-30

Cisco In Service Software Upgrade 3-31

Online Insertion and Removal 3-31

Summary 3-31

CHAPTER 4

Medianet QoS Design Considerations 4-1

Drivers for QoS Design Evolution 4-1

New Applications and Business Requirements 4-1

The Evolution of Video Applications 4-2

The Transition to High-Definition Media 4-4

The Explosion of Media 4-5

The Phenomena of Social Networking 4-6

The Emergence of Bottom-Up Media Applications 4-6

The Convergence Within Media Applications 4-7

The Globalization of the Workforce	4-8
The Pressures to be Green	4-8
New Industry Guidance and Best Practices	4-8
RFC 2474 Class Selector Code Points	4-9
RFC 2597 Assured Forwarding Per-Hop Behavior Group	4-10
RFC 3246 An Expedited Forwarding Per-Hop Behavior	4-11
RFC 3662 A Lower Effort Per-Domain Behavior for Differentiated Services	4-11
Cisco's QoS Baseline	4-12
RFC 4594 Configuration Guidelines for DiffServ Classes	4-13
New Platforms and Technologies	4-16
Cisco QoS Toolset	4-16
Classification and Marking Tools	4-16
Policing and Markdown Tools	4-19
Shaping Tools	4-20
Queuing and Dropping Tools	4-21
CBWFQ	4-21
LLQ	4-22
1PxQyT	4-23
WRED	4-24
Link Efficiency Tools	4-24
Hierarchical QoS	4-25
AutoQoS	4-26
QoS Management	4-27
Admission Control Tools	4-28
Enterprise Medianet Strategic QoS Recommendations	4-29
Enterprise Medianet Architecture	4-30
Enterprise Medianet QoS Application Class Recommendations	4-31
VoIP Telephony	4-32
Broadcast Video	4-33
Realtime Interactive	4-33
Multimedia Conferencing	4-33
Network Control	4-33
Signaling	4-33
Operations, Administration, and Management (OAM)	4-34
Transactional Data and Low-Latency Data	4-34
Bulk Data and High-Throughput Data	4-34
Best Effort	4-34
Scavenger and Low-Priority Data	4-34
Media Application Class Expansion	4-35
Cisco QoS Best Practices	4-36

Hardware versus Software QoS	4-36
Classification and Marking Best Practices	4-36
Policing and Markdown Best Practices	4-36
Queuing and Dropping Best Practices	4-37
QoS for Security Best Practices	4-39
Summary	4-45
References	4-46
White Papers	4-46
IETF RFCs	4-46
Cisco Documentation	4-47

CHAPTER 5

Medianet Security Design Considerations 5-1

An Introduction to Securing a Medianet	5-1
Medianet Foundation Infrastructure	5-1
Medianet Collaboration Services	5-2
Cisco SAFE Approach	5-2
Security Policy and Procedures	5-3
Security of Medianet Foundation Infrastructure	5-3
Security Architecture	5-3
Network Foundation Protection	5-4
Endpoint Security	5-5
Web Security	5-6
E-mail Security	5-6
Network Access Control	5-7
User Policy Enforcement	5-7
Secure Communications	5-7
Firewall Integration	5-8
IPS Integration	5-8
Telemetry	5-9
Security of Medianet Collaboration Services	5-9
Security Policy Review	5-10
Architecture Integration	5-10
Application of Cisco SAFE Guidelines	5-10
Medianet Security Reference Documents	5-12

CHAPTER 6

Medianet Management and Visibility Design Considerations 6-1

Network-Embedded Management Functionality	6-2
NetFlow	6-5
NetFlow Strategies Within an Enterprise Medianet	6-6

NetFlow Collector Considerations	6-7
NetFlow Export of Multicast Traffic Flows	6-9
NetFlow Configuration Example	6-10
Cisco Network Analysis Module	6-12
NAM Analysis of Chassis Traffic	6-13
NAM Analysis of NetFlow Traffic	6-15
NAM Analysis of SPAN/RSPAN Traffic	6-22
Cisco IP Service Level Agreements	6-24
IPSLAs as a Pre-Assessment Tool	6-24
IPSLA as an Ongoing Performance Monitoring Tool	6-32
Router and Switch Command-Line Interface	6-35
Traceroute	6-37
show interface summary and show interface Commands	6-43
Platform Specific Queue-Level Commands	6-45
Simple Network Management Protocol	6-63
Application-Specific Management Functionality	6-66
Cisco TelePresence	6-66
Cisco TelePresence Manager	6-70
Cisco Unified Communications Manager	6-73
Cisco TelePresence Multipoint Switch	6-75
Cisco TelePresence System Endpoint	6-78
Cisco TelePresence SNMP Support	6-80
IP Video Surveillance	6-81
Digital Media Systems	6-81
Desktop Video Collaboration	6-81
Summary	6-82

CHAPTER 7

Medianet Auto Configuration 7-1

Auto Smartports	7-1
Platform Support	7-2
Switch Configuration	7-3
ASP Macro Details	7-7
Medianet Devices with Built-in ASP Macros	7-9
Cisco IPVS Cameras	7-9
Cisco Digital Media Players (DMPs)	7-12
Medianet Devices without Built-in ASP Macros	7-13
Cisco TelePresence (CTS) Endpoints	7-13
Other Video Conferencing Equipment	7-14
Overriding Built-in Macros	7-14

Macro-of-Last-Resort	7-18
Custom Macro	7-20
Security Considerations	7-22
Authenticating Medianet Devices	7-23
CDP Fallback	7-24
Guest VLANs and LAST_RESORT Macro	7-24
Verifying the VLAN Assignment on an Interface	7-25
ASP with Multiple Attached CDP Devices	7-25
Deployment Considerations	7-26
Location Services	7-26
Summary	7-28
References	7-29



CHAPTER 1

Medianet Architecture Overview

Executive Summary

Media applications—particularly video-oriented media applications—are exploding over corporate networks, exponentially increasing bandwidth utilization and radically shifting traffic patterns. There are several business drivers behind media application growth, including a globalized workforce, the pressure to go “green,” the transition to high-definition media (both in consumer and corporate markets) and social networking phenomena that are crossing over into the workplace. As a result, media applications are fueling a new wave of IP convergence, necessitating a fresh look at the network architecture.


Converging media applications onto an IP network is much more complex than converging VoIP alone; this is not only because media applications are generally bandwidth-intensive and bursty (as compared to VoIP), but also because there are so many different types of media applications: beyond IP Telephony, these can include live and on-demand streaming media applications, digital signage applications, high-definition room-based conferencing applications as well as an infinite array of data-oriented applications. By embracing media applications as the next cycle of convergence, IT departments can think holistically about their network architecture and its readiness to support the coming tidal wave of media applications and develop a network-wide strategy to ensure high quality end-user experiences.

Furthermore, thinking about your media application strategy now can help you take the first steps toward the next IP convergence wave and give your business competitive advantages, including the ability to harness the collective creativity and knowledge of your employees and to fundamentally change the experience your customers receive, all through the availability, simplicity and effectiveness of media applications.

Additionally, media applications featuring video are quickly taking hold as the de facto medium for communication, supplementing virtually every other communication media. As a result, a significant portion of know-how and intellectual property is migrating into video mediums. It is critical to get ahead of this trend in order to maintain control of company assets and intellectual property.

Offering both compelling media applications, like TelePresence and WebEx, as well as an end-to-end network design to support this next convergence wave, Cisco is in a unique position to provide a medianet architecture which can ensure the experience well into the collaborative workforce, enabling strategic and competitive advantage.

High-level requirements of medianets are addressed, including availability and quality requirements, bandwidth and optimization requirements, and access control and security requirements. Following these, specific strategic recommendations in designing campus, WAN and branch, and data center medianets are presented.

Figure 1-1 Media Applications

Business Drivers for Media Applications

There are several business drivers behind media application growth, including a globalized workforce, the pressure to go green, the transition to high-definition media (both in consumer and corporate markets), and social networking phenomena that are crossing over into the workplace. These and other business drivers are discussed in additional detail below.

Global Workforce and the Need for Real-Time Collaboration

The first stage of productivity for most companies is acquiring and retaining the skilled and talented individuals in a single or few geographic locations. More recently the focus has been on finding technology solutions to enable a geographically-distributed workforce to collaborate together as a team, enabling companies more flexibly to harness talent “where it lives.” While this approach has been moderately successful, there is a new wave of productivity on the horizon: harnessing collective and collaborative knowledge.

Future productivity gains will be achieved by creating collaborative teams that span corporate boundaries, national boundaries, and geographies. Employees will collaborate with partners, research and educational institutions, and customers to create a new level of collective knowledge.

To do so, real-time multimedia collaboration applications will be absolutely critical to the success of these virtual teams. Video offers a unique medium which streamlines the effectiveness of communications between members of such teams. For this reason, real-time interactive video will become increasingly prevalent, as will media integrated with corporate communications systems.

Pressures to be “Green”

For many reasons, companies are seeking to reduce employee travel. Travel creates expenses to the bottom-line, as well as significant productivity impacts while employees are in-transit and away from their usual working environments. Many solutions have emerged to assist with productivity while traveling, including Wireless LAN hotspots, remote access VPNs, and softphones, all attempting to keep the employee connected while traveling.

More recently companies are under increasing pressures to demonstrate environmental responsibility, often referred to as being “green.” On the surface such initiatives may seem like a pop-culture trend, but lacking in tangible corporate returns. However, it is entirely possible to pursue “green” initiatives, while simultaneously increasing productivity and lowering expenses.

Media applications, such as Cisco TelePresence, offer real solutions to remote collaboration challenges and have demonstrable savings as well. For example, during the first year of deployment, Cisco measured its usage of TelePresence in direct comparison to the employee travel that would otherwise have taken place and found that over 80,000 hours of meetings were held by TelePresence instead of physical travel, avoiding \$100 million of travel expenses, as well as over 30,000 tons of carbon emissions.

Being “green” does not have to be a “tax;” it can improve productivity and reduce corporate expenses, offering many dimensions of return on investment, while at the same time sending significant messages to the global community of environmental responsibility.

New Opportunities for IP Convergence

Many advantages were achieved through the convergence of voice onto IP networks. In addition to cost savings, new communications applications were made possible by the integration of VoIP with other media applications on the IP network.

There is a new wave of IP convergence emerging for media applications. One source of convergence is from applications historically having dedicated video transmission and broadcast networks. For example, high-definition video collaboration, video surveillance systems, and video advertising signage typically had dedicated private systems for the creation and dissemination of video content. Increasingly, companies are further leveraging the investment in their corporate network by converging these video applications onto a single IP network. Cisco TelePresence, Cisco IP video surveillance, and Cisco Digital Media System products all make this convergence a reality.

A second source of convergence is the integration of video as a medium into many other forms of corporate communications. For example, video cameras integrated with the VoIP system (such as Cisco Unified Personal Communicator) provide an easy way to add video to existing VoIP calling patterns. Further, collaboration tools such as Cisco MeetingPlace and Cisco WebEx add video mediums as a capability for simple conferencing and real-time collaboration.

Transition to High-Definition Media

One of the reasons traditional room-to-room video conferencing and desktop webcam-style video conferencing are sometimes questioned as less than effective communications systems is the reliance on low-definition audio and video formats.

On the other hand, high-definition interactive media applications, like Cisco TelePresence, demonstrate how high-definition audio and video can create an experience where meeting participants feel like they are in the same meeting room, enabling a more effective remote collaboration experience. IP video surveillance cameras are migrating to high-definition video in order to have digital resolutions needed for new functions, such as pattern recognition and intelligent event triggering based on motion and visual characteristics. Cisco fully expects other media applications to migrate to high-definition in the near future, as people become accustomed to the format in their lives as consumers, as well as the experiences starting to appear in the corporate environment.

High-definition media formats transmitted over IP networks create unique challenges and demands on the network that need to be planned for. Demands including not only bandwidth, but also transmission reliability and low delay become critical issues to address.

Media Explosion

Another factor driving the demand for video on IP networks is a sheer explosion of media content. The barriers to media production, distribution, and viewing have been dramatically lowered. For example, five to ten years ago video cameras became so affordable and prevalent that just about everyone bought one and became an amateur video producer. Additionally, video cameras are so common that almost every cell phone, PDA, laptop, and digital still camera provide a relatively high-quality video capture capability. However, until recently, it was not that easy to be a distributor of video content, as distribution networks were not common.

Today, social networking sites like YouTube, MySpace and many others appearing every day have dramatically lowered the barrier to video publishing to the point where anyone can do it. Video editing software is also cheap and easy to use. Add to that a free, global video publishing and distribution system, and essentially anyone, anywhere can be a film studio. With little or no training, people are making movie shorts that rival those of dedicated video studios.

The resulting explosion of media content is now the overwhelming majority of consumer network traffic, and is quickly “crossing over” to corporate networks. The bottom line is there are few barriers left to inhibit video communication, and so this incredibly effective medium is appearing in new and exciting applications every day.

Social Networking—Not Just For Consumers Anymore

Social networking started as a consumer phenomenon, with every day people producing and sharing rich media communications such as blogs, photos, and videos. When considering the affect it may have on corporate networks, some IT analysts believed social networking would stay as a consumer trend, while others believed the appearance in corporate networks was inevitable.

Skeptics look at social networking sites like Myspace, YouTube and others and see them as fads primarily for the younger population. However, looking beyond the sites themselves it is important to understand the new forms of communication and information sharing they are enabling. For example, with consumer social networking typically people are sharing information about themselves, about subjects they have experience in, and interact with others in real-time who have similar interests. In the workplace, we already see the parallels happening, because the same types of communication and information sharing are just as effective.

The corporate directory used to consist of employee names, titles, and phone numbers. Companies embracing social networking are adding to that skillsets and experience, URL links to shared work spaces, blogs, and other useful information. The result is a more productive and effective workforce that can adapt and find the skillsets and people needed to accomplish dynamic projects.

Similarly, in the past information was primarily shared via text documents, E-mail, and slide sets. Increasingly, we see employees filming short videos to share best practices with colleagues, provide updates to peers and reports, and provide visibility into projects and initiatives. Why have social networking trends zeroed in on video as the predominant communication medium? Simple: video is the most effective medium. People can show or demonstrate concepts much more effectively and easily using video than any other medium.

Just as a progression occurred from voice exchange to text, to graphics, and to animated slides, video will start to supplant those forms of communications. Think about the time it would take to create a good set of slides describing how to set up one of your company’s products. Now how much easier would it be just to film someone actually doing it? That’s just one of many examples where video is supplanting traditional communication formats.

At Cisco, we have seen the cross-over with applications like Cisco Vision (C-Vision). Started as an ad-hoc service by several employees, C-Vision provides a central location for employees to share all forms of media with one another, including audio and video clips. Cisco employees share information on projects, new products, competitive practices, and many other subjects. The service was used by so many employees, Cisco's IT department assumed ownership and scaled the service globally within Cisco. The result is a service where employees can become more effective and productive, quickly tapping into each other's experience and know-how, all through the effectiveness and simplicity of video.

Bottom-Up versus Top-Down Media Application Deployments

Closely-related to the social-networking aspect of media applications is that users have increasingly driven certain types of media application deployments within the enterprise from the "bottom-up" (i.e., the user base either demands or just begins to use a given media application with or without formal management or IT support). Such bottom-up deployments are illustrated by the Cisco C-Vision example mentioned in the previous section. Similar bottom-up deployment patterns have been noted for other Web 2.0 and multimedia collaboration applications.

In contrast, company-sponsored video applications are pushed from the "top-down" (i.e., the management team decides and formally directs IT to support a given media application for their user-base). Such top-down media applications may include Cisco TelePresence, digital signage, video surveillance, and live broadcast video meetings.

The combination of top-down and bottom-up media application proliferation places a heavy burden on the IT department as it struggles to cope with officially-supported and officially-unsupported, yet highly-proliferated, media applications.

Multimedia Integration with Communications Applications

Much like the integration of rich text and graphics into documentation, audio and video media will continue to be integrated into many forms of communication. Sharing of information with E-mailed slide sets will start to be replaced with video clips. The audio conference bridge will be supplanted with the video-enabled conference bridge. Collaboration tools designed to link together distributed employees will increasingly integrate desktop video to bring teams closer together.

Cisco WebEx is a prime example of such integration, providing text, audio, instant messaging, application sharing, and desktop video conferencing easily to all meeting participants, regardless of their location. Instead of a cumbersome setup of a video conference call, applications such as Cisco Unified Personal Communicator and Cisco WebEx greatly simplify the process, and video capability is added to the conference just as easily as any other type of media, like audio.

Demand for Universal Media Access

Much like the mobile phone and wireless networking, people want to extend communications everywhere they want to use them. The mobile phone unwired audio, making voice communications accessible virtually anywhere on the planet. Wireless networking untethered the laptop and PDA, extending high-speed data communications to nearly everywhere and many different devices.

Media applications will follow the same model. As multimedia applications become increasingly utilized and integrated, the demands from users will be to access these applications wherever they are, and on their device of choice. These demands will drive the need for new thinking about how employees work and how to deliver IT services to them.

Today employees extend the workplace using mobile phones and wireless networking to home offices, airports, hotels, and recreation venues. But, for example, with increased reliance on video as a communication medium, how will video be extended to these same locations and with which devices? We already see the emergence of video clips filmed with mobile phones and sent to friends and colleagues. Participation in video conferencing, viewing the latest executive communications, and collaborating with co-workers will need to be accessible to employees, regardless of their work location.

Challenges of Medianets

There are a number of challenges in designing an IP network with inherent support for the limitless number of media applications, both current and future. The typical approach followed is to acquire a media application, like IP Video Conferencing, make the network improvements and upgrades needed to deliver that specific application, and then monitor the user feedback. While a good way to implement a single application, the next media application will likely require the same process, and repeated efforts, and often another round of network upgrades and changes.

A different way to approach the challenge is to realize up-front that there are going to be a number of media applications on the network, and that these applications are likely to start consuming the majority of network resources in the future. Understanding the collection of these applications and their common requirements on the network can lead to a more comprehensive network design, better able to support new media applications as they are added. This design is what we term the *medianet*.

Considerations for the medianet include media delivery, content management, client access and security, mobility, as well as integration with other communications systems and applications.

Understanding Different Media Application Models

Different media applications will behave differently and put different requirements on the network. For example, Cisco TelePresence has relatively high bandwidth requirements (due to the HD video streams being transmitted) and tight tolerances for delivery. Traffic patterns are somewhat predictable, due to the room-to-room calling characteristics. In contrast, Cisco Digital Signage typically has less stringent delivery tolerances, and the traffic flows are from a central location (or locations) out towards several or many endpoints (see [Figure 1-2](#)).

Figure 1-2 Understanding Media Application Behavior Models

	Model	Direction of Flows	Traffic Trends
Interactive	TelePresence	Many to Many Client ← → Client MCU ← → Client	High-def video requires up to 4-12Mbps per location Expansion down to the individual user
	Desktop Multimedia Conferencing	Many to Many Client ← → Client MCU ← → Client	Collaboration across geographies Growing peer -to-peer model driving higher on -demand bandwidth
	Video Surveillance	Many to Few Source → Storage Storage → Client Source → Client	IP convergence opening up usage and applications Higher quality video requirements driving higher bandwidth (up to 3-4Mbps per camera)
Streaming	Desktop Streaming Media and Digital Signage	Few to Many Storage → Client Source → Client	Tremendous increase in applications driving more streams Demand for higher quality video increases each stream

224514

The four media applications shown in [Figure 1-2](#) cover a significant cross-section of models of media application behavior. To include additional applications in the inventory, critical questions to consider include:

- Is the media stored and viewed (streaming) or real-time (interactive)?
- Where are the media sources and where are the viewers?
- Which direction do the media flows traverse the network?
- How much bandwidth does the media application require? And how much burst?
- What are the service level tolerances (in terms of latency, jitter and loss)?
- What are the likely media application usage patterns?
- Are there requirements to connect to other companies (or customers)?
- In what direction is the media application likely to evolve in the future?

With a fairly straightforward analysis, it is possible to gain tremendous understanding into the network requirements various media applications.

One important consideration is: where is/are the media source(s) and where is/are the consumer(s)? For example, with desktop multimedia conferencing, the sources and consumers are both the desktop; therefore, the impacts to the network are very likely to be within the campus switching network, across the WAN/VPN, and the branch office networks. Provisioning may be challenging, as the ad-hoc conference usage patterns may be difficult to predict; however, voice calling patterns may lend insight into likely media conferencing calling patterns.

To contrast, the sources of on-demand media streams are typically within the data center, from high-speed media servers. Because viewers can be essentially any employee, this will affect the campus switching network, the WAN/VPN, the branch offices, and possibly even remote teleworkers. Since there will may be many simultaneous viewers, it would be inefficient to duplicate the media stream to each viewer; so wherever possible, we would like to take advantage of broadcast optimization technologies.

In these simplistic examples, you can see why it is important to understand how different media applications behave in order to understand how they are likely to impact your network. Start by making a table with (at least) the above questions in mind and inventory the various media applications in use today, as well as those being considered for future deployments. Common requirements will emerge, such as the need to meet “tight” service levels, the need to optimize bandwidth, and the need to optimize broadcasts, which will be helpful in determining media application class groupings (discussed in more detail later).

Delivery of Media Applications

A critical challenge the converged IP network needs to address is delivery of media application traffic, in a reliable manner, while achieving the service levels required by each application. Media applications inherently consume significant amounts of network resources, including bandwidth. A common tendency is to add network bandwidth to existing IP networks and declare them ready for media applications; however, bandwidth is just one factor in delivering media applications.

Media applications, especially those which are real-time or interactive, require reliable networks with maximum up-time. For instance, consider the loss sensitivities of VoIP compared to high-definition media applications, such as HD video. For a voice call, a packet loss percentage of even 1% can be effectively concealed by VoIP codecs; whereas, the loss of two consecutive VoIP packets will cause an audible “click” or “pop” to be heard by the receiver. In stark contrast, however, video-oriented media applications generally have a much greater sensitivity to packet loss, especially HD video applications, as these utilize highly-efficient compression techniques, such as H.264. As a result, a tremendous amount of visual information is represented by a relatively few packets, which if lost, immediately become visually apparent in the form of screen pixelization. With such HD media applications, such as Cisco TelePresence, the loss of even one packet in 10,000 can be noticed by the end user. This represents a hundred-fold increase in loss sensitivity when VoIP is compared to HD video.

Therefore, for each media application, it is important to understand the delivery tolerances required in order to deliver a high-quality experience to the end user.

Prioritizing the Right Media Applications, Managing the Rest


With the first stage of IP convergence, the Cisco Architecture for Voice, Video, and Integrated Data (AVVID) provided the foundation for different applications to effectively and transparently share the same IP network. One of the challenges to overcome with converged networks is to be able to simultaneously meet different application requirements, prioritizing network resources accordingly. Quality of Service (QoS) continues to be a critical set of functions relied upon in the network to provide differentiated service levels, assuring the highest priority applications can meet their delivery requirements.

The AVVID model defined best practices for adding Voice-over-IP (VoIP) and Video over IP applications to the existing data IP network. Most QoS implementations assume a number of data applications, a single or few VoIP applications, and a single or few video applications.

Today there is a virtual explosion of media applications on the IP network with many different combinations of audio, video and data media. For example, VoIP streams can be standard IP telephony, high-definition audio, internet VoIP, or others. Video streams can range from relatively low-definition webcams to traditional video-over-IP room-to-room conferencing to or high-definition Cisco TelePresence systems. Additionally, there are new IP convergence opportunities occurring which further expand the number of media applications and streams on the IP network (see [Figure 1-3](#)).

Another source of new media streams on the network is “unmanaged” media applications; namely, applications which are considered primarily for consumers, but are also used by corporate employees. Many of these unmanaged media applications may fall into a gray area for some companies in terms of usage policies. For instance, at first glance, consumer media sharing sites such as YouTube may appear to be clearly consumer-only applicability; however, many of these same services also contain videos that can provide considerable know-how and information that are useful to employees as well.

Figure 1-3 Media Explosion Driving New Convergence Evolution



Beyond the current “media explosion” which is driving a new wave of IP convergence, new and exciting applications targeted at collaboration are integrating numerous types of streams and media into end-user applications. Cisco TelePresence is one example, combining HD video streams, HD audio, application sharing, and some level of interoperability with traditional video conferencing, into an overall collaboration tool and near in-person meeting experience. Cisco WebEx is another example, combining many types of media sharing for web-based meetings. Such applications provide new challenges for prioritizing media application streams.

The explosion of media content, types and applications—both managed and unmanaged—requires network architects to take a new look at their media application provisioning strategy. Without a clear strategy, the number and volume of media applications on the IP network could very well exceed the ability of the network administrator to provision and manage.

Media Application Integration

As media applications increase on the IP network, integration will play a key role in two ways: first, media streams and endpoints will be increasingly leveraged by multiple applications. For example, desktop video endpoints may be leveraged for desktop video conferencing, web conferencing, and for viewing stored streaming video for training and executive communications.

Second, many media applications will require common sets of functions, such as transcoding, recording, and content management. To avoid duplication of resources and higher implementation costs, common media services need to be integrated into the IP network so they can be leveraged by multiple media applications.

Securing Media Applications

Because of the effectiveness of multimedia communication and collaboration, the security of media endpoints and communication streams becomes an important part of the media-ready strategy. Access controls for endpoints and users, encryption of streams, and securing content files stored in the data center are all part of a required comprehensive media application security strategy.

Other specialized media applications, such as IP video surveillance and digital signage, may warrant additional security measures due to their sensitivity and more restricted user group. Placing such media applications within private logical networks within the IP network can offer an additional layer of security to keep their endpoints and streams confidential.

Finally, as the level of corporate intellectual property migrates into stored and interactive media, it is critical to have a strategy to manage the media content, setting and enforcing clear policies, and having the ability to protect intellectual property in secure and managed systems. Just as companies have policies and processes for handling intellectual property in document form, they also must develop and update these policies and procedures for intellectual property in media formats.

Solution

The Need for a Comprehensive Media Network Strategy

It is possible to pursue several different strategies for readying the IP network for media applications. One strategy is to embrace media applications entirely, seeing these technologies as driving the next wave of productivity for businesses. Another strategy is to adopt a stance to manage and protect select media applications on the network. Still another strategy would be to not manage media applications at all. Which strategy should you pursue?

If we have learned anything from past technology waves which enable productivity, it is this: if corporate IT does not deploy (or lags significantly in deployment) users will try to do it themselves... and usually poorly. For example, several years ago, some IT departments were skeptical of the need to deploy Wireless LANs (WLANS) or questioned-and rightly so-their security. As a result, many WLAN deployments lagged. Users responded by purchasing their own consumer-grade WLAN access-points and plugging them into corporate networks, creating huge holes in the network security strategy. Such “rogue” access-points in the corporate network, lacking proper WLAN security, not only represented critical security vulnerabilities to the network as a whole, but also were difficult for network administrators to locate and shut down.

The coming media application wave will be no different and is already happening. IT departments lacking a media application strategy may find themselves in the future trying to regain control of traffic on the network. It is advantageous to define a comprehensive strategy now for how media applications will be managed on the network. Key questions the strategy should answer include:

- Which are the business-critical media applications? And what service levels must be ensured for these applications?
- Which media applications will be managed or left unmanaged?


- What will the usage policies be and how will they be enforced?

As mentioned earlier, one approach to planning the network is to assess the network upgrades and changes required for each new media application deployed by the company. This approach could lead to a lot of repeated effort and change cycles by the IT staff and potentially incompatible network designs. A more efficient and far-sighted approach would be to consider all the types of media applications the company is currently using—or may use in the future—and design a network-wide architecture with media services in mind.

Architecture of a Medianet

A medianet is built upon an architecture that supports the different models of media applications and optimizes their delivery, such as those shown in the architectural framework in [Figure 1-4](#).

Figure 1-4 Architectural Framework of a Medianet



A medianet framework starts with an end-to-end network infrastructure designed and built to achieve high availability, including the data center, campus, WAN, and branch office networks. The network provides a set of services to video applications, including:

- Access services—Provide access control and identity of video clients, as well as mobility and location services
- Transport services—Provide packet delivery, ensuring the service levels with QoS and delivery optimization
- Bridging services—Transcoding, conferencing, and recording services
- Storage services—Content capture, storage, retrieval, distribution, and management services

- Session control services—Signaling and control to setup and tear-down sessions, as well as gateways

When these media services are made available within the network infrastructure, endpoints can be multi-purpose and rely upon these common media services to join and leave sessions for multiple media applications. Common functions such as transcoding and conferencing different media codecs within the same session can be deployed and leveraged by multiple applications, instead of being duplicated for each new media application.

Where these different services are deployed within the network can also be customized for different business models or media applications. For example, it may be advantageous to store all IP video surveillance feeds centrally in the data center, or for some companies it may be preferable to have distributed storage in branch office networks.


Common Requirements and Recommendations

After understanding the behavior of the different media applications in the network, there are common threads of requirements that can be derived. The top recommendations based on these common requirements are discussed in the follow subsections.

Network Design for High Availability

Data applications are tolerant of multi-second interruptions, while VoIP and video applications require tighter delivery requirements in order to achieve high quality experiences for the end users. Networks that have already implemented higher availability designs with VoIP convergence in mind are a step ahead.

Loss of packets, whether due to network outage or other cause, necessitates particular attention for media applications, especially those that require extreme compression. For example, HD video, would require billions of bytes to be transmitted over the IP network and is not practically deployable without efficient compression schemes like MPEG4 or H.264. To illustrate this point, consider a high-definition 1080p30 video stream, such as used by Cisco TelePresence systems. The first parameter “1080” refers to 1080 lines of horizontal resolution, which are matrixed with 1920 lines of vertical resolution (as per the 16:9 Widescreen Aspect Ratio used in High Definition video formatting), resulting in 2,073,600 pixels per screen. The second parameter “p” indicates a progressive scan, which means that every line of resolution is refreshed with each frame (as opposed to an interlaced scan, which would be indicated with an “i” and would mean that every other line is refreshed with each frame). The third parameter “30” refers to the transmission rate of 30 frames per second. While video sampling techniques may vary, each pixel has approximately 3 Bytes of color and/or luminance information. When all of this information is factored together (2,073,600 pixels x 3 Bytes x 8 bits per Byte x 30 frames per second), it results in approximately 1.5 Gbps of information. However, H.264-based Cisco TelePresence codecs transmit this information at approximately 5 Mbps (maximum), which translates to over 99% compression. Therefore, the overall effect of packet loss is proportionally magnified, such that dropping even one packet in 10,000 (0.01% packet loss) is noticeable to end users in the form of minor pixelization. This is simply because a single packet represents a hundred or more packets’ worth of information, due to the extreme compression ratios applied, as illustrated in [Figure 1-5](#).

Figure 1-5 Compression Ratios for HD Video Applications

Traditional network designs supporting data applications may have targeted packet loss at less than 1-2%. For VoIP, network designs were tightened to have only 0.5-1% of packet loss. For media-ready networks, especially those supporting high-definition media applications, network designs need to be tightened again by an order of magnitude, targeting 0-0.05% packet loss.

However, an absolute target for packet loss is not the only consideration in HA network design. Loss, during normal network operation, should effectively be 0% on a properly-designed network. In such a case, it is generally only during network events, such as link failures and/or route-flaps, that packet loss would occur. Therefore, it is usually more meaningful to express availability targets not only in absolute terms, such as <0.05%, but also in terms of convergence targets, which are sometimes also referred to as the Mean-Time-to-Repair (MTRR) targets.

Statistical analysis on speech and communications have shown that overall user satisfaction with a conversation (whether voice or interactive video) begins to drop when latency exceeds 200 ms¹. This is because 200 ms is about the length of time it takes for one party to figure out that the other person has stopped talking and thus, it is their turn to speak. This value (200 ms) provides a subjective “conversation disruption” metric. Put another way, a delay in excess of 200 ms—whether network transmission delay or network convergence delay—would impact the naturalness of a voice or video conversation. This is not to say that a loss of packets for 200 ms is unnoticeable to end users (as already mentioned, a loss of a single packet in 10,000 may be noticeable as minor pixelization in some HD video applications); however, a temporary interruption in a media application of 200 ms would likely not be considered intolerable, should it happen, and would not significantly impact a conversation.

Therefore, a network convergence target for highly-available campus and data center networks supporting media applications is 200 ms. On other network topologies, such as WAN and branch networks, this target is more likely unattainable, given the technologies and constraints involved, in which case the network should be designed to converge in the lowest achievable amount of time.

1. ITU G.114 (E-Model)—Note: The primary application of the ITU G.114 E-Model is to target one-way transmission latency; however, these observations and metrics can also be applied to target network convergence latency.

To summarize: the targets for media-ready campus and data center networks in terms of packet loss is 0.05% with a network convergence target of 200 ms; on WAN and branch networks, loss should still be targeted to 0.05%, but convergence targets will be higher depending on topologies, service providers, and other constraints. Finally, it should be noted that by designing the underlying network architecture for high availability, all applications on the converged network benefit.

Bandwidth and Burst

There is no way around the fact that media applications require significant network bandwidth. An important step to implement a medianet is to assess current and future bandwidth requirements across the network. Consider current bandwidth utilization and add forecasts for media applications, especially for video-oriented media applications. Because video is in a relatively early stage of adoption, use aggressive estimates of possible bandwidth consumption. Consider bandwidth of different entry and transit points in the network. What bandwidth is needed at network access ports both in the campus as well as branch offices? What are the likely media streams needing transport across the WAN?

It is important to consider all types of media applications. For example, how many streaming video connections will be utilized for training and communications? These typically flow from a central point, such as the data center, outward to employees in campus and branch offices. As another example, how many IP video surveillance cameras will exist on the network? These traffic flows are typically from many sources at the edges of the network inward toward central monitoring and storage locations.

Map out the media applications that will be used, considering both managed and un-managed applications. Understand the bandwidth required by each stream and endpoint, as well as the direction(s) in which the streams will flow. Mapping those onto the network can lead to key bandwidth upgrade decisions at critical places in the network architecture, including campus switching as well as the WAN.


Another critical bandwidth-related concern is burst. So far, we have discussed bandwidth in terms of bits per second (i.e., how much traffic is sent over a one second interval); however, when provisioning bandwidth, burst must also be taken into account. Burst is defined as the amount of traffic (generally measured in Bytes) transmitted per millisecond which exceeds the per-second average.

For example, a Cisco TelePresence 3000 system may average 15 Megabits per second, which equates to an average per millisecond rate of 1,875 Bytes ($15 \text{ Mbps} \div 1,000 \text{ milliseconds} \div 8 \text{ bits per Byte}$). Cisco TelePresence operates at 30 frames per second, which means that every 33 ms a video frame is transmitted. Each frame consists of several thousand Bytes of video payload, and therefore each frame interval consists of several dozen packets, with an average packet size of 1,100 bytes per packet. However, because video is variable in size (due to the variability of motion in the encoded video), the packets transmitted by the codec are not spaced evenly over each 33 ms frame interval, but rather are transmitted in bursts measured in shorter intervals. Therefore, while the overall bandwidth (maximum) averages out to 15 Mbps over one second, when measured on a per millisecond basis, the packet transmission rate is highly variable, and the number of Bytes transmitted per millisecond for a 15 Mbps stream can burst well above the 1,875 Bytes per millisecond average. Therefore, adequate burst tolerance must be accommodated by all switch and router interfaces in the path.

Given these considerations, it can be noted that converging voice onto a common IP-based network is a significantly simpler exercise than converging video onto the same network. The principle reason is that VoIP is a very well-behaved application, from a networking perspective. For instance, each VoIP packet size is known and constant (for example, G.711 codecs generate packets that are always 160 Bytes [+ Layer 2 overhead]); similarly, VoIP packetization rates are known and constant (the default packetization rate for VoIP is 50 packet-per-second, which produces a packet every 20 ms). Furthermore, VoIP has very light bandwidth requirements (as compared to video and data) and these requirements can be very cleanly calculated by various capacity planning formulas (such as Erlang and Endset formulas).

In contrast, video is a completely different type of application in almost every way. Video packet sizes vary significantly and video packetization rates also vary significantly (both in proportion to the amount of motion in the video frames being encoded and transmitted); furthermore, video applications are generally quite bursty—especially during sub-second intervals—and can wreak havoc on underprovisioned network infrastructures. Additionally, there are no clean formulas for provisioning video, as there are with VoIP. This contrast—from a networking perspective—between voice and video traffic is illustrated in [Figure 1-6](#)

Figure 1-6 Sub-Second Bandwidth Analysis—Voice versus Video



Summing up, converging media applications—especially video-based media applications—onto the IP network is considerably more complex than converging voice and data, due to the radically different bandwidth and burst requirements of video compared to voice. While deployment scenarios will vary, in most cases, capacity planning exercises will indicate that Campus and Data Center medianets will require GigabitEthernet (GE) connections at the edge and 10 GigabitEthernet (10GE) connections—or multiples thereof—in the core; additionally, medianets will likely have a minimum bandwidth requirement of 45 Mbps/DS3 circuits. Furthermore, network administrators not only have to consider the bandwidth requirements of applications as a function of bits-per-second, but also they must consider the burst requirements of media, such as video, as a function of Bytes-per-millisecond, and ensure that the routers and switches have adequate buffering capacity to handle bursts.

Latency and Jitter

Media applications, particularly interactive media applications, have strict requirements for network latency. Network latency can be broken down further into fixed and variable components:

- Serialization (fixed)

- Propagation (fixed)
- Queuing (variable)


Serialization refers to the time it takes to convert a Layer 2 frame into Layer 1 electrical or optical pulses onto the transmission media. Therefore, serialization delay is fixed and is a function of the line rate (i.e., the clock speed of the link). For example, a 45 Mbps DS3 circuit would require 266 μ s to serialize a 1500 byte Ethernet frame onto the wire. At the circuit speeds required for medianets (generally speaking DS3 or higher), serialization delay is not a significant factor in the overall latency budget.

The most significant network factor in meeting the latency targets for video is propagation delay, which can account for over 95% of the network latency budget. Propagation delay is also a fixed component and is a function of the physical distance that the signals have to travel between the originating endpoint and the receiving endpoint. The gating factor for propagation delay is the speed of light: 300,000 km/s or 186,000 miles per second. Roughly speaking, the speed of light in an optical fiber is about one-sixth the speed of light in a vacuum. Thus, the propagation delay works out to be approximately 4-6 μ s per km (or 6.4-9.6 μ s per mile)¹.

Another point to keep in mind when calculating propagation delay is that optical fibers and coaxial cables are not always physically placed over the shortest path between two geographic points, especially over transoceanic links. Due to installation convenience, circuits may be hundreds or thousands of miles longer than theoretically necessary.

The network latency target specified in the ITU G.114 specification for voice and video networks is 150 ms. This budget allows for nearly 24,000 km (or 15,000 miles) worth of propagation delay (which is approximately 60% of the earth's circumference); the theoretical worst-case scenario (exactly half of the earth's circumference) would require 120 ms of latency. Therefore, this latency target (of 150 ms) should be achievable for virtually any two locations on the planet, given relatively direct transmission paths. Nonetheless, it should be noted that overall quality does not significantly degrade for either voice or video calls until latency exceeds 200 ms, as shown in [Figure 1-7](#) (taken from ITU G.114).

1. Per ITU G.114 Table 4.1: The transmission speeds of terrestrial coaxial cables is 4 μ s /km, of optical fiber cable systems with digital transmission is 5 μ s / km, and of submarine coaxial cable systems is 6 μ s /km (allowing for delays in repeaters and regenerators).

Figure 1-7 Network Latency versus Call Quality

The final network latency component to be considered is queuing delay, which is variable. Variance in network latency is also known as jitter. For instance, if the average latency is 100 ms and packets are arriving between 95 ms and 105 ms, the peak-to-peak jitter is defined as 10 ms. Queuing delay is the primary cause of jitter and is a function of whether a network node is congested or not, and if it is, what scheduling policies (if any) have been configured to manage congestion. For interactive media applications, packets that are excessively late (due to network jitter) are no better than packets that have been lost. Media endpoints usually have a limited amount of playout-buffering capacity to offset jitter. However, in general, it is recommended that jitter for real-time interactive media applications not exceed 10 ms peak-to-peak.

To recap: the one-way latency target for interactive media applications is 150 ms (with a threshold limit of 200 ms). Additionally, since the majority of factors contributing to the latency budget are fixed, careful attention has to be given to queuing delay, as this is the only latency/jitter factor that is directly under the network administrator's control (via QoS queuing policies, which are discussed in the next section, [Application Intelligence and Quality of Service](#)).

Application Intelligence and Quality of Service

Implementation of a comprehensive QoS strategy requires the ability to identify the business critical media applications and set a QoS service policy to mark and service such traffic. With the dramatic increase in types of media applications and streams, it becomes increasingly difficult to identify the critical media application streams from those that are considered unimportant. Streams using similar codecs may have similar packet construction and be difficult to classify using IP packet header information alone.

Therefore, packet classification needs to evolve to utilize deeper packet inspection technologies in order to have the granularity needed to distinguish between different types of media streams. Developing additional application intelligence within the network infrastructure is a crucial requirement to build a medianet, especially at the edges of the network where media endpoints first handoff packets into the network for transport.

Additionally, there are advantages of being able to perform media application sub-component separation, such that data components of a media application receive one level of service, whereas the audio and video components of the same application receive a different level of service¹. Such separation can simplify bandwidth provisioning, admission control, and capacity planning. That being said, media application sub-component separation more often than not requires deep packet inspection technologies, especially for media applications that are transported entirely within HTTP.

An alternative approach that presents another consideration is whether to trust media endpoints to mark their own traffic or not. Typically such endpoints can mark at Layer 2 (via 802.1Q/p CoS) or at Layer 3 (DSCP). Key factors the administrator needs to consider is how secure is the marking? Is it the marking centrally administered or locally set? Can it be changed or exploited by the end users? While trusting the endpoints to correctly mark themselves may simplify the network edge policies, it could present security vulnerabilities that could be inadvertently or deliberately exploited. In general, hardware-based media endpoints (such as dedicated servers, cameras, codecs, and gateways) are more “trustworthy,” whereas software-based media endpoints (such as PCs) are usually less “trustworthy.”

Nonetheless, whether media applications are explicitly classified and marked or are implicitly trusted, the question still remains of how should media applications be marked and serviced? As previously discussed, different media applications have different traffic models and different service level requirements. Ultimately, each class of media applications that has unique traffic patterns and service level requirements will need a dedicated service class in order to provision and guarantee these service level requirements. There is simply no other way to make service level guarantees. Thus, the question “how should media applications be marked and serviced?” becomes “how many classes of media applications should be provisioned and how should these individual classes be marked and serviced?”

To this end, Cisco continues to advocate following relevant industry standards and guidelines whenever possible, as this extends the effectiveness of your QoS policies beyond your direct administrative control. For example, if you (as a network administrator) decide to mark a realtime application, such as VoIP, to the industry standard recommendation (as defined in RFC 3246, “An Expedited Forwarding Per-Hop Behavior”), then you will no doubt provision it with strict priority servicing at every node within your enterprise network. Additionally, if you handoff to a service provider following this same industry standard, they also will similarly provision traffic marked Expedited Forwarding (EF-or DSCP 46) in a strict priority manner at every node within their cloud. Therefore, even though you do not have direct administrative control of the QoS policies within the service provider's cloud, you have extended the influence of your QoS design to include your service provider's cloud, simply by jointly following the industry standard recommendations.

That being said, it may be helpful to overview a guiding RFC for QoS marking and provisioning, namely RFC 4594, “Configuration Guidelines for DiffServ Service Classes.” The first thing to point out is that this RFC is not in the standards track, meaning that the guidelines it presents are not mandatory but rather it is in the informational track of RFCs, meaning that these guidelines are to be viewed as industry best practice recommendations. As such, enterprises and service providers are encouraged to adopt these marking and provisioning recommendations, with the aim of improving QoS consistency, compatibility, and interoperability. However, since these guidelines are not standards, modifications can be made to these recommendations as specific needs or constraints require. To this end, Cisco has made a minor modification to its adoption of RFC 4594, as shown in [Figure 1-8](#)².

1. However, it should be noted that in general it would not be recommended to separate audio components from video components within a media application and provision these with different levels of service, as this could lead to loss of synchronization between audio and video.

Figure 1-8 Cisco Media QoS Recommendations (RFC 4594-based)

Application Class	Per-Hop Behavior	Admission Control	Queuing and Dropping	Media Application Examples
VoIP Telephony	EF	Required	Priority Queue (PQ)	Cisco IP Phones (G.711, G.729)
Broadcast Video	CS5	Required	(Optional) PQ	Cisco IP Video Surveillance/Cisco Enterprise TV
Real-Time Interactive	CS4	Required	(Optional) PQ	Cisco TelePresence
Multimedia Conferencing	AF4	Required	BW Queue + DSCP WRED	Cisco Unified Personal Communicator
Multimedia Streaming	AF3	Recommended	BW Queue + DSCP WRED	Cisco Digital Media System (VoDs)
Network Control	CS6		BW Queue	EIGRP, OSPF, BGP, HSRP, IKE
Signaling	CS3		BW Queue	SCCP, SIP, H.323
Ops/Admin/Mgmt (OAM)	CS2		BW Queue	SNMP, SSH, Syslog
Transactional Data	AF2		BW Queue + DSCP WRED	Cisco WebEx/MeetingPlace/ERP Apps
Bulk Data	AF1		BW Queue + DSCP WRED	E-mail, FTP, Backup Apps, Content Distribution
Best Effort	DF		Default Queue + RED	Default Class
Scavenger	CS1		Min BW Queue	YouTube, iTunes, BitTorrent, Xbox Live

224550

RFC 4594 outlines twelve classes of media applications that have unique service level requirements:

- **VoIP Telephony**—This service class is intended for VoIP telephony (bearer-only) traffic (VoIP signaling traffic is assigned to the “Call Signaling” class). Traffic assigned to this class should be marked EF (DSCP 46). This class is provisioned with an Expedited Forwarding (EF) Per-Hop Behavior (PHB). The EF PHB—defined in RFC 3246—is a strict-priority queuing service, and as such, admission to this class should be controlled. Example traffic includes G,711 and G,729a.
- **Broadcast Video**—This service class is intended for broadcast TV, live events, video surveillance flows, and similar “inelastic” streaming media flows (“inelastic” flows refer to flows that are highly drop sensitive and have no retransmission and/or flow-control capabilities). Traffic in this class should be marked Class Selector 5 (CS5/DSCP 40) and may be provisioned with an EF PHB; as such, admission to this class should be controlled (either by an explicit admission control mechanisms or by explicit bandwidth provisioning). Examples traffic includes live Cisco Digital Media System (DMS) streams to desktops or to Cisco Digital Media Players (DMPs), live Cisco Enterprise TV (ETV) streams, and Cisco IP Video Surveillance (IPVS).
- **Real-time Interactive**—This service class is intended for (inelastic) room-based, high-definition interactive video applications and is intended primarily for audio and video components of these applications. Whenever technically possible and administratively feasible, data sub-components of this class can be separated out and assigned to the “Transactional Data” traffic class. Traffic in this class should be marked CS4 (DSCP 32) and may be provisioned with an EF PHB; as such, admission to this class should be controlled. An example application is Cisco TelePresence.

2. RFC 4594 recommends marking Call Signaling traffic to CS5. Cisco has recently completed a lengthy and expensive marking migration for Call Signaling from AF31 to CS3, and as such, have no plans to embark on another marking migration in the near future. RFC 4594 is an informational RFC (i.e., an industry best practice) and not a standard. Therefore, lacking a compelling business case at the time of writing, Cisco plans to continue marking Call Signaling as CS3 until future business requirements may arise that necessitate another marking migration. Therefore, the modification in [Figure 1-8](#) is that Call Signaling is marked CS3 and Broadcast Video (recommended to be marked CS3 in RFC 4594) is marked CS5.

- **Multimedia Conferencing**—This service class is intended for desktop software multimedia collaboration applications and is intended primarily for audio and video components of these applications. Whenever technically possible and administratively feasible, data sub-components of this class can be separated out and assigned to the “Transactional Data” traffic class. Traffic in this class should be marked Assured Forwarding¹ Class 4 (AF41/DSCP 34) and should be provisioned with a guaranteed bandwidth queue with DSCP-based Weighted-Random Early Detect (DSCP-WRED) enabled. Admission to this class should be controlled; additionally, traffic in this class may be subject to policing and re-marking². Example applications include Cisco Unified Personal Communicator, Cisco Unified Video Advantage, and the Cisco Unified IP Phone 7985G.
- **Multimedia Streaming**—This service class is intended for Video-on-Demand (VoD) streaming media flows which, in general, are more elastic than broadcast/live streaming flows. Traffic in this class should be marked Assured Forwarding Class 3 (AF31/DSCP 26) and should be provisioned with a guaranteed bandwidth queue with DSCP-based WRED enabled. Admission control is recommended on this traffic class (though not strictly required) and this class may be subject to policing and re-marking. Example applications include Cisco Digital Media System Video-on-Demand streams to desktops or to Digital Media Players.
- **Network Control**—This service class is intended for network control plane traffic, which is required for reliable operation of the enterprise network. Traffic in this class should be marked CS6 (DSCP 48) and provisioned with a (moderate, but dedicated) guaranteed bandwidth queue. WRED should not be enabled on this class, as network control traffic should not be dropped (if this class is experiencing drops, then the bandwidth allocated to it should be re-provisioned). Example traffic includes EIGRP, OSPF, BGP, HSRP, IKE, etc.
- **Call-Signaling**—This service class is intended for signaling traffic that supports IP voice and video telephony; essentially, this traffic is control plane traffic for the voice and video telephony infrastructure. Traffic in this class should be marked CS3 (DSCP 24) and provisioned with a (moderate, but dedicated) guaranteed bandwidth queue. WRED should not be enabled on this class, as call-signaling traffic should not be dropped (if this class is experiencing drops, then the bandwidth allocated to it should be re-provisioned). Example traffic includes SCCP, SIP, H.323, etc.
- **Operations/Administration/Management (OAM)**—This service class is intended for—as the name implies—network operations, administration, and management traffic. This class is important to the ongoing maintenance and support of the network. Traffic in this class should be marked CS2 (DSCP 16) and provisioned with a (moderate, but dedicated) guaranteed bandwidth queue. WRED should not be enabled on this class, as OAM traffic should not be dropped (if this class is experiencing drops, then the bandwidth allocated to it should be re-provisioned). Example traffic includes SSH, SNMP, Syslog, etc.
- **Transactional Data (or Low-Latency Data)**—This service class is intended for interactive, “foreground” data applications (“foreground” applications refer to applications that users are expecting a response—via the network—in order to continue with their tasks; excessive latency in response times of foreground applications directly impacts user productivity). Traffic in this class should be marked Assured Forwarding Class 2 (AF21 / DSCP 18) and should be provisioned with a dedicated bandwidth queue with DSCP-WRED enabled. This traffic class may be subject to policing and re-marking. Example applications include data components of multimedia collaboration applications, Enterprise Resource Planning (ERP) applications, Customer Relationship Management (CRM) applications, database applications, etc.
- **Bulk Data (or high-throughput data)**—This service class is intended for non-interactive “background” data applications (“background” applications refer to applications that the users are not awaiting a response—via the network—in order to continue with their tasks; excessive latency

1. The Assured Forwarding Per-Hop Behavior is defined in RFC 2597.

2. These policers may include Single-Rate Three Color Policers or Dual-rate Three Color Policers, as defined in RFC 2697 and 2698, respectively.

in response times of background applications does not directly impact user productivity. Furthermore, as most background applications are TCP-based file-transfers, these applications—if left unchecked—could consume excessive network resources away from more interactive, foreground applications). Traffic in this class should be marked Assured Forwarding Class 1 (AF11/DSCP 10) and should be provisioned with a dedicated bandwidth queue with DSCP-WRED enabled. This traffic class may be subject to policing and re-marking. Example applications include E-mail, backup operations, FTP/SFTP transfers, video and content distribution, etc.

- **Best Effort (or default class)**—This service class is the default class. As only a relative minority of applications will be assigned to priority, preferential, or even to deferential service classes, the vast majority of applications will continue to default to this best effort service class; as such, this default class should be adequately provisioned¹. Traffic in this class is marked Default Forwarding² (DF or DSCP 0) and should be provisioned with a dedicated queue. WRED is recommended to be enabled on this class. Although, since all the traffic in this class is marked to the same “weight” (of DSCP 0), the congestion avoidance mechanism is essentially Random Early Detect (RED).
- **Scavenger (or Low-Priority Data)**—This service class is intended for non-business related traffic flows, such as data or media applications that are entertainment-oriented. The approach of a less-than best effort service class for non-business applications (as opposed to shutting these down entirely) has proven to be a popular, political compromise: these applications are permitted on enterprise networks, as long as resources are always available for business-critical voice, video, and data applications. However, as soon the network experiences congestion, this class is the first to be penalized and aggressively dropped. Furthermore, the scavenger class can be utilized as part of an effective strategy for DoS and worm attack mitigation³. Traffic in this class should be marked CS1⁴ (DSCP 8) and should be provisioned with a minimal bandwidth queue that is the first to starve should network congestion occur. Example traffic includes YouTube, Xbox Live/360 Movies, iTunes, BitTorrent, etc.

Admission Control



Note

The reason “Admission Control” is used in this document, rather than “Call Admission Control,” is that not all media applications are call-oriented (e.g., IPVS and streaming video). Nonetheless, these non-call-oriented flows can also be controlled by administrative policies and mechanisms, in conjunction with bandwidth provisioning.

Bandwidth resources dedicated to strict-priority queuing need to be limited in order to prevent starvation of non-priority (yet business critical) applications. As such, contention for priority queues needs to be strictly controlled by higher-layer mechanisms.

Admission control solutions are most effective when built on top of a DiffServ-enabled infrastructure, that is, a network that has Differentiated Services (QoS policies for marking, queuing, policing, and dropping) configured and activated, as illustrated in [Figure 1-9](#).

The first level of admission control is simply to enable mechanisms to protect voice-from-voice and/or video-from-video on a first-come, first-serve basis. This functionality provides a foundation on which higher-level policy-based decisions can be built.

1. Cisco recommends provisioning no less than 25% of a link’s bandwidth for the default best effort class.
2. Default Forwarding is defined in RFC 2474.
3. See the QoS SRND at www.cisco.com/go/srnd for more details.
4. A Lower-Effort Per-Domain Behavior that defines a less than best effort or scavenger level of service—along with the marking recommendation of CS1—is defined in RFC 3662.

The second level of admission control factors in dynamic network topology and bandwidth information into a real-time decision of whether or not a media stream should be admitted.

The third level of admission control introduces the ability to preempt existing flows in favor of “higher-priority” flows.

The fourth level of admission control contains policy elements and weights to determine what exactly constitutes a “higher-priority” flow, as defined by the administrative preferences of an organization. Such policy information elements may include-but are not limited to-the following:


- Scheduled versus Ad Hoc—Media flows that have been scheduled in advance would likely be granted priority over flows that have been attempted ad hoc.
- Users/Groups—Certain users or user groups may be granted priority for media flows.
- Number of participants—Multipoint media calls with larger number of participants may be granted priority over calls with fewer participants.
- External versus internal participants—Media sessions involving external participants, such as customers, may be granted priority over sessions comprised solely of internal participants.
- Business critical factor—Additional subjective elements may be associated with media streams, such as a business critical factor. For instance, a live company meeting would likely be given a higher business critical factor than a live training session. Similarly, a media call to close a sale or to retain a customer may be granted priority over regular, ongoing calls.



Note

It should be emphasized this is not an exhaustive list of policy information elements that could be used for admission control, but rather is merely a sample list of possible policy information elements. Additionally, each of these policy information elements could be assigned administratively-defined weights to yield an overall composite metric to calculate and represent the final admit/deny admission control decision for the stream.

And finally, the fifth level of admission control provides graceful conflict resolution, such that-should preemption of a media flow be required-existing flow users are given a brief message indicating that their flow is about to be preempted (preferably including a brief reason as to why) and a few seconds to make alternate arrangements (as necessary).

Figure 1-9 Levels of Admission Control Options

Broadcast Optimization

Several media applications which utilize streaming, such as corporate broadcast communications, live training sessions, and video surveillance, have a traffic model with a single or few media sources transmitting to many simultaneous viewers. With such media applications present on the network, it is advantageous to optimize these broadcasts so that preferably a single (or few) packet streams are carried on the network that multiple viewers can join, instead of each viewer requiring their own dedicated packet stream.

IP multicast (IPmc) is a proven technology that can be leveraged to optimize such media applications. Stream “splitting” is an alternative starting to appear in products. Stream splitting behaves in a similar fashion as IP multicast, only instead of a real multicast packet stream in the network, usually a proxy device receives the stream, then handles “join” requests, much like a rendezvous point in IPmc. Cisco’s Wide Areas Application Services (WAAS) product line is an example product that has an integrated stream splitting capability for certain types of media streams.

Securing Media Communications

There are a number of threats to media communications that network administrators would want to be aware of in their medianet designs, including:

- **Eavesdropping**—The unauthorized listening/recording of media conversations, presenting the risk of privacy loss, reputation loss, and regulatory non-compliance.
- **Denial of Service**—The loss of media applications or services, presenting the risk of lost productivity and/or business.
- **Compromised video clients**—Hacker control of media clients, such as cameras, displays, and conferencing units, presenting the risk of fraud, data theft, and damaged reputations.

- Compromised system integrity—Hacker control of media application servers or the media control infrastructure, presenting similar risks as compromised clients, but on a significantly greater scale, as well as major productivity and business loss.

When it comes to securing a medianet, there is no silver-bullet technology that protects against all forms of attacks and secures against all types of vulnerabilities. Rather, a layered approach to security, with security being integral to the overall network design, presents the most advantages in terms of protection, operational efficiency, and management.

Visibility and Monitoring Service Levels


It is more important than ever to understand the media applications running on your network, what resources they are consuming, and how they are performing. Whether you are trying to ensure a high-quality experience for video conferencing users or trying to understand how YouTube watchers may be impacting your network, it is important to have visibility into the network.

Tools like Cisco NetFlow can be essential to understanding what portion of traffic flows on the network are critical data applications, VoIP applications, “managed” media applications, and the “unmanaged” media (and other) applications. For example, if you discover that YouTube watchers are consuming 50% of the WAN bandwidth to your branch offices, potentially squeezing out other business critical applications, network administrators may want to put usage policies into place or even more drastic measures such as network-based policing.

Another important aspect is to understand how the media applications deemed business critical are performing? What kind of experience are users receiving? One way to proactively monitor such applications are using network-based tools such as IP Service Level Assurance (IP SLA), which can be programmed to send periodic probes through the network to measure critical performance parameters such as latency, jitter, and loss. It can be helpful to discover trouble spots with long-latency times, for example, and take actions with the service provider (or other root cause) to correct them before users get a bad experience and open trouble reports.

Campus Medianet Architecture

Deploying the medianet in the campus takes place on the standard hierarchical campus design recommendations, following the access, distribution, and core architecture model (see [Figure 1-10](#)). The subsections that follow provide the top design recommendations for the campus switching architecture.

Figure 1-10 Campus Medianet Architecture

Design for Non-Stop Communications in the Campus

As previously discussed, the campus switching network must be designed with high-availability in mind, with the design targets of 0-0.05% packet loss and network convergence within 200 ms.

Designs to consider for the campus include those that include the Cisco Virtual Switching System (VSS), which dramatically simplifies the core and distribution design, implementation, and management. VSS is network system virtualization technology that pools multiple Cisco Catalyst 6500 Series Switches into one virtual switch, increasing operational efficiency by simplifying management to a single virtual device with a single configuration file, boosting nonstop communications by provisioning interchassis stateful failover, and scaling system bandwidth capacity to 1.4 Tbps.

Additionally, Cisco Non-Stop Forwarding (NSF) with Stateful Switchover (SSO) is another feature to consider deploying in the campus switching network to increase network up-time and more gracefully handle failover scenarios if they occur.

Cisco Catalyst switching product lines have industry-leading high-availability features including VSS and NSF/SSO. When deployed with best practices network design recommendations, including routed access designs for the campus switching network, media applications with even the strictest tolerances can be readily supported.

Bandwidth, Burst, and Power

As discussed earlier, provisioning adequate bandwidth is a key objective when supporting many types of media applications, especially interactive real-time media applications such as Cisco TelePresence.

In the access layer of the campus switching network, consider upgrading switch ports to Gigabit Ethernet (GE). This provides sufficient bandwidth for high-definition media capable endpoints. In the distribution and core layers of the campus switching network, consider upgrading links to 10 Gigabit Ethernet (10GE), allowing aggregation points and the core switching backbone to handle the traffic loads as the number of media endpoints and streams increases.

Additionally, ensure that port interfaces have adequate buffering capacity to handle the burstiness of media applications, especially video-oriented media applications. The amount of buffering needed depends on the number and type of media applications traversing the port.

Finally, the campus infrastructure can also supply Power-over-Ethernet to various media endpoints, such as IP video surveillance cameras and other devices.

Application Intelligence and QoS

Having a comprehensive QoS strategy can protect critical media applications including VoIP and video, as well as protect the campus switching network from the effects of worm outbreaks. The Cisco Catalyst switching products offer industry-leading QoS implementations, accelerated with low-latency hardware ASICs, which are critical for ensuring the service level for media applications.

QoS continues to evolve to include more granular queuing, as well as additional packet identification and classification technologies. One advance is the Cisco Programmable Intelligent Services Adapter (PISA), which employs deeper packet inspection techniques mappable to service policies. Intelligent features like PISA will continue to evolve at the network edge to allow application intelligence, enabling the network administrator to prioritize critical applications while at the same time control and police unmanaged or unwanted applications which may consume network resources.


Once traffic has been classified and marked, then queuing policies must be implemented on every node where the possibility of congestion could occur (regardless of how often congestion scenarios actually do occur). This is an absolute requirement to guarantee service levels. In the campus, queuing typically occurs in very brief bursts, usually only lasting a few milliseconds. However, due to the speeds of the links used within the campus, deep buffers are needed to store and re-order traffic during these bursts. Additionally, within the campus, queuing is performed in hardware, and as such, queuing models will vary according to hardware capabilities. Obviously, the greater the number of queues supported, the better, as this presents more policy flexibility and granularity to the network administrator. Four queues would be considered a minimum (one strict-priority queue, one guaranteed bandwidth queue, one default queue, and one deferential queue). Similarly, Catalyst hardware that supports DSCP-to-Queue mappings would be preferred, as these (again) present the most granular QoS options to the administrator.

Consider an example, the Catalyst 6500 WS-X6708-10G, which provides a 1P7Q4T queuing model, where:

- 1P represents a single, strict-priority queue
- 7Q represents 7 non-priority, guaranteed-bandwidth queues
- 4T represents 4 dropping thresholds per queue

Additionally, the WS-X6708-10G supports DSCP-to-Queue mapping, providing additional policy granularity. With such a linecard, voice, video, and data applications could be provisioned as shown in Figure 1-11.

Figure 1-11 *Campus Medianet Queuing Model Example*



Broadcast Optimization with IP Multicast

IP multicast is an important part of many campus switching network designs, optimizing the broadcast of one-to-many streams across the network. Cisco Catalyst switching products provide industry-leading IP multicast proven in business critical network implementations. The IPmc foundation offers further value in networks in optimizing broadcast streaming.

Leveraging Network Virtualization for Restricted Video Applications

The objective of many media applications is to improve effectiveness of communication and collaboration between groups of people. These applications typically have a fairly open usage policy, meaning that they are accessible by and available to a large number of employees in the company.

Other media applications have more restrictive access requirements, and are only available to a relatively small number of well defined users. For example, IP video surveillance is typically available to the Safety and Security department. Access to Digital Signage may only be needed by the few content programmers and the sign endpoints themselves. Additionally, it would generally be prudent to restrict visiting guests from on-demand or streaming content that is confidential to the company.

For these restricted access video scenarios, network virtualization technologies can be deployed to isolate the endpoints, servers, and corresponding media applications within a logical network partition, enhancing the security of the overall solution. Cisco Catalyst switching products offer a range of network virtualization technologies, including Virtual Routing and Forwarding (VRF) Lite and Generic Route Encapsulation (GRE), that are ideal for logical isolation of devices and traffic.

Securing Media in the Campus

As previously discussed, a layered and integrated approach to security provides the greatest degree of protection, while at the same time increases operational and management efficiency. To this end, campus network administrators are encouraged to use the following tactics and tools to secure the Campus medianet:

Basic security tactics and tools:

- Access-lists to restrict unwanted traffic
- Separate voice/video VLANs from data VLANs
- Harden software media endpoints with Host-based Intrusion Protection Systems (HIPS), like Cisco Security Agent (CSA)
- Disable gratuitous ARP
- Enable AAA and roles based access control (RADIUS/TACACS+) for the CLI on all devices
- Enable SYSLOG to a server; collect and archive logs
- When using SNMP, use SNMPv3
- Disable unused services
- Use SSH to access devices instead of Telnet
- Use FTP or SFTP (SSH FTP) to move images and configurations around and avoid TFTP when possible
- Install VTY access-lists to limit which addresses can access management and CLI services
- Apply basic protections offered by implementing RFC 2827 filtering on external edge inbound interfaces

Intermediate security tactics and tools:

- Deploy firewalls with stateful inspection
- Enable control plane protocol authentication where it is available (EIGRP, OSPF, HSRP, VTP, etc.)
- Leverage the Cisco Catalyst Integrated Security Feature (CISF) set, including:
 - Dynamic Port Security
 - DHCP Snooping
 - Dynamic ARP Inspection
 - IP Source Guard

Advanced security tactics and tools:

- Deploy Network Admission Control (NAC) and 802.1x
- Encrypt all media calls with IPSec
- Protect the media control plane with Transport Layer Security (TLS)
- Encrypt configuration files

- Enable Control Plane Policing (CoPP)
- Deploy scavenger class QoS (data plane policing)

WAN and Branch Office Medianet Architecture

Many employees in the typical large company now work in satellite or branch offices away from the main headquarters. These employees expect access to the same set of media applications as your HQ employees. In fact, they may rely on them even more because of the need to communicate effectively and productively with corporate.

Deploying the medianet in the WAN and branch office networks takes place on the standard design recommendations, following the services aggregation edge, service provider, and branch office architecture model (see [Figure 1-12](#) and [Figure 1-13](#)). The subsections that follow provide the top design recommendations for the WAN and branch office architecture.

Figure 1-12 WAN/MAN Medianet Architecture





Figure 1-13 Branch Medianet Architecture

Design for Non-Stop Communications over the WAN

For reasons previously discussed, the WAN and branch office networks must be designed with high availability in mind. The target for packet loss on the WAN and branch networks is the same as for campus networks: 0-0.05%. However, the convergence target of 200 ms for campus networks is most likely unachievable over the WAN and as such, WAN convergence times should be designed to the minimum achievable times.

Because branch offices need to stay consistently and reliably connected to the regional hub or central site, it is highly recommended that each branch office have dual WAN connections, using diverse SP circuits. In the event of an outage on one WAN connection, the secondary WAN provides survivability. Designs for the WAN and branch office should deploy Cisco Performance Routing (PfR), which provides highly-available utilization of the dual WAN connections, as well as fast convergence and rerouting in the event of lost connectivity. At the branch office, consider designs with dual Cisco Integrated Services Routers (ISR) to offer redundancy in the event of an equipment failure.

Additionally, at the services aggregation edge, deploy designs based on highly-available WAN aggregation, including Stateful Switchover (SSO). The Cisco Aggregation Services Router (ASR) product line has industry-leading high-availability features including built-in hardware and processor redundancy, In-Service Software Upgrade (ISSU) and NSF/SSO. When deployed with best practices network design recommendations for the WAN edge, video applications with even the strictest tolerances can be readily supported.

Bandwidth Optimization over the WAN

When not properly planned and provisioned, the WAN may raise the largest challenge to overcome in terms of delivering simultaneous converged network services for media applications. Video-oriented media applications in particular consume significant WAN resources and understanding application requirements and usage patterns at the outset is critical.

Starting with a survey of current WAN speeds can assist in decisions regarding which branch offices need to be upgraded to higher speed and secondary WAN connections. Some quick calculations based on the number of seats in a branch office can provide a quick indicator about bandwidth needs. For example, suppose there are 20 employees in a branch office and the company relies on TelePresence and desktop multimedia conferencing for collaboration, streaming media for training and corporate communications broadcasts, and plans to install IP video surveillance cameras at all branches for security. Let us further assume a 5:1 over-subscription on desktop multimedia conferencing. A quick calculation might look similar to the following:

- VoIP: 5 simultaneous calls over the WAN to HQ @ 128 kbps each
- Video Surveillance: 2 camera feeds @ 512 kbps each
- Cisco TelePresence: 1 call @ 15 Mbps
- Desktop Multimedia Conferencing: 4 simultaneous calls over the WAN to HQ @ 512 kbps each
- Training VoDs: 2 simultaneous viewers @ 384 kbps each
- Data Applications: 1 Mbps x 20 employees

With simple estimates, it is possible to see that this Branch Office may need 45 Mbps or more of combined WAN bandwidth.

One technology which can aid the process is to “harvest” bandwidth using WAN optimization technologies such as Cisco Wide Area Application Services (WAAS). Using compression and optimization, Cisco WAAS can give back 20-50% or more of our current WAN bandwidth, without sacrificing application speed. WAAS or any other WAN optimization technology is unlikely to save bandwidth in video applications themselves, because of the high degree of compression already “built-in” to most video codecs. But rather, the point of implementing WAN optimization is to “clear” bandwidth from other applications to be re-used by newer or expanding media applications, such as video.

The question whether to optimize the WAN or upgrade the WAN bandwidth is often raised. The answer when adding significant video application support is *both*. Optimizing the WAN typically allows the most conservative WAN upgrade path.

Application Intelligence and QoS

Having a comprehensive QoS strategy can protect critical media applications as well as protect the WAN and branch office networks from the effects of worm outbreaks.


Cisco ISR and ASR product families offer industry-leading QoS implementations, accelerated with low-latency hardware ASICs, that are critical for ensuring the service level for video applications. QoS continues to evolve to include more granular queuing, as well as additional packet identification and classification technologies.

Another critical aspect of the overall QoS strategy is the Service Level Assurance (SLA) contracted with the service provider (or providers) for the WAN connectivity. In general, for video applications an SLA needs to specify the lowest practical latency (such as less than 60 milliseconds one-way SP edge-to-edge

latency; however, this value would be greater for intercontinental distances), low jitter (such as less than 5 ms peak-to-peak jitter within the SP network), and lowest practical packet loss (approaching 0-0.05%). SP burst allowances and capabilities are also factors to consider.

When selecting SPs, the ability to map the company's QoS classes to those offered by the SP is also essential. The SP service should be able to preserve Layer 3 DSCP markings and map as many classes as practical across the SP network. An example enterprise edge medianet mapping to a 6-class MPLS VPN SP is illustrated in Figure 1-14.

Figure 1-14 Enterprise to 6-Class MPLS VPN Service Provider Mapping Model Example



Broadcast Optimization for Branch Offices

IP multicast is supported by the Cisco ISR and ASR product families. Certain SP WAN services may or may not support the capability to use IPmc over the WAN. For example, if using an MPLS service, typically the provider must be able to offer a multicast VPN service to allow IPmc to continue to operate over the MPLS WAN topology.

Similarly, certain WAN topologies and integrated security designs also may preclude the use of IPmc. For example, IPsec VPNs cannot transport multicast packets natively. Cisco IPsec VPN WANs combined with Cisco GRE, Cisco Virtual Tunnel Interface (VTI), and Cisco Dynamic Multipoint VPN (DMVPN) do support multicast traffic.


Scalability of WANs with encryption-enabled can suffer from multicast traffic due to the requirements to encrypt the same packet numerous times, once for each branch office connection. The Cisco Group Encrypted Transport VPN (GETVPN) offers a solution, allowing many branch office connections to share the same encryption key. This is an ideal solution for maintaining the secure connectivity that VPNs offer, while not compromising scalability when IP multicast is required to be broadcast over the WAN.

Finally, for situations where multicast of the WAN is not possible, the Cisco WAAS product line also offers a stream “splitting” capability as an alternative to IPmc. The WAAS device in the branch office network acts as a proxy device, allowing multiple users to join the single media stream received over the WAN connection.

Data Center Medianet Architecture

Deploying the medianet in the data center takes place on the standard design recommendations, following the data center architecture model (see [Figure 1-15](#)). The subsections that follow provide the top design recommendations for the data center architecture.

Figure 1-15 Data Center Medianet Architecture



224520

Design for Non-Stop Communications in the Data Center

As with the campus network, the data center network must be designed with high-availability in mind, with the design targets of 0-0.05% packet loss and network convergence within 200 ms.

Designs to consider for the data center include those that include Cisco Non-Stop Forwarding (NSF) with Stateful Switchover (SSO) to increase network up-time and more gracefully handle failover scenarios if they occur.

Cisco Catalyst switching product lines, including the Catalyst 6000 family, and the Cisco Nexus family have industry-leading high-availability features. When deployed with best practices network design recommendations for the data center switching network, video applications with even the strictest tolerances can be readily supported.

High-Speed Media Server Access

As discussed earlier, minimizing latency is a key objective when supporting many types of media applications, especially interactive real-time media applications such as desktop multimedia conferencing and Cisco TelePresence. If conferencing resources are located in the data center, it is important to provide high-speed, low-latency connections to minimize unnecessary additions to the latency budget.

In the aggregation layer of the data center switching network, consider upgrading links to 10 Gigabit Ethernet (10GE), allowing aggregation points and the core switching backbone to handle the traffic loads as the number of media endpoints and streams increases.

In the access layer of the data center switching network, consider upgrading targeted server cluster ports to 10 Gigabit Ethernet (10GE). This provides sufficient speed and low-latency for storage and retrieval needed for streaming intensive applications, including Cisco IP Video Surveillance (IPVS) and Cisco Digital Media System (DMS).

Media Storage Considerations

Several media applications need access to high-speed storage services in the data center, including IP video surveillance, digital signage, and desktop streaming media. It is important to recognize that video as a media consumes significantly more storage than many other types of media. Factor video storage requirements into data center planning. As the number and usage models of video increases, the anticipated impact to storage requirements is significant.


Another consideration is how to manage the increasing volume of video media that contain proprietary, confidential, or corporate intellectual property. Policies and regulatory compliance planning must be in place to manage video content as a company would manage any of its sensitive financial or customer information.

Conclusions

Media applications are increasing exponentially on the IP network. It is best to adopt a comprehensive and proactive strategy to understand how these media applications will affect your network now and in the future. By taking an inventory of video-enabled applications and understanding the new and changing requirements they will place on the network, it is possible to successfully manage through this next evolution of IP convergence, and take steps to enable your network to continue to be the converged platform for your company's communications and collaborations.

By designing the deployment of an end-to-end medianet architecture, it is possible to enable faster adoption of new media applications, while providing IT staff with the tools to proactively manage network resources and ensure the overall user experience (see [Figure 1-16](#)). Enterprises that lack a comprehensive network architecture plan for media applications may find themselves in a difficult situation, as the proportion of media application traffic consumes the majority of network resources.

Figure 1-16 *Bringing it All Together*



Cisco is uniquely positioned to provide medianets, offering a comprehensive set of products for the network infrastructure designed with built-in media support, as well as being a provider of industry leading media applications, including Cisco TelePresence, Cisco WebEx, and Cisco Unified Communications. Through this unique portfolio of business media solutions and network platforms, Cisco leads the industry in the next wave of IP convergence and will lead the media revolution as companies move to the next wave of productivity and collaboration.

Terms and Acronyms

Acronyms	Definition
10GE	10 Gigabit Ethernet
AVVID	Architecture for Voice, Video, and Integrated Data
Codec	Coder/Decoder
DC	Data Center
DMS	Digital Media System
DMVPN	Dynamic Multipoint VPN
DPI	Deep Packet Inspection
GE	Gigabit Ethernet
GETVPN	Group Encrypted Transport VPN
GRE	Generic Route Encapsulation
H.264	Video Compression standard, also known as MPEG4
HA	High Availability
HD	High Definition video resolution

HDTV	High-Definition Television
IPmc	IP Multicast
IP SLA	IP Service Level Assurance
IPTV	IP Television
IPVS	IP Video Surveillance
LD	Low Definition video resolution
MPEG4	Moving Pictures Expert Group 4 standard
NSF	Non-Stop Forwarding
NV	Network Virtualization
PfR	Performance Routing
PISA	Programmable Intelligent Services Adapter
QoS	Quality of Service
SLA	Service Level Agreement
SP	Service Provider
SSO	Stateful-Switchover
SVC	Scalable Video Coding
UC	Unified Communications
VoD	Video On Demand
VoIP	Voice over IP
VPN	Virtual Private Network
VRF	Virtual Routing and Forwarding
VRN	Video Ready Network
VSS	Virtual Switching System
WAN	Wide Area Network
WLAN	Wireless LAN
WAAS	Wide Area Application Services

Related Documents

White Papers

- [The Exabyte Era](http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/net_implementation_white_paper0900aecd806a81a7.pdf)
http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/net_implementation_white_paper0900aecd806a81a7.pdf
- [Global IP Traffic Forecast and Methodology, 2006-2011](http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/net_implementation_white_paper0900aecd806a81aa.pdf)
http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/net_implementation_white_paper0900aecd806a81aa.pdf
- [Video: Improving Collaboration in the Enterprise Campus](#)

http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns431/solution_overview_c22_468222.pdf

System Reference Network Designs

- Enterprise 3.0 Campus Architecture Overview and Framework
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/campover.html>
- WAN Transport Diversity Design Guide
http://www.cisco.com/application/pdf/en/us/guest/netso/ns483/c649/ccmigration_09186a008094
- Branch Office Architecture Overview
http://www.cisco.com/application/pdf/en/us/guest/netso/ns171/c649/ccmigration_09186a00807593b7.pdf
- Data Center Infrastructure Design Guide
http://www.cisco.com/application/pdf/en/us/guest/netso/ns107/c649/ccmigration_09186a008073377d.pdf
- End-to-End Quality of Service (QoS) Design Guide
http://www.cisco.com/application/pdf/en/us/guest/netso/ns432/c649/ccmigration_09186a008049b062.pdf
- Telepresence Network System Design Guide
http://www.cisco.com/en/US/docs/solutions/TelePresence_Network_Systems_1.1_DG.pdf
- IP Video Surveillance Stream Manager Design Guide
http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns656/net_design_guidance0900aecd805ee51d.pdf
- Branch Wide Area Application Services (WAAS) Design Guide
http://www.cisco.com/application/pdf/en/us/guest/netso/ns477/c649/ccmigration_09186a008081c7d5.pdf
- Network Virtualization Path Isolation Design Guide
http://www.cisco.com/application/pdf/en/us/guest/netso/ns171/c649/ccmigration_09186a0080851cc6.pdf

Websites

- Campus Solutions
http://www.cisco.com/en/US/netso/ns340/ns394/ns431/networking_solutions_packages_list.html
- WAN and Aggregation Services Solutions
http://www.cisco.com/en/US/netso/ns483/networking_solutions_packages_list.html
- Branch Office Solutions
http://www.cisco.com/en/US/netso/ns477/networking_solutions_packages_list.html
- Data Center 3.0 Solutions
http://www.cisco.com/en/US/netso/ns708/networking_solutions_solution_segment_home.html
- Video Solutions

http://www.cisco.com/en/US/netsol/ns340/ns394/ns158/networking_solutions_packages_list.html

- Telepresence Solutions

http://www.cisco.com/en/US/netsol/ns669/networking_solutions_solution_segment_home.html

- Unified Communications Solutions

http://www.cisco.com/en/US/netsol/ns340/ns394/ns165/ns152/networking_solutions_package.html

- Wide Area Application Services Solutions

http://www.cisco.com/en/US/products/ps5680/Products_Sub_Category_Home.html



CHAPTER 2

Medianet Bandwidth and Scalability

This chapter discusses the bandwidth requirements for different types of video on the network, as well as scalability techniques that allow additional capacity to be added to the network.

Bandwidth Requirements

Video is the ultimate communications tool. People naturally use visual clues to help interpret the spoken word. Facial expression, hand gestures, and other clues form a large portion of the messaging that is normal conversation. This information is lost on traditional voice-only networks. If enough of this visual information can be effectively transported across the network, potential productivity gains can be realized. However, if the video is restricted by bandwidth constraints, much of the visual information is lost. In the case of video conferencing, the user community does not significantly reduce travel. In the case of video surveillance, small distinguishing features may be lost. Digital media systems do not produce engaging content that draws in viewers. In each case, the objectives that motivated the video deployment cannot be met if the video is restricted by bandwidth limitations.

Quantifying the amount of bandwidth that a video stream consumes is a bit more difficult than other applications. Specifying an attribute in terms of bits per second is not sufficient. The per-second requirements result from other more stringent requirements. To fully understand the bandwidth requirements, the packet distribution must be fully understood. This is covered in [Chapter 1, “Medianet Architecture Overview,”](#) and briefly revisited here.

The following video attributes affect how much bandwidth is consumed:

- **Resolution**—The number of rows and columns in a given frame of video in terms of pixel count. Often resolution is specified as the number of rows. Row counts of 720 or greater are generally accepted as high definition (HD) video. The number of columns can be derived from the number of rows by using the aspect ratio of the video. Most often HD uses an aspect ratio of 16:9, meaning 16 columns for every 9 rows. As an example, a resolution of 720 and an aspect ratio of 16:9 gives a screen dimension of 1280 x 720 pixels. The same 720 resolution at a common 4:3 aspect ratio gives a screen dimension of 960 x 720 pixels. Resolution has a significant effect on bandwidth requirements as well as the productivity gains of video on the network. Resolution is a second-degree term when considering network load. If the aspect ratio is held at 16:9 and the resolution is increased from 720 to 1080, the number of pixels per frame jumps from 921,600 to 2,073,600, which is significant. If the change is in terms of percent, a 50 percent increase in resolution results in a 125 percent increase in pixel count. Resolution is also a key factor influencing the microburst characteristics of video. A microburst results when an encoded frame of video is sliced into packets and placed on the outbound queue of the encoder network interface card (NIC). This is discussed in more detail later in this chapter.

- **Encoding implementation**—Encoding is the process of taking the visual image and representing it in terms of bytes. Encoders can be distinguished by how well they compress the information. Two factors are at work. One is the algorithm that is used. Popular encoding algorithms are H.264 and MPEG-4. Other older encoders may use H.263 or MPEG-2. The second factor is how well these algorithms are implemented. Multiple hardware digital signal processors (DSPs) are generally able to encode the same video in less bytes than a battery operated camera using a low power CPU. For example, a flip camera uses approximately 8 Mbps to encode H.264 at a 720 resolution. A Cisco TelePresence CTS-1000 can encode the same resolution at 2 Mbps. The algorithm provides the encoder flexibility to determine how much effort is used to optimize the compression. This in turn gives vendors some latitude when trying to meet other considerations, such as cost and power.
- **Quality**—Encoding video uses a poor compression. This means that some amount of negligible visual information can be discarded without having a detrimental impact of the viewer experience. Examples are small variations in color at the outer edges of the visual spectrum of red and violet. As more detail is omitted from the encoded data, small defects in the rendered video begin to become apparent. The first noticeable impact is color banding. This is when small color differences are noticed in an area of common color. This is often most pronounced at the edge of the visible spectrum, such as a blue sky.
- **Frame rate**—This is the number of frames per second (fps) used to capture the motion. The higher the frame rate, the smoother and more life-like is the resulting video. At frame rates less than 5 fps, the motion becomes noticeably jittery. Typically, 30 fps is used, although motion pictures are shot at 24 fps. Video sent at more than 30 fps offers no substantial gain in realism. Frame rates have a linear impact on bandwidth. A video stream of 15 fps generates approximately half as much network traffic as a stream of 30 fps.
- **Picture complexity**—Encoders must take a picture and encode it in as few bytes as possible without noticeably impacting the quality. As the image becomes more complex, it takes more bytes to describe the scene. Video of a blank wall does not consume as much bandwidth as a scene with a complex image, such as a large room of people. The impact on bandwidth is not substantial but does have some influence.
- **Motion**—Just like picture complexity, the amount of motion in a video has some influence over how much bandwidth is required. The exception is Motion JPEG (M-JPEG). The reason is that all other encoding techniques involve temporal compression, which capitalizes on savings that can be made by sending only the changes from one frame to the next. As a result, video with little motion compresses better than video with a great deal of motion. Usually, this means that video shot outside, where a breeze may be moving grass or leaves, often requires more network bandwidth than video shot indoors. Temporal compression is discussed in more detail in [Chapter 1, “Medianet Architecture Overview.”](#)

It is possible to have some influence on the bandwidth requirements by changing the attributes of the video stream. A 320x240 video at 5 fps shot in a dark closet requires less bandwidth than a 1080x1920 video at 30 fps shot outside on a sunny, breezy day. The attributes that have the most influence on network bandwidth are often fully configurable. These are resolution, frame rate, and quality settings. The remaining attributes are not directly controlled by the administrator.

Measuring Bandwidth

Network bandwidth is often measured in terms of bits per second. This is adequate for capacity planning. If a video stream is expected to run at 4 megabits per second (Mbps), a 45 Mbps circuit can theoretically carry 11 of these video streams. The number is actually less, because of the sub-second bandwidth requirements. This is referred to as the *microburst* requirements, which are always greater than the one second smoothed average.

Consider the packet distribution of video. First remember that frames are periodic around the frame rate. At 30 fps, there is a frame every 33 msec. The size of this frame can vary. Video is composed of two basic frame types, I-frames and P-frames. I-frames are also referred to as full reference frames. They are larger than P-frame but occur much less frequently. It is not uncommon to see 128 P-frames for every 1 I-frame. Some teleconference solutions send out even fewer I-frames. When they are sent, they look like a small burst on the network when compared to the adjacent P-frames. It may take as many as 80 packets or more to carry an I-frame of high definition 1080 video. These 80+ packets show up on the outbound NIC of the encoder in one chunk. The NIC begins to serialize the packet onto the Ethernet wire. During this time, the network media is being essentially used at 100 percent; the traffic bursts to line rate for the duration necessary to serialize an I-frame. If the interface is a Gigabit interface, the duration of this burst is one-tenth as long as the same burst on a 100 Mbs interface. A microburst entails the concept that the NIC is 100 percent used during the time it takes to serialize all the packets that compose the entire frame. The more packets, the longer duration required to serialize them.

It is best to conceive of a microburst as either the serialization delay of the I-frame or the total size of the frame. It is not very useful to characterize an I-frame in terms of a rate such as Kbps, although this is fairly common. On closer examination, all bursts, and all packets, are sent at line rate. Interfaces operate only at a single speed. The practice of averaging all bits sent over a one-second interval is somewhat arbitrary. At issue is the network ability to buffer packets, because multiple inbound streams are in contention for the same outbound interface. A one-second measurement interval is too long to describe bandwidth requirements because very few devices can buffer one second worth of line rate data. A better interval is 33 msec, because this is the common frame rate.

There are two ways to consider this time interval. First, the serialization delay of any frame should be less than 33 msec. Second, any interface in the network should be able to buffer the difference in serialization delay between the ingress and egress interface over a 33-msec window. During congestion, the effective outbound serialization delay for a given stream may fall to zero. In this case, the interface may have to queue the entire frame. If queue delays of above 33 msec are being experienced, the video packets are likely to arrive late. Network shapers and policers are typical points of concern when talking about transporting I-frames. These are discussed in more detail in [Chapter 4, “Medianet QoS Design Considerations,”](#) and highlighted later in this chapter.

Video Transports


Several classifications can be used to describe video, from real-time interactive streaming video on the high end to prerecorded video on the low end. Real-time video is being viewed and responded to as it is occurring. This type of stream has the highest network requirements. Remote medicine is an example of an application that uses this type of video. TelePresence is a more common application. In all cases, packets that are dropped by the network cannot be re-sent because of the time-sensitive nature. Real-time decoders are built with the smallest de-jitter buffers possible. On the other extreme is rebroadcasting previously recorded video. This is usually done over TCP and results in a network load similar to large FTP file transfers. Dropped packets are easily retransmitted. [Chapter 4, “Medianet QoS Design Considerations”](#) expands on this concept and discusses the various types of video and the service levels required of the network.

Packet Flow Malleability

Video packets are constrained by the frame rate. Each frame consists of multiple packets, which should arrive within the same frame window. There are I-frames and P-frames. The network is not aware of what type of frame has been sent, or that a group of packets are traveling together as a frame. The network considers each packet only as a member of a flow, without regard to packet distribution. When tools such


as policers and shapers are deployed, some care is required to accommodate the grouping of packets into frames, and the frame rate. The primary concern is the I-frame, because it can be many times larger than a P-frame, because of the way video encoders typically place I-frames onto the network. (See [Figure 2-1](#).)

Figure 2-1 P-frames and I-frames



When an I-frame is generated, the entire frame is handed to the network abstraction layer (NAL). This layer breaks the frame into packets and sends them on to the IP stack for headers. The processor on the encoder can slice the frame into packets much faster than the Ethernet interface can serialize packets onto the wire. As a result, video frames generate a large number of packets that are transmitted back-to-back with only the minimum interpacket gap (IPG). (See [Figure 2-2](#).)


Figure 2-2 I-frame Serialization



The service provider transport and video bandwidth requirements set the limit to which video streams can be shaped and recast. Natural network video is sent at a variable bit rate. However, many transports have little tolerance for traffic flows that exceed a predetermined contract committed information rate (CIR). Although [Chapter 4, “Medianet QoS Design Considerations”](#) discusses this in further details, some overview of shapers and policers is warranted as part of the discussion of bandwidth requirements.

Any interface can transmit only at line rate. Interfaces use a line encoding scheme to ensure that both the receiver and transmitter are bit synchronized. When a user states that an interface is running at x Mbps, that is an average rate over 1 second of time. The interface was actually running at 100 percent utilization while those packets were being transmitted, and idle at all the other times. [Figure 2-3](#) illustrates this concept:

Figure 2-3 Interface Load/Actual Load



228638

Microbursts

In video, frames are sent as a group of packets. These packets are packed tightly because they are generated at the same time. The larger the frame, the longer the duration of the microburst that results when the frame is serialized. It is not uncommon to find microbursts measured in terms of bits per second. Typically the rate is normalized over a frame. For example, if an I-frame is 80 Kb, and must be sent within a 33 msec window, it is tempting to say the interface is running at 4 Mbps but bursting to $(80 \times 1000 \times 8) / 0.033 = 19.3$ Mbps. In actuality, the interface is running at line rate long enough to serialize the entire frame. The interface speed and buffers are important in determining whether there will be drops. The normalized 33 msec rate gives some useful information when setting shapers. If the line rate in the example above is 100 Mbps, you know that the interface was idle for 80.7 percent of the time during the 33 msec window. Shapers can help distribute idle time. However, this does not tell us whether the packets were evenly distributed over the 33 msec window, or whether they arrived in sequence during the first 6.2 msec. The encoders used by TelePresence do some level of self shaping so that packets are better distributed over a 33 msec window, while the encoders used by the IP video surveillance cameras do not.


Shapers

Shapers are the most common tool used to try to mitigate the effect of bursty traffic. Their operation should be well understood so that other problems are not introduced. Shapers work by introducing delay. Ideally, the idle time is distributed between each packet. Only hardware-based shapers such as those found in the Cisco Catalyst 3750 Metro device can do this. Cisco IOS shapers use a software algorithm to enforce packets to delay. Cisco IOS-based shapers follow the formula $Bc = CIR * Tc$. The target bandwidth (*CIR*) is divided into fixed time slices (*Tc*). Each *Tc* can send only *Bc* bytes worth of data. Additional traffic must wait for the next available time slice. This algorithm is generally effective, but keep in mind some details. First, IOS shapers can meter only traffic with time slices of at least 4 msec. This means that idle time cannot be evenly distributed between all packets. Within a time slice, the interface still sends packets at line rate. If the queue of packets waiting is deeper than *Bc* bytes, all the packets are sent in sequence at the start of each *Tc*, followed by an idle period. In effect, if the offered rate exceeds the *CIR* rate for an extended period, the shaper introduces microbursts that are limited to *Bc* in size. Each time slice is independent of the previous time slice. A burst of packets may arrive at the shaper and completely fill a *Bc* at the very last moment, followed immediately by a new time slice with another *Bc* worth of available bandwidth. This means that although the interface routinely runs at line rate for each *Bc* worth of data, it is possible that it will run at line rate for $2*Bc$ worth of bytes. When a shaper first becomes active, the traffic alignment in the previous *Tc* is not considered.

Partial packets are another feature of shapers to consider. Partial packets occur when a packet arrives whose length exceeds the remaining *Bc* bits available in the current time slice. There are two possible approaches to handle this. First, delay the packet until there are enough bits available in the bucket. The down side of this approach is twofold. First, the interface is not able to achieve *CIR* rate because time slices are expiring with bits still left in the *Bc* bucket. Secondly, while there may not be enough *Bc* bits for a large packet, there could be enough bits for a much smaller packet in queue behind the large packet. There are problems with trying to search the queue looking for the best use of the remaining *Bc* bits. Instead, the router allows the packet to transmit by borrowing some bits from the next time slice.

Figure 2-4 shows the impact of using shapers.

Figure 2-4 Shaper Impact



228639

Choosing the correct values for T_c , B_c , and CIR requires some knowledge of the traffic patterns. The CIR must be above the sustained rate of the traffic load; otherwise, traffic continues to be delayed until shaper drops occur. In addition, the shaper should delay as few packets as possible. Finally, if the desire is to meet a service level enforced by a policer, the shaper should not send bursts (B_c) larger than the policer allows. The attributes of the upstream policer are often unknown, yet these values are a dominant consideration when configuring the shaper. It might be tempting to set the shaper B_c to its smallest possible value. However, as T_c falls below $2 * 33$ msec, the probability of delaying packets increases, as does the jitter. Jitter is at its worst when only one or two packets are delayed by a large T_c . As T_c approaches 0, jitter is reduced and delay is increased. In the limit as T_c approaches 0, the introduced delay equals the serialization delay if the circuit can be clocked at a rate equal to CIR. With TelePresence, the shaper T_c should be 20 msec or less to get the best balance between delay and jitter. If the service provider cannot accept bursts, the shaper can be set as low as 4 msec.

With shapers, if packets continue to arrive at a rate that exceeds the CIR of the shaper, the queue depth continues to grow and eventually saturates. At this point, the shaper begins to discard packets. Normally, a theoretically ideal shaper has infinite queue memory and does not discard packets. In practice, it is actually desirable to have shapers begin to look like policers if the rate exceeds CIR for a continued duration. The result of drops is that the sender throttles back its transmission rate. In the case of TCP flows, window sizes are reduced. In the case of UDP, lost transmissions cause upper layers such as TFTP, LDAP, or DNS to pause for the duration of a response timeout. UDP flows in which the session layer has no feedback mechanism can overdrive a shaper. Denial-of-service (DoS) attacks are in this class. Some Real-Time Protocol (RTP)/UDP video may also fall in this class where Real-Time Control Protocol (RTCP) is not used. Real-Time Streaming Protocol (RTSP)-managed RTP flows are an example of this type of video. In these cases, it is very important to ensure that the shaper CIR is adequately configured. When a shaper queue saturates, all non-priority queuing (PQ) traffic can be negatively impacted.

Shapers versus Policers

Policers and shapers are related methods that are implemented somewhat differently. Typically, a shaper is configured on customer equipment to ensure that traffic is not sent out of contract. The service provider uses a policer to enforce a contracted rate. The net effect is that shapers are often used to prevent upstream policers from dropping packets. Typically, the policer is set in place without regard to customer shapers. If the customer knows what the parameters of the policer are, this knowledge can be used to correctly configure a shaper.

Understanding the difference between policers and shapers helps in understanding the difference in implementation. First, a policer does not queue any packets. Any packets that do not conform are dropped. The shaper is the opposite. No packets are dropped until all queue memory is starved. Policers do not require the router to perform an action; instead, the router only reacts. Shaping is an active process. Queues must be managed. Events are triggered based on the fixed T_c timer. The algorithm for shaping is to maintain a token bucket. Each T_c seconds, B_c tokens are added to the bucket. When a packet arrives, the bucket is checked for available tokens. If there are enough tokens, the packet is allowed onto the TxRing and the token bucket is debited by the size of the packet. If the bucket does not have enough tokens, the packet must wait in queue. At each T_c interval, B_c tokens are credited to the bucket. If there are packets waiting in queue, these packets can be processed until either the queue is empty or the bucket is again depleted of tokens.

By contrast, policing is a passive process. There is no time constant and no queue to manage. A simple decision is made to pass or drop a packet. With policing, the token bucket initially starts full with B_c tokens. When a packet arrives, the time interval since the last packet is calculated. The time elapsed is multiplied by the CIR to determine how many tokens should be added to the bucket. After these tokens have been credited, the size of the packet is compared with the token balance in the bucket. If there are available tokens, the packet is placed on the TxRing and the size of the packet is subtracted from the token bucket. If the bucket does not have enough available tokens, the packet is dropped. As the policed rate approaches the interface line rate, the size of the bucket become less important. When $CIR = Line\ Rate$, the bucket refills at the same rate that it drains.

Because tokens are added based on packet arrival times, and not as periodic events as is done with shapers, there is no time constant (T_c) when discussing policers. The closest equivalent is the time required for an empty bucket to completely refill if no additional packets arrive. In an ideal case, a shaper sends B_c bytes at line rate, which completely drains the policer B_c bucket. The enforced idle time of the shaper for the remaining T_c time then allows the B_c bucket of the policer to completely refill. The enforced idle time of the shaper is $T_c * (1 - CIR / Line_Rate)$. In practice, it is best to set the shaper so that the policer B_c bucket does not go below half full. This is done by ensuring that when the shaped CIR equals the policed CIR, the shaper B_c should be half of the policer B_c .


It is not always possible to set the shaper B_c bucket to be smaller than the policer B_c bucket, because shapers implemented in software have a minimum configurable T_c value of 4 msec. The shaper T_c is not directly configured; instead, B_c and CIR are configured and T_c is derived from the equation $T_c = B_c / CIR$. This means that the shaper B_c cannot be set to less than $0.004 * CIR$. If the policer does not allow bursts of this size, some adjustments must be made. Possible workarounds are as follows:

- Place a hardware-based shaper inline (see [Figure 2-5](#)).

Examples of devices that support hardware based shaping are the Cisco Catalyst 3750 Metro Series Switches. However, the Cisco Catalyst 3750 Metro supports hardware shaping only on 1 Gigabit uplink interfaces. These interfaces do not support any speed other than 1 Gigabit. This can be a problem if the service provider is not using a 1 Gigabit interface to hand off the service. In this case, if the Cisco Catalyst 3750 Metro is to be used, the hardware shaping must occur before the customer edge (CE) router. The Cisco Catalyst 3750 Metro would attach to a router instead of directly with the service provider. The router would handle any Border Gateway Protocol (BGP) peering, security, encryption, and so on. The Cisco Catalyst 3750 Metro would provide wiring closet access and the

shaping. This works only if the router is being fed by a single Metro device. Of course, if more than 48 ports are needed, additional switches must be fed through the Cisco Catalyst 3750 Metro such that the hardware shaper is metering all traffic being fed into the CE router.

Figure 2-5 Hardware-Based Shaper Inline



- Contract a higher CIR from the service provider.

As the contracted CIR approaches the line rate of the handoff circuit, the policer bucket refill rate begins to approach the drain rate. The shaper does not need to inject as much idle time. When the contracted CIR equals the line rate of the handoff circuit, shaping is no longer needed because the traffic never bursts above CIR. Testing in the lab resulted in chart shown in [Figure 2-6](#), which can be used to determine the contracted service provider CIR necessary when shaping is required but the shapers Bc cannot be set below the maximum burst allowed by the service provider. This is often a concern when the service provider is offering a constant bit rate (CBR) service. Video is generally thought of as a variable bit rate, real-time (VBR-RT) service.

Figure 2-6 Higher CIR




Figure 2-6 shows validation results and gives some guidance about the relationship between policers and shapers. The plots are the result of lab validation where a shaper was fixed with a CIR of 15.3 Mbps and a Bc of 7650 bytes. The plot shows the resulting policer drops as the policer CIR values are changed. The Y-Axis shows the drops that were reported by the service provider policer after two minutes of traffic. The X-Axis shows the configured CIR on the policer. This would be the equivalent bandwidth purchased from the provider. Six plots are displayed, each at a unique Policer Bc. This represents how tolerant the service provider is of bursts above CIR. The objective is to minimize the drops to zero at the smallest policer CIR possible. The plot that represents a Bc of 7650 bytes is of particular interest because this is the case where the policer Bc equals the shaper Bc.

The results show that the policed CIR should be greater than twice that of the shaped CIR. Also note that at a policed Bc of 12 KB. This represents the smallest policer Bc that allows the policed CIR to equal to the shaped CIR. As a best practice, it is recommended that the policer Bc be at least twice large as the shaper Bc if the CIR is set to the same value. As this chart shows, if this best practice cannot be met, additional CIR must be purchased from the service provider.

Key points are as follows:

- Shapers do not change the speed at which packets are sent, but rather introduce idle times.
- Policers allow traffic at line rate until the Bc bucket is empty. Policers do not enforce a rate, but rather a maximum burst beyond a rate.
- Shapers that feed upstream policers should use a Bc that is half of the policer Bc.

In the case of TelePresence, the validation results plotted above can be used to derive the following recommendations:

- The shaper Tc should be 20 msec or less. At 20 msec, the number of delayed P-frames is minimized.
- The cloud should be able to handle a burst of at least two times the shaper Bc value. At 20 msec Tc and 15.3 MB CIR, this would be buffer space or an equivalent policer Bc of at least 76.5 KB.
- If the burst capabilities of the cloud are reduced, the shaper Tc must be reduced to maintain the 2:1 relationship (policer Bc twice that of the shaper Bc).
- The minimum shaper Tc is 4 msec on most platforms. If the resulting Bc is too large, additional bandwidth can be purchased from the service provider using the information in Table 2-1.



Note Table 2-1 applies to the Cisco TelePresence System 3000.

Table 2-1 CIR Guidelines

Policed Bc or interface buffer (Kbyte)		CIR (Mbit/sec)
Less than	But more than	
15	12	20
12	11	25
11	10	30
10	8.8	40
8.8	7.65	50
7.65	6.50	75


Table 2-1 CIR Guidelines (continued)

Policed Bc or interface buffer (Kbyte)		CIR (Mbit/sec)
Less than	But more than	
15	12	20
12	11	25
11	10	30
10	8.8	40
6.50	3.0	100
3.0	0.0	N/A

Because shapers can send Bc bytes at the beginning of each Tc time interval, and because shapers feed indirectly into the TxRing of the interface, it is possible to tune the TxRing to accommodate this traffic.

TxRing

The TxRing and RxRings are memory structures shared by the main processor and the interface processor (see [Figure 2-7](#)). This memory is arranged as a first in, first out (FIFO) queue. The ring can be thought of as a list of memory pointers. For each ring, there is a read pointer and a write pointer. The main processor and interface process each manage the pair of pointers appropriate to their function. The pointers move independently of one another. The difference between the write and read pointers gives the depth of the queue. Each pointer links a particle of memory. Particles are an efficient means of buffering packets of all different sizes within a pool of memory. A packet can be spread over multiple particles depending on the size of the packet. The pointers of a single packet form a linked list.

Figure 2-7 TxRings and RxRings

The rest of the discussion on Cisco IOS architecture is out of scope for this section, but some key points should be mentioned. Because a shaper can deposit Bc bytes of traffic onto an interface at the beginning of each Tc time period, the TxRing should be at least large enough to handle this traffic. The exact number of particles required depends on the average size of the packets to be sent, and the average number of particles that a packet may link across. It may not be possible to know these values in all cases.

But some worst case assumptions can be made. For example, video flows typically use larger packets of approximately 1100 bytes (average). Particles are 256 bytes. An approximate calculation for a shaper configured with a CIR of 15 Mb and a Tc of 20 msec would yield a Bc of 37.5 Kb. If that much traffic is placed on the TxRing at once, it requires 146 particles.

However, there are several reasons the TxRing should not be this large. First, a properly configured shaper is not active most of the time. QoS cannot re-sequence packets already on the TxRing. A smaller TxRing size is needed to allow QoS to properly prioritize traffic. Second, the downside of a TxRing that is too small is not compelling. In the case where the shaper is active and a Bc worth of data is being moved to the output interface, packets that do not fit onto the ring wait on the interface queue. Third, in a converged network with voice and video, the TxRing should be kept as small as possible. A ring size of 10 is adequate in a converged environment if a slow interface such as a DS-3 is involved. This provides the needed back-pressure for the interface queueing. In most other cases, the default setting provides the best balance between the competing objects.

The default interface hold queue may not be adequate for video. There are several factors such as the speed of the link, other types of traffic that may be using the link as well as the QoS service policy. In most cases, the default value is adequate, but it can be adjusted if output drops are being reported.

Converged Video

Mixing video with other traffic, including other video, is possible. [Chapter 4, “Medianet QoS Design Considerations”](#) discusses techniques to mark and service various types of video.

In general terms, video classification follows the information listed in [Table 2-2](#).

Table 2-2 Video Classification

Application Class	Per-Hop Behavior	Media Application Example
Broadcast Video	CS5	IP video surveillance/enterprise TV
Real-Time Interactive	CS4	TelePresence
Multi-media Conferencing	AF4	Unified Personal Communicator
Multi-media Streaming	AF3	Digital media systems (VoDs)
HTTP Embedded Video	DF	Internal video sharing
Scavenger	CS1	YouTube, iTunes, Xbox Live, and so on

Queueing is not video frame-aware. Each packet is treated based solely on its markings, and the capacity of the associated queue. This means that it is possible, if not likely, that video frames can be interleaved with other types of packets. Consider a P-frame that is 20 packets deep. The encoder places those twenty packets in sequence, with the inter-packet gaps very close to the minimum 9.6 usec allowed. As these twenty packets move over congested interfaces, packets from other queues may be interleaved on the interface. This is the normal QoS function. If the video flow does not cross any interface where there is congestion, queueing is not active and the packet packing is not be disturbed.

Congestion is not determined by the one-second average of the interface. Congestion occurs any time an interface has to hold packets in queue because the TxRing is full. Interfaces with excessively long TxRings are technically less congested than the same interface with the same traffic flows, but with a smaller TxRing. As mentioned above, congestion is desirable when the objective is to place priority traffic in front of non-priority traffic. When a video frame is handled in a class-based queue structure, the result at the receiving codec is additional gaps in packet spacing. The more often this occurs, the greater the fanout of the video frame. The result is referred to as application jitter. This is slightly

different than packet jitter. Consider again the P-frame of video. At 30 fps, the start of each frame is aligned on 33 msec boundaries; this means the initial packet of each frame also aligns with this timing. If all the interfaces along the path are empty, this first packet arrives at the decoding station spaced exactly 33 msec apart. The delay along the path is not important, but as the additional packets of the frame transit the interface, some TxRings may begin to fill. When this happens, the probability that a non-video packet (or video from a different flow) will be interleaved increases. The result is that even though each frame initially had zero jitter, the application cannot decode the frame until the last packet arrives.

Measuring application jitter is somewhat arbitrary because not all frames are the same size. The decoder may process a small frame fairly quickly, but then have to decode a large frame. The end result is the same, the frame decode completion time is not a consistent 33 msec. Decoders employ playout buffers to address this situation. If the decoder knows the stream is not real-time, the only limit is the tolerance of the user to the initial buffering delay. Because of this, video that is non-real-time can easily be run on a converged network. The Internet is a perfect example. Because the stream is non-real-time, the video is sent as a bulk transfer. Within HTML, this usually a progressive load. The data transfer may complete long before the video has played out. What this means is that a video that was encoded at 4 Mbps flows over the network as fast as TCP allows, and can easily exceed the encoded rate. Many players make an initial measurement of TCP throughput and then buffer enough of the video such that the transfer completes just as the playout completes. If the video is real-time, the playout buffers must be as small as possible. In the case of TelePresence, a dynamic playout buffer is implemented. The duration of any playout has a direct impact on the time delay of the video. Running real-time flows on a converged network takes planning to ensure that delay and jitter are not excessive. Individual video applications each have unique target thresholds.

As an example, assume a network with both real-time and non-real-time video running concurrently with data traffic. Real-time video is sensitive to application jitter. This type of jitter can occur any time there is congestion along that path. Congestion is defined as a *TxRing* that is full. RxRings can also saturate, but the result is more likely a drop. Traffic shapers can cause both packet jitter and application jitter. Jitter can be reduced by placing real-time video in the PQ. TxRings should be fairly small to increase the effectiveness of the PQ. The PQ should be provisioned with an adequate amount of bandwidth, as shown by [Table 2-1](#). This is discussed in more depth in [Chapter 4, “Medianet QoS Design Considerations.”](#)

**Note**

TxRings and RxRings are memory structures found primarily in IOS-based routers.

Bandwidth Over Subscription

Traditionally, the network interfaces were oversubscribed in the voice network. The assumption is that not everyone will be on the phone at the same time. The ratio of oversubscription was often determined by the type of business and the expected call volumes as a percent of total handset. Oversubscription was possible because of Call Admission Control (CAC), which was an approach to legacy time-division multiplexing (TDM) call blocking. This ensured that new connections were blocked to preserve the quality of the existing connections. Without this feature, all users are negatively impacted when call volumes approached capacity.


With medianet, there is not a comparable feature for video. As additional video is loaded onto a circuit, all user video experience begins to suffer. The best method is to ensure that aggregation of all real-time video does not exceed capacity, through provisioning. This is not always a matter of dividing the total bandwidth by the per-flow usage because frames are carried in grouped packets. For example, assume that two I-frames from two different flows arrive on the priority queue at the same time. The router places all the packets onto the outbound interface queue, where they drain off onto the TxRing for serialization

on the wire. The next device upstream sees an incoming microburst twice as large as normal. If the RxRing saturates, it is possible to begin dropping packets at very modest 1 second average loads. As more video is added, the probability that multiple frames will converge increases. This can also load Tx Queues, especially if the multiple high speed source interfaces are bottlenecking into a single low-speed WAN link.

Another concern is when service provider policers cannot accept large, or back-to-back bursting. Video traffic that may naturally synchronize frame transmission is of particular concern and is likely to experience drops well below 90 percent circuit utilization. Multipoint TelePresence is a good example of this type of traffic. The Cisco TelePresence Multipoint Switch replicates the video stream to each participant by swapping IP headers. Multicast interfaces with a large fanout are another example. These types of interfaces are often virtual WAN links such as Dynamic Multipoint Virtual Private Network (DMVPN), or virtual interfaces such as Frame Relay. In both cases, multipoint flows fanout at the bandwidth bottleneck. The same large packet is replicated many times and packed on the wire close to the previous packet.

Buffer and queue depths of the Tx interface can be overrun. Knowing the queue buffer depth and maximum expected serialization delay is a good method to determine how much video an interface can handle before drops. When multiple video streams are on a single path, consider the probability that one frame will overlap or closely align with another frame. Some switches allow the user some granularity when allocated shared buffer space. In this case, it is wise to ensure buffers that can be expected to process long groups of real-time packets and have an adequate pool of memory. This can mean reallocating memory away from queues where packets are very periodic and groups of packets are generally small.

For now, some general guidelines are presented as the result of lab verification of multipoint TelePresence. [Figure 2-8](#) shows both the defaults and tuned buffer allocation on a Cisco Catalyst 3750G Switch. Additional queue memory has been allocated to queues where tightly spaced packets are expected. By setting the buffer allocation to reflect the anticipated packet distribution, the interface can reach a higher utilization as a percent of line speed.

Figure 2-8 Default and Tuned Buffer Allocation

It may take some fine tuning to discover the values most appropriate to the load placed on the queues. Settings depend on the exact mix of applications using the interface.

Capacity Planning

Capacity planning involves determining the following:

- How much video is currently running over the network
- How much future video is expected on the network
- The bandwidth requirements for each type of video
- The buffer requirements for each type of video

The first item above is discussed in [Chapter 6, “Medianet Management and Visibility Design Considerations.”](#) Tools available in the network such as NetFlow can help understand the current video loads.

The future video requirements can be more subjective. The recent trend is for more video and for that video to be HD. Even if the number of video streams stays the same, but is updated from SD to HD, the video load on the network will grow substantially.

The bandwidth requirements for video as a 1 second smoothed average are fairly well known. Most standard definition video consumes between 1–3 MB of bandwidth. High definition video takes between 4–6 Mbps, although it can exceed this with the highest quality settings. There are some variances because of attributes such as frame rate (fps) and encoding in use. Table 2-3 lists the bandwidth requirements of common video streams found on a medianet.

Table 2-3 *BAndwidth Requirements of Common Video Streams*

Video Source	Transport	Encoder	Frame Rate	Resolution	Typical Load ¹
Cisco TelePresence System 3000		H.264	30 fps	1080p	12.3 Mbps
Cisco TelePresence System 3000		H.264	30 fps	720p	6.75 Mbps
Cisco TelePresence System 1000		H.264	30 fps	1080p	4.1 Mbps
Cisco TelePresence System 1000		H.264	30 fps	720p	2.25 Mbps
Cisco 2500 Series Video Surveillance IP Camera		MPEG-4	D1 (720x480)	15 fps	1 Mbps
Cisco 2500 Series Video Surveillance IP Camera		MPEG-4	D1 (720x480)	30 fps	2 Mbps
Cisco 2500 Series Video Surveillance IP Camera		M-JPEG	D1 (720x480)	5 fps	2.2 Mbps
Cisco 4500 Series Video Surveillance IP Camera		H.264	1080p	30 fps	4–6 Mbps
Cisco Digital Media System (DMS)—Show and Share VoD		WMV	720x480	30 fps	1.5 Mbps
Cisco Digital Media System (DMS)—Show and Share Live		WMV	720x480	30 fps	1.5 Mbps
Cisco DMS—Digital Sign SD (HTTP)		MPEG-2	720x480	30 fps	3–5 Mbps
Cisco DMS—Digital Sign HD (HTTP)		MPEG-2	1080p	30 fps	13–15 Mbps
Cisco DMS—Digital Sign SD (HTTP)		H.264	720x480	30 fps	1.5–2.5 Mbps
Cisco DMS—Digital Sign HD (HTTP)		H.264	1080p	30 fps	8–12 Mbps
Cisco Unified Video Advantage	UDP/5445	H.264	CIF	variable	768 Kbps
Cisco WebEx	TCP/HTTPS		CIF	variable	128K per small thumbnail
YouTube	TCP/HTTP	MPEG-4	320x240		768 Kbps
YouTube HD	TCP/HTTP	H.264	720p		2 Mbps

1. This does not include audio or auxiliary channels.

The one second smoothed average is influenced by the stream of P-frame. Although I-frames do not occur often enough to have a substantive influence over the average load, they do influence the burst size. From an overly simplified planning capacity standpoint, if a 10 percent overhead is added to the one second load, and the high end of the range is used, the planning numbers become 3.3 MB for standard definition and 6.6 MB for HD video. If you allow 25 percent as interface headroom, [Table 2-4](#) provides some guidance for common interface speeds.

Table 2-4 Common Interface Speeds

Interface	Provisioned Rate	HD	SD
10 Gbps	7.5 Gbps	1136	2272
1 Gbps	750 Mbps	113	226
155 Mbps	11633 Mbps	17	34
100 Mbps	75 Mbps	11	22
45 Mbps	33 Mbps	5	10



Note


These values are based on mathematical assumptions about the frame distribution. They give approximate guidance where only video is carried on the link. These values, as of this writing, have not yet been validated, with the exception of TelePresence, where the numbers modeled above are appropriate. In cases where the encoder setting results in larger video streams, the values given here are not appropriate.

Load Balancing

Inevitably, there are multiple paths between a sender and receiver. The primary goal of multiple paths is to provide an alternate route around a failure in the network. If this is done at each hop, and the metrics are equal to promote load balancing, the total number of paths can grow to be quite large. The resulting binary tree is $2^{(\text{hop count})}$. If the hop count is 4, the number of possible paths is 16 (2^4). If there were three next hops for each destination, the total number of paths is $3^{(\text{hop count})}$. (See [Figure 2-9](#)).

Support and troubleshooting issues arise as the number of possible paths increases. These are covered in [Chapter 6, “Medianet Management and Visibility Design Considerations.”](#)

Figure 2-9 Load Balancing



Although not the primary purpose of redundant links, most designs attempt to use all bandwidth when it is available. In this case, it is prudent to remember that the load should still be supported if any link fails. The more paths available, the higher utilization each path can be provisioned for. If a branch has two circuits, each circuit should run less than 50 percent load to allow failover capacity. If there are three paths available, each circuit can be provisioned to 66 percent capacity. At four paths, each is allowed to run at 75 percent total capacity, and still mathematically allow the load of any single failed path to be distributed to the remaining circuits. In the extreme case, the total bandwidth can be distributed onto so many circuits that a large size flow would congest a particular path.

The exercise can easily be applied to upstream routers in addition to the feeder circuits. As is often the case, there are competing objectives. If there are too many paths, troubleshooting difficulties can extend outages and decrease overall availability. If there are too few paths, expensive bandwidth must be purchased that is rarely used, and some discipline must be employed to ensure the committed load does not grow beyond the single path capacity. Port channels or Multilink PPP are methods to provide L1 redundancy without introducing excessive Layer 3 complexity. These each introduce other complexities and will be discussed in more detail in a future version of this document.

Another approach is to restrict some load that can be considered non-mission critical, such as the applications in the scavenger class. This is valid if you are willing to accept that not all applications will be afforded an alternate path. There are various ways to achieve this, from simple routing to more advanced object tracking.


Consider the following guidelines when transporting video with multi-path routing:

- Ensure that load balancing is per-flow and not per-packet—This helps prevent out-of-order packets. Per-flow also minimizes the chance of a single congested or failing link ruining the video. Because each frame is composed of multiple packets, in a per-packet load balancing scenario, each frame is spread over every path. If any one path has problems, the entire frame can be destroyed.
- Determine a preferred path—With equal cost metrics, the actual path may be difficult to discover. Tools such as trace cannot effectively discover the path a particular flow has taken over equal cost routes, because Cisco Express Forwarding considers both the source and destination address in the hash to determine the next hop. The source address used by trace may not hash to the same path as the stream experiencing congestion. If there are problems, it takes longer to isolate the issue to a particular link if the path is not deterministic. Enhanced Interior Gateway Routing Protocol (EIGRP) provides an offset list that can be used to influence the metric of a particular route by changing its delay. To make use of this feature, mission-critical video such as TelePresence needs to be on

dedicated subnets. The specific routes need to be allowed through any summary boundaries. Offset lists are used at each hop to prefer one path over another for just that subnet (or multiple subnets, as defined in the access control list). This method is useful only to set a particular class of traffic on a determined route, while all other traffic crossing the interface is using the metric of the interface. Offset lists do take additional planning, but can be useful to manage a balanced load in a specific and known way.

- When possible, load balance multiple circuits such that similar traffic is flowing together, and competing traffic is kept apart. For example, video and VoIP should both be handled in the priority queue as real-time RTP traffic. This can be done with the dual-PQ algorithm, or by setting each to prefer a unique path. Without special handling, it is possible that the large packets packed tightly in a video frame can inject jitter into the much smaller and periodic VoIP packets, especially on lower speed links where serialization delay can be a concern.
- Hot Standby Routing Protocol (HSRP), Virtual Router Redundancy Protocol (VRRP), and Gateway Load Balancing Protocol (GLBP)—These are all gateway next-hop protocols. They can be used to direct media traffic off of the LAN into the routed domain. HSRP and VRRP are very similar. VRRP is an open standards protocol while HSRP is found in Cisco products. HSRP does not provide load balancing natively but does allow multiple groups to serve the same LAN. The Dynamic Host Configuration Protocol (DHCP) pool is then broken into two groups, each with its gateway address set to match one of the two HSRP standby addresses. GLBP does support native load balancing. It has only a single address, but the participating devices take turns responding to Address Resolution Protocol (ARP) requests. This allows the single IP address to be load balanced over multiple gateways on a per-client basis. This approach also allows a single DHCP pool to be used.

Both HSRP and GLBP can be used in a video environment. Ideally a given LAN is mission-specific for tasks such as video surveillance, digital media signage, or TelePresence. These tasks should not be on the same LAN as other types of default traffic. This allows unique subnets to be used. The design should consider the deterministic routing in the network as discussed above. Often multiple VLANs are used for data, voice, video, and so on. In this case, it may make operational sense to set the active address of each VLAN on a predetermined path that aligns with the routing. For example, real-time voice would use box A as the active gateway, while real-time video would use box B. In the example shown in [Figure 2-10](#), voice and video are both treated with priority handling. Data and other class-based traffic can be load balanced over both boxes.

Figure 2-10 Load Balancing Example

Trace Route. Video path is predetermined and known. Troubleshooting will focus on this path only.

A-A-A-A

Trace Route. Data path, determined by CEF hash based on IP address source and destination. A different source address could take another path. In this example the path is:

B-A-B-B

228645

These design considerations attempt to reduce the time required to troubleshoot a problem because the interfaces in the path are known. The disadvantage of this approach is that the configurations are more complicated and require more administrative overhead. This fact can offset any gains from a predetermined path, depending on the discipline of the network operations personnel. The worst case would be a hybrid, where a predetermined path is thought to exist, but in fact does not or is not the expected path. Processes and procedures should be followed consistently to ensure that troubleshooting does not include false assumptions. In some situations, load balancing traffic may be a non-optimal but less error-prone approach. It may make operational sense to use a simplified configuration. Each situation is unique.

EtherChannel

It is possible to bond multiple Ethernet interfaces together to form an EtherChannel. This effectively increases the bandwidth because parallel paths are allowed at Layer 2 without spanning tree blocking any of the redundant paths. EtherChannel is documented by the IEEE as 802.ad. Although EtherChannels do effectively increase the bandwidth to the aggregation of the member interfaces, there are a few limitations worth noting. First, packets are not split among the interfaces as they are with Multilink PPP. In addition, packets from the flow will use the same interface based on a hash of that flow. There are some advantages of this approach. Packets will arrive in the same order they were sent. If a flow was sent over multiple interfaces, some resolution is needed to reorder any out of order packets. However, this also means that the bandwidth available for any single flow is still restricted to a single member interface. If many video flows hash to the same interface, then it is possible that the buffer space of that physical interface will be depleted.

Bandwidth Conservation

There are two fundamental approaches to bandwidth management. The first approach is to ensure there is more bandwidth provisioned than required. This is easier in the campus where high speed interfaces can be used to attach equipment located in the same physical building. This may not always be possible where distance is involved, such as the WAN. In this case, it may be more cost-effective to try to minimize the bandwidth usage.

Multicast

Broadcast video is well suited for the bandwidth savings that can be realized with multicast. In fact, IPTV was the original driver for multicast deployments in many enterprise environments. Multicasts allow a single stream to be split by the network as it fans out to receiving stations. In a traditional point-to-point topology, the server must generate a unique network stream for every participant. If everyone is essentially getting the same video stream in real-time, the additional load on both the server and shared portions of the network can negatively impact the scalability. With a technology such as Protocol Independent Multicast (PIM), listeners join a multicast stream. If hundreds or even thousands of users have joined the stream from the same location, only one copy of that stream needs to be generated by the server and sent over the network.

There are some specific cases that can benefit from multicast, but practical limitations warrant the use of unicast. Of the all various types of video found on a medianet, Cisco DMS is the best suited for multicast. This is because of the one-to-many nature of signage. The benefits are best suited when several displays are located at the same branch. The network savings are not significant when each branch has only a single display, because the fanout occurs at the WAN aggregation router, leaving the real savings for the high speed LAN interface.

Aside from DMS, other video technologies have some operational restrictions that limit the benefits of multicast. For example, TelePresence does support multipoint conference calls. However, this is accomplished with a Multipoint Conferencing Unit (MCU), which allows for unique control plane activity to manage which stations are sending, and which stations are the receivers. The MCU serves as a central control device. It also manipulates information in the packet header to control screen placement. This helps ensure that participants maintain a consistent placement when a conference call has both one screen and three screen units.

IP Virtual Server (IPVS) is another technology that can benefit from multicast in very specific situations. However, in most cases, the savings are not realized. Normally, the UDP/RTP streams from the camera terminate on a media server, and not directly on a display station. The users use HTTP to connect to the media server and view various cameras at the discretion of the user. Video surveillance is a many-to-one as opposed to one-to-many. Many cameras transmit video to a handful of media servers, which then serve unicast HTTP clients.

For a more detailed look at video over multicast, see the Multicast chapter in the *Cisco Digital Media System 5.1 Design Guide for Enterprise Medianet* at the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/DMS_DG/DMS_DG.html.


Cisco Wide Area Application Services

Cisco Wide Area Application Services (WAAS) is another technique that can be used to more efficiently use limited bandwidth. Cisco WAAS is specifically geared for all applications that run over the WAN. A typical deployment has a pair or more of Cisco Wide Area Application Engines (WAEs) on either side

of the WAN. The WAEs sit in the flow of the path, and replace the segment between the WAEs with an optimized segment. Video is only one service that can benefit from WAAS. Other WAAS components include the following:

- **TCP Flow Optimization (TFO)**—This feature can help video that is transported in TCP sessions (see Figure 2-11). The most common TCP transport for video is HTTP or HTTPS. There are also video control protocols such as RTSP and RTP Control Protocol (RTCP) that use TCP and benefit from WAAS, such as Cisco DMS. TFO can shield the TCP session from WAN conditions such as loss and congestion. TFO is able to better manage the TCP windowing function. Whereas normal TCP cuts the window size in half and then slowly regains windowing depth, TFO uses an sophisticated algorithm to set window size and recover from lost packets. Video that is transported over TCP can benefit from WAAS, including Adobe Flash, Quicktime, and HTTP, which commonly use TCP. RTP or other UDP flows do not benefit from TFO.

Figure 2-11 TFO



- **Data Redundancy Elimination**—WAAS can discover repeating patterns in the data. The pattern is then replaced with an embedded code that the paired device recognizes and replaces with the pattern. Depending on the type of traffic, this can represent a substantial savings in bandwidth. This feature is not as useful with video, because the compression used by the encoders tends to eliminate any redundancy in the data. There may still be gains in the control plane being used by video. Commonly these are Session Initiation Protocol (SIP) or RTSP.
- **Persistent LZ Compression**—This is a compression technique that also looks for mutual redundancy, but in the bit stream, outside of byte boundaries. The video codecs have already compressed the bit stream using one of two techniques, context-adaptive binary arithmetic coding (CABAC) or context-adaptive variable-length coding (CAVLC). LZ Compression and CABAC/CAVLC are both forms of entropy encoding. By design, these methods eliminate any mutual redundancy. This means that compressing a stream a second time does not gain any appreciable savings. This is the case with LZ compression of a video stream. The gains are modest at best.

Cisco Application and Content Network Systems


Cisco Application and Content Network Systems (ACNS) is another tool that can better optimize limited WAN bandwidth. Cisco ACNS runs on the Cisco WAE product family as either a content engine or content distribution manager. Cisco ACNS saves WAN bandwidth by caching on-demand content or prepositioning content locally. When many clients in a branch location request this content, ACNS can fulfill the request locally, thereby saving repeated requests over the WAN. Of the four technologies that form a medianet, ACNS is well suited for Cisco DMS and desktop broadcast video. For more information, see the *Cisco Digital Media System 5.1 Design Guide for Enterprise Medianet* at the following URL:

http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/DMS_DG/DMS_dgbk.pdf.

Cisco Performance Routing

Cisco Performance Routing (PfR) is a feature available in Cisco routers that allows the network to make routing decisions based on network performance. This tool can be used to ensure that the WAN is meeting specific metrics such as loss, delay, and jitter. PfR can operate in either a passive or active mode. One or more border routers is placed at the edge of the WAN. A master controller collects performance information from the border routers and makes policy decisions. These decisions are then distributed to the border routers for implementation. Figure 2-12 shows a typical topology.

Figure 2-12 Typical Topology using PfR



Multiprotocol Environments

In the early days of networking, it was common to see multiple protocols running simultaneously. Many networks carried IP, Internetwork Packet Exchange (IPX), Systems Network Architecture (SNA), and perhaps AppleTalk or DEC. It was not uncommon for an IPX Service Advertising Protocol (SAP) update to occasionally cause 3270 sessions to clock. Modern networks are increasingly IP only, yet convergence remains a concern for the same reason: large blocks of packets are traveling together with small time-sensitive packets. The difference now is that the large stream is also time-sensitive. QoS is the primary tool currently used to ensure that bandwidth is used as efficiently as possible. This feature allows UDP RTP video to be transported on the same network as TCP-based non-real-time video and mission-critical data applications. In addition to many different types of video along with traditional data and voice, new sophisticated features are being added to optimize performance, including those discussed here: Cisco WAAS, multicast, Cisco ACNS, PfR, and so on, as well as other features to support QoS, security, and visibility. New features are continuously being developed to further improve network performance. The network administrator is constantly challenged to ensure that the features are working together to obtain the desired result. In most cases, features are agnostic and do not interfere with one another.



Note

Future revisions to this chapter will include considerations where this is not the case. For example, security features can prevent WAAS from properly functioning.

Summary

Bandwidth is an essential base component of a medianet architecture. Other features can help to maximize the utilization of the circuit in the network, but do not replace the need to adequately provisioned links. Because CAC-like functionality is not yet available for video, proper planning should accommodate the worst-case scenario when many HD devices are present. When network bandwidth saturates, all video suffers. Near-perfect HD video is necessary to maximize the potential in productivity gains. Bandwidth is the foundational component of meeting this requirement, but not the only service needed. Other functionality such as QoS, availability, security, management, and visibility are also required. These features cannot be considered standalone components, but all depend on each other. Security requires good management and visibility. QoS requires adequate bandwidth. Availability depends on effective security. Each feature must be considered in the context of an overall medianet architecture.



CHAPTER 3

Medianet Availability Design Considerations

The goal of network availability technologies is to maximize network uptime such that the network is always ready and able to provide needed services to critical applications, such as TelePresence or other critical network video.

Network video has varying availability requirements. At one extreme, if a single packet is lost, the user likely notices an artifact in the video. On the other extreme, video is a unidirectional session; the camera always sends packets and the display always receives packets. When an outage occurs, the camera may not recognize it, and continue to send video packets. Upper layer session control protocols, such as Session Initiation Protocol (SIP) and Real-Time Streaming Protocol (RTSP), are responsible to validate the path. Video applications may respond differently to session disruptions. In all cases, the video on the display initially freezes at the last received frame, and looks to the session control for some resolution. If the packet stream is restored, quite often the video recovers without having to restart the session. TelePresence can recover after a 30-second network outage before SIP terminates the call. Broadcast video may be able to go longer. Availability techniques should be deployed such that the network converges faster than the session control protocol hello interval. The user notices that the video has frozen, but in most cases, the stream recovers without having to restart the media.

Network Availability

Network availability is the cornerstone of network design, on which all other services depend.

The three primary causes of network downtime are as follows:

- Hardware failures, which can include system and sub-component failures, as well as power failures and network link failures
- Software failures, which can include incompatibility issues and bugs
- Operational processes, which mainly include human error; however, poorly-defined management and upgrading processes may also contribute to operational downtime

To offset these types of failures, the network administrator attempts to provision the following types of resiliency:

- Device resiliency—Deploying redundant hardware (including systems, supervisors, line cards, and power-supplies) that can failover in the case of hardware and/or software failure events
- Network resiliency—Tuning network protocols to detect and react to failure events as quickly as possible
- Operational resiliency—Examining and defining processes to maintain and manage the network, leveraging relevant technologies that can reduce downtime, including provisioning for hardware and software upgrades with minimal downtime (or optimally, with no downtime)



Note

Because the purpose of this overview of availability technologies is to provide context for the design chapters to follow, this discussion focuses on device and network resiliency, rather than operational resiliency.

Network availability can be quantitatively measured by using the formula shown in [Figure 3-1](#), which correlates the mean time between failures (MTBF) and the mean time to repair (MTTR) such failures.

Figure 3-1 Availability Formula

$$Availability = \frac{MTBF}{MTBF + MTTR}$$


For example, if a network device has an MTBF of 10,000 hours and an MTTR of 4 hours, its availability can be expressed as 99.96 percent[(10,000)/(10,000 + 4), converted to a percentage].

Therefore, from this formula it can be seen that availability can be improved by either increasing the MTBF of a device (or network), or by decreasing the MTTR of the same.

The most effective way to increase the MTBF of a device (or network) is to design with redundancy. This can be mathematically proven by comparing the availability formula of devices connected in serial (without redundancy) with the formula of devices connected in parallel (with redundancy).

The availability of devices connected in series is shown in [Figure 3-2](#).

Figure 3-2 Availability Formula for Devices Connected in Serial



S₁, S₂ - Series Components

System is available when both components are available:

$$A_{series} = A_1 \times A_2$$


223825

S1 and S2 represent two separate systems (which may be individual devices or even networks). A1 and A2 represent the availability of each of these systems, respectively. Aseries represents the overall availability of these systems connected in serial (without redundancy).

For example, if the availability of the first device (S1) is 99.96 percent and the availability of the second device (S2) is 99.98 percent, the overall system availability, with these devices connected serially, is 99.94 percent (99.96% x 99.98%).

Therefore, connecting devices in serial actually reduces the overall availability of the network.

In contrast, consider the availability of devices connected in parallel, as shown in [Figure 3-3](#).

Figure 3-3 Availability Formula for Devices Connected in Parallel

S₃, S₄ - Parallel Components

System is unavailable when both components are unavailable:

$$A_{\text{parallel}} = 1 - (1 - A_1) \times (1 - A_2)$$

223826


S₃ and S₄ represent two separate systems (devices or networks). A₃ and A₄ represent the availability of each of these systems, respectively. A_{parallel} represents the overall availability of these systems connected in parallel (with redundancy).

Continuing the example, using the same availability numbers for each device as before yields an overall system availability, with these devices connected in parallel, of 99.99992 percent [1-(1-99.96%) * (1-99.98%)].


Therefore, connecting devices in parallel significantly increases the overall availability of the combined system. This is a foundational principle of available network design, where individual devices as well as networks are designed to be fully redundant, whenever possible. Figure 3-4 illustrates applying redundancy to network design and its corresponding effect on overall network availability.

Figure 3-4 Impact of Redundant Network Design on Network Availability


Reliability = 99.938% with Four Hour MTTR (325 Minutes/Year)



Reliability = 99.961% with Four Hour MTTR (204 Minutes/Year)



Reliability = 99.9999% with Four Hour MTTR (30 Seconds/Year)



223690

A *five nines* network (a network with 99.999 percent availability) has been considered the hallmark of excellent enterprise network design for many years. However, a five nines network allows for only five minutes of downtime per year.

Another commonly used metric for measuring availability is defects per million (DPM). Measuring the probability of failure of a network and establishing the service-level agreement (SLA) that a specific design is able to achieve is a useful tool, but DPM takes a different approach by measuring the impact of defects on the service from the end-user perspective. This is often a better metric for determining the availability of the network because it better reflects the user experience relative to event effects. DPM is calculated based on taking the total affected user minutes for each event, total users affected, and the duration of the event, as compared to the total number of service minutes available during the period in question. The sum of service downtime minutes is divided by the total service minutes and multiplied by 1,000,000, as shown in [Figure 3-5](#).

Figure 3-5 Defects Per Million Calculation

$$DPM = \frac{\sum(\text{number of users affected} * \text{Outage Minutes})}{\text{Total Users} * \text{Total Service Minutes}} \times 1,000,000$$

For example, if a company of 50 employees suffers two separate outages during the course of a year, with the first outage affecting 12 users for 4 hours and the second outage affecting 25 users for 2 hours, the total DPM is 224 [[(12 users x 240 min)+(25 users x 120 min)]/(50 users x 525,960 min/year)]x 1,000,000, rounded].



Note

The benefit of using a “per-million” scale in a defects calculation is that it allows the final ratio to be more readable, given that this ratio becomes extremely small as availability improves.

DPM is useful because it is a measure of the observed availability and considers the impact to the end user as well as the network itself. Adding this user experience element to the question of network availability is very important to understand, and is becoming a more important part of the question of what makes a highly available network.

[Table 3-1](#) summarizes the availability targets, complete with their DPM and allowable downtime/year.

Table 3-1 Availability, DPM, and Downtime

Availability (Percent)	DPM	Downtime/Year
99.000	10,000	3 days, 15 hours, 36 minutes
99.500	5,000	1 day, 19 hours, 48 minutes
99.900	1,000	8 hours, 46 minutes
99.950	500	4 hours, 23 minutes
99.990	100	53 minutes
99.999	10	5 minutes
99.9999	1	0.5 minutes

Having reviewed these availability principles, metrics, and targets, the next section discusses some of the availability technologies most relevant for systems and networks supporting TelePresence systems.

Device Availability Technologies

Every network design has single points of failure, and the overall availability of the network might depend on the availability of a single device. The access layer of a campus network is a prime example of this. Every access switch represents a single point of failure for all the attached devices (assuming that the endpoints are single-homed; this does not apply to endpoint devices that are dual-homed). Ensuring the availability of the network services often depends on the resiliency of the individual devices.

Device resiliency, as with network resiliency, is achieved through a combination of the appropriate level of physical redundancy, device hardening, and supporting software features. Studies indicate that most common failures in campus networks are associated with Layer 1 failures, from components such as power supplies, fans, and fiber links. The use of diverse fiber paths with redundant links and line cards, combined with fully redundant power supplies and power circuits, are the most critical aspects of device resiliency. The use of redundant power supplies becomes even more critical in access switches with the introduction of power over Ethernet (PoE) devices such as IP phones. Multiple devices now depend on the availability of the access switch and its ability to maintain the necessary level of power for all the attached end devices. After physical failures, the most common cause of device outage is often related to the failure of supervisor hardware or software. The network outages caused by the loss or reset of a device because of supervisor failure can be addressed through the use of supervisor redundancy. Cisco Catalyst switches provide the following mechanisms to achieve this additional level of redundancy:

- Cisco StackWise and Cisco StackWise-Plus
- Cisco non-stop forwarding (NSF) with stateful switchover (SSO)

Both these mechanisms, which are discussed in the following sections, provide for a hot active backup for the switching fabric and control plane, thus ensuring that data forwarding and the network control plane seamlessly recover (with sub-second traffic loss, if any) during any form of software or supervisor hardware crash.

Cisco StackWise and Cisco StackWise Plus

Cisco StackWise and Cisco StackWise Plus technologies are used to create a unified, logical switching architecture through the linkage of multiple, fixed configuration Cisco Catalyst 3750G and/or Cisco Catalyst 3750E switches.

Cisco Catalyst 3750G switches use StackWise technology and Cisco Catalyst 3750E switches can use either StackWise or StackWise Plus. StackWise Plus is used only if all switches within the group are 3750E switches; whereas, if some switches are 3750E and others are 3750G, StackWise technology is used.



Note

“StackWise” is used in this section to refer to both Cisco StackWise and Cisco StackWise Plus technologies, with the exception of explicitly pointing out the differences between the two at the end of this section.

Cisco StackWise technology intelligently joins individual switches to create a single switching unit with a 32-Gbps switching stack interconnect. Configuration and routing information is shared by every switch in the stack, creating a single switching unit. Switches can be added to and deleted from a working stack without affecting availability.

The switches are united into a single logical unit using special stack interconnect cables that create a bidirectional closed-loop path. This bidirectional path acts as a switch fabric for all the connected switches. Network topology and routing information are updated continuously through the stack interconnect. All stack members have full access to the stack interconnect bandwidth. The stack is managed as a single unit by a master switch, which is elected from one of the stack member switches.

Each switch in the stack has the capability to behave as a master in the hierarchy. The master switch is elected and serves as the control center for the stack. Each switch is assigned a number. Up to nine separate switches can be joined together.

Each stack of Cisco Catalyst 3750 Series Switches has a single IP address and is managed as a single object. This single IP management applies to activities such as fault detection, VLAN creation and modification, security, and quality of service (QoS) controls. Each stack has only one configuration file, which is distributed to each member in the stack. This allows each switch in the stack to share the same network topology, MAC address, and routing information. In addition, this allows for any member to immediately take over as the master, in the event of a master failure.

To efficiently load balance the traffic, packets are allocated between two logical counter-rotating paths. Each counter-rotating path supports 16 Gbps in both directions, yielding a traffic total of 32 Gbps bidirectionally. When a break is detected in a cable, the traffic is immediately wrapped back across the single remaining 16-Gbps path (within microseconds) to continue forwarding.

Switches can be added and deleted to a working stack without affecting stack availability. However, adding additional switches to a stack may have QoS performance implications, as is discussed in more in [Chapter 4, “Medianet QoS Design Considerations.”](#) Similarly, switches can be removed from a working stack with no operational effect on the remaining switches.

Stacks require no explicit configuration, but are automatically created by StackWise when individual switches are joined together with stacking cables, as shown in [Figure 3-6](#). When the stack ports detect electromechanical activity, each port starts to transmit information about its switch. When the complete set of switches is known, the stack elects one of the members to be the master switch, which becomes responsible for maintaining and updating configuration files, routing information, and other stack information.

Figure 3-6 Cisco Catalyst 3750G StackWise Cabling



Each switch in the stack can serve as a master, creating a 1:N availability scheme for network control. In the unlikely event of a single unit failure, all other units continue to forward traffic and maintain operation. Furthermore, each switch is initialized for routing capability and is ready to be elected as master if the current master fails. Subordinate switches are not reset so that Layer 2 forwarding can continue uninterrupted.

The following are the three main differences between StackWise and StackWise Plus:

- StackWise uses source stripping and StackWise Plus uses destination stripping (for unicast packets). Source stripping means that when a packet is sent on the ring, it is passed to the destination, which copies the packet, and then lets it pass all the way around the ring. When the packet has traveled all

the way around the ring and returns to the source, it is stripped off the ring. This means bandwidth is used up all the way around the ring, even if the packet is destined for a directly attached neighbor. Destination stripping means that when the packet reaches its destination, it is removed from the ring and continues no further. This leaves the rest of the ring bandwidth free to be used. Thus, the throughput performance of the stack is multiplied to a minimum value of 64 Gbps bidirectionally. This ability to free up bandwidth is sometimes referred to as spatial reuse.

**Note**

Even in StackWise Plus, broadcast and multicast packets must use source stripping because the packet may have multiple targets on the stack.

- StackWise Plus can locally switch, whereas StackWise cannot. Furthermore, in StackWise, because there is no local switching and there is source stripping, even locally destined packets must traverse the entire stack ring.
- StackWise Plus supports up to two Ten Gigabit Ethernet ports per Cisco Catalyst 3750-E.

Finally, both StackWise and StackWise Plus can support Layer 3 non-stop forwarding (NSF) when two or more nodes are present in a stack.

Non-Stop Forwarding with Stateful Switch Over

Stateful switchover (SSO) is a redundant route- and/or switch-processor availability feature that significantly reduces MTTR by allowing extremely fast switching between the main and backup processors. SSO is supported on routers (such as the Cisco 7600, 10000, and 12000 Series) and switches (such as the Cisco Catalyst 4500 and 6500 Series).

Before discussing the details of SSO, a few definitions may be helpful. For example, *state* in SSO refers to maintaining between the active and standby processors, among many other elements, the protocol configurations and current status of the following:

- Layer 2 (L2)
- Layer 3 (L3)
- Multicast
- QoS policy
- Access list policy
- Interface

Also, the adjectives *cold*, *warm*, and *hot* are used to denote the readiness of the system and its components to assume the network services functionality and the job of forwarding packets to their destination. These terms appear in conjunction with Cisco IOS verification command output relating to NSF/SSO, as well as with many high availability feature descriptions. These terms are generally defined as follows:

- **Cold**—The minimum degree of resiliency that has been traditionally provided by a redundant system. A redundant system is cold when no state information is maintained between the backup or standby system and the system to which it offers protection. Typically, a cold system must complete a boot process before it comes online and is ready to take over from a failed system.
- **Warm**—A degree of resiliency beyond the cold standby system. In this case, the redundant system has been partially prepared, but does not have all the state information known by the primary system to take over immediately. Additional information must be determined or gleaned from the traffic flow or the peer network devices to handle packet forwarding. A warm system is already booted and needs to learn or generate only the state information before taking over from a failed system.

- Hot—The redundant system is fully capable of handling the traffic of the primary system. Substantial state information has been saved, so the network service is continuous, and the traffic flow is minimally or not affected.

To better understand SSO, it may be helpful to consider its operation in detail within a specific context, such as within a Cisco Catalyst 6500 with two supervisors per chassis.

The supervisor engine that boots first becomes the active supervisor engine. The active supervisor is responsible for control plane and forwarding decisions. The second supervisor is the standby supervisor, which does not participate in the control or data plane decisions. The active supervisor synchronizes configuration and protocol state information to the standby supervisor, which is in a hot standby mode. As a result, the standby supervisor is ready to take over the active supervisor responsibilities if the active supervisor fails. This take-over process from the active supervisor to the standby supervisor is referred to as a *switchover*.

Only one supervisor is active at a time, and supervisor engine redundancy does not provide supervisor engine load balancing. However, the interfaces on a standby supervisor engine are active when the supervisor is up and thus can be used to forward traffic in a redundant configuration.


NSF/SSO evolved from a series of progressive enhancements to reduce the impact of MTTR relating to specific supervisor hardware/software network outages. NSF/SSO builds on the earlier work known as Route Processor Redundancy (RPR) and RPR Plus (RPR+). Each of these redundancy modes of operation incrementally improves on the functions of the previous mode.

- RPR-RPR is the first redundancy mode of operation introduced in Cisco IOS Software. In RPR mode, the startup configuration and boot registers are synchronized between the active and standby supervisors, the standby is not fully initialized, and images between the active and standby supervisors do not need to be the same. Upon switchover, the standby supervisor becomes active automatically, but it must complete the boot process. In addition, all line cards are reloaded and the hardware is reprogrammed. Because the standby supervisor is *cold*, the RPR switchover time is two or more minutes.
- RPR+-RPR+ is an enhancement to RPR in which the standby supervisor is completely booted and line cards do not reload upon switchover. The running configuration is synchronized between the active and the standby supervisors. All synchronization activities inherited from RPR are also performed. The synchronization is done before the switchover, and the information synchronized to the standby is used when the standby becomes active to minimize the downtime. No link layer or control plane information is synchronized between the active and the standby supervisors. Interfaces may bounce after switchover, and the hardware contents need to be reprogrammed. Because the standby supervisor is *warm*, the RPR+ switchover time is 30 or more seconds.
- NSF with SSO-NSF works in conjunction with SSO to ensure Layer 3 integrity following a switchover. It allows a router experiencing the failure of an active supervisor to continue forwarding data packets along known routes while the routing protocol information is recovered and validated. This forwarding can continue to occur even though peering arrangements with neighbor routers have been lost on the restarting router. NSF relies on the separation of the control plane and the data plane during supervisor switchover. The data plane continues to forward packets based on pre-switchover Cisco Express Forwarding information. The control plane implements graceful restart routing protocol extensions to signal a supervisor restart to NSF-aware neighbor routers, reform its neighbor adjacencies, and rebuild its routing protocol database (in the background) following a switchover. Because the standby supervisor is *hot*, the NSF/SSO switchover time is 0–3 seconds.

As previously described, neighbor nodes play a role in NSF function. A node that is capable of continuous packet forwarding during a route processor switchover is NSF-capable. Complementing this functionality, an NSF-aware peer router can enable neighbor recovery without resetting adjacencies, and support routing database re-synchronization to occur in the background. [Figure 3-5](#) illustrates the difference between NSF-capable and NSF-aware routers. To gain the greatest benefit from NSF/SSO deployment, NSF-capable routers should be peered with NSF-aware routers

(although this is not absolutely required for implementation), because only a limited benefit is achieved unless routing peers are aware of the ability of the restarting node to continue packet forwarding and assist in restoring and verifying the integrity of the routing tables after a switchover.

Figure 3-7 NSF-Capable Compared to NSF-Aware Routers



Cisco Nonstop Forwarding and Stateful Switchover are designed to be deployed together. NSF relies on SSO to ensure that links and interfaces remain up during switchover, and that the lower layer protocol state is maintained. However, it is possible to enable SSO with or without NSF, because these are configured separately.

The configuration to enable SSO is very simple, as follows:

```
Router(config)#redundancy
Router(config-red)#mode sso
```

NSF, on the other hand, is configured within the routing protocol itself, and is supported within Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), Intermediate System to Intermediate System (IS-IS), and (to an extent) Border Gateway Protocol (BGP). Sometimes NSF functionality is also called *graceful-restart*.

To enable NSF for EIGRP, enter the following commands:

```
Router(config)# router eigrp 100
Router(config-router)# nsf
```

Similarly, to enable NSF for OSPF, enter the following commands:

```
Router(config)# router ospf 100
Router(config-router)# nsf
```

Continuing the example, to enable NSF for IS-IS, enter the following commands:

```
Router(config)#router isis level2
Router(config-router)#nsf cisco
```

And finally, to enable NSF/graceful-restart for BGP, enter the following commands:

```
Router(config)#router bgp 100
Router(config-router)#bgp graceful-restart
```

You can see from the example of NSF that the line between device-level availability technologies and network availability technologies is sometimes uncertain. A discussion of more network availability technologies follows.

Network Availability Technologies

Network availability technologies, which include link integrity protocols, link bundling protocols, loop detection protocols, first-hop redundancy protocols (FHRPs) and routing protocols, are used to increase the resiliency of devices connected within a network. Network resiliency relates to how the overall design implements redundant links and topologies, and how the control plane protocols are optimally configured to operate within that design. The use of physical redundancy is a critical part of ensuring the availability of the overall network. In the event of a network device failure, having a path means that the overall network can continue to operate. The control plane capabilities of the network provide the ability to manage the way in which the physical redundancy is leveraged, the network load balances traffic, the network converges, and the network is operated.

The following basic principles can be applied to network availability technologies:

- Wherever possible, leverage the ability of the device hardware to provide the primary detection and recovery mechanism for network failures. This ensures both a faster and a more deterministic failure recovery.
- Implement a defense-in-depth approach to failure detection and recovery mechanisms. Multiple protocols, operating at different network layers, can complement each other in detecting and reacting to network failures.
- Ensure that the design is self-stabilizing. Use a combination of control plane modularization to ensure that any failures are isolated in their impact and that the control plane prevents flooding or thrashing conditions from arising.

These principles are intended to complement the overall structured modular design approach to the network architecture and to re-enforce good resilient network design practices.



Note

A complete discussion of all network availability technologies and best practices could easily fill an entire volume. Therefore, this discussion introduces only an overview of the most relevant network availability technologies for TelePresence enterprise network deployments.

The following sections discuss L2 and L3 network availability technologies.

L2 Network Availability Technologies

L2 network availability technologies that particularly relate to TelePresence network design include the following:

- Unidirectional Link Detection (UDLD)
- IEEE 802.1d Spanning Tree Protocol (STP)
- Cisco Spanning Tree Enhancements
- IEEE 802.1w Rapid Spanning Tree Protocol (RSTP)
- Trunks, Cisco Inter-Switch Link, and IEEE 802.1Q
- EtherChannels, Cisco Port Aggregation Protocol, and IEEE 802.3ad

- Cisco Virtual Switching System (VSS)


Each of these L2 technologies are discussed in the following sections.

UniDirectional Link Detection

UDLD protocol is a Layer 2 protocol, which uses a keep-alive to test that the switch-to-switch links are connected and operating correctly. Enabling UDLD is a prime example of how a defense-in-depth approach to failure detection and recovery mechanisms can be implemented, because UDLD (an L2 protocol) acts as a backup to the native Layer 1 unidirectional link detection capabilities provided by IEEE 802.3z (Gigabit Ethernet) and 802.3ae (Ten Gigabit Ethernet) standards.

The UDLD protocol allows devices connected through fiber optic or copper Ethernet cables connected to LAN ports to monitor the physical configuration of the cables and detect when a unidirectional link exists. When a unidirectional link is detected, UDLD shuts down the affected LAN port and triggers an alert. Unidirectional links, such as shown in [Figure 3-8](#), can cause a variety of problems, including spanning tree topology loops.

Figure 3-8 Unidirectional Link Failure



You can configure UDLD to be globally enabled on all fiber ports by entering the following command:

```
Switch(config)#udld enable
```


Additionally, you can enable UDLD on individual LAN ports in interface mode, by entering the following commands:

```
Switch(config)#interface GigabitEthernet8/1
Switch(config-if)#udld port
```

Interface configurations override global settings for UDLD.

IEEE 802.1D Spanning Tree Protocol


IEEE 802.1D STP prevents loops from being formed when switches are interconnected via multiple paths. STP implements the spanning tree algorithm by exchanging Bridge Protocol Data Unit (BPDU) messages with other switches to detect loops, and then removes the loop by shutting down selected switch interfaces. This algorithm guarantees that there is only one active path between two network devices, as illustrated in [Figure 3-9](#).

Figure 3-9 STP-Based Redundant Topology

STP prevents a loop in the topology by transitioning all (STP-enabled) ports through four STP states:

- **Blocking**—The port does not participate in frame forwarding. STP can take up to 20 seconds (by default) to transition a port from blocking to listening.
- **Listening**—The port transitional state after the blocking state when the spanning tree determines that the interface should participate in frame forwarding. STP takes 15 seconds (by default) to transition between listening and learning.
- **Learning**—The port prepares to participate in frame forwarding. STP takes 15 seconds (by default) to transition from learning to forwarding (provided such a transition does not cause a loop; otherwise, the port is be set to blocking).
- **Forwarding**—The port forwards frames.

Figure 3-10 illustrates the STP states, including the disabled state.

Figure 3-10 STP Port States

You can enable STP globally on a per-VLAN basis, using Per-VLAN Spanning-Tree (PVST), by entering the following command:

```
Switch(config)# spanning-tree vlan 100
```

The two main availability limitations for STP are as follows:

- To prevent loops, redundant ports are placed in a blocking state and as such are not used to forward frames/packets. This significantly reduces the advantages of redundant network design, especially with respect to network capacity and load sharing.
- Adding up all the times required for STP port-state transitions shows that STP can take up to 50 seconds to converge on a loop-free topology. Although this may have been acceptable when the protocol was first designed, it is certainly unacceptable today.

Both limitations are addressable using additional technologies. The first limitation can be addressed by using the Cisco Virtual Switching System (VSS), discussed later in this section; and the second limitation can be addressed by various enhancements that Cisco developed for STP, as is discussed next.

Cisco Spanning Tree Enhancements

To improve STP convergence times, Cisco has made a number of enhancements to 802.1D STP, including the following:

- PortFast (with BPDU Guard)
- UplinkFast
- BackboneFast


STP PortFast causes a Layer 2 LAN port configured as an access port to enter the forwarding state immediately, bypassing the listening and learning states. PortFast can be used on Layer 2 access ports connected to a single workstation or server to allow those devices to connect to the network immediately, instead of waiting for STP to converge, because interfaces connected to a single workstation or server should not receive BPDUs. Because the purpose of PortFast is to minimize the time that access ports must wait for STP to converge, it should only be used on access ports. Optionally, for an additional level of security, PortFast may be enabled with BPDU Guard, which immediately shuts down a port that has received a BPDU.

You can enable PortFast globally (along with BPDU Guard), or on a per-interface basis, by entering the following commands:

```
Switch(config)# spanning-tree portfast default
Switch(config)# spanning-tree portfast bpduguard default
```

UplinkFast provides fast convergence after a direct link failure and achieves load balancing between redundant Layer 2 links, as shown in Figure 3-11. If a switch detects a link failure on the currently active link (a direct link failure), UplinkFast unblocks the blocked port on the redundant link port and immediately transitions it to the forwarding state without going through the listening and learning states. This switchover takes approximately one to five seconds.

Figure 3-11 UplinkFast Recovery Example After Direct Link Failure




UplinkFast is enabled globally, as follows:

```
Switch(config)# spanning-tree uplinkfast
```

In contrast, BackboneFast provides fast convergence after an indirect link failure, as shown in Figure 3-12. This switchover takes approximately 30 seconds (yet improves on the default STP convergence time by 20 seconds).

Figure 3-12 BackboneFast Recovery Example After Indirect Link Failure



BackboneFast is enabled globally, as follows:

```
Switch(config)# spanning-tree backbonefast
```

These Cisco-proprietary enhancements to 802.1D STP were adapted and adopted into a new standard for STP, IEEE 802.1w or Rapid Spanning-Tree Protocol (RSTP), which is discussed next.

IEEE 802.1w-Rapid Spanning Tree Protocol

RSTP is an evolution of the 802.1D STP standard. RSTP is a Layer 2 loop prevention algorithm like 802.1D; however, RSTP achieves rapid failover and convergence times, because RSTP is not a timer-based spanning tree algorithm (STA) like 802.1D; but rather a handshake-based spanning tree algorithm. Therefore, RSTP offers an improvement of over 30 seconds or more, as compared to 802.1D, in transitioning a link into a forwarding state.

There are the following three port states in RSTP:

- Learning
- Forwarding
- Discarding

The disabled, blocking, and listening states from 802.1D have been merged into a unique 802.1w discarding state.

Rapid transition is the most important feature introduced by 802.1w. The legacy STA passively waited for the network to converge before moving a port into the forwarding state. Achieving faster convergence was a matter of tuning the conservative default timers, often sacrificing the stability of the network.

RSTP is able to actively confirm that a port can safely transition to forwarding without relying on any timer configuration. There is a feedback mechanism that operates between RSTP-compliant bridges. To achieve fast convergence on a port, the RSTP relies on two new variables: edge ports and link type.

The edge port concept basically corresponds to the PortFast feature. The idea is that ports that are directly connected to end stations cannot create bridging loops in the network and can thus directly transition to forwarding (skipping the 802.1D listening and learning states). An edge port does not generate topology changes when its link toggles. Unlike PortFast, however, an edge port that receives a BPDU immediately loses its edge port status and becomes a normal spanning tree port.

RSTP can achieve rapid transition to forwarding only on edge ports and on point-to-point links. The link type is automatically derived from the duplex mode of a port. A port operating in full-duplex is assumed to be point-to-point, while a half-duplex port is considered as a shared port by default. In switched networks today, most links are operating in full-duplex mode and are therefore treated as point-to-point links by RSTP. This makes them candidates for rapid transition to forwarding.

Like STP, you can enable RSTP globally on a per-VLAN basis, also referred to as Rapid-Per-VLAN-Spanning Tree (Rapid-PVST) mode, using the following command:

```
Switch(config)# spanning-tree mode rapid-pvst
```

Beyond STP, there are many other L2 technologies that also play a key role in available network design, such as trunks, which are discussed in the following section.

Trunks, Cisco Inter-Switch Link, and IEEE 802.1Q

A trunk is a point-to-point link between two networking devices (switches and/or routers) capable of carrying traffic from multiple VLANs over a single link. VLAN frames are encapsulated with trunking protocols to preserve logical separation of traffic while transiting the trunk.

There are two trunking encapsulations available to Cisco devices:

- Inter-Switch Link (ISL)—ISL is a Cisco-proprietary trunking encapsulation.
- IEEE 802.1Q—802.1Q is an industry-standard trunking encapsulation. Trunks may be configured on individual links or on EtherChannel bundles (discussed in the following section). ISL encapsulates the original Ethernet frame with both a header and a field check sequence (FCS) trailer,


for a total of 30 bytes of encapsulation. ISL trunking can be configured on a switch port interface, as shown in [Example 3-1](#). The trunking mode is set to ISL, and the VLANs permitted to traverse the trunk are explicitly identified; in this example, VLANs 2 and 102 are permitted over the ISL trunk.

Example 3-1 ISL Trunk Example

```
Switch(config)#interface GigabitEthernet8/3
Switch(config-if)# switchport
Switch(config-if)# switchport trunk encapsulation isl
Switch(config-if)# switchport trunk allowed 2, 102
```

In contrast with ISL, 802.1Q does not actually encapsulate the Ethernet frame, but rather inserts a 4-byte tag after the source address field, as well as recomputes a new FCS, as shown in [Figure 3-13](#). This tag not only preserves VLAN information, but also includes a 3-bit field for class of service (CoS) priority (which is discussed in more detail in [Chapter 4, “Medianet QoS Design Considerations”](#)).

Figure 3-13 IEEE 802.1Q Tagging



IEEE 802.1Q also supports the concept of a native VLAN. Traffic sourced from the native VLAN is not tagged, but is rather simply forwarded over the trunk. As such, only a single native VLAN can be configured for an 802.1Q trunk, to preserve logical separation.



Note

Because traffic from the native VLAN is untagged, it is important to ensure that the same native VLAN be specified on both ends of the trunk. Otherwise, this can cause a routing blackhole and potential security vulnerability.

IEEE 802.1Q trunking is likewise configured on a switch port interface, as shown in [Example 3-2](#). The trunking mode is set to 802.1Q, and the VLANs permitted to traverse the trunk are explicitly identified; in this example, VLANs 3 and 103 are permitted over the 802.1Q trunk. Additionally, VLAN 103 is specified as the native VLAN.

Example 3-2 IEEE 802.1Q Trunk Example


```
Switch(config)# interface GigabitEthernet8/4
Switch(config-if)# switchport
Switch(config-if)# switchport trunk encapsulation dot1q
Switch(config-if)# switchport trunk allowed 3, 103
Switch(config-if)# switchport trunk native vlan 103
```

Trunks are typically, but not always, configured in conjunction with EtherChannels, which allow for network link redundancy, and are described next.

EtherChannels, Cisco Port Aggregation Protocol, and IEEE 802.3ad

EtherChannel technologies create a single logical link by bundling multiple physical Ethernet-based links (such as Gigabit Ethernet or Ten Gigabit Ethernet links) together, as shown in Figure 3-14. As such, EtherChannel links can provide for increased redundancy, capacity, and load balancing. To optimize the load balancing of traffic over multiple links, Cisco recommends deploying EtherChannels in powers of two (two, four, or eight) physical links. EtherChannel links can operate at either L2 or L3.

Figure 3-14 EtherChannel Bundle



EtherChannel links can be created using Cisco Port Aggregation Protocol (PAgP), which performs a negotiation before forming a channel, to ensure compatibility and administrative policies.

PAgP can be configured in four channeling modes:

- **On**—This mode forces the LAN port to channel unconditionally. In the On mode, a usable EtherChannel exists only when a LAN port group in the On mode is connected to another LAN port group in the On mode. Ports configured in the On mode do not negotiate to form EtherChannels; they simply do or do not, depending on the configuration of the other port.
- **Off**—This mode precludes the LAN port from channeling unconditionally.
- **Desirable**—This PAgP mode places a LAN port into an active negotiating state, in which the port initiates negotiations with other LAN ports to form an EtherChannel by sending PAgP packets. A port in this mode forms an EtherChannel with a peer port that is in either auto or desirable PAgP mode.
- **Auto**—This (default) PAgP mode places a LAN port into a passive negotiating state, in which the port responds to PAgP packets it receives but does not initiate PAgP negotiation. A port in this mode forms an EtherChannel with a peer port that is in desirable PAgP mode (only).

PAgP, when enabled as an L2 link, is enabled on the physical interface (only). Optionally, you can change the PAgP mode from the default autonegotiation mode, as follows:

```
Switch(config)# interface GigabitEthernet8/1
Switch(config-if)# channel-protocol pagp
Switch(config-if)# channel-group 15 mode desirable
```

Alternatively, EtherChannels can be negotiated with the IEEE 802.3ad Link Aggregation Control Protocol (LACP). LACP similarly allows a switch to negotiate an automatic bundle by sending LACP packets to the peer. LACP supports two channel negotiation modes:

- **Active**—This LACP mode places a port into an active negotiating state, in which the port initiates negotiations with other ports by sending LACP packets. A port in this mode forms a bundle with a peer port that is in either active or passive LACP mode.
- **Passive**—This (default) LACP mode places a port into a passive negotiating state, in which the port responds to LACP packets it receives but does not initiate LACP negotiation. A port in this mode forms a bundle with a peer port that is in active LACP mode (only).

Similar to PAgP, LACP requires only a single command on the physical interface when configured as an L2 link. Optionally, you can change the LACP mode from the default passive negotiation mode, as follows:

```
Switch(config)# interface GigabitEthernet8/2
Switch(config-if)# channel-protocol lacp
```

```
Switch(config-if)# channel-group 16 mode active
```

However, note that PAgP and LACP do not interoperate with each other; ports configured to use PAgP cannot form EtherChannels with ports configured to use LACP, nor can ports configured to use LACP form EtherChannels with ports configured to use PAgP.

EtherChannel plays a critical role in provisioning network link redundancy, especially at the campus distribution and core layers. Furthermore, an evolution of EtherChannel technology plays a key role in Cisco VSS, which is discussed in the following section.

Cisco Virtual Switching System

The Cisco Catalyst 6500 Virtual Switching System (VSS) represents a major leap forward in device and network availability technologies, by combining many of the technologies that have been discussed thus far into a single, integrated system. VSS allows for the combination of two switches into a single, logical network entity from the network control plane and management perspectives. To the neighboring devices, the VSS appears as a single, logical switch or router.

Within the VSS, one chassis is designated as the active virtual switch and the other is designated as the standby virtual switch. All control plane functions, Layer 2 protocols, Layer 3 protocols, and software data path are centrally managed by the active supervisor engine of the active virtual switch chassis. The supervisor engine on the active virtual switch is also responsible for programming the hardware forwarding information onto all the distributed forwarding cards (DFCs) across the entire Cisco VSS as well as the policy feature card (PFC) on the standby virtual switch supervisor engine.

From the data plane and traffic forwarding perspectives, both switches in the VSS actively forward traffic. The PFC on the active virtual switch supervisor engine performs central forwarding lookups for all traffic that ingresses the active virtual switch, whereas the PFC on the standby virtual switch supervisor engine performs central forwarding lookups for all traffic that ingresses the standby virtual switch.

The first step in creating a VSS is to define a new logical entity called the virtual switch domain, which represents both switches as a single unit. Because switches can belong to one or more switch virtual domains, a unique number must be used to define each switch virtual domain, as [Example 3-3](#) demonstrates.

Example 3-3 VSS Virtual Domain Configuration

```
VSS-sw1(config)#switch virtual domain 100
```

Domain ID 100 config will take effect only
after the exec command 'switch convert mode virtual' is issued

```
VSS-sw1(config-vs-domain)#switch 1
```



Note

A corresponding set of commands must be configured on the second switch, with the difference being that *switch 1* becomes *switch 2*. However, the switch virtual domain number must be identical (in this example, 100).

Additionally, to bond the two chassis together into a single, logical node, special signaling and control information must be exchanged between the two chassis in a timely manner. To facilitate this information exchange, a special link is needed to transfer both data and control traffic between the peer chassis. This link is referred to as the virtual switch link (VSL). The VSL, formed as an EtherChannel interface, can comprise links ranging from one to eight physical member ports, as shown by [Example 3-4](#).

Example 3-4 VSL Configuration and VSS Conversion

```

VSS-sw1(config)#interface port-channel 1
VSS-sw1(config-if)#switch virtual link 1
VSS-sw1(config-if)#no shut
VSS-sw1(config-if)#exit
VSS-sw1(config)#interface range tenGigabitEthernet 5/4 - 5
VSS-sw1(config-if-range)#channel-group 1 mode on
VSS-sw1(config-if-range)#no shut
VSS-sw1(config-if-range)#exit
VSS-sw1(config)#exit
VSS-sw1#switch convert mode virtual

```

This command converts all interface names to naming convention *interface-type switch-number/slot/port*, saves the running configuration to the startup configuration, and reloads the switch.

```

Do you want to proceed? [yes/no]: yes
Converting interface names
Building configuration...
[OK]
Saving converted configurations to bootflash ...
[OK]


```

**Note**

As previously discussed, a corresponding set of commands must be configured on the second switch, with the difference being that *switch virtual link 1* becomes *switch virtual link 2*. Additionally, *port-channel 1* becomes *port-channel 2*.


VSL links carry two types of traffic: the VSS control traffic and normal data traffic. [Figure 3-15](#) illustrates the virtual switch domain and the VSL.

Figure 3-15 Virtual Switch Domain and Virtual Switch Link




Furthermore, VSS allows for an additional addition to EtherChannel technology: multi-chassis EtherChannel (MEC). Before VSS, EtherChannels were restricted to reside within the same physical switch. However, in a VSS environment, the two physical switches form a single logical network entity, and therefore EtherChannels can be extended across the two physical chassis, forming an MEC.


Thus, MEC allows for an EtherChannel bundle to be created across two separate physical chassis (although these two physical chassis are operating as a single, logical entity), as shown in [Figure 3-16](#).

Figure 3-16 Multi-Chassis EtherChannel Topology


Therefore, MEC allows all the dual-homed connections to and from the upstream and downstream devices to be configured as EtherChannel links, as opposed to individual links. From a configuration standpoint, the commands to form a MEC are the same as a regular EtherChannel; they are simply applied to interfaces that reside on two separate physical switches, as shown in [Figure 3-17](#).

Figure 3-17 MEC--Physical and Logical Campus Network Blocks

As a result, MEC links allow for implementation of network designs where true Layer 2 multipathing can be implemented without the reliance on Layer 2 redundancy protocols such as STP, as shown in [Figure 3-18](#).

Figure 3-18 STP Topology and VSS Topology

The advantage of VSS over STP is highlighted further by comparing [Figure 3-19](#), which shows a full campus network design using VSS, with [Figure 3-9](#), which shows a similar campus network design using STP.

Figure 3-19 VSS Campus Network Design

The ability to remove physical loops from the topology, and no longer be dependent on spanning tree, is one of the significant advantages of the virtual switch design. However, it is not the only difference. The virtual switch design allows for a number of fundamental changes to be made to the configuration and operation of the distribution block. By simplifying the network topology to use a single virtual distribution switch, many other aspects of the network design are either greatly simplified or, in some cases, no longer necessary.

Furthermore, network designs using VSS can be configured to converge in under 200 ms, which is 250 times faster than STP.

L3 Network Availability Technologies

L3 network availability technologies that particularly relate to TelePresence network design include the following:

- Hot Standby Router Protocol (HSRP)
- Virtual Router Redundancy Protocol (VRRP)

- Gateway Load Balancing Protocol (GLBP)
- IP Event Dampening

Hot Standby Router Protocol

Cisco HSRP is the first of three First Hop Redundancy Protocols (FHRPs) discussed in this chapter (the other two being VRRP and GLBP). A FHRP provides increased availability by allowing for transparent failover of the first-hop IP router, also known as the default gateway (for endpoint devices).

HSRP is used in a group of routers for selecting an active router and a standby router. In a group of router interfaces, the active router is the router of choice for routing packets; the standby router is the router that takes over when the active router fails or when preset conditions are met.

Endpoint devices, or IP hosts, have an IP address of a single router configured as the default gateway. When HSRP is used, the HSRP virtual IP address is configured as the host default gateway instead of the actual IP address of the router.

When HSRP is configured on a network segment, it provides a virtual MAC address and an IP address that is shared among a group of routers running HSRP. The address of this HSRP group is referred to as the virtual IP address. One of these devices is selected by the HSRP to be the active router. The active router receives and routes packets destined for the MAC address of the group.


HSRP detects when the designated active router fails, at which point a selected standby router assumes control of the MAC and IP addresses of the hot standby group. A new standby router is also selected at that time.

HSRP uses a priority mechanism to determine which HSRP configured router is to be the default active router. To configure a router as the active router, you assign it a priority that is higher than the priority of all the other HSRP-configured routers. The default priority is 100, so if just one router is configured to have a higher priority, that router is the default active router.

Devices that are running HSRP send and receive multicast UDP-based hello messages to detect router failure and to designate active and standby routers. When the active router fails to send a hello message within a configurable period of time, the standby router with the highest priority becomes the active router. The transition of packet forwarding functions between routers is completely transparent to all hosts on the network.

Multiple hot standby groups can be configured on an interface, thereby making fuller use of redundant routers and load sharing.

[Figure 3-20](#) shows a network configured for HSRP. By sharing a virtual MAC address and IP address, two or more routers can act as a single virtual router. The virtual router does not physically exist but represents the common default gateway for routers that are configured to provide backup to each other. All IP hosts are configured with IP address of the virtual router as their default gateway. If the active router fails to send a hello message within the configurable period of time, the standby router takes over and responds to the virtual addresses and becomes the active router, assuming the active router duties.

Figure 3-20 HSRP Topology

HSRP also supports object tracking, such that the HSRP priority of a router can dynamically change when an object that is being tracked goes down. Examples of objects that can be tracked are the line protocol state of an interface or the reachability of an IP route. If the specified object goes down, the HSRP priority is reduced.

Furthermore, HSRP supports SSO awareness, such that HSRP can alter its behavior when a router with redundant route processors (RPs) are configured in SSO redundancy mode. When an RP is active and the other RP is standby, SSO enables the standby RP to take over if the active RP fails.

With this functionality, HSRP SSO information is synchronized to the standby RP, allowing traffic that is sent using the HSRP virtual IP address to be continuously forwarded during a switchover without a loss of data or a path change. Additionally, if both RPs fail on the active HSRP router, the standby HSRP router takes over as the active HSRP router.

**Note**

SSO awareness for HSRP is enabled by default when the redundancy mode of operation of the RP is set to SSO, as was shown in [Non-Stop Forwarding with Stateful Switch Over, page 3-7](#).

Example 3-5 demonstrates the HSRP configuration that can be used on the LAN interface of the active router from [Figure 3-20](#). Each HSRP group on a given subnet requires a unique number; in this example, the HSRP group number is set to 10. The IP address of the virtual router (which is what each IP host on the network uses as a default gateway address) is set to 172.16.128.3. The HSRP priority of this router has been set to 105 and preemption has been enabled on it; preemption allows for the router to immediately take over as the virtual router (provided it has the highest priority on the segment). Finally, object tracking has been configured, such that should the line protocol state of interface Serial0/1 go down (the WAN link for the active router, which is designated as object-number 110), the HSRP priority for this interface dynamically decrements (by a value of 10, by default).

Example 3-5 HSRP Example

```

track 110 interface Serial0/1 line-protocol
!
interface GigabitEthernet0/0
ip address 172.16.128.1 255.255.255.0
standby 10 ip 172.16.128.3
standby 10 priority 105 preempt
standby 10 track 110
!

```

Because HSRP was the first FHRP and because it was invented by Cisco, it is Cisco-proprietary. However, to support multi-vendor interoperability, aspects of HSRP were standardized in the Virtual Router Redundancy Protocol (VRRP), which is discussed next.


Virtual Router Redundancy Protocol

VRRP, defined in RFC 2338, is an FHRP very similar to HSRP, but is able to support multi-vendor environments. A VRRP router is configured to run the VRRP protocol in conjunction with one or more other routers attached to a LAN. In a VRRP configuration, one router is elected as the virtual router master, with the other routers acting as backups in case the virtual router master fails.

VRRP enables a group of routers to form a single virtual router. The LAN clients can then be configured with the virtual router as their default gateway. The virtual router, representing a group of routers, is also known as a VRRP group.

Figure 3-21 shows a LAN topology in which VRRP is configured. In this example, two VRRP routers (routers running VRRP) comprise a virtual router. However, unlike HSRP, the IP address of the virtual router is the same as that configured for the LAN interface of the virtual router master; in this example, 172.16.128.1.

Figure 3-21 VRRP Topology



228679

Router A assumes the role of the virtual router master and is also known as the IP address owner, because the IP address of the virtual router belongs to it. As the virtual router master, Router A is responsible for forwarding packets sent to this IP address. Each IP host on the subnet is configured with the default gateway IP address of the virtual route master, in this case 172.16.128.1.

Router B, on the other hand, functions as a virtual router backup. If the virtual router master fails, the router configured with the higher priority becomes the virtual router master and provides uninterrupted service for the LAN hosts. When Router A recovers, it becomes the virtual router master again.

Additionally, like HSRP, VRRP supports object tracking, preemption, and SSO awareness.

**Note**

SSO awareness for VRRP is enabled by default when the redundancy mode of operation of the RP is set to SSO, as was shown in [Non-Stop Forwarding with Stateful Switch Over, page 3-7](#).

[Example 3-6](#) shows a VRRP configuration that can be used on the LAN interface of the virtual router master from [Figure 3-21](#). Each VRRP group on a given subnet requires a unique number; in this example, the VRRP group number is set to 10. The virtual IP address is set to the actual LAN interface address, designating this router as the virtual router master. The VRRP priority of this router has been set to 105. Unlike HSRP, preemption for VRRP is enabled by default. Finally, object tracking has been configured, such that should the line protocol state of interface Serial0/1 go down (the WAN link for this router, which is designated as object-number 110), the VRRP priority for this interface dynamically decrements (by a value of 10, by default).

Example 3-6 VRRP Example

```
!  
track 110 interface Serial0/1 line-protocol  
!  
interface GigabitEthernet0/0  
ip address 172.16.128.1 255.255.255.0  
vrrp 10 ip 172.16.128.1  
vrrp 10 priority 105  
vrrp 10 track 110  
!
```

A drawback to both HSRP and VRRP is that the standby/backup router is not used to forward traffic, and as such wastes both available bandwidth and processing capabilities. This limitation can be worked around by provisioning two complementary HSRP/VRRP groups on each LAN subnet, with one group having the left router as the active/master and the other group having the right router as the active/master router. Then, approximately half of the hosts are configured to use the virtual IP address of one HSRP/VRRP group, and the remaining hosts are configured to use the virtual IP address of the second group. This requires additional operational and management complexity. To improve the efficiency of these FHRP models without such additional complexity, GLBP can be used, which is discussed next.

Gateway Load Balancing Protocol


Cisco GLBP improves the efficiency of FHRP protocols by allowing for automatic load balancing of the default gateway. The advantage of GLBP is that it additionally provides load balancing over multiple routers (gateways) using a single virtual IP address and multiple virtual MAC addresses per GLBP group (in contrast, both HSRP and VRRP used only one virtual MAC address per HSRP/VRRP group). The forwarding load is shared among all routers in a GLBP group rather than being handled by a single router while the other routers stand idle. Each host is configured with the same virtual IP address, and all routers in the virtual router group participate in forwarding packets.

Members of a GLBP group elect one gateway to be the active virtual gateway (AVG) for that group. Other group members provide backup for the AVG in the event that the AVG becomes unavailable. The function of the AVG is that it assigns a virtual MAC address to each member of the GLBP group. Each gateway assumes responsibility for forwarding packets sent to the virtual MAC address assigned to it by the AVG. These gateways are known as active virtual forwarders (AVFs) for their virtual MAC address.

The AVG is also responsible for answering Address Resolution Protocol (ARP) requests for the virtual IP address. Load sharing is achieved by the AVG replying to the ARP requests with different virtual MAC addresses (corresponding to each gateway router).

In [Figure 3-22](#), *Router A* is the AVG for a GLBP group, and is primarily responsible for the virtual IP address 172.16.128.3; however, *Router A* is also an AVF for the virtual MAC address 0007.b400.0101. *Router B* is a member of the same GLBP group and is designated as the AVF for the virtual MAC address 0007.b400.0102. All hosts have their default gateway IP addresses set to the virtual IP address of 172.16.128.3. However, when these use ARP to determine the MAC of this virtual IP address, *Host A* and *Host C* receive a gateway MAC address of 0007.b400.0101 (directing these hosts to use *Router A* as their default gateway), but *Host B* and *Host D* receive a gateway MAC address 0007.b400.0102 (directing these hosts to use *Router B* as their default gateway). In this way, the gateway routers automatically load share.

Figure 3-22 GLBP Topology



If *Router A* becomes unavailable, *Hosts A* and *C* do not lose access to the WAN because *Router B* assumes responsibility for forwarding packets sent to the virtual MAC address of *Router A*, and for responding to packets sent to its own virtual MAC address. *Router B* also assumes the role of the AVG for the entire GLBP group. Communication for the GLBP members continues despite the failure of a router in the GLBP group.

Additionally, like HSRP and VRRP, GLBP supports object tracking, preemption, and SSO awareness.

**Note**

SSO awareness for GLBP is enabled by default when the route processor's redundancy mode of operation is set to SSO, as was shown in [Non-Stop Forwarding with Stateful Switch Over, page 3-7](#).

However, unlike the object tracking logic used by HSRP and VRRP, GLBP uses a weighting scheme to determine the forwarding capacity of each router in the GLBP group. The weighting assigned to a router in the GLBP group can be used to determine whether it forwards packets and, if so, the proportion of hosts in the LAN for which it forwards packets. Thresholds can be set to disable forwarding when the weighting for a GLBP group falls below a certain value; when it rises above another threshold, forwarding is automatically re-enabled.

GLBP group weighting can be automatically adjusted by tracking the state of an interface within the router. If a tracked interface goes down, the GLBP group weighting is reduced by a specified value. Different interfaces can be tracked to decrement the GLBP weighting by varying amounts.

[Example 3-7](#) shows a GLBP configuration that can be used on the LAN interface of the AVG from [Figure 3-22](#). Each GLBP group on a given subnet requires a unique number; in this example, the GLBP group number is set to 10. The virtual IP address for the GLBP group is set to 172.16.128.3. The GLBP priority of this interface has been set to 105, and like HSRP, preemption for GLBP must be explicitly enabled (if desired). Finally, object tracking has been configured, such that should the line protocol state of interface Serial0/1 go down (the WAN link for this router, which is designated as object-number 110), the GLBP priority for this interface dynamically decrements (by a value of 10, by default).

Example 3-7 GLBP Example

```
!  
track 110 interface Serial0/1 line-protocol  
!  
interface GigabitEthernet0/0  
ip address 172.16.128.1 255.255.255.0  
glbp 10 ip 172.16.128.3  
glbp 10 priority 105  
glbp 10 preempt  
glbp 10 weighting track 110  
!
```

Having concluded an overview of these FHRPs, a discussion of another type of L3 network availability feature, IP Event Dampening, follows.

IP Event Dampening

Whenever the line protocol of an interface changes state, or flaps, routing protocols are notified of the status of the routes that are affected by the change in state. Every interface state change requires all affected devices in the network to recalculate best paths, install or remove routes from the routing tables, and then advertise valid routes to peer routers. An unstable interface that flaps excessively can cause other devices in the network to consume substantial amounts of system processing resources and cause routing protocols to lose synchronization with the state of the flapping interface.

The IP Event Dampening feature introduces a configurable exponential decay mechanism to suppress the effects of excessive interface flapping events on routing protocols and routing tables in the network. This feature allows the network administrator to configure a router to automatically identify and selectively dampen a local interface that is flapping. Dampening an interface removes the interface from the network until the interface stops flapping and becomes stable.

Configuring the IP Event Dampening feature improves convergence times and stability throughout the network by isolating failures so that disturbances are not propagated, which reduces the use of system processing resources by other devices in the network and improves overall network stability.

IP Event Dampening uses a series of administratively-defined thresholds to identify flapping interfaces, to assign penalties, to suppress state changes (if necessary), and to make stabilized interfaces available to the network. These thresholds are as follows:

- **Suppress threshold**—The value of the accumulated penalty that triggers the router to dampen a flapping interface. The flapping interface is identified by the router and assigned a penalty for each up and down state change, but the interface is not automatically dampened. The router tracks the penalties that a flapping interface accumulates. When the accumulated penalty reaches the default or preconfigured suppress threshold, the interface is placed in a dampened state. The default suppress threshold value is 2000.
- **Half-life period**—Determines how fast the accumulated penalty can decay exponentially. When an interface is placed in a dampened state, the router monitors the interface for additional up and down state changes. If the interface continues to accumulate penalties and the interface remains in the suppress threshold range, the interface remains dampened. If the interface stabilizes and stops flapping, the penalty is reduced by half after each half-life period expires. The accumulated penalty is reduced until the penalty drops to the reuse threshold. The default half-life period timer is five seconds.
- **Reuse threshold**—When the accumulated penalty decreases until the penalty drops to the reuse threshold, the route is unsuppressed and made available to the other devices on the network. The default value is 1000 penalties.
- **Maximum suppress time**—The maximum suppress time represents the maximum amount of time an interface can remain dampened when a penalty is assigned to an interface. The default maximum penalty timer is 20 seconds.

IP Event Dampening is configured on a per-interface basis (where default values are used for each threshold) as follows:

```
interface FastEthernet0/0
dampening
```

IP Event Dampening can be complemented with the use of route summarization, on a per-routing protocol basis, to further compartmentalize the effects of flapping interfaces and associated routes.

Operational Availability Technologies

As has been shown, the predominant way that availability of a network can be improved is to improve its MTBF by using devices that have redundant components and by engineering the network itself to be as redundant as possible, leveraging many of the technologies discussed in the previous sections.

However, glancing back to the general availability formula from [Figure 3-1](#), another approach to improving availability is to reduce MTTR. Reducing MTTR is primarily a factor of operational resiliency.

MTTR operations can be significantly improved in conjunction with device and network redundant design. Specifically, the ability to make changes, upgrade software, and replace or upgrade hardware in a production network is extensively improved because of the implementation of device and network redundancy. The ability to upgrade individual devices without taking them out of service is based on having internal component redundancy complemented with the system software capabilities. Similarly,

by having dual active paths through redundant network devices designed to converge in sub-second timeframes, it is possible to schedule an outage event on one element of the network and allow it to be upgraded and then brought back into service with minimal or no disruption to the network as a whole.

MTTR can also be improved by reducing the time required to perform any of the following operations:

- Failure detection
- Notification
- Fault diagnosis
- Dispatch/arrival
- Fault repair

Technologies that can help automate and streamline these operations include the following:

- Cisco General Online Diagnostics (GOLD)
- Cisco IOS Embedded Event Manager (EEM)
- Cisco In Service Software Upgrade (ISSU)
- Online Insertion and Removal (OIR)

This section briefly introduces each of these technologies.

Cisco Generic Online Diagnostics

Cisco GOLD defines a common framework for diagnostic operations for Cisco IOS Software-based products. GOLD has the objective of checking the health of all hardware components and verifying the proper operation of the system data plane and control plane at boot time, as well as run-time.

GOLD supports the following:

- Bootup tests (includes online insertion)
- Health monitoring tests (background non-disruptive)
- On-demand tests (disruptive and non-disruptive)
- User scheduled tests (disruptive and non-disruptive)
- Command-line interface (CLI) access to data via a management interface

GOLD, in conjunction with several of the technologies previously discussed, can reduce device failure detection time.

Cisco IOS Embedded Event Manager

The Cisco IOS EEM offers the ability to monitor device hardware, software, and operational events and take informational, corrective, or any desired action, including sending an e-mail alert, when the monitored events occur or when a threshold is reached.

EEM can notify a network management server and/or an administrator (via e-mail) when an event of interest occurs. Events that can be monitored include the following:

- Application-specific events
- CLI events
- Counter/interface-counter events

- Object-tracking events
- Online insertion and removal events
- Resource events
- GOLD events
- Redundancy events
- Simple Network Management Protocol (SNMP) events
- Syslog events
- System manager/system monitor events
- IOS watchdog events
- Timer events

Capturing the state of network devices during such situations can be helpful in taking immediate recovery actions and gathering information to perform root-cause analysis, reducing fault detection and diagnosis time. Notification times are reduced by having the device send e-mail alerts to network administrators. Furthermore, availability is also improved if automatic recovery actions are performed without the need to fully reboot the device.

Cisco In Service Software Upgrade

Cisco ISSU provides a mechanism to perform software upgrades and downgrades without taking a switch out of service. ISSU leverages the capabilities of NSF and SSO to allow the switch to forward traffic during supervisor IOS upgrade (or downgrade). With ISSU, the network does not re-route and no active links are taken out of service. ISSU thereby expedites software upgrade operations.

Online Insertion and Removal

OIR allows line cards to be added to a device without affecting the system. Additionally, with OIR, line cards can be exchanged without losing the configuration. OIR thus expedites hardware repair and/or replacement operations.

Summary

Availability was shown to be a factor of two components: the mean time between failures (MTBF) and the mean time to repair (MTTR) such failures. Availability can be improved by increasing MTBF (which is primarily a function of device and network resiliency/redundancy), or by reducing MTTR (which is primarily a function of operational resiliency).

Device availability technologies were discussed, including Cisco Catalyst StackWise/StackWise Plus technologies, which provide 1:N control plane redundancy to Cisco Catalyst 3750G/3750E switches, as well as NSF with SSO, which similarly provides hot standby redundancy to network devices with multiple route processors.

Network availability technologies were also discussed, beginning with Layer 2 technologies, such as spanning tree protocols, trunking protocols, EtherChannel protocols, and Cisco VSS. Additionally, Layer 3 technologies, such as HSRP, VRRP, GLBP, and IP Event dampening were introduced.

Finally, operational availability technologies were introduced to show how availability can be improved by automating and streamlining MTTR operations, including GOLD, EEM, ISSU, and OIR.



CHAPTER 4

Medianet QoS Design Considerations

This document provides an overview of Quality of Service (QoS) tools and design recommendations relating to an enterprise medianet architecture and includes high-level answers to the following:

- Why is Cisco providing new QoS design guidance at this time?
- What is Cisco's Quality of Service toolset?
- How can QoS be optimally deployed for enterprise medianets?

QoS has proven itself a foundational network infrastructure technology required to support the transparent convergence of voice, video, and data networks. Furthermore, QoS has also been proven to complement and strengthen the overall security posture of a network. However, business needs continue to evolve and expand, and as such, place new demands on QoS technologies and designs. This document examines current QoS demands and requirements within an enterprise medianet and presents strategic design recommendations to address these needs.

Drivers for QoS Design Evolution

There are three main sets of drivers pressuring network administrators to reevaluate their current QoS designs (each is discussed in the following sections):

- New applications and business requirements
- New industry guidance and best practices
- New platforms and technologies

New Applications and Business Requirements

Media applications—particularly video-oriented media applications—are exploding over corporate networks, exponentially increasing bandwidth utilization and radically shifting traffic patterns. For example, according to recent studies, global IP traffic will nearly double every two years through 2012¹ and the sum of all forms of video will account for close to 90 percent of consumer traffic by 2012².

1. Cisco Visual Networking Index—Forecast and Methodology, 2007-2012

http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360_ns827_Networking_Solutions_White_Paper.html

2. Approaching the Zettabyte Era

http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481374_ns827_Networking_Solutions_White_Paper.html

Businesses recognize the value that media applications—particularly video-based collaborative applications—bring to the enterprise, including:

- Increasing productivity
- Improving the quality of decision making
- Speeding time-to-market
- Facilitating knowledge sharing
- Fueling innovation
- Reducing travel time and expenses
- Protecting the environment

Corresponding to these values and benefits of media applications, there are several business drivers behind media application growth, including:

- Evolution of video applications
- Transition to high-definition media
- Explosion of media
- Phenomena of social networking
- Emergence of “bottoms-up” media applications
- Convergence within media applications
- Globalization of the workforce
- Pressures to go green


These business drivers briefly described in the following sections.

The Evolution of Video Applications

When the previous Cisco Enterprise QoS Design Guide was published (in 2003), there were basically only two broad types of video applications deployed over enterprise networks:

- Interactive video—Generally describes H.323-based collaborative video applications (typically operating at 384 kbps or 768 kbps); video flows were bi-directional and time-sensitive.
- Streaming video—Generally describes streaming or video-on-demand (VoD) applications; video flows were unidirectional (either unicast or multicast) and were not time-sensitive (due to significant application buffering).

However, at the time of writing this document (2009), video applications have evolved considerably, as illustrated in [Figure 4-1](#).

Figure 4-1 Video Application Evolution

Consider first the streaming video branch—the earliest sets of video applications were VoD streams to the desktop. VoD streams can include pre-recorded content such as employee communications, training, e-learning, and social-interaction content. Today, due to the ease of content creation, on-demand content may either be professionally-produced (top-down) or self-produced (bottom-up). It is important to also note that not all VoD content is necessarily business-related, as non-business, entertainment-oriented content is often widely available for on-demand video viewing.

VoD applications soon expanded to include the development and support of “live” or “broadcast” video streams to the desktop. Broadcast streams may include company meetings, special events, internal corporate announcements or similar content. As such, broadcast streaming video content is typically professionally-produced, top-down content.

Thereafter, with the proliferation of flat-screen digital displays, it became increasingly apparent that the desktop is not the only display option for streaming video. Thus, digital signage began to emerge as another streaming video application (for both on-demand and broadcast video streams). Digital signage refers to centrally-managed publishing solutions for delivering digital media to networked displays. For example, Cisco offers a Digital Media Player and an enterprise TV solution that works in conjunction with its Digital Media System to support a comprehensive digital signage solution. Digital signage can be used to broadcast internal information, such as sharing up-to-date schedules and news where people need it most or providing realtime location and directional guidance. Additionally, digital signage is an effective tool for marketing, helping companies to promote products and services directly to customers.

Around the same time that digital signage was being developed, the advantages that IP brought to video were being gradually being applied to the video surveillance market. These advantages include the ability to forward live video streams to local or remote control centers for observation and efficient processing. Cisco offers comprehensive IP surveillance (IPVS) solutions, including IP cameras, hybrid analog-to-digital video gateways (to facilitate transitioning from closed-circuit TV surveillance solutions to IPVS), and IPVS management applications. Interestingly, video surveillance has a unique degree of interactivity not found in any other streaming video application, namely, that of having an observer “interact” with the video stream by sending control information to the transmitting video camera, for instance, to track an event-in-progress.

On the interactive video side of the video application hemisphere, there has also been considerable application evolution. Basic video conferencing applications, which were initially dedicated room-based units, evolved into software-based PC applications. The factors behind this shift from room-based hardware to PC-based software were two-fold:

- The convenience of immediate desktop collaboration (rather than having to book or hunt for an available video-conferencing enabled room).

- The availability of inexpensive Webcams. Desktop video conferencing may be utilized on a one-to-one basis or may support a few participants simultaneously.

Once video conferencing moved to software, a whole new range of communication possibilities opened up, which morphed desktop video conferencing applications into multimedia collaboration applications. Multimedia collaboration applications, including Cisco Unified Personal Communicator (CUPC) and Cisco WebEx, share not only voice and video, but also data applications, such as instant messaging, document and presentation sharing, application sharing, and other integrated multimedia features.

However, not all interactive video migrated to the desktop. Room-based video conferencing solutions continued to evolve and leveraged advances in high-definition video and audio, leading to solutions like Cisco TelePresence. Additionally, application sharing capabilities—borrowed from multimedia conferencing applications—were added to these high-definition room-based video conferencing solutions.


And video application evolution doesn't end here, but will continue to expand and morph over time as new demands and technologies emerge.

The Transition to High-Definition Media

One of the reasons traditional room-to-room video conferencing and desktop Webcam-style video conferencing are sometimes questioned as less than effective communications systems is the reliance on low-definition audio and video formats.

On the other hand, high-definition interactive media applications, like Cisco TelePresence, demonstrate how high-definition audio and video can create an more effective remote collaboration experience, where meeting participants actually feel like they are in the same meeting room. Additionally, IP video surveillance cameras are migrating to high-definition video in order to have the digital resolutions needed for new functions, such as pattern recognition and intelligent event triggering based on motion and visual characteristics. Cisco fully expects other media applications to migrate to high-definition in the near future, as people become accustomed to the format in their lives as consumers, as well as the experiences starting to appear in the corporate environment.

High-definition media formats transmitted over IP networks create unique challenges and demands on the network that need to be planned for. For example, [Figure 4-2](#) contrasts the behavior of VoIP as compared to high definition video at the packet level.

Figure 4-2 VoIP versus High-Definition Video—At the Packet Level

The network demands of high-definition video include not only radically more bandwidth, but also significantly higher transmission reliability, as compared to standard-definition video applications.

The Explosion of Media

Another factor driving the demand for video on IP networks is the sheer explosion of media content. The barriers to media production, distribution, and viewing have been dramatically lowered. For example, five to ten years ago video cameras became so affordable and prevalent that just about anyone could buy one and become an amateur video producer. Additionally, video cameras are so common that almost every cell phone, PDA, laptop, and digital still camera provides a relatively high-quality video capture capability. However, until recently, it was not that easy to be a distributor of video content, as distribution networks were not common.

Today, social networking sites like YouTube, MySpace and many others appearing every day and have dramatically lowered the barrier to video publishing to the point where anyone can do it. Video editing software is also cheap and easy to use. Add to that a free, global video publishing and distribution system and essentially anyone, anywhere can be a film studio. With little or no training, people are making movie shorts that rival those of dedicated video studios.

The resulting explosion of media content is now the overwhelming majority of consumer network traffic and is quickly crossing over to corporate networks. The bottom line is there are few barriers left to inhibit video communication and so this incredibly effective medium is appearing in new and exciting applications every day.

The Phenomena of Social Networking

Social networking started as a consumer phenomenon, with people producing and sharing rich media communications such as blogs, photos, and videos. When considering the effect it may have on corporate networks, some IT analysts believed social networking would remain a consumer trend, while others believed the appearance in corporate networks was inevitable.

Skeptics look at social networking sites like YouTube, MySpace, and others and see them as fads primarily for the younger population. However, looking beyond the sites themselves, it is important to understand the new forms of communication and information sharing they are enabling. For example, with consumer social networking, typically people are sharing information about themselves, about subjects they have experience in, and interact with others in real-time who have similar interests. In the workplace, we already see parallel activities, because the same types of communication and information sharing are just as effective.

The corporate directory used to consist of employee names, titles, and phone numbers. Companies embracing social networking are adding to that skillsets and experience, URL links to shared work spaces, blogs, and other useful information. The result is a more productive and effective workforce that can adapt and find the skillsets and people needed to accomplish dynamic projects.

Similarly, in the past information was primarily shared via text documents, E-mail, and slide sets. Increasingly, we see employees filming short videos to share best practices with colleagues, provide updates to peers and reports, and provide visibility into projects and initiatives. Why have social networking trends zeroed in on video as the predominant communication medium? Simple: video is the most effective medium. People can show or demonstrate concepts much more effectively and easily using video than any other medium.

Just as a progression occurred from voice exchange to text, to graphics, and to animated slides, video will start to supplant those forms of communications. Think about the time it would take to create a good set of slides describing how to set up or configure a product. Now how much easier would it be just to film someone actually doing it? That's just one of many examples where video is supplanting traditional communication formats.

Internally, Cisco has witnessed the cross-over of such social networking applications into the workplace, with applications like Cisco Vision (C-Vision). C-Vision started as an ad hoc service by several employees, providing a central location for employees to share all forms of media with one another, including audio and video clips. Cisco employees share information on projects, new products, competitive practices, and many other subjects. The service was used by so many employees that Cisco's IT department had to assume ownership and subsequently scaled the service globally within Cisco. The result is a service where employees can become more effective and productive, quickly tapping into each other's experiences and know-how, all through the effectiveness and simplicity of media.

The Emergence of Bottom-Up Media Applications

As demonstrated in the C-Vision example, closely related to the social-networking aspect of media applications is the trend of users driving certain types of media application deployments within the enterprise from the bottom-up (in other words, the user base either demands or just begins to use a given media application with or without formal management or IT support). Bottom-up deployment patterns have been noted for many Web 2.0 and multimedia collaboration applications.

In contrast, company-sponsored video applications are pushed from the top-down (in other words, the management team decides and formally directs the IT department to support a given media application for their user base). Such top-down media applications may include Cisco TelePresence, digital signage, video surveillance, and live broadcast video meetings.

The combination of top-down and bottom-up media application proliferation places a heavy burden on the IT department as it struggles to cope with officially-supported and officially-unsupported, yet highly proliferated, media applications.


The Convergence Within Media Applications

Much like the integration of rich text and graphics into documentation, audio and video media continue to be integrated into many forms of communication. Sharing of information with E-mailed slide sets will gradually be replaced with E-mailed video clips. The audio conference bridge will be supplanted with the video-enabled conference bridge. Collaboration tools designed to link together distributed employees will increasingly integrate desktop video to bring teams closer together.

Cisco WebEx is a prime example of such integration, providing text, audio, instant messaging, application sharing, and desktop video conferencing easily to all meeting participants, regardless of their location. Instead of a cumbersome setup of a video conference call, applications such as CUPC and WebEx greatly simplify the process and video capability is added to the conference just as easily as any other type of media, such as audio.

The complexity that application presents to the network administrator relates to application classification: as media applications include voice, video, and data sub-components, the question of how to mark and provision a given media application becomes more difficult and blurry, as illustrated in Figure 4-3.

Figure 4-3 Media Application Convergence—Voice, Video, and Data Within an Application



For example, since Cisco WebEx has voice, video, and data sub-components, how should it be classified? As a voice application? As a video application? As a data application? Or is an altogether new application-class model needed to accommodate multimedia applications?

The Globalization of the Workforce

In the past, most companies focused on acquiring and retaining skilled and talented individuals in a single or few geographic locations. More recently, this focus has shifted to finding technology solutions to enable a geographically-distributed workforce to collaborate together as a team. This new approach enables companies to more flexibly harness talent “where it lives.”

Future productivity gains will be achieved by creating collaborative teams that span corporate boundaries, national boundaries, and geographies. Employees will collaborate with partners, research and educational institutions, and customers to create a new level of collective knowledge.

To do so, real-time multimedia collaboration applications are absolutely critical to the success of these virtual teams. Video offers a unique medium which streamlines the effectiveness of communications between members of such teams. For this reason, real-time interactive video will become increasingly prevalent, as will media integrated with corporate communications systems.

The Pressures to be Green

For many reasons, companies are seeking to reduce employee travel. Travel creates bottom line expenses, as well as significant productivity impacts while employees are in-transit and away from their usual working environments. Many solutions have emerged to assist with productivity while traveling, including wireless LAN hotspots, remote access VPNs, and softphones, all designed to keep the employee connected while traveling.

More recently companies are under increasing pressures to demonstrate environmental responsibility, often referred to as being “green.” On the surface, such initiatives may seem like a pop-culture trend that lacks tangible corporate returns. However, it is entirely possible to pursue green initiatives while simultaneously increasing productivity and lowering expenses.

Media applications, such as Cisco TelePresence, offer real solutions to remote collaboration challenges and have demonstrable savings as well. For example, during the first year of deployment, Cisco measured its usage of TelePresence in direct comparison to the employee travel that would otherwise have taken place. Cisco discovered that over 80,000 hours of meetings were held by TelePresence instead of physical travel, avoiding \$100 million of travel expenses, as well as over 30,000 tons of carbon emissions, the equivalent of removing over 10,000 vehicles off the roads for a period of one year.

Being green does not have to be a “tax;” but rather can improve productivity and reduce corporate expenses, offering many dimensions of return on investment, while at the same time sending significant messages to the global community of environmental responsibility.

Thus, having reviewed several key business drivers for evolving QoS designs, relevant industry guidance and best practices are discussed next.

New Industry Guidance and Best Practices

A second set of drivers behind QoS design evolution are advances in industry standards and guidance. Cisco has long advocated following industry standards and recommendations—whenever possible—when deploying QoS, as this simplifies QoS designs, extends QoS policies beyond an administrative domain, and improves QoS between administrative domains.

To the first point of simplifying QoS, there are 64 discrete Differentiated Services Code Point (DSCP) values to which IP packets can be marked. If every administrator were left to their own devices to arbitrarily pick-and-choose DSCP markings for applications, there would be a wide and disparate set of

marking schemes that would likely vary from enterprise to enterprise, perhaps even within an enterprise (such as department to department). However if industry standard marking values are used, then marking schemes become considerably simplified and consistent.

To the second point of extending QoS policies beyond an administrative domain, if an enterprise administrator wishes a specific type of Per-Hop Behavior (PHB)—which is the manner in which a packet marked to a given DSCP value is to be treated at each network node—they mark the packet according to the industry recommended marking value that corresponds to the desired PHB. Then, as packets are handed off to other administrative domains, such as service provider networks or partner networks, these packets continue to receive the desired PHB (provided that the SP or partner network is also following the same industry standards). Therefore, the PHB treatment is extended beyond the original administrative domain and thus the overall quality of service applied to the packet end-to-end is improved.

To the third point of improving QoS between administrative domains, as networks pass packets to adjacent administrative domains, sometimes their QoS policies differ. Nonetheless, the differences are likely to be minor, as compared to the scenario in which every administrator handled packets in an arbitrary, locally-defined fashion. Thus, the mapping of QoS policies is much easier to handle between domains, as these ultimately use many—if not most—of the same industry-defined PHBs.


However, there may be specific constraints, either financial, technical, or otherwise, that may preclude following industry standards 100% of the time. In such cases, administrators need to make careful decisions as to when and how to deviate from these standards and recommendations to best meet their specific objectives and constraints and to allow them maximum flexibility and consistency in the end-to-end scenarios described above.

Therefore, in line with the principle of following industry standards and recommendations whenever possible, it would be beneficial to briefly review some of the standards and recommendations most relevant to QoS design.

RFC 2474 Class Selector Code Points

The IETF RFC 2474 standard defines the use of 6 bits in the IPv4 and IPv6 Type of Service (ToS) byte, termed Differentiated Services Code Points (DSCP). Additionally, this standard introduces Class Selector codepoints to provide backwards compatibility for legacy (RFC 791) IP Precedence bits, as shown in [Figure 4-4](#).

Figure 4-4 The IP ToS Byte—IP Precedence Bits and DiffServ Extensions



Class Selectors, defined in RFC 2474, are not Per-Hop Behaviors per se, but rather were defined to provide backwards compatibility to IP Precedence. Each Class Selector corresponds to a given IP Precedence value, with its three least-significant bits set to 0. For example, IP Precedence 1 is referred to as Class Selector 1 (or DSCP 8), IP Precedence 2 is referred to as Class Selector 2 (or DSCP 16), and so on. [Table 4-1](#) shows the full table of IP Precedence to Class Selector mappings.

Table 4-1 *IP Precedence to Class Selector/DSCP Mappings*

IP Precedence Value	IP Precedence Name	IPP Binary Equivalent	Class Selector	CS Binary Equivalent	DSCP Value (Decimal)
0	Normal	000	CS0 ¹ /DF	000 000	0
1	Priority	001	CS1	001 000	8
2	Immediate	010	CS2	010 000	16
3	Flash	011	CS3	011 000	24
4	Flash-Override	100	CS4	100 000	32
5	Critical	101	CS5	101 000	40
6	Internetwork Control	110	CS6	110 000	48
7	Network Control	111	CS7	111 000	56


1. Class Selector 0 is a special case, as it represents the default marking value (defined in RFC 2474-Section 4.1); as such, it is not typically called Class Selector 0, but rather Default Forwarding or DF.

RFC 2597 Assured Forwarding Per-Hop Behavior Group

RFC 2597 defines four Assured Forwarding groups, denoted by the letters “AF” followed by two digits:

- The first digit denotes the AF class number and can range from 1 through 4 (these values correspond to the three most-significant bits of the codepoint or the IPP value that the codepoint falls under). Incidentally, the AF class number does not in itself represent hierarchy (that is, AF class 4 does not necessarily get any preferential treatment over AF class 1).
- The second digit refers to the level of drop precedence within each AF class and can range from 1 (lowest drop precedence) through 3 (highest drop precedence).

Figure 4-5 shows the Assured Forwarding PHB marking scheme.

Figure 4-5 *Assured Forwarding PHB Marking Scheme*

The three levels of drop precedence are analogous to the three states of a traffic light:

- Drop precedence 1, also known as the “conforming” state, is comparable to a green traffic light.
- Drop precedence 2, also known as the “exceeding” state, is comparable to a yellow traffic light (where a moderate allowance in excess of the conforming rate is allowed to prevent erratic traffic patterns).
- Drop precedence 3, also known as the “violating” state, is comparable to a red traffic light.

Packets within an AF class are always initially marked to drop precedence of 1 and can only be remarked to drop precedence 2 or 3 by a policer, which meters traffic rates and determines if the traffic is exceeding or violating a given traffic contract.

Then, for example, during periods of congestion on an RFC 2597-compliant node, packets remarked AF33 (representing the highest drop precedence for AF class 3) would be dropped more often than packets remarked AF32; in turn, packets remarked AF32 would be dropped more often than packets marked AF31.

The full set of AF PHBs are detailed in [Figure 4-6](#).

Figure 4-6 Assured Forwarding PHBs with Decimal and Binary Equivalents

	AF PHB			DSCP		
	Conforming DP	Exceeding DP	Violating DP			
AF Class 1	AF11	AF12	AF13	10 001 010	12 001 100	14 001 110
AF Class 2	AF21	AF22	AF23	18 010 010	20 010 100	22 010 110
AF Class 3	AF31	AF32	AF33	26 011 010	28 011 100	30 011 110
AF Class 4	AF41	AF42	AF43	34 100 010	36 100 100	38 100 110

226605

RFC 3246 An Expedited Forwarding Per-Hop Behavior

The Expedited Forwarding PHB is defined in RFC 3246. In short, the definition describes a strict-priority treatment for packets that have been marked to a DSCP value of 46 (101110), which is also termed Expedited Forwarding (or EF). Any packet marked 46/EF that encounters congestion at a given network node is to be moved to the front-of-the-line and serviced in a strict priority manner. It doesn't matter how such behavior is implemented—whether in hardware or software—as long as the behavior is met for the given platform at the network node.



Note

Incidentally, the RFC 3246 does not specify which application is to receive such treatment; this is open to the network administrator to decide, although the industry norm over the last decade has been to use the EF PHB for VoIP.

The EF PHB provides an excellent case-point of the value of standardized PHBs. For example, if a network administrator decides to mark his VoIP traffic to EF and service it with strict priority over his networks, he can extend his policies to protect his voice traffic even over networks that he does not have direct administrative control. He can do this by partnering with service providers and/or extranet partners who follow the same standard PHB and who thus continue to service his (EF marked) voice traffic with strict priority over their networks.

RFC 3662 A Lower Effort Per-Domain Behavior for Differentiated Services

While most of the PHBs discussed so far represent manners in which traffic may be treated preferentially, there are cases it may be desired to treat traffic deferentially. For example, certain types of non-business traffic, such as gaming, video-downloads, peer-to-peer media sharing, and so on might dominate network links if left unabated.

To address such needs, a Lower Effort Per-Domain Behavior is described in RFC 3662 to provide a less than Best Effort service to undesired traffic. Two things should be noted about RFC 3662 from the start:

- RFC 3662 is in the “informational” category of RFCs (not the standards track) and as such is not necessary to implement in order to be DiffServ standard-compliant.
- A Per-Domain Behavior (PDB) has a different and larger scope than a Per-Hop Behavior (PHB). A PDB does not require that undesired traffic be treated within a “less than Best Effort service” at necessarily every network node (which it would if this behavior were defined as a Per-Hop Behavior); rather, as long as one (or more) nodes within the administrative domain provide a “less than best effort service” to this undesired traffic class, the Per-Domain Behavior requirement has been met.

The reason a PDB is sufficient to provision this behavior, as opposed to requiring a PHB, is that the level of service is deferential, not preferential. To expand, when dealing with preferential QoS policies, sometimes it is said that “a chain of QoS policies is only as strong as the weakest link.” For example, if provisioning an EF PHB for voice throughout a network and only one node in the path does not have EF properly provisioned on it, then the overall quality of voice is (potentially) ruined. On the other hand, if the objective is to provide a deferential level of service, all one needs is a single weak link in the path in order to lower the overall quality of service for a given class. Thus, if only a single weak link is required per administrative domain, then a Per-Domain Behavior-rather than a Per-Hop Behavior-better suits the requirement.

The marking value recommended in RFC 3662 for less than best effort service (sometimes referred to as a Scavenger service) is Class Selector 1 (DSCP 8). This marking value is typically assigned and constrained to a minimally-provisioned queue, such that it will be dropped the most aggressively under network congestion scenarios.

Cisco's QoS Baseline

While the IETF DiffServ RFCs (discussed thus far) provided a consistent set of per-hop behaviors for applications marked to specific DSCP values, they never specified which application should be marked to which DiffServ Codepoint value. Therefore, considerable industry disparity existed in application-to-DSCP associations, which led Cisco to put forward a standards-based application marking recommendation in their strategic architectural QoS Baseline document (in 2002). Eleven different application classes were examined and extensively profiled and then matched to their optimal RFC-defined PHBs. The application-specific marking recommendations from Cisco's QoS Baseline of 2002 are summarized in [Figure 4-7](#).

Figure 4-7 Cisco's QoS Baseline Marking Recommendations

Application	L3 Classification		IETF
	PHB	DSCP	RFC
Routing	CS6	48	RFC 2474
Voice	EF	46	RFC 3246
Interactive Video	AF41	34	RFC 2597
Streaming Video	CS4	32	RFC 2474
Mission-Critical Data	AF31	26	RFC 2597
Call Signaling	CS3	24	RFC 2474
Transactional Data	AF21	18	RFC 2597
Network Management	CS2	16	RFC 2474
Bulk Data	AF11	10	RFC 2597
Best Effort	0	0	RFC 2474
Scavenger	CS1	8	RFC 2474

220199

**Note**

The previous Cisco Enterprise QoS SRND (version 3.3 from 2003) was based on Cisco's QoS Baseline; however, as will be discussed, newer RFCs have since been published that improve and expand on the Cisco QoS Baseline.

RFC 4594 Configuration Guidelines for DiffServ Classes

More than four years after Cisco put forward its QoS Baseline document, RFC 4594 was formally accepted as an informational RFC (in August 2006).

Before getting into the specifics of RFC 4594, it is important to comment on the difference between the IETF RFC categories of informational and standard. An informational RFC is an industry recommended best practice, while a standard RFC is an industry requirement. Therefore RFC 4594 is a set of formal DiffServ QoS configuration best practices, not a requisite standard.

RFC 4594 puts forward twelve application classes and matches these to RFC-defined PHBs. These application classes and recommended PHBs are summarized in [Figure 4-8](#).

Figure 4-8 RFC 4594 Marking Recommendations

Application	L3 Classification		IETF
	PHB	DSCP	RFC
Network Control	CS6	48	RFC 2474
VoIP Telephony	EF	46	RFC 3246
Call Signaling	CS5	40	RFC 2474
Multimedia Conferencing	AF41	34	RFC 2597
Real-Time Interactive	CS4	32	RFC 2474
Multimedia Streaming	AF31	26	RFC 2597
Broadcast Video	CS3	24	RFC 2474
Low-Latency Data	AF21	18	RFC 2597
OAM	CS2	16	RFC 2474
High-Throughput Data	AF11	10	RFC 2597
Best Effort	DF	0	RFC 2474
Low-Priority Data	CS1	8	RFC 3662

220200

It is fairly obvious that there are more than a few similarities between Cisco's QoS Baseline and RFC 4594, as there should be, since RFC 4594 is essentially an industry-accepted evolution of Cisco's QoS Baseline. However, there are some differences that merit attention.

The first set of differences is minor, as they involve mainly nomenclature. Some of the application classes from the QoS Baseline have had their names changed in RFC 4594. These changes in nomenclature are summarized in [Table 4-2](#).

Table 4-2 Nomenclature Changes from Cisco QoS Baseline to RFC 4594

Cisco QoS Baseline Class Names	RFC 4594 Class Names
Routing	Network Control
Voice	VoIP Telephony
Interactive Video	Multimedia Conferencing
Streaming Video	Multimedia Streaming
Transactional Data	Low-Latency Data
Network Management	Operations/Administration/Management (OAM)
Bulk Data	High-Throughput Data
Scavenger	Low-Priority Data

The remaining changes are more significant. These include one application class deletion, two marking changes, and two new application class additions:

- The QoS Baseline Locally-Defined Mission-Critical Data class has been deleted from RFC 4594.
- The QoS Baseline marking recommendation of CS4 for Streaming Video has been changed in RFC 4594 to mark Multimedia Streaming to AF31.

- The QoS Baseline marking recommendation of CS3 for Call Signaling has been changed in RFC 4594 to mark Call Signaling to CS5.
- A new application class has been added to RFC 4594, Real-Time Interactive. This addition allows for a service differentiation between elastic conferencing applications (which would be assigned to the Multimedia Conferencing class) and inelastic conferencing applications (which would include high-definition applications, like Cisco TelePresence, in the Realtime Interactive class). Elasticity refers to the applications ability to function despite experiencing minor packet loss. Multimedia Conferencing uses the AF4 class and is subject to markdown (and potential dropping) policies, while the Realtime Interactive class uses CS4 and is not subject neither to markdown nor dropping policies.
- A second new application class has been added to RFC 4594, Broadcast video. This addition allows for a service differentiation between elastic and inelastic streaming media applications. Multimedia Streaming uses the AF3 class and is subject to markdown (and potential dropping) policies, while Broadcast Video uses the CS3 class and is subject neither to markdown nor dropping policies.

The most significant of the differences between Cisco's QoS Baseline and RFC 4594 is the RFC 4594 recommendation to mark Call Signaling to CS5. Cisco has completed a lengthy and expensive marking migration for Call Signaling from AF31 to CS3 (as per the original QoS Baseline of 2002) and, as such, there are no plans to embark on another marking migration in the near future. It is important to remember that RFC 4594 is an informational RFC (in other words, an industry best-practice) and not a standard. Therefore, lacking a compelling business case at the time of writing, Cisco plans to continue marking Call Signaling as CS3 until future business requirements arise that necessitate another marking migration.

Therefore, for the remainder of this document, RFC 4594 marking values are used throughout, with the one exception of swapping Call-Signaling marking (to CS3) and Broadcast Video (to CS5). These marking values are summarized in [Figure 4-9](#).

Figure 4-9 Cisco-Modified RFC 4594-based Marking Values (Call-Signaling is Swapped with Broadcast Video)

Application	L3 Classification		IETF RFC
	PHB	DSCP	
Network Control	CS6	48	RFC 2474
VoIP Telephony	EF	46	RFC 3246
Broadcast Video	CS5	40	RFC 2474
Multimedia Conferencing	AF41	34	RFC 2597
Real-Time Interactive	CS4	32	RFC 2474
Multimedia Streaming	AF31	26	RFC 2597
Call Signaling	CS3	24	RFC 2474
Low-Latency Data	AF21	18	RFC 2597
OAM	CS2	16	RFC 2474
High-Troughput Data	AF11	10	RFC 2597
Best Effort	DF	0	RFC 2474
Low-Priority Data	CS1	8	RFC 3662

221258

A final note regarding standards and RFCs is that there are other RFCs relating to DiffServ design that are currently in draft status (as of the time of writing). One such example is RFC 5127, “Aggregation of Diffserv Service Classes.” As such drafts are finalized, these will correspondingly impact respective areas of QoS design.

Having reviewed various relevant industry guidance and best practices relating to QoS evolution, a final driver—namely advances in QoS technologies—is briefly introduced.

New Platforms and Technologies

As network hardware and software technologies evolve, so do their QoS capabilities and features. New switches and linecards boast advanced classification engines or queuing structures, new routers support sophisticated QoS tools that scale with greater efficiency, and new IOS software features present entirely new QoS options to solve complex scenarios. Therefore, a third set of drivers behind QoS design evolution are the advances in QoS technologies, which are discussed in detail in their respective Place-in-the-Network (PIN) QoS design chapters.

As can be noted from the discussion to this point, all of the drivers behind QoS design evolution are in a constant state of evolution themselves—business drivers will continue to expand and change, as will relevant industry standards and guidance, and so too will platforms and technologies. Therefore, while the strategic and detailed design recommendations presented in this document are as forward-looking as possible, these will no doubt continue to evolve over time.

Before discussing current strategic QoS design recommendations, it may be beneficial to set a base context by first overviewing Cisco’s QoS toolset.

Cisco QoS Toolset

This section describes the main categories of the Cisco QoS toolset and includes these topics:

- Admission control tools
- Classification and marking tools
- Policing and markdown tools
- Scheduling tools
- Link-efficiency tools
- Hierarchical QoS
- AutoQoS
- QoS management

Classification and Marking Tools

Classification tools serve to identify traffic flows so that specific QoS actions may be applied to the desired flows. Often the terms classification and marking are used interchangeably (yet incorrectly so); therefore, it is important to understand the distinction between classification and marking operations:

- Classification refers to the inspection of one or more fields in a packet (the term packet is being used loosely here, to include all Layer 2 to Layer 7 fields, not just Layer 3 fields) to identify the type of traffic that the packet is carrying. Once identified, the traffic is directed to the applicable

policy-enforcement mechanism for that traffic type, where it receives predefined treatment (either preferential or deferential). Such treatment can include marking/remarking, queuing, policing, shaping, or any combination of these (and other) actions.


- Marking, on the other hand, refers to changing a field within the packet to preserve the classification decision that was reached. Once a packet has been marked, a “trust-boundary” is established on which other QoS tools later depend. Marking is only necessary at the trust boundaries of the network and (as with all other QoS policy actions) cannot be performed without classification. By marking traffic at the trust boundary edge, subsequent nodes do not have to perform the same in-depth classification and analysis to determine how to treat the packet.

Cisco IOS software performs classification based on the logic defined within the class map structure within the Modular QoS Command Line Interface (MQC) syntax. MQC class maps can perform classification based on the following types of parameters:


- Layer 1 parameters—Physical interface, sub-interface, PVC, or port
- Layer 2 parameters—MAC address, 802.1Q/p Class of Service (CoS) bits, MPLS Experimental (EXP) bits
- Layer 3 parameters—Differentiated Services Code Points (DSCP), IP Precedence (IPP), IP Explicit Congestion Notification (IP ECN), source/destination IP address
- Layer 4 parameters—TCP or UDP ports
- Layer 7 parameters—Application signatures and URLs in packet headers or payload via Network Based Application Recognition (NBAR)

NBAR is the most sophisticated classifier in the IOS tool suite. NBAR can recognize packets on a complex combination of fields and attributes. NBAR deep-packet classification engine examines the data payload of stateless protocols and identifies application-layer protocols by matching them against a Protocol Description Language Module (PDLM), which is essentially an application signature. NBAR is dependent on Cisco Express Forwarding (CEF) and performs deep-packet classification only on the first packet of a flow. The rest of the packets belonging to the flow are then CEF-switched. However, it is important to recognize that NBAR is merely a classifier, nothing more. NBAR can identify flows by performing deep-packet inspection, but it is up to the MQC policy-map to define what action should be taken on these NBAR-identified flows.


Marking tools change fields within the packet, either at Layer 2 or at Layer 3, such that in-depth classification does not have to be performed at each network QoS decision point. The primary tool within MQC for marking is Class-Based Marking (though policers-sometimes called markers-may also be used, as is discussed shortly). Class-Based Marking can be used to set the CoS fields within an 802.1Q/p tag (as shown in [Figure 4-10](#)), the Experimental bits within a MPLS label (as shown in [Figure 4-11](#)), the Differentiated Services Code Points (DSCPs) within an IPv4 or IPv6 header (as shown in [Figure 4-12](#)), the IP ECN Bits (also shown in [Figure 4-12](#)), as well as other packet fields. Class-Based Marking, like NBAR, is CEF-dependant.

Figure 4-10 **Figure 1-10 802.1Q/p CoS Bits**

226606

Figure 4-11 **Figure 1-11 MPLS EXP Bits**

226607


Figure 4-12 **Figure 1-12 IP ToS Bits: DSCP and IP ECN**

226608

Policing and Markdown Tools


Policers are used to monitor traffic flows and to identify and respond to traffic violations. Policers achieve these objectives by performing ongoing, instantaneous checks for traffic violations and taking immediate prescribed actions when such violations occur. For example, a policer can determine if the offered load is in excess of the defined traffic rate and then drop the out-of-contract traffic, as illustrated in Figure 4-13.

Figure 4-13 A Policing Action




Alternatively, policers may be used to remark excess traffic instead of dropping it. In such a role, the policer is called a marker. Figure 4-14 illustrates a policer functioning as a marker.

Figure 4-14 A Policer as a Marker



The rate at which the policer is configured to either drop or remark traffic is called the Committed Information Rate (CIR). However, policers may police to multiple rates, such as the dual rate policer defined in RFC 2698. With such a policer, the CIR is the principle rate to which traffic is policed, but an upper limit, called the Peak Information Rate (PIR), is also set. The action of a dual-rate policer is analogous to a traffic light, with three conditional states—green light, yellow light, and red light. Traffic equal to or below the CIR (a green light condition) is considered to conform to the rate. An allowance for moderate amounts of traffic above this principal rate is permitted (a yellow light condition) and such traffic is considered to exceed the rate. However, a clearly-defined upper-limit of tolerance (the PIR) is also set (a red light condition), beyond which traffic is considered to violate the rate. As such, a dual-rate RFC 2698 policer performs the traffic conditioning for RFC 2597 Assured Forwarding PHBs, as previously discussed. The actions of such a dual-rate policer (functioning as a three-color marker) are illustrated in Figure 4-15.

Figure 4-15 A Dual-Rate Policer as a Three-Color Marker


A RFC 2698 “Two Rate Three Color Marker” can:

- mark *conforming* traffic to one value (such as AF31)
- remark *exceeding* traffic to another value (such as AF32)
- remark *violating* traffic to yet another value (such as AF33)

226611

Shaping Tools

Shapers operate in a manner similar to policers, in that they meter traffic rates. However, the principle difference between a policer and a shaper is that where a policer remarks or drops traffic as a policy action, a shaper merely delays traffic. [Figure 4-16](#) illustrates generic traffic shaping.

Figure 4-16 Traffic Shaping

Shapers are particularly useful when traffic must conform to a specific rate of traffic in order to meet a service level agreement (SLA) or to guarantee that traffic offered to a service provider is within a contracted rate. Traditionally, shapers have been associated with Non-Broadcast Multiple-Access (NBMA) Layer 2 WAN topologies, like ATM and Frame-Relay, where potential speed-mismatches exist. However, shapers are becoming increasingly necessary on Layer 3 WAN access circuits, such as Ethernet-based handoffs, in order to conform to sub-line access-rates.

Queuing and Dropping Tools

Normally, over uncongested interfaces, packets are transmitted in order on a First-In-First-Out (FIFO) basis. However, if packets arrive at an interface faster than they can be transmitted out the interface, then excess packets may be buffered. When packets are buffered, they may be reordered prior to transmission according to administratively-defined algorithms, which are generally referred to as queuing policies. It is important to recognize that queuing policies are engaged only when the interface is experiencing congestion and are deactivated shortly after the interface congestion clears.


Queuing may be performed in software or in hardware. Within Cisco IOS Software there are two main queuing algorithms available, Class-Based Weighted-Fair Queuing (CBWFQ) and Low-Latency Queuing (LLQ). Within Cisco Catalyst hardware, queuing algorithms fall under a 1PxQyT model, which are overviewed in the following sections.

CBWFQ

Regardless of what queuing policy is applied to an interface within Cisco IOS, there is always an underlying queuing mechanism in place called the Tx-Ring, which is a final (FIFO) output buffer. The Tx-Ring serves the purpose of always having packets ready to be placed onto the wire so that link utilization can be driven to 100%. The Tx-Ring also serves to indicate congestion to the IOS software; specifically, when the Tx-Ring fills to capacity, then the interface is known to be congested and a signal is sent to engage any LLQ/CBWFQ policies that have been configured on the interface.

Class-Based Weighted-Fair Queuing (CBWFQ) is a queuing algorithm that combines the ability to guarantee bandwidth with the ability to dynamically ensure fairness to other flows within a class of traffic. Each queue is serviced in a weighted-round-robin (WRR) fashion based on the bandwidth assigned to each class. The operation of CBWFQ is illustrated in [Figure 4-17](#).

Figure 4-17 CBWFQ Operation



In [Figure 4-17](#), a router interface has been configured with a 4-class CBWFQ policy, with an explicit CBWFQ defined for Network Control, Transactional Data, and Bulk Data respectively, as well as the default CBWFQ queue, which has a Fair-Queuing (FQ) pre-sorter assigned to it.




Note

CBWFQ is a bit of a misnomer because the pre-sorter that may be applied to certain CBWFQs, such as class-default, is not actually a Weighted-Fair Queuing (WFQ) pre-sorter, but rather a Fair-Queuing (FQ) pre-sorter. As such, it ignores any IP Precedence values when calculating bandwidth allocations traffic flows. To be more technically precise, this queuing algorithm would be more accurately named Class-Based Fair-Queuing or CBFQ.

LLQ

Low-Latency Queuing (LLQ) is essentially CBWFQ combined with a strict priority queue. In fact, the original name for the LLQ scheduling algorithm was PQ-CBWFQ. While this name was technically more descriptive, it was obviously clumsy from a marketing perspective and hence the algorithm was renamed LLQ. LLQ operation is illustrated in [Figure 4-18](#).


Figure 4-18 LLQ/CBWFQ Operation



In [Figure 4-18](#), a router interface has been configured with a 5-class LLQ/CBWFQ policy, with voice assigned to a 100 kbps LLQ, three explicit CBWFQs are defined for Call-Signaling, Transactional Data, and Bulk Data respectively, as well as a default queue that has a Fair-Queuing pre-sorter assigned to it. However, an underlying mechanism that doesn't appear within the IOS configuration, but is shown in [Figure 4-18](#), is an implicit policer attached to the LLQ.

The threat posed by any strict priority-scheduling algorithm is that it could completely starve lower priority traffic. To prevent this, the LLQ mechanism has a built-in policer. This policer (like the queuing algorithm itself) engages only when the LLQ-enabled interface is experiencing congestion. Therefore, it is important to provision the priority classes properly. In this example, if more than 100 kbps of voice traffic was offered to the interface, and the interface was congested, the excess voice traffic would be discarded by the implicit policer. However, traffic that is admitted by the policer gains access to the strict priority queue and is handed off to the Tx-Ring ahead of all other CBWFQ traffic.

Not only does the implicit policer for the LLQ protect CBWFQs from bandwidth-starvation, but it also allows for sharing of the LLQ. TDM of the LLQ allows for the configuration and servicing of multiple LLQs, while abstracting the fact that there is only a single LLQ "under-the-hood," so to speak. For example, if both voice and video applications required realtime service, these could be provisioned to two separate LLQs, which would not only protect voice and video from data, but also protect voice and video from interfering with each other, as illustrated in [Figure 4-19](#).

Figure 4-19 Dual-LLQ/CBWFQ Operation

In Figure 4-19, a router interface has been configured with a 6-class LLQ/CBWFQ policy, with voice assigned to a 100 kbps LLQ, video assigned to a “second” 400 kbps LLQ, three explicit CBWFQs are defined for Call-Signaling, Transactional Data, and Bulk Data respectively, as well as a default queue that has a Fair-Queuing pre-sorter assigned to it.

Within such a dual-LLQ policy, two separate implicit policers have been provisioned, one each for the voice class (to 100 kbps) and another for the video class (to 400 kbps), yet there remains only a single strict-priority queue, which is provisioned to the sum of all LLQ classes, in this case to 500 kbps (100 kbps + 400 kbps). Traffic offered to either LLQ class is serviced on a first-come, first-serve basis until the implicit policer for each specific class has been invoked. For example, if the video class attempts to burst beyond its 400 kbps rate then it is dropped. In this manner, both voice and video are serviced with strict-priority, but do not starve data flows, nor do they interfere with each other.

1PxQyT

In order to scale QoS functionality to campus speeds (like GigabitEthernet or Ten GigabitEthernet), Catalyst switches must perform QoS operations within hardware. For the most part, classification, marking, and policing policies (and syntax) are consistent in both Cisco IOS Software and Catalyst hardware; however, queuing (and dropping) policies are significantly different when implemented in hardware. Hardware queuing across Catalyst switches is implemented in a model that can be expressed as 1PxQyT, where:

- 1P represents the support of a strict-priority hardware queue (which is usually disabled by default).
- xQ represents x number of non-priority hardware queues (including the default, Best-Effort queue).
- yT represents y number of drop-thresholds per non-priority hardware queue.

For example, consider a Catalyst 6500 48-port 10/100/1000 RJ-45 Module, the WS-X6748-GE-TX, which has a 1P3Q8T egress queuing structure, meaning that it has:

- One strict priority hardware queue
- Three additional non-priority hardware queues, each with:
 - Eight configurable Weighted Random Early Detect (WRED) drop thresholds per queue

Traffic assigned to the strict-priority hardware queue is treated with an Expedited Forwarding Per-Hop Behavior (EF PHB). That being said, it bears noting that on some platforms there is no explicit limit on the amount of traffic that may be assigned to the PQ and as such, the potential to starve non-priority

queues exists. However, this potential for starvation may be effectively addressed by explicitly configuring input policers that limit—on a per-port basis—the amount of traffic that may be assigned to the priority queue (PQ). Incidentally, this is the recommended approach defined in RFC 3246 (Section 3).

Traffic assigned to a non-priority queue is provided with bandwidth guarantees, subject to the PQ being either fully-serviced or bounded with input policers.

WRED


Selective dropping of packets when the queues are filling is referred to as congestion avoidance. Congestion avoidance mechanisms work best with TCP-based applications because selective dropping of packets causes the TCP windowing mechanisms to “throttle-back” and adjust the rate of flows to manageable rates.

Congestion avoidance mechanisms are complementary to queueing algorithms; queueing algorithms manage the front of a queue, while congestion avoidance mechanisms manage the tail of the queue. Congestion avoidance mechanisms thus indirectly affect scheduling.

The principle congestion avoidance mechanism is WRED, which randomly drops packets as queues fill to capacity. However, the randomness of this selection can be skewed by traffic weights. The weight can either be IP Precedence values, as is the case with default WRED which drops lower IPP values more aggressively (for example, IPP 1 would be dropped more aggressively than IPP 6) or the weights can be AF Drop Precedence values, as is the case with DSCP-Based WRED which drops higher AF Drop Precedence values more aggressively (for example, AF23 is dropped more aggressively than AF22, which in turn is dropped more aggressively than AF21). WRED can also be used to set the IP ECN bits to indicate that congestion was experienced in transit.

The operation of DSCP-based WRED is illustrated in [Figure 4-20](#).

Figure 4-20 DSCP-Based WRED Example Operation



Link Efficiency Tools

Link Efficiency Tools are typically relevant only on link speeds ≤ 768 kbps, and come in two main types:

- Link Fragmentation and Interleaving (LFI) tools—With slow-speed WAN circuits, large data packets take an excessively long time to be placed onto the wire. This delay, called serialization delay, can easily cause a VoIP packet to exceed its delay and/or jitter threshold. There are two LFI tools to mitigate serialization delay on slow speed (≤ 768 kbps) links, Multilink PPP Link Fragmentation and Interleaving (MLP LFI) and Frame Relay Fragmentation (FRF.12).
- Compression tools—Compression techniques, such as compressed Real-Time Protocol (cRTP), minimize bandwidth requirements and are highly useful on slow links. At 40 bytes total, the header portion of a VoIP packet is relatively large and can account for up to two-thirds or the entire VoIP packet (as in the case of G.729 VoIP). To avoid the unnecessary consumption of available bandwidth, cRTP can be used on a link-by-link basis. cRTP compresses IP/UDP/RTP headers from 40 bytes to between two and five bytes (which results in a bandwidth savings of approximately 66% for G.729 VoIP). However, cRTP is computationally intensive, and therefore returns the best bandwidth-savings value vs. CPU-load on slow speed (≤ 768 kbps) links.


This document is intended to address network designs for today's media networks and, as such, link speeds that are ≤ 768 kbps are unsuitable in such a context. Therefore, little or no mention is given to link efficiency tools. For networks that still operate at or below 768 kbps, refer to design recommendations within the Enterprise QoS SRND version 3.3 at http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND/QoS-SRND-Book.html

Hierarchical QoS

Cisco IOS MQC-based tools may be combined in a hierarchical fashion, meaning QoS policies may contain other “nested” QoS policies within them. Such policy combinations are commonly referred to as Hierarchical QoS policies or HQoS policies.

Consider a couple of examples where HQoS policies may be useful. In the first case, there may be scenarios where some applications require policing at multiple levels. Specifically, it might be desirable to limit all TCP traffic to 5 Mbps while, at the same time, limiting FTP traffic (which is a subset of TCP traffic) to no more than 1.5 Mbps. To achieve this nested policing requirement, Hierarchical Policing can be used. The policer at the second level in the hierarchy acts on packets transmitted or marked by the policer at the first level, as illustrated in Figure 4-21. Therefore, any packets dropped by the first level are not seen by the second level. Up to three nested levels are supported by the Cisco IOS Hierarchical Policing feature.

Figure 4-21 Hierarchical Policing Policy Example



226615


Additionally, it is often useful to combine shaping and queuing policies in a hierarchical manner, particularly over sub-line rate access scenarios. As previously discussed, queuing policies only engage when the physical interface is congested (as is indicated to IOS software by a full Tx-Ring). This means that queuing policies never engage on media that has a contracted sub-line rate of access, whether this media is Frame Relay, ATM, or Ethernet. In such a scenario, queuing can only be achieved at a sub-line rate by introducing a two-part HQoS policy wherein:

- Traffic is shaped to the sub-line rate.
- Traffic is queued according to the LLQ/CBWFQ policies within the sub-line rate.

With such an HQoS policy, it is not the Tx-Ring that signals IOS software to engage LLQ/CBWFQ policies, but rather it is the Class-Based Shaper that triggers software queuing when the shaped rate has been reached.

Consider a practical example in which a service provider offers an enterprise subscriber a GigabitEthernet handoff, but with a (sub-line rate) contract for only 60 Mbps, over which he wants to deploy IP Telephony and TelePresence, as well as applications. Normally, queuing policies only engage on this GE interface when the offered traffic rate exceeds 1000 Mbps. However, the enterprise administrator wants to ensure that traffic within the 60 Mbps contracted rate is properly prioritized prior to the handoff so that both VoIP and TelePresence are given the highest levels of service. Therefore, the administrator configures an HQoS policy, such that the software shapes all traffic to the contracted 60 Mbps rate and attaches a nested LLQ/CBWFQ queuing policy within the shaping policy, such that traffic is properly prioritized within this 60 Mbps sub-line rate. Figure 4-22 illustrates the underlying mechanisms for this HQoS policy.

Figure 4-22 Hierarchical Shaping and Queuing Policy Example



AutoQoS

The richness of the Cisco QoS toolset inevitably increases its deployment complexity. To address customer demand for simplification of QoS deployment, Cisco has developed the Automatic QoS (AutoQoS) features. AutoQoS is an intelligent macro that allows an administrator to enter one or two simple AutoQoS commands to enable all the appropriate features for the recommended QoS settings for an application on a specific interface.

AutoQoS VoIP, the first release of AutoQoS, provides best-practice QoS designs for VoIP on Cisco Catalyst switches and Cisco IOS routers. By entering one global and/or one interface command (depending on the platform), the AutoQoS VoIP macro expands these commands into the recommended VoIP QoS configurations (complete with all the calculated parameters and settings) for the platform and interface on which the AutoQoS is being applied.

In the second release, AutoQoS Enterprise, this feature consists of two configuration phases, completed in the following order:

- Auto Discovery (data collection)—Uses NBAR-based protocol discovery to detect the applications on the network and performs statistical analysis on the network traffic.
- AutoQoS template generation and installation—Generates templates from the data collected during the Auto Discovery phase and installs the templates on the interface. These templates are then used as the basis for creating the class maps and policy maps for the network interface. After the class maps and policy maps are created, they are then installed on the interface.

Some may naturally then ask, Why should I read this lengthy and complex QoS design document when I have AutoQoS? It is true that AutoQoS-VoIP is an excellent tool for customers with the objective of enabling QoS for VoIP (only) on their campus and WAN infrastructures. It is also true that AutoQoS-Enterprise is a fine tool for enabling basic branch-router WAN-Edge QoS for voice, video, and multiple classes of data. And as such, customers that have such basic QoS needs and/or do not have the time or desire to do more with QoS, AutoQoS is definitely the way to go.

However, it is important to remember where AutoQoS came from. AutoQoS tools are the result of Cisco QoS feature development coupled with Cisco QoS design guides based on large-scale lab-testing. AutoQoS VoIP is the product of the first QoS design guide (published in 1999). AutoQoS Enterprise is based on the second QoS design guide (published in 2002) and the AutoQoS feature has not been updated since. Therefore, if the business requirements for QoS are quite basic, then—as mentioned—AutoQoS would be an excellent tool to expedite the QoS deployment. If, on the other hand, there are more advanced requirements of QoS—such as those presented in this document—then the configurations presented herein would be recommended over AutoQoS.

QoS Management

Cisco offers a variety of applications to manage quality of service, including

- Cisco QoS Policy Manager (QPM)—QPM supports centralized management of network QoS by providing comprehensive QoS provisioning and monitoring capabilities to deploy, tune, monitor, and optimize the performance characteristics of the network. QPM leverages intelligent network services such as NBAR and other QoS features to identify and monitor networked applications and control their behavior throughout the network.
- Cisco Bandwidth Quality Manager (BQM)—BQM provides end-to-end network service quality monitoring with unique visibility and analysis of traffic, bandwidth, and service quality on IP access networks. BQM can be used to monitor, troubleshoot, and assure end-to-end network performance objectives for converged application traffic. BQM provides micro-level visibility into the network and the network service quality events compromising user experience.
- Cisco Network Analysis Modules (NAM)—Available as Cisco router network modules or as Cisco Catalyst 6500 linecard modules, NAMs can perform extensive voice quality monitoring, intelligent application performance analytics, QoS analysis, and advanced troubleshooting.

Such tools can enable administrators to more efficiently baseline, deploy, monitor, and manage QoS policies over their network infrastructure.

Admission Control Tools

Interactive applications—particularly voice and video applications—often require realtime services from the network. As these resources are finite, they must be managed efficiently and effectively. If the number of flows contending for such priority resources were not limited, then as these resources become oversubscribed, the quality of all realtime flows would degrade—eventually to the point of unusability.

**Note**

Admission Control (AC) is sometimes also referred to as Call Admission Control (CAC); however, as applications evolve, not all applications requiring priority services are call-oriented, and as such AC is a more encompassing designation.

Admission control functionality is most effectively controlled an application-level, such as is the case with Cisco Unified CallManager, which controls VoIP and IP video and/or TelePresence flows. As such, admission control design is not discussed in detail in this document, but will be deferred to application-specific design guides, such as the Cisco Unified Communications design guides and/or the Cisco TelePresence design guides at www.cisco.com/go/designzone.

As discussed previously, media applications are taxing networks as never before. To that end, current admission control tools are not sufficient to make the complex decisions that many collaborative media applications require. Thus, admission control continues to be an field for extended research and development in the coming years, with the goal of developing multi-level admission control solutions, as described below:

- The first level of admission control is simply to enable mechanisms to protect voice-from-voice and/or video-from-video on a first-come, first-serve basis. This functionality provides a foundation on which higher-level policy-based decisions can be built.
- The second level of admission control factors in dynamic network topology and bandwidth information into a real-time decision of whether or not a media stream should be admitted. These decisions could be made by leveraging intelligent network protocols, such as Resource Reservation Protocol (RSVP).
- The third level of admission control introduces the ability to preempt existing flows in favor of “higher-priority” flows.
- The fourth level of admission control contains policy elements and weights to determine what exactly constitutes a “higher-priority” flow, as defined by the administrative preferences of an organization. Such policy information elements may include—but are not limited to—the following:
 - Scheduled versus ad hoc—Media flows that have been scheduled in advance would likely be granted priority over flows that have been attempted ad hoc.
 - Users and groups—Certain users or user groups may be granted priority for media flows.
 - Number of participants—Multipoint media calls with larger number of participants may be granted priority over calls with fewer participants.
 - External versus internal participants—Media sessions involving external participants, such as customers, may be granted priority over sessions comprised solely of internal participants.
 - Business critical factor—Additional subjective elements may be associated with media streams, such as a business critical factor. For instance, a live company meeting would likely be given a higher business critical factor than a live training session. Similarly, a media call to close a sale or to retain a customer may be granted priority over regular, ongoing calls.


**Note**

It should be emphasized this is not an exhaustive list of policy information elements that could be used for admission control, but rather is merely a sample list of possible policy information elements. Additionally, each of these policy information elements could be assigned administratively-defined weights to yield an overall composite metric to calculate and represent the final admit/deny admission control decision for the stream.

- The fifth level of admission control provides graceful conflict resolution, such that—should preemption of a media flow be required—existing flow users are given a brief message indicating that their flow is about to be preempted (preferably including a brief reason as to why) and a few seconds to make alternate arrangements (as necessary).

A five-level admission control model, deployed over a DiffServ-enabled infrastructure is illustrated in Figure 4-23.

Figure 4-23 *Five-Level Admission Control Model Deployed Over a DiffServ Infrastructure*



Thus, having laid a foundational context by reviewing QoS technologies, let us turn our attention to Cisco's strategic QoS recommendations for enterprise medianets.

Enterprise Medianet Strategic QoS Recommendations

As media applications increase on the IP network, QoS will play a progressively vital role to ensure the required service level guarantees to each set of media applications, all without causing interference to each other. Therefore, the QoS strategies must be consistent at each PIN, including the campus, data center, branch WAN/MAN/VPN, and branch.

Also, integration will play a key role in two ways. First, media streams and endpoints will be increasingly leveraged by multiple applications. For example, desktop video endpoints may be leveraged for desktop video conferencing, Web conferencing, and for viewing stored streaming video for training and executive communications.

Additionally, many media applications will require common sets of functions, such as transcoding, recording, and content management. To avoid duplication of resources and higher implementation costs, common media services need to be integrated into the IP network so they can be leveraged by multiple media applications.

Furthermore, because of the effectiveness of multimedia communication and collaboration, the security of media endpoints and communication streams becomes an important part of the media-ready strategy. Access controls for endpoints and users, encryption of streams, and securing content files stored in the data center are all part of a required comprehensive media application security strategy.


Finally, as the level of corporate intellectual property migrates into stored and interactive media, it is critical to have a strategy to manage the media content, setting and enforcing clear policies, and having the ability to protect intellectual property in secure and managed systems. Just as companies have policies and processes for handling intellectual property in document form, they also must develop and update these policies and procedures for intellectual property in media formats.

Therefore, to meet all these media application requirements, Cisco recommends—not to reengineer networks to support each wave of applications—but rather to utilize an architectural approach, namely a medianet architecture.

Enterprise Medianet Architecture

A medianet is built upon an architecture that supports the different models of media applications and optimizes their delivery, such as those shown in the architectural framework in [Figure 4-24](#).

Figure 4-24 Enterprise Medianet Architectural Framework



226617

An enterprise medianet framework starts with an end-to-end QoS-enabled network infrastructure designed and built to achieve high availability, including the data center, campus, WAN, and branch office networks. The network provides a set of services to video applications, including:

- Access services—Provide access control and identity of video clients, as well as mobility and location services.
- Transport services—Provide packet delivery, ensuring the service levels with QoS and delivery optimization.
- Bridging services—Transcoding, conferencing, and recording services.
- Storage services—Content capture, storage, retrieval, distribution, and management services.
- Session control services—Signaling and control to setup and tear-down sessions, as well as gateways.

When these media services are made available within the network infrastructure, endpoints can be multi-purpose and rely upon these common media services to join and leave sessions for multiple media applications. Common functions such as transcoding and conferencing different media codecs within the same session can be deployed and leveraged by multiple applications, instead of being duplicated for each new media application.

With this architectural framework in mind, let us take a closer look at the strategic QoS recommendations for a medianet.

Enterprise Medianet QoS Application Class Recommendations

As mentioned previously, Cisco has slightly modified its implementation of (informational) RFC 4594 (as shown in [Figure 4-9](#)). With Admission Control recommendations added to this model, these combined recommendations are summarized in [Figure 4-25](#).

Figure 4-25 Enterprise Medianet QoS Recommendations

Application Class	Per-Hop Behavior	Admission Control	Queuing and Dropping	Media Application Examples
VoIP Telephony	EF	Required	Priority Queue (PQ)	Cisco IP Phones (G.711, G.729)
Broadcast Video	CS5	Required	(Optional) PQ	Cisco IP Video Surveillance/Cisco Enterprise TV
Real-Time Interactive	CS4	Required	(Optional) PQ	Cisco TelePresence
Multimedia Conferencing	AF4	Required	BW Queue + DSCP WRED	Cisco Unified Personal Communicator
Multimedia Streaming	AF3	Recommended	BW Queue + DSCP WRED	Cisco Digital Media System (VoDs)
Network Control	CS6		BW Queue	EIGRP, OSPF, BGP, HSRP, IKE
Signaling	CS3		BW Queue	SCCP, SIP, H.323
Ops/Admin/Mgmt (OAM)	CS2		BW Queue	SNMP, SSH, Syslog
Transactional Data	AF2		BW Queue + DSCP WRED	Cisco WebEx/MeetingPlace/ERP Apps
Bulk Data	AF1		BW Queue + DSCP WRED	E-mail, FTP, Backup Apps, Content Distribution
Best Effort	DF		Default Queue + RED	Default Class
Scavenger	CS1		Min BW Queue	YouTube, iTunes, BitTorrent, Xbox Live

224550

The 12 classes of applications within this enterprise medianet QoS model—which have unique service level requirements and thus require explicit QoS PHBs—are outlined as follows:

- VoIP Telephony
- Broadcast Video
- Realtime Interactive
- Multimedia Conferencing
- Network Control
- Signaling
- Operations, Administration, and Management (OAM)
- Transactional Data and Low-Latency Data
- Bulk Data and High-Throughput Data
- Best Effort
- Scavenger and Low-Priority Data

VoIP Telephony

This service class is intended for VoIP telephony (bearer-only) traffic (VoIP signaling traffic is assigned to the Call Signaling class). Traffic assigned to this class should be marked EF (DSCP 46) and should be admission controlled. This class is provisioned with an Expedited Forwarding Per-Hop Behavior. The EF PHB—defined in RFC 3246—is a strict-priority queuing service and as such, admission to this class should be controlled. Example traffic includes G.711 and G.729a.

Broadcast Video

This service class is intended for broadcast TV, live events, video surveillance flows, and similar “inelastic” streaming media flows (“inelastic” flows refer to flows that are highly drop sensitive and have no retransmission and/or flow-control capabilities). Traffic in this class should be marked Class Selector 5 (CS5/DSCP 40) and may be provisioned with an EF PHB; as such, admission to this class should be controlled (either by an explicit admission control mechanisms or by explicit bandwidth provisioning). Examples traffic includes live Cisco Digital Media System (DMS) streams to desktops or to Cisco Digital Media Players (DMPs), live Cisco Enterprise TV (ETV) streams, and Cisco IP Video Surveillance (IPVS).

Realtime Interactive

This service class is intended for inelastic high-definition interactive video applications and is intended primarily for audio and video components of these applications. Whenever technically possible and administratively feasible, data sub-components of this class can be separated out and assigned to the Transactional Data traffic class. Traffic in this class should be marked CS4 (DSCP 32) and may be provisioned with an EF PHB; as such, admission to this class should be controlled. An example application is Cisco TelePresence.

Multimedia Conferencing

This service class is intended for desktop software multimedia collaboration applications and is intended primarily for audio and video components of these applications. Whenever technically possible and administratively feasible, data sub-components of this class can be separated out and assigned to the Transactional Data traffic class. Traffic in this class should be marked Assured Forwarding Class 4 (AF41/DSCP 34) and should be provisioned with a guaranteed bandwidth queue with DSCP-based Weighted-Random Early Detect (DSCP-WRED) enabled. Admission to this class should be controlled; additionally, traffic in this class may be subject to policing and re-marking. Example applications include Cisco Unified Personal Communicator, Cisco Unified Video Advantage, and the Cisco Unified IP Phone 7985G.

Network Control

This service class is intended for network control plane traffic, which is required for reliable operation of the enterprise network. Traffic in this class should be marked CS6 (DSCP 48) and provisioned with a (moderate, but dedicated) guaranteed bandwidth queue. WRED should not be enabled on this class, as network control traffic should not be dropped (if this class is experiencing drops, then the bandwidth allocated to it should be re-provisioned). Example traffic includes EIGRP, OSPF, BGP, HSRP, IKE, etc.

Signaling

This service class is intended for signaling traffic that supports IP voice and video telephony; essentially, this traffic is control plane traffic for the voice and video telephony infrastructure. Traffic in this class should be marked CS3 (DSCP 24) and provisioned with a (moderate, but dedicated) guaranteed bandwidth queue. WRED should not be enabled on this class, as signaling traffic should not be dropped (if this class is experiencing drops, then the bandwidth allocated to it should be re-provisioned). Example traffic includes SCCP, SIP, H.323, etc.

Operations, Administration, and Management (OAM)

This service class is intended for—as the name implies—network operations, administration, and management traffic. This class is important to the ongoing maintenance and support of the network. Traffic in this class should be marked CS2 (DSCP 16) and provisioned with a (moderate, but dedicated) guaranteed bandwidth queue. WRED should not be enabled on this class, as OAM traffic should not be dropped (if this class is experiencing drops, then the bandwidth allocated to it should be re-provisioned). Example traffic includes SSH, SNMP, Syslog, etc.

Transactional Data and Low-Latency Data

This service class is intended for interactive, “foreground” data applications (“foreground” applications refer to applications from which users are expecting a response—via the network—in order to continue with their tasks. Excessive latency in response times of foreground applications directly impacts user productivity). Traffic in this class should be marked Assured Forwarding Class 2 (AF21 / DSCP 18) and should be provisioned with a dedicated bandwidth queue with DSCP-WRED enabled. This traffic class may be subject to policing and re-marking. Example applications include data components of multimedia collaboration applications, Enterprise Resource Planning (ERP) applications, Customer Relationship Management (CRM) applications, database applications, etc.

Bulk Data and High-Throughput Data

This service class is intended for non-interactive “background” data applications (“background” applications refer to applications from which users are not awaiting a response—via the network—in order to continue with their tasks. Excessive latency in response times of background applications does not directly impact user productivity. Furthermore, as most background applications are TCP-based file-transfers, these applications—if left unchecked—could consume excessive network resources away from more interactive, foreground applications). Traffic in this class should be marked Assured Forwarding Class 1 (AF11/DSCP 10) and should be provisioned with a moderate, but dedicated bandwidth queue with DSCP-WRED enabled. This traffic class may be subject to policing and re-marking. Example applications include E-mail, backup operations, FTP/SFTP transfers, video and content distribution, etc.

Best Effort

This service class is the default class. As only a relative minority of applications are assigned to priority, guaranteed-bandwidth, or even to deferential service classes, the vast majority of applications continue to default to this best effort service class; as such, this default class should be adequately provisioned (a minimum bandwidth recommendation, for this class is 25%). Traffic in this class is marked Default Forwarding (DF or DSCP 0) and should be provisioned with a dedicated queue. WRED is recommended to be enabled on this class. Although, since all the traffic in this class is marked to the same “weight” (of DSCP 0), the congestion avoidance mechanism is essentially Random Early Detect (RED).

Scavenger and Low-Priority Data


This service class is intended for non-business related traffic flows, such as data or media applications that are entertainment-oriented. The approach of a less-than best effort service class for non-business applications (as opposed to shutting these down entirely) has proven to be a popular, political compromise. These applications are permitted on enterprise networks, as long as resources are always available for business-critical media applications. However, as soon the network experiences congestion, this class is the first to be penalized and aggressively dropped. Furthermore, the scavenger class can be

utilized as part of an effective strategy for DoS and worm attack mitigation (discussed later in this chapter). Traffic in this class should be marked CS1 (DSCP 8) and should be provisioned with a minimal bandwidth queue that is the first to starve should network congestion occur. Example traffic includes YouTube, Xbox Live/360 Movies, iTunes, BitTorrent, etc.

Media Application Class Expansion

While there are merits to adopting a 12-class model, as outlined in the previous section, Cisco recognizes that not all enterprises are ready to do so, whether this be due to business reasons, technical constraints, or other reasons. Therefore, rather than considering these medianet QoS recommendations as an all-or-nothing approach, Cisco recommends considering a phased approach to media application class expansion, as illustrated in [Figure 4-26](#).

Figure 4-26 **Media Application Class Expansion**



Utilizing such a phased approach to application class expansion, enterprise administrators can incrementally implement QoS policies across their infrastructures in a progressive manner, inline with their business needs and technical constraints. Familiarity with this enterprise medianet QoS model can assist in the smooth expansion of QoS policies to support additional media applications as future requirements arise. Nonetheless, at the time of QoS deployment, the enterprise needs to clearly define their business objectives with QoS, which correspondingly determines how many traffic classes will be required at each phase of deployment.

Cisco QoS Best Practices

With an overall application PHB strategy in place, end-to-end QoS policies can be designed for each device and interface, as determined by their roles in the network infrastructure. These are detailed in the various PIN-specific QoS design chapters that follow. However, because the Cisco QoS toolset provides many QoS design and deployment options, a few succinct design principles can help simplify strategic QoS deployments.

Hardware versus Software QoS

A fundamental QoS design principle is to **always enable QoS policies in hardware—rather than software—whenever a choice exists**. Cisco IOS routers perform QoS in software, which places incremental loads on the CPU, depending on the complexity and functionality of the policy. Cisco Catalyst switches, on the other hand, perform QoS in dedicated hardware ASICs on Ethernet-based ports and as such do not tax their main CPUs to administer QoS policies. This allows complex policies to be applied at line rates at even Gigabit or Ten-Gigabit speeds.

Classification and Marking Best Practices

When classifying and marking traffic, a recommended design principle is to **classify and mark applications as close to their sources as technically and administratively feasible**. This principle promotes end-to-end Differentiated Services and PHBs.

In general, it is not recommended to trust markings that can be set by users on their PCs or other similar devices, because users can easily abuse provisioned QoS policies if permitted to mark their own traffic. For example, if an EF PHB has been provisioned over the network, a PC user can easily configure all their traffic to be marked to EF, thus hijacking network priority queues to service non-realtime traffic. Such abuse could easily ruin the service quality of realtime applications throughout the enterprise. On the other hand, if enterprise controls are in place that centrally administer PC QoS markings, then it may be possible and advantageous to trust these.

Following this rule, it is further recommended to **use DSCP markings whenever possible**, because these are end-to-end, more granular, and more extensible than Layer 2 markings. Layer 2 markings are lost when media changes (such as a LAN-to-WAN/VPN edge). There is also less marking granularity at Layer 2. For example, 802.1Q/p CoS supports only three bits (values 0-7), as does MPLS EXP. Therefore, only up to eight classes of traffic can be supported at Layer 2 and inter-class relative priority (such as RFC 2597 Assured Forwarding Drop Preference markdown) is not supported. On the other hand, Layer 3 DSCP markings allow for up to 64 classes of traffic, which is more than enough for most enterprise requirements for the foreseeable future.

As the line between enterprises and service providers continues to blur and the need for interoperability and complementary QoS markings is critical, you should **follow standards-based DSCP PHB markings to ensure interoperability and future expansion**. Because the enterprise medianet marking recommendations are standards-based—as has been previously discussed—enterprises can easily adopt these markings to interface with service provider classes of service. Network mergers—whether the result of acquisitions, mergers, or strategic alliances—are also easier to manage when using standards-based DSCP markings.

Policing and Markdown Best Practices

There is little reason to forward unwanted traffic only to police and drop it at a subsequent node, especially when the unwanted traffic is the result of DoS or worm attacks. Furthermore, the overwhelming volume of traffic that such attacks can create can cause network outages by driving

network device processors to their maximum levels. Therefore, it is recommended to **police traffic flows as close to their sources as possible**. This principle applies also to legitimate flows, as worm-generated traffic can masquerade under legitimate, well-known TCP/UDP ports and cause extreme amounts of traffic to be poured onto the network infrastructure. Such excesses should be monitored at the source and marked down appropriately.

Whenever supported, markdown should be done according to standards-based rules, such as RFC 2597 (AF PHB). For example, excess traffic marked to AFx1 should be marked down to AFx2 (or AFx3 whenever dual-rate policing—such as defined in RFC 2698—is supported). Following such markdowns, congestion management policies, such as DSCP-based WRED, should be configured to drop AFx3 more aggressively than AFx2, which in turn should be dropped more aggressively than AFx1.

Queuing and Dropping Best Practices

Critical media applications require service guarantees regardless of network conditions. **The only way to provide service guarantees is to enable queuing at any node that has the potential for congestion**, regardless of how rarely this may occur. This principle applies not only to campus-to-WAN/VPN edges, where speed mismatches are most pronounced, but also to campus interswitch links, where oversubscription ratios create the potential for congestion. There is simply no other way to guarantee service levels than by enabling queuing wherever a speed mismatch exists.

Additionally, **because each medianet application class has unique service level requirements, each should optimally be assigned a dedicated queue**. However, on platforms bounded by a limited number of hardware or service provider queues, no fewer than four queues would be required to support medianet QoS policies, specifically:

- Realtime queue (to support a RFC 3246 EF PHB service)
- Guaranteed-bandwidth queue (to support RFC 2597 AF PHB services)
- Default queue (to support a RFC 2474 DF service)
- Bandwidth-constrained queue (to support a RFC 3662 Scavenger service)

Additional queuing recommendations for these classes are discussed next.

Strict-Priority Queuing Recommendations—The 33 Percent LLQ Rule

The Realtime or Strict Priority class corresponds to the RFC 3246 EF PHB. The amount of bandwidth assigned to the realtime queuing class is variable. However, if the majority of bandwidth is provisioned with strict priority queuing (which is effectively a FIFO queue), then the overall effect is a dampening of QoS functionality, both for latency and jitter sensitive realtime applications (contending with each other within the FIFO priority queue) and also for non-realtime applications (as these may periodically receive wild bandwidth allocation fluctuations, depending on the instantaneous amount of traffic being serviced by the priority queue). Remember the goal of convergence is to enable voice, video, and data applications to transparently co-exist on a single IP network. When realtime applications dominate a link, then non-realtime applications fluctuate significantly in their response times, destroying the transparency of the converged network.

For example, consider a (45 Mbps) DS3 link configured to support two TelePresence (CTS-3000) calls with an EF PHB service. Assuming that both systems are configured to support full high definition, each such call requires 15 Mbps of strict-priority queuing. Prior to TelePresence calls being placed, non-realtime applications have access to 100% of the bandwidth on the link (to simplify the example, assume there are no other realtime applications on this link). However, once these TelePresence calls are established, all non-realtime applications would suddenly be contending for less than 33% of the link.

TCP windowing would take effect and many applications hang, time-out, or become stuck in a non-responsive state, which usually translates into users calling the IT help desk complaining about the network (which happens to be functioning properly, albeit in a poorly-configured manner).

To obviate such scenarios, Cisco Technical Marketing has done extensive testing and has found that a significant decrease in non-realtime application response times occurs when realtime traffic exceeds one-third of link bandwidth capacity. Extensive testing and customer deployments have shown that a general best queuing practice is to **limit the amount of strict priority queuing to 33% of link bandwidth capacity**. This strict priority queuing rule is a conservative and safe design ratio for merging realtime applications with data applications.



Note

As previously discussed, Cisco IOS software allows the abstraction (and thus configuration) of multiple strict priority LLQs. In such a multiple LLQ context, this design principle would apply to the sum of all LLQs to be within one-third of link capacity.

It is vitally important to understand that **this strict priority queuing rule is simply a best practice design recommendation and is not a mandate**. There may be cases where specific business objectives cannot be met while holding to this recommendation. In such cases, enterprises must provision according to their detailed requirements and constraints. However, it is important to recognize the tradeoffs involved with over-provisioning strict priority traffic and its negative performance impact both on other realtime flows and also on non-realtime-application response times.

And finally, **any traffic assigned to a strict-priority queue should be governed by an admission control mechanism**.


Best Effort Queuing Recommendation

The Best Effort class is the default class for all traffic that has not been explicitly assigned to another application-class queue. Only if an application has been selected for preferential/deferential treatment is it removed from the default class. Because most enterprises have several thousand applications running over their networks, adequate bandwidth must be provisioned for this class as a whole in order to handle the sheer number and volume of applications that default to it. Therefore, it is recommended to **reserve at least 25 percent of link bandwidth for the default Best Effort class**.

Scavenger Class Queuing Recommendations

Whenever Scavenger queuing class is enabled, it should be assigned a minimal amount of bandwidth, such as 1% (or whatever the minimal bandwidth allocation that the platform supports). On some platforms, queuing distinctions between Bulk Data and Scavenger traffic flows cannot be made, either because queuing assignments are determined by CoS values (and both of these application classes share the same CoS value of 1) or because only a limited amount of hardware queues exist, precluding the use of separate dedicated queues for each of these two classes. In such cases, the Scavenger/Bulk queue can be assigned moderate amount of bandwidth, such as 5%.

These queuing rules are summarized in [Figure 4-27](#), where the inner pie chart represents a hardware or service provider queuing model that is limited to four queues and the outer pie chart represents a corresponding, more granular queuing model that is not bound by such constraints.

Figure 4-27 Compatible 4-Class and 12-Class Medianet Queuing Models

226619

QoS for Security Best Practices

While the primary objective of most QoS deployments is to provision preferential—and sometimes deferential—service to various application classes, QoS policies can also provide an additional layer of security to the network infrastructure, especially in the case of mitigating Denial-of-Service (DoS) and worm attacks.


There are two main classes of DoS attacks:

- **Spoofing attacks**—The attacker pretends to provide a legitimate service, but provides false information to the requester (if any).
- **Slamming attacks**—The attacker exponentially generates and propagates traffic until service resources (servers and/or network infrastructure) are overwhelmed.

Spoofing attacks are best addressed by authentication and encryption technologies. Slamming (also known as “flooding”) attacks, on the other hand, can be effectively mitigated through QoS technologies.

In contrast, worms exploit security vulnerabilities in their targets and disguisedly carry harmful payloads that usually include a self-propagating mechanism. Network infrastructure usually is not the direct target of a worm attack, but can become collateral damage as worms exponentially self-propagate. The rapidly multiplying volume of traffic flows eventually drowns the CPU/hardware resources of routers and switches in their paths, indirectly causing Denial of Service to legitimate traffic flows, as shown in Figure 4-28.

Figure 4-28 Direct and Indirect Collateral Damage from DoS/Worm Attacks



A *reactive* approach to mitigating such attacks is to reverse-engineer the worm and set up intrusion detection mechanisms and/or ACLs and/or NBAR policies to limit its propagation. However, the increased sophistication and complexity of worms make them harder and harder to separate from legitimate traffic flows. This exacerbates the finite time lag between when a worm begins to propagate and when the following can take place:

- Sufficient analysis has been performed to understand how the worm operates and what its network characteristics are.
- An appropriate patch, plug, or ACL is disseminated to network devices that may be in the path of worm; this task may be hampered by the attack itself, as network devices may become unreachable for administration during the attacks.

These time lags may not seem long in absolute terms, such as in minutes, but the relative window of opportunity for damage is huge. For example, in 2003, the number of hosts infected with the Slammer worm (a Sapphire worm variant) doubled every 8.5 seconds on average, infecting over 75,000 hosts in just 11 minutes and performing scans of 55 million more hosts within the same time period.

A *proactive* approach to mitigating DoS/worm attacks within enterprise networks is to have control plane policing and data plane policing policies in place within the infrastructure which immediately respond to out-of-profile network behavior indicative of DoS or worm attacks. Control plane policing serves to protect the CPU of network devices—such as switches and routers—from becoming bogged down with interruption-handling and thus not having enough cycles to forward traffic. Data plane policing—also referred to as Scavenger-class QoS—serves to protect link bandwidth from being consumed by forwarding DoS/worm traffic to the point of having no room to service legitimate, in-profile flows.

Control Plane Policing

A router or switch can be logically divided into four functional components or planes:

- Data plane
- Management plane
- Control plane
- Services plane


The vast majority of traffic travels through the router via the data plane. However the route processor must handle certain packets, such as routing updates, keepalives, and network management. This is often referred to as control and management plane traffic.

Because the route processor is critical to network operations, any service disruption to the route processor or the control and management planes can result in business-impacting network outages. A DoS attack targeting the route processor, which can be perpetrated either inadvertently or maliciously, typically involves high rates of punted traffic (traffic that results in a processor-interruption) that results in excessive CPU utilization on the route processor itself. This type of attack, which can be devastating to network stability and availability, may display the following symptoms:

- High route processor CPU utilization (near 100%)
- Loss of line protocol keepalives and routing protocol updates, leading to route flaps and major network transitions
- Interactive sessions via the Command Line Interface (CLI) are slow or completely unresponsive due to high CPU utilization
- Route processor resource exhaustion—resources such as memory and buffers are unavailable for legitimate IP data packets
- Packet queue backup, which leads to indiscriminate drops (or drops due to lack of buffer resources) of other incoming packets

Control Plane Policing (CPP for Cisco IOS routers or CoPP for Cisco Catalyst Switches) addresses the need to protect the control and management planes, ensuring routing stability, availability, and packet delivery. It uses a dedicated control plane configuration via the Modular QoS CLI (MQC) to provide filtering and rate limiting capabilities for control plane packets.

[Figure 4-29](#) illustrates the flow of packets from various interfaces. Packets destined to the control plane are subject to control plane policy checking, as depicted by the control plane services block.

Figure 4-29 Packet Flow Within a Switch/Router

By protecting the route processor, CPP/CoPP helps ensure router and network stability during an attack. For this reason, **a best practice recommendation is to deploy CPP/CoPP as a key protection mechanism on all routers and switches that support this feature.**

To successfully deploy CPP, the existing control and management plane access requirements must be understood. While it can be difficult to determine the exact traffic profile required to build the filtering lists, the following summarizes the recommended steps necessary to properly define a CPP policy:


1. Start the deployment by defining liberal policies that permit most traffic.
2. Monitor traffic pattern statistics collected by the liberal policy.
3. Use the statistics gathered in the previous step to tighten the control plane policies.

Data Plane Policing/Scavenger-Class QoS

The logic applied to protecting the control plane can also be applied to the data plane. Data plane policing has two components:


- Campus access-edge policers that meter traffic flows from endpoint devices and remark “abnormal” flows to CS1 (the Scavenger marking value).
- Queuing policies on all nodes that include a deferential service class for Scavenger traffic.

These two components of data plane policing/Scavenger-class QoS are illustrated in [Figure 4-30](#).

Figure 4-30 Data Plane Policing/Scavenger-class QoS Components

Most endpoint devices have fairly predictable traffic patterns and, as such, can have metering policers to identify "normal" flows (the volume of traffic that represents 95% of the typically-generated traffic rates for the endpoint device) vs. "abnormal" flows (the remainder). For instance, it would be "abnormal" for a port that supposedly connects to an IP phone to receive traffic in excess of 128 kbps. Similarly, it would be "abnormal" for a port that supposedly connects to a Cisco TelePresence system to receive traffic in excess of 20 Mbps. Both scenarios would be indicative of network abuse—either intentional or inadvertent. Endpoint PCs also have traffic patterns that can be fairly accurately baselined with statistical analysis.

For example, for users of Windows-based systems, the Windows Task Manager (which can be selected by simultaneously pressing CTRL-ALT-DEL) can graphically display networking statistics (available from the networking tab). Most users are generally surprised at how low the average network utilization rates of PCs are during everyday use, as compared to their link speed capacities. Such a graphical display of network utilization is shown in Figure 4-31, where the radical and distinctive difference in network utilization rates after worm-infection is highlighted.

Figure 4-31 Sample PC Network Utilization Rates—Before and After Infection by a Worm

These access edge metering policers are relatively unintelligent. They do not match specific network characteristics of specific types of attacks, but simply meter traffic volumes and respond to abnormally high volumes as close to the source as possible. The simplicity of this approach negates the need for the policers to be programmed with knowledge of the specific details of how the attack is being generated or propagated. It is precisely this unintelligence of such access layer metering policers that allow them

to maintain relevancy as worms mutate and become more complex. The policers do not care how the traffic was generated or what it looks like; they care only how much traffic is being put onto the wire. Therefore, they continue to police even advanced worms that continually change their tactics of traffic-generation.

For example, in most enterprises it is quite abnormal (within a 95% statistical confidence interval) for PCs to generate sustained traffic in excess of 5% of link capacity. In the case of a GigabitEthernet access switch port, this means that it would be unusual in most organizations for an end user PC to generate more than 50 Mbps of uplink traffic on a sustained basis.

**Note**

It is important to recognize that this value (5%) for normal endpoint utilization by PC endpoints is just an example value. This value would likely vary from enterprise to enterprise, as well as within a given enterprise (such as by departmental functions).

It is very important to recognize that what is being recommended by data plane policing/Scavenger class QoS is not to police all traffic to 50 Mbps and automatically drop the excess. Should that be the case, there would not be much reason to deploy GigabitEthernet switch ports to endpoint devices. But rather, these campus access-layer policers do not drop traffic at all; they only perform remarking (if traffic rates appear abnormal). These policers are coupled with queuing policies on all network nodes that include a deferential service class for traffic marked as Scavenger (CS1). Queuing policies only engage when links are congested; as such, if links capacity exists, then traffic is never dropped. It is only in scenarios where offered traffic flows exceed link capacity—forcing queuing policies to engage and queuing buffers to fill to capacity—that drops may occur. In such scenarios, dropping can either occur indiscriminately (on a last-come-first-dropped basis) or with a degree of intelligence (as would be the case if abnormal traffic flows were previously identified).

Let's illustrate how this might work for both legitimate excess traffic and also the case of illegitimate excess traffic resulting from a DoS or worm attack.

In the former case, assume that the PC generates over 50 Mbps of traffic, perhaps because of a large file transfer or backup. Congestion (under normal operating conditions) is rarely if ever experienced within the campus because there is generally abundant capacity to carry the traffic. Uplinks to the distribution and core layers of the campus network are typically GigabitEthernet or Ten Gigabit Ethernet, which would require 1,000 or 10,000 Mbps of traffic (respectively) from the access layer switch to congest. If the traffic is destined to the far side of a WAN/VPN link, queuing and dropping typically occurs even without the access layer policer, because of the bottleneck caused by the typical campus-to-WAN/VPN speed mismatch. In such a case, the TCP sliding windows mechanism would eventually find an optimal speed (under 50 Mbps) for the file transfer. Access layer policers that mark out-of-profile traffic to Scavenger (CS1) would thus not affect legitimate traffic, aside from the obvious remarking. *No reordering or dropping would occur on such flows as a result of these data plane policers that would not have occurred anyway.*

In the latter case, the effect of access layer policers on traffic caused by DoS or worm attacks is quite different. As hosts become infected and traffic volumes multiply, congestion may be experienced even within the campus. For example, if just 11 end user PCs on a single access switch begin spawning worm flows to their maximum GigabitEthernet link capacities, even Ten Gigabit Ethernet uplinks/core links will congest and queuing and dropping policies will engage. At this point, VoIP and media applications, and even best effort applications, would gain priority over worm-generated traffic (as Scavenger traffic would be dropped the most aggressively). Furthermore, network devices would remain accessible for administration of the patches/plugs/ACLs/NBAR policies required to fully neutralize the specific attack. WAN/VPN links would also be similarly protected, which is a huge advantage, as generally WAN/VPN links are the first to be overwhelmed by DoS/worm attacks. Scavenger class policies thus *significantly mitigate network traffic generated by DoS or worm attacks.*

Therefore, for network administrators to implement data plane policing/Scavenger class QoS, they need to first **profile applications to determine what constitutes normal as opposed to abnormal flows, within a 95 percent confidence interval**. Thresholds demarking normal/abnormal flows vary from enterprise to enterprise and from application to application. Beware of over-scrutinizing traffic behavior because this could exhaust time and resources and could easily change daily. Remember, legitimate traffic flows that temporarily exceed thresholds are not penalized by the presented Scavenger class QoS strategy. Only sustained, abnormal streams generated simultaneously by multiple hosts (highly indicative of DoS/worm attacks) are subject to aggressive dropping only after legitimate traffic has been serviced.

To contain such abnormal flows, **deploy campus access edge policers to remark abnormal traffic to Scavenger (CS1)**. Additionally, whenever possible, **deploy a second line of policing defense at the distribution layer**. And to complement these remarking policies, it is necessary to **enforce deferential Scavenger class queuing policies** throughout the network.

A final word on this subject—it is important to recognize the distinction between mitigating an attack and preventing it entirely. **Control plane policing and data plane policing policies do not guarantee that no DoS or worm attacks will ever happen, but serve only to reduce the risk and impact that such attacks could have on the network infrastructure. Therefore, it is vital to overlay a comprehensive security strategy over the QoS-enabled network infrastructure.**

Summary

This chapter began by discussing the reasons driving the QoS design updates presented in this document by examining three sets of drivers behind QoS design evolution, including:

- New applications and business requirements
- New industry guidance and best practices
- New platforms and technologies

Business drivers—including the evolution of video applications, the phenomena of social networking, the convergence within media applications, the globalization of the workforce, and the pressures to be “green”—were examined to determine how these impact new QoS designs over enterprise media networks. Next, developments in industry standards and best practices—with particular emphasis on RFC 4594 Configuration Guidelines for DiffServ Classes—were discussed, as were developments in QoS technologies and their respective impacts on QoS design.

Cisco’s QoS toolset was overviewed to provide a foundational context for the strategic best practices that followed. Classification and marking tools, policing and markdown tools, shaping, queuing, and dropping tools were all reviewed, as were AutoQoS and QoS management tools.

An enterprise medianet architecture was then presented, along with strategic QoS design recommendations. These recommendations included an RFC 4594-based application class model and an application expansion class model (for enterprises not yet ready to deploy a 12-class QoS model). Additionally, QoS best practices for classification, marking, policing, and queuing were presented, including:

- Always deploy QoS in hardware (over software) whenever possible.
- Mark as close to the source as possible with standards-based DSCP values.
- Police as close to the source as possible.
- Markdown according to standards-based rules.
- Deploy queuing policies on all network nodes.

- Optimally assign each medianet class a dedicated queue.
- Limit strict-priority queuing to 33% of link-capacity whenever possible.
- Provision at least 25% of a link's capacity for best effort applications.
- Provision a minimal queue (such as 1%) for the Scavenger applications class.
- Enable control plane policing on platforms that support this feature.
- Deploy data plane policing/Scavenger class QoS policies whenever possible.

These strategic design recommendations will serve to make the PIN-specific designs that follow more cohesive, complementary, and consistent.

References

White Papers

- Cisco Visual Networking Index—Forecast and Methodology, 2007-2012
http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360_ns827_Networking_Solutions_White_Paper.html
- Approaching the Zettabyte Era
http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481374_ns827_Networking_Solutions_White_Paper.html
- Cisco Enterprise QoS Solution Reference Design Guide, version 3.3
http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND/QoS-SRND-Book.html
- Overview of a Medianet Architecture
<http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/vrn.html>

IETF RFCs

- RFC 791 Internet Protocol
<http://www.ietf.org/rfc/rfc791>
- RFC 2474 Definition of the Differentiated Services Field
<http://www.ietf.org/rfc/rfc2474>
- RFC 2597 Assured Forwarding PHB Group
<http://www.ietf.org/rfc/rfc2597>
- RFC 3246 An Expedited Forwarding PHB
<http://www.ietf.org/rfc/rfc3246>
- RFC 3662 A Lower Effort Per-Domain Behavior for Differentiated Services
<http://www.ietf.org/rfc/rfc3662>
- RFC 4594 Configuration Guidelines for DiffServ Service Classes
<http://www.ietf.org/rfc/rfc4594>
- RFC 5187 [Draft] Aggregation of Diffserv Service Classes
<http://tools.ietf.org/html/rfc5127>

Cisco Documentation

- Cisco IOS Quality of Service Solutions Configuration Guide, Release 12.4
http://www.cisco.com/en/US/docs/ios/qos/configuration/guide/12_4/qos_12_4_book.html



CHAPTER 5

Medianet Security Design Considerations

A medianet is the foundation for media-rich collaboration across borderless networks. The availability and overall security of a medianet is thus critical to global business operations.

The security challenge is enabling an enterprise to confidently embrace and deliver these rich global collaboration services without compromising the overall security posture of the company.

The chapter illustrates the key strategies for enabling secure collaboration by employing a defense-in-depth approach that extends and integrates consistent, end-to-end security policy enforcement, and system-wide intelligence, across an enterprise medianet.

An Introduction to Securing a Medianet

The security of a medianet is addressed as two broad categories:

- Medianet foundation infrastructure

This consists of the end-to-end network infrastructure and services that are fundamental to a medianet, including switches, routers, wireless infrastructure, network clients, servers, baseline network services, as well as the WAN and other elements that enable pervasive access to medianet services.

- Medianet collaboration services

This consists of the media-rich collaboration and communication services that a medianet may support, such as TelePresence, Digital Media Systems (DMS), IP Video surveillance (IPVS), Unified Communications, desktop video and WebEx conferencing, along with their associated infrastructure and clients.

In order to secure a medianet, Cisco SAFE guidelines are applied to these two broad categories of a medianet. The security of both being critical to the delivery of pervasive secure collaboration.

Medianet Foundation Infrastructure

The network infrastructure and clients of a medianet are its fundamental elements. Security of these medianet clients and infrastructure thus provides the secure foundation for all the collaboration services that a medianet enables. Without the security of this foundational element, secure collaboration is impossible to deliver and any additional security measures are futile.

The Cisco SAFE guidelines must be applied to this fundamental area and each of its elements in order to provide a secure medianet foundation.

Medianet Collaboration Services

Each of the collaboration and communication services deployed on a medianet must each be assessed and secured in accordance with security policy and by applying the Cisco SAFE guidelines. This requires detailed analysis of the platforms and protocols used, the traffic flows and communication points, as well as possible attack vectors. The extension and integration of current, and possibly new, security techniques to each of these services can then be developed and deployed.


The implementation details may vary but the Cisco SAFE guidelines provide a consistent blueprint of the security considerations that need to be addressed.

Cisco SAFE Approach

Cisco SAFE provides a reference guide, an architecture and design blueprints for consistent, end-to-end security policy enforcement and system-wide intelligence. We will apply Cisco SAFE to a medianet in order to extend this approach to all elements of a medianet.

The Cisco SAFE approach includes proactive techniques to provide protection from initial compromise. This includes Network Foundation Protection, endpoint security, web and E-mail security, virtualization and network access control, as well as secure communications. These are complemented by reactive techniques that provide the ability to identify anomalous activity on the network and, where necessary, mitigate their impact. This includes telemetry, event correlation, firewall, IPS, data loss prevention and switching security.

Figure 5-1 Cisco SAFE



For more information about Cisco SAFE, see the link referenced in [Medianet Security Reference Documents, page 5-12](#).

Security Policy and Procedures

Every organization should have defined security policies and procedures that form the basis of a strong security framework. These policies concisely define the required security actions and may, in turn, specify associated standards and guidelines. Procedures define how these policy goals are to be accomplished.

Security policies and procedures must be in place in order to achieve consistent, effective network security. The security guidelines provided in this chapter can be leveraged to enforce these policies, according to the specific policy requirements.

For more information on developing and implementing a security policy, the SANS Technology Institute offers some excellent resources including training, guidelines and sample security policies, see [Medianet Security Reference Documents, page 5-12](#).


Security of Medianet Foundation Infrastructure

The security of this foundational element of a medianet is critical to the security of all services that a medianet enables. If the medianet itself is vulnerable, fundamental network services are vulnerable and thus, all additional services are vulnerable. If the clients that access a medianet are vulnerable, any hosts, devices or services they have access to are vulnerable.

To address this area, we can leverage the Cisco SAFE reference guide to provide the fundamental security guidelines. This chapter provides a brief overview of the key elements of Cisco SAFE; for the complete *Cisco SAFE Reference Guide* and additional Cisco SAFE collateral, see the link referenced in [Medianet Security Reference Documents, page 5-12](#).

Security Architecture

The Cisco SAFE architecture features a modular design with the overall network represented by functional modules, including campus, branch, data center, Internet edge, WAN edge, and core. This enables the overall security design, as well as the security guidelines for each individual module to be leveraged, applied, and integrated into a medianet architecture.

Figure 5-2 Cisco SAFE Architecture

226659

The Cisco SAFE architecture features virtualization and segmentation to enable different functional and security domains, secure communications for data in transit, centralized management and control for ease of operations and consistent policy enforcement, along with fundamental design principles such as the Cisco Security Control Framework and the architecture lifecycle.

Network Foundation Protection

The focus of Network Foundation Protection (NFP) is security of the network infrastructure itself, primarily protecting the control and management planes of a medianet. NFP mitigates unauthorized access, denial-of-service (DoS) and local attacks such as man-in-the-middle (MITM) attacks that can be used to perform eavesdropping, sniffing, and data stream manipulation.

The key areas NFP addresses include the following:

- Secure Device Access
- Service Resiliency
- Network Policy Enforcement
- Routing Security
- Switching Security

Integration of these elements is critical to medianet security and, unless implemented, renders any more advanced techniques futile. For instance, if a malicious user can access the local LAN switch using a simple password, they will have access to all traffic flowing through that switch, can reconfigure the device and mount a vast array of attacks.

Endpoint Security

Endpoints are exposed to a wide range of threats, including malware, botnets, worms, viruses, trojans, spyware, theft of information, and unauthorized access. Hardening these endpoints is thus critical to overall network security, protecting both the endpoint itself, the data they host and any network to which they connect.

Endpoint security includes the following:

- Operating system and application hardening

It is critical that the operating system and applications running on an endpoint are hardened and secured in order to reduce the attack surface and render the endpoint as resilient as possible to attacks. This involves implementing a secure initial configuration, as well as the regular review of vulnerabilities and the timely application of any necessary updates and security patches.

- User education and training

End-users should receive ongoing education and training to make them aware of the critical role they play in mitigating existing and emerging threats, including security awareness, protection of corporate data, acceptable use policy and minimizing risk exposure. This should be presented in a simple, collaborative way to reinforce corporate policies.

- Host-based IPS (HIPS)

HIPS provides endpoints with protection against both known and zero-day or unpatched attacks, whichever network they may be connected to. This is achieved through both signature- and behavior-based threat detection and mitigation that are key features of HIPS. This functionality is offered by the Cisco Security Agent (CSA), along with the ability to enforce policy and perform data loss prevention on the endpoint itself. Some of this functionality may also be available in the host operating system.

- Cisco Security Services Client (CSSC)

The CSSC is a software supplicant that enables identity-based access and policy enforcement on a client, across both wired and wireless networks. This includes the ability to enforce secure network access controls, such as requiring the use of WPA2 for wireless access and automatically starting a VPN connection when the endpoint is connected to a non-corporate network.

For more information about Cisco CSA and CSSC, see [Medianet Security Reference Documents](#), page 5-12.

Web Security

The web is increasingly being used to distribute malware and, whilst malicious sites continue to operate as one key delivery method, the majority of today's web-based threats are delivered through legitimate websites that have been compromised. Add to this the threats posed by spyware, traffic tunneling, client usage of unauthorized sites and services, and the sharing of unauthorized data, and it is easy to see why web security is critical to any organization.

Cisco offers four web security options:

- Cisco Ironport S-Series Web Security Appliance (WSA)
An on-premise, dedicated appliance offering high performance web-based threat mitigation and security policy enforcement. The WSA provides web usage controls, known and unknown malware protection through multiple scanning engines and reputation filtering, data loss prevention, URL filtering, protocol tunneling protection and malware activity monitoring.
- Cisco ScanSafe
Hosted web security (SaaS) offering web-based malware protection in the cloud. ScanSafe provides real-time scanning of inbound and outbound web traffic for known and unknown malware, as well as monitoring of malware activity.
- Cisco ASA 5500 Series Content Security and Control Security Services Module (CSC-SSM)
Service module for the Cisco ASA 5500 Series providing comprehensive antivirus, anti-spyware, file blocking, anti-spam, anti-phishing, URL blocking and filtering, and content filtering.
- Cisco IOS Content Filtering
Integrated web security in Cisco IOS platforms offering whitelist and blacklist URL filtering, keyword blocking, security rating, and category filtering

For more information about Cisco Ironport WSA, ScanSafe, and Cisco IOS security, see [Medianet Security Reference Documents, page 5-12](#).

E-mail Security

E-mail is one of the primary malware distribution methods, be it through broad phishing attacks, malware in attachments or more sophisticated, targeted E-mail attacks. E-mail spam is a major revenue generator for the miscreant community, and E-mail is one of the most common methods for unauthorized data exchange. Consequently, E-mail security is critical to an enterprise.

Cisco offers E-mail security through the Ironport C-Series E-mail Security Appliance (ESA), providing spam filtering, malware filtering, reputation filtering, data loss prevention (DLP) and E-mail encryption. This is available in three deployment options:

- On-premise appliance enforcing both inbound and outbound policy controls.
- Hybrid Hosted service offering an optimal design that features inbound filtering in the cloud for spam and malware filtering, and an on-premise appliance performing outbound control for DLP and encryption.
- Dedicated hosted E-mail security service (SaaS) offering the same rich E-mail security features but with inbound and outbound policy enforcement being performed entirely in the cloud.

For more information on Cisco Ironport ESA, see [Medianet Security Reference Documents, page 5-12](#).

Network Access Control

With the pervasiveness of networks, controlling who has access and what they are subsequently permitted to do are critical to network and data security. Consequently, identity, authentication and network policy enforcement are key elements of network access control.

Cisco Trusted Security (TrustSec) is a comprehensive solution that offers policy-based access control, identity-aware networking, and data confidentiality and integrity protection in the network. Key Cisco technologies integrated in this solution include:

- Cisco Catalyst switches providing rich infrastructure security features such as 802.1X, web authentication, MAC authentication bypass, MACSec, Security Group Tags (SGT), and a selection of dynamic policy enforcement mechanisms and deployment modes.
- Cisco Secure Access Control System (ACS) as a powerful policy server for centralized network identity and access control.
- Cisco Network Access Control (NAC) offering appliance-based network access control and security policy enforcement, as well as posture assessment.

For more information about Cisco TrustSec, see [Medianet Security Reference Documents, page 5-12](#).

User Policy Enforcement

User policy enforcement is a broad topic and, based on the defined security policy, may include:

- Acceptable Use Policy (AUP) Enforcement

For example, restricting web access and application usage, such as P2P applications and adult content. This can be achieved through Cisco IOS Content Filtering and Ironport WSA Web Usage Controls (WUC).

- Data Loss Prevention (DLP)

DLP is often required for regulatory purposes and refers to the ability to control the flow of certain data, as defined by security policy. For example, this may include credit card numbers or medical records. DLP can be enforced at multiple levels, including on a host, through the use of Cisco Security Agent (CSA), in E-mail through integration of the Ironport ESA and via web traffic through integration of the Ironport WSA.

Secure Communications

The confidentiality, integrity, and availability of data in transit is critical to business operations and is thus a key element of network security. This encompasses the control and management, as well as data planes. The actual policy requirements will typically vary depending on the type of data being transferred and the network and security domains being transited. This is a reflection of the risk and vulnerabilities to which data may be subject, including unauthorized access, and data loss and manipulation from sniffing or man-in-the-middle (MITM) attacks.

For example, credit card processing over the Internet is governed by regulatory requirements that require it to be in an isolated security domain and encrypted. A corporate WLAN may require the use of WPA2 for internal users and segmented wireless access for guests.

Secure communications is typically targeted at securing data in transit over WAN and Internet links that are exposed to external threats, but the threats posed by compromised internal hosts is not to be overlooked. Similarly, sensitive data or control and management traffic transiting internal networks may also demand additional security measures.

Cisco offers a range of VPN technology options for securing WAN access, either site-to-site or for remote access, along with PKI for secure, scalable, and manageable authentication. Cisco VPN technologies include MPLS, IPSec VPN, SSL VPN, GRE, GETVPN, DMVPN.

For more information about Cisco VPN technologies, see [Medianet Security Reference Documents, page 5-12](#).

Firewall Integration

Firewall integration enables extended segmentation and network policy enforcement of different security policy domains. For example, to isolate and secure servers that store highly sensitive data or segment users with different access privileges.

In addition, firewall integration offers more advanced, granular services, such as stateful inspection and application inspection and control on Layer 2 through Layer 7. These advanced firewall services are highly effective of detecting and mitigating TCP attacks and application abuse in HTTP, SMTP, IM/P2P, voice, and other protocols.

Cisco offers the following two key firewall integration options:

- Adaptive Security Appliance (ASA) 5500 Series

Dedicated firewall enabling a highly scalable, high performance, high availability and fully featured deployment that is available on a range of platforms. The ASA 5500 Series also features the Cisco ASA Botnet Traffic Filter, providing real-time traffic monitoring, anomalous traffic detection, and reputation-based control that enables the mitigation of botnets and other malware that shares phone-home communication patterns.

- Cisco IOS Firewall

Cost-effective, integrated firewall offered as a classic, interface-based firewall or as a zone-based firewall (ZBFW) that enables the application of policies to defined security zones.

For more information about the Cisco ASA 5500 Series and Cisco IOS Firewall, see [Medianet Security Reference Documents, page 5-12](#).

IPS Integration

The integration of network IPS provides the ability to accurately identify, classify, and stop malicious traffic on the network, including worms, spyware, adware, attacks, exploits, network viruses, and application abuse. Cisco IPS offers dynamic and flexible signature, vulnerability, exploit, behavioral and reputation-based threat detection and mitigation, as well as protocol anomaly detection.

In addition, the collaboration of Cisco IPS with other Cisco devices provides enhanced visibility and control through system-wide intelligence. This includes host-based IPS collaboration with Cisco Security Agent (CSA), reputation-based filtering and global correlation using SensorBase, automated threat mitigation with the WLAN controller (WLC), multi-vendor event correlation and attack path identification using Cisco Security Monitoring, Analysis, and Response System (CS-MARS), and common policy management using Cisco Security Manager (CSM).

Cisco IPS is available in a wide range of network IPS deployment options, including:

- Cisco IPS 4200 Series Appliances

Dedicated high scalability, high availability hardware appliances.

- Integrated modules for ISR, ASA and Catalyst 6500

Offering flexible deployment options but consistent rich signature set and policy enforcement

- Cisco IOS IPS

Cost-effective integrated IPS with sub-set of common signatures.

For more information about the Cisco IPS offerings, see [Medianet Security Reference Documents](#), page 5-12.

Telemetry

Visibility into the status of a medianet and the identification of any anomalous activity is critical to overall network security. Security monitoring, analysis & correlation is thus essential to the timely and accurate detection and mitigation of anomalies.

The baseline elements of telemetry are very simple and inexpensive to implement, and include:

- Time Synchronization

Synchronize all network devices to the same network clock by using Network Time Protocol (NTP) to enable accurate and effective event correlation.

- Monitoring of System Status Information

Maintain visibility into overall device health by monitoring CPU, memory and processes.

- Implementation of CDP Best Common Practices

Enable CDP on all infrastructure interfaces for operational purposes but disable CDP on any interfaces where CDP may pose a risk, such as external-facing interfaces.

- Remote Monitoring

Leverage syslog, SNMP and additional telemetry techniques, such as Netflow, to a centralized server, such as CS-MARS, for cross-network data aggregation. This enables detailed and behavioral analysis of the data which is key to traffic profiling, anomaly-detection and attack forensics, as well as general network visibility and routine troubleshooting.

For more information about management and visibility in a medianet, see [Chapter 6, “Medianet Management and Visibility Design Considerations.”](#)

Security of Medianet Collaboration Services

Once the foundational elements of a medianet are secured, the next step is to address the security of each of the collaboration and communication services that a medianet is being used to deliver, whether it is TelePresence, DMS, IPVS, Unified Communications, desktop video, WebEx conferencing, or any other collaboration and communication service.

As each collaboration service is deployed, the service must be well-researched and understood, security policy must be reviewed and applied, and network security measures extended to encompass it. To achieve this, the same Cisco SAFE guidelines are applied to each medianet collaboration service and their associated infrastructure, enabling consistent, end-to-end, security policy enforcement.

Security Policy Review

Prior to deployment of a new service, it is critical to review it in relation to security policy. This will initially require detailed analysis of the service itself, including the protocols it uses, the traffic flows and traffic profile, the type of data involved, as well as its associated infrastructure devices and platforms. This enables a security threat and risk assessment to be generated that identifies possible attack vectors and their associated risk. In addition, there may be regulatory requirements to take into consideration.

The service can then be reviewed in relation to security policy in order to determine how to enforce the security policy and, if necessary, what changes are required to the policy. This is generally referred to as a policy impact assessment.

Reviewing a new service in relation to the security policy enables consistent enforcement that is critical to overall network security.

Architecture Integration

Integration of a new service into a medianet requires an assessment of the traffic flows, the roles of its associated infrastructure and the communications that take place, as well as an understanding of the current corporate network design. This enables the most appropriate deployment model to be adopted, including the appropriate segmentation of security domains.

For example, a WebEx Node resides on the corporate network, but communicates with the external WebEx cloud as well as internal clients. Consequently, the logical placement for this device, performing an intermediary role between internal clients and an external service, is the DMZ. For more information about WebEx Node integration, see [Medianet Security Reference Documents, page 5-12](#).

Application of Cisco SAFE Guidelines

For each medianet collaboration service, we will apply the Cisco SAFE guidelines to enable the consistent enforcement of security policy. Taking each of the Cisco SAFE security areas, we will assess if and how they apply to this service and its associated infrastructure, and what additions or changes may need to be made to the current security measures. The Cisco SAFE security areas we will apply include:

- **Network Foundation Protection (NFP)**
Hardening of each of the service infrastructure components and services, including secure device access and service resiliency. QoS and Call Admission Control (CAC) being two key features of service resiliency for media-rich communication services.
- **Endpoint Security**
Hardening of each of the service endpoints and a review of current endpoint security policies. For instance, if the CSA Trusted QoS feature is currently employed, this may need to be modified to reflect the requirements of a new desktop video deployment.
- **Web Security**
Extension of web security policies to the service, including perhaps the modification of web usage controls, DLP policies, and URL filtering. For instance, a WebEx Node should only connect to the WebEx Cloud and so corporate URL filtering policies may be modified to enforce this.
- **E-mail Security**
A review of E-mail security policies may be required if the service involves the use of E-mail, either as an integral part of the service itself or as part of its monitoring and management.

- Network Access Control (NAC)

Extension of network access control to the service, including identification, authentication and network policy enforcement of users and devices. This may involve the extension of policies to include service-specific policy enforcement, such as to restrict the authorized users, devices, protocols and flows of a particular service, thereby only granting minimum access privileges and reducing the risk exposure of the service endpoints.

- User Policy Enforcement

A review of user policies may be required to reflect the new service offerings. For instance, to define the data sharing policy for external Cisco WebEx Connect Spaces.

- Secure Communications

The path and risk exposure of data in transit must be assessed in order to deploy the most appropriate security solution. This may include the security of control and management planes, as well as the data plane. For example, the encryption of TelePresence media flows may be required if data traverses an insecure security domain or the media content is sensitive.

- Firewall Integration

Firewall policies may need to be modified to allow firewall traversal for the service. For instance, if you wish to provide secure access to your UC infrastructure from external softphones, you may enable the ASA Phone Proxy feature.

- IPS Integration

IPS integration and signature tuning may be required to ensure the accurate and timely detection and mitigation of anomalies in these new services. For instance, to identify SIP attacks or DoS attacks against UC servers.

- Telemetry

Extension of monitoring to the new service in order to provide visibility into its operational status, to enable the detection of anomalous activity that may be indicative of an incident, as well as to record activity for detailed analysis and forensics.

Implementation involves leveraging the available security features on the service infrastructure devices themselves and those offered within the service, as well as extending existing or new network security techniques to these new services.

Since the actual implementation of security for each service is very specific and often very different, it should be addressed as an integral part of the overall design and deployment of each service. For more information on securing each of the collaboration services, see [Medianet Security Reference Documents](#), page 5-12 for additional collateral.

Medianet Security Reference Documents

- ASA 5500 Series
<http://www.cisco.com/go/asa>
- Cisco Data Center Security
http://www.cisco.com/en/US/netsol/ns750/networking_solutions_sub_program_home.html
- Cisco IOS Content Filtering
<http://www.cisco.com/en/US/products/ps6643/index.html>
- Cisco IOS Firewall
<http://www.cisco.com/en/US/products/sw/secursw/ps1018/index.html>
- Cisco IOS NetFlow
http://www.cisco.com/en/US/products/ps6601/products_ios_protocol_group_home.html
- Cisco IP Video Surveillance (IPVS)
http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns819/landing_vid_surveillance.html
- Cisco IronPort C-Series E-mail Security Appliance (ESA)
http://www.ironport.com/products/email_security_appliances.html
- Cisco IronPort S-Series Web Security Appliance (WSA)
http://www.ironport.com/products/web_security_appliances.html
- Cisco Medianet
<http://www.cisco.com/web/solutions/medianet/index.html>
- Cisco Network Admission Control (NAC)
http://cisco.com/en/US/netsol/ns466/networking_solutions_package.html
- Cisco SAFE
<http://www.cisco.com/go/safe>
- Cisco SAFE WebEx Node Integration
http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/WebEx_wpf.html
- Cisco ScanSafe Web Security
<http://www.scansafe.com/>
- Cisco Secure Services Client (CSSC)
<http://cisco.com/en/US/products/ps7034/index.html>
- Cisco Security Portfolio
<http://www.cisco.com/go/security>
- Cisco Security Agent (CSA)
<http://www.cisco.com/go/csa>
- Cisco Trust and Identity Management Solutions
http://cisco.com/en/US/netsol/ns463/networking_solutions_sub_solution_home.html
- Cisco Trusted Security (TrustSec)
http://www.cisco.com/en/US/netsol/ns774/networking_solutions_package.html

- Cisco Unified Communications (UC) Security
http://www.cisco.com/en/US/netsol/ns340/ns394/ns165/ns391/networking_solutions_package.html
- Cisco VPN
http://cisco.com/en/US/products/ps5743/Products_Sub_Category_Home.html
- Cisco WebEx Security Overview
http://www.cisco.com/en/US/prod/collateral/ps10352/cisco_webex_security_overview.pdf
- SANS Policy Resources
<http://www.sans.org/security-resources/policies/>



CHAPTER 6

Medianet Management and Visibility Design Considerations

This chapter provides a high-level overview of various functionalities that can be used to provide management and visibility into video flows within an enterprise medianet. This functionality can be divided into the following two broad categories:

- **Network-embedded**—Management functionality embedded within the IP network infrastructure itself (that is, routers, switches, and so on). Network-embedded management functionality may benefit a single video application solution, or may benefit all video application solutions, depending on the specific functionality.
- **Application-specific**—Management functionality embedded within the components that comprise individual video application solutions, such as Cisco TelePresence, Cisco Digital Media Systems, Cisco IP Video Surveillance, and Cisco Desktop Video Collaboration. Although individual video application solutions co-exist over a converged IP network infrastructure, the application-specific management functionality may be unique to the video solution.



Note

Management applications that make use of the functionality embedded within both the IP network infrastructure and/or individual components of video application solutions to provide a centralized point of monitoring, control, and reporting within the medianet infrastructure, may be considered a third category of functionality. Examples of such applications are the Cisco QoS Policy Manager (QPM), which provides centralized QoS provisioning and monitoring for Cisco router platforms. Future revisions of this design chapter may include discussion around these applications.

In this design guide, management functionality is presented using the International Organization for Standardization (ISO)/International Telecommunications Union (ITU) Fault, Configuration, Accounting, Performance, and Security (FCAPS) model. The five major categories of network management defined within the FCAPS model are as follows:

- **Fault management**—Detection and correction of problems within the network infrastructure or end device.
- **Configuration management**—Configuration of network infrastructure components or end devices, including initial provisioning and ongoing scheduled changes.
- **Accounting management**—Scheduling and allocation of resources among end users, as well as billing back for that use if necessary.
- **Performance management**—Performance of the network infrastructure or end devices, including maintaining service level agreements (SLAs), quality of service (QoS), network resource allocation, and long-term trend analysis.

- Security management—Maintaining secure authorization and access to network resources and end devices, as well as maintaining confidentiality of information crossing the network infrastructure.

**Note**

Note that the security management aspects of the medianet infrastructure are only briefly discussed in this chapter. A separate chapter of this design guide deals with medianet security design considerations.

Network-Embedded Management Functionality

The following sections highlight functionality embedded within network infrastructure devices that can be used to provide visibility and management of video flows within an enterprise medianet. Although specific examples within each section discuss the use of a particular functionality for a specific video application solution (Cisco TelePresence, Cisco Digital Media Systems, Cisco IP Video Surveillance, or Cisco Desktop Video Collaboration), the features discussed can generally provide benefit across multiple video application solutions. A complete list of network-embedded management functionality is outside the scope of this document. Instead, for brevity, only specific features relevant to medianet management and visibility are discussed. [Table 6-1](#) provides a high level summarization of the functionality discussed in following sections.

Table 6-1 Summary of Network-Embedded Management Functionality

Management Product /Tool	Management Functionality	Description
NetFlow	Performance and security management	- NetFlow services embedded within Cisco router and Cisco Catalyst switch platforms provide the ability to collect and export flow information that can be used to determine the amount of video traffic crossing key points within a medianet. Flow information collected at a NetFlow collector, such as the Cisco Network Analysis Module (NAM) can be used to provide ongoing monitoring and/or reports that may be used to determine whether adequate bandwidth is provisioned per service class to support the video traffic applications. - NetFlow export version 9 provides the ability to export multicast flows as well, providing some visibility into the amount of multicast traffic crossing key points within the medianet infrastructure. - Netflow can also be used to identify anomalous flows within the medianet infrastructure, alerting security operations staff of potential worm propagation or a DDoS attack. For further information, see the <i>Cisco SAFE Reference Guide</i> at the following URL: http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg.html.
Cisco Network Analysis Module (NAM)	Performance management	- The Cisco Catalyst 6500 Series Network Analysis Module (NAM) provides the ability to monitor and generate reports regarding data flows within a medianet. Data flows from the supervisor of a Cisco Catalyst 6500 switch platform, SPAN/RSPAN ports, or NetFlow Data Export (NDE) from other routers and switches within the medianet infrastructure can be analyzed. - The NAM provides the ability to monitor and generate reports on traffic flows aggregated by Differentiated Services Code Point (DSCP) marking. This can assist in providing visibility into the amount of traffic per service class crossing key points within the medianet, and can aid in provisioning adequate bandwidth per service class across the network infrastructure.
IP service level agreements (IPSLAs)	Performance management	- IPSLA functionality embedded within Cisco Catalyst switches, Cisco IOS routers, and Cisco TelePresence endpoints can be used as a pre-assessment tool, to determine whether the medianet infrastructure has the capability to support additional video flows before becoming production resources. - IPSLAs may be used cautiously to perform ongoing performance monitoring of the medianet infrastructure to determine whether a particular video class is experiencing degradation because of packet loss and/or jitter.

Table 6-1 Summary of Network-Embedded Management Functionality (continued)

Management Product /Tool	Management Functionality	Description
Router and switch command-line interface	Performance management and fault management	- The traceroute utility can be used to determine the Layer 3 hop path of video flows through a medianet infrastructure. - After the path has been determined, high-level CLI commands such as show interface summary and show interface can be used on each router and switch along the path to determine quickly whether drops or errors are occurring on relevant interfaces. - Other platform-specific commands can be used to display packet drops per queue on Cisco Catalyst switch platforms. When separate traffic service classes (corresponding to different video applications) are mapped to different queues, network administrators can use these commands to determine whether particular video applications are experiencing degradation because of packet loss within the medianet infrastructure. - When policy maps are used to map specific traffic service classes (corresponding to different video applications) to software queues within Cisco router platforms, or hardware queues within certain Cisco Catalyst switch platforms, the show policy-map command can be used to display the amount of traffic per service class as well as drops experienced by the particular service class. Network administrators can use this command to determine whether adequate bandwidth is provisioned, as well as to determine whether particular video applications are experiencing degradation because of packet loss within the medianet infrastructure.
Syslog	Security management and fault management	- Telemetry using syslog can be used to provide some key fault management information on network infrastructure devices within a medianet, such as CPU utilization, memory utilization, and link status. - For further information regarding network security best practices, see the <i>Cisco SAFE Reference Guide</i> at the following URL: http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg.html.

Table 6-1 Summary of Network-Embedded Management Functionality (continued)

Management Product /Tool	Management Functionality	Description
Simple Network Management Protocol (SNMP)	Security management, fault management, and performance management	• Telemetry using SNMP can also be used to provide key fault management information on network infrastructure devices within a medianet. • SNMP can be used to collect statistics from network infrastructure devices for performance management purposes. • SNMP traps can be generated for authentication failures to devices, providing an additional layer of security management.
AAA services	Security management	• AAA services can be used to provide centralized access control for security management, as well as an audit trail providing visibility into access of network infrastructure devices. • For further information regarding network security best practices, see the <i>Cisco SAFE Reference Guide</i> at the following URL: http://www.cisco.com/en/US/docs/solutions/Enterprise/Security/SAFE_RG/SAFE_rg.html.

NetFlow


NetFlow services provide network administrators access to information regarding IP flows within their networks. IP flows are unidirectional streams of packets flowing through a network device. They share common properties such as source address, destination address, protocol, port, DSCP value, and so on. Network devices, such as switches and routers, can collect and store flow data in the form of flow records within a NetFlow table or cache. Flow records can then be periodically exported from the NetFlow cache to one or more NetFlow management collectors located centrally within a data center or campus service module. NetFlow collectors aggregate exported NetFlow records to provide monitoring and reporting information regarding the IP traffic flows within the network.

NetFlow provides a means of gaining additional visibility into the various video flows within an enterprise medianet. From an FCAPS perspective, this visibility can be used for either performance management purposes or for accounting management purposes. More specifically, NetFlow data can assist in determining whether sufficient bandwidth has been provisioned across the network infrastructure to support existing video applications. NetFlow data records can be exported in various formats depending on the version. The most common formats are versions 1, 5, 7, 8, and 9. NetFlow export version 9 is the latest version, which has been submitted to the IETF as informational RFC 3954, providing a model for the IP Flow Information Export (IPFIX) working group within the IETF. NetFlow version 9 provides a flexible and extensible means of exporting NetFlow data, based on the use of templates that are sent along with the flow record. Templates contain structural information about the flow record fields, allowing the NetFlow collector to interpret the flow records even if it does not understand the semantics of the fields. For more information regarding NetFlow version 9, see the following URL: <http://www.ietf.org/rfc/rfc3954.txt>.

NetFlow Strategies Within an Enterprise Medianet

Simply enabling NetFlow on every interface on every network device, exporting all the flow data to a central NetFlow collector, and then aggregating the flow data into a single set of information across the entire enterprise medianet, is generally considered only marginally useful for anything but small networks. This strategy typically results in information overload, in which a lot of statistics are collected, yet the network administrator has no idea where traffic is flowing within the network infrastructure. An alternative strategy is to collect flow information based on specific requirements for the flow data itself. One such strategy is to selectively enable NetFlow to collect traffic flows on certain interfaces at key points within the enterprise medianet. The data from each collection point in the network can then be kept as separate information sets, either at a single NetFlow collector or in multiple NetFlow collectors, rather than aggregated together. This can be used to provide a view of what traffic is flowing through the different points within the enterprise medianet. Depending on the capabilities of the NetFlow collector, this can be done in various ways. Some NetFlow collectors allow different UDP port numbers to be used for flows from different devices. This allows the aggregation of NetFlow information from multiple interfaces on a single router or switch to appear as a single data set or source. It also allows the flows from a redundant set of routers or switches to appear as a single data set or source. Other NetFlow collectors, such as the Cisco Network Analysis Module (NAM), use a fixed port (UDP 3000) for flows from devices. Flows from multiple interfaces on the same device can be aggregated into a single custom data source. Flows from multiple devices, such as a redundant pair of routers or switches, appear as separate data sources. However, the use of Virtual Switching System (VSS) on a pair of Cisco Catalyst 6500 Series switches allows flows from multiple interfaces on the redundant switch pair to appear as a single data source on the NAM. [Figure 6-1](#) shows an example of some key network points within an enterprise medianet where NetFlow collection can be enabled. Note that pairs of Cisco Catalyst 6500 Series Switches can be VSS-enabled, although not specifically shown.

Figure 6-1 Example of Collecting NetFlow Data at Key Network Points



This example is not the only recommended model for enabling NetFlow within an enterprise medianet, but is an example of a methodology for collecting NetFlow data to gain some useful insight regarding video flows at various points within the network infrastructure. You can choose to selectively enable NetFlow collection at one or more strategic aggregation points in the network, such as the distribution layer within different modules of a campus, depending on the desired visibility for video flows. For example, NetFlow statistics can be collected at the ingress interfaces of the distribution layer switch pairs at each module within a campus. In other words, statistics can be collected for traffic flows exiting the core and entering each campus module. Statistics gathered from this type of NetFlow deployment can be used to determine the following video traffic flows:

- Aggregated flows outbound across the corporate WAN to all the branch locations
- Flows into each building within the campus
- Aggregated flows outbound toward the Internet

This model can be useful because many video flows emanate from a central point within a campus data center or campus service module, and flow out to users within each campus building or each branch location. For example, unicast or broadcast enterprise TV as well as video-on-demand (VoD) flows to desktop devices often follow this flow pattern. Likewise, because of the nature of TelePresence video, the majority of the video flows within a multipoint meeting are from a centralized Cisco TelePresence Multipoint Switch, potentially located within a data center or campus service module, out to the Cisco TelePresence System endpoints located within the campus buildings and branch locations. Additional flow information can be gathered by implementing NetFlow bidirectionally at the distribution layer of each module. Note that this can preferably be done by enabling NetFlow statistics collection in an ingress direction on other interfaces. Although video broadcasts, VoD, and multipoint TelePresence tend to follow a flow model where the majority of traffic emanates from a central point outward to the endpoints, Cisco IP video surveillance follows the opposite model. The majority of traffic in a Cisco IP video surveillance deployment flows from cameras deployed within the campus buildings back to the Video Surveillance Operations Manager (VSOM) server potentially deployed within a data center or campus service module. However, note that implementing NetFlow collection bidirectionally can result in some duplication of flow information when multiple collection points exist within the network infrastructure.

Additional flow information can also be gathered by implementing NetFlow at the branch router itself, to gain insight into the flows into and out of individual branch locations, if that level of detail is needed. Keep in mind, however, that the NetFlow data export uses some of the available branch bandwidth. Also, NetFlow in Cisco IOS router platforms is performed in software, potentially resulting in somewhat higher CPU utilization depending on the platform and the amount of flow statistics collected and exported. The use of flow filters and/or sampling may be necessary to decrease both CPU utilization and bandwidth usage because of NetFlow flow record exports. Even with the campus distribution switches, it may be desirable to implement flow filters and/or sampling to decrease CPU and bandwidth usage. Note that data sampling may distort statistics regarding how much traffic is flowing across a single point in the network. However, the relative percentages of the flows can still be useful from a bandwidth allocation perspective. An alternative strategy may be to SPAN the flow traffic from the Cisco Catalyst switch to a separate device, such as the Cisco Service Control Engine (SCE), which can then perform analysis of the flows and export records to a centralized NetFlow collector for monitoring and reporting.


NetFlow Collector Considerations

The aggregation capabilities of the NetFlow collector determine to a large extent the usefulness of the NetFlow data from a medianet perspective. Most NetFlow collectors provide monitoring and historical reporting of aggregate bit rates, byte counts, and packet counts of overall IP data. Typically, this can be further divided into TCP, UDP, and other IP protocols, such as Internet Control Message Protocol (ICMP). However, beyond this level of analysis, some NetFlow collectors simply report Real-Time

Transport Protocol (RTP) traffic as “Other UDP” or “VoIP”, because RTP can use a range of UDP ports. Further, the ability to drill down to monitor and generate reports that show the specific hosts and flows that constitute RTP video and/or VoIP traffic, versus all UDP flows, may be limited. Further, both VoD, and in some cases video surveillance traffic, can be sent using HTTP instead of RTP. Therefore, generating useful reports showing medianet-relevant information, such as how much video data (RTP- and/or HTTP-based) is crossing a particular point within the network, may not be straightforward.

For devices such as TelePresence endpoints and IP video surveillance cameras, you can often simply assume that most of the data generated from the device is video traffic, and therefore use the overall amount of IP traffic from the device as a good estimate of the overall amount of video traffic generated by the device. Figure 6-2 shows a sample screen capture from a generic NetFlow collector, showing flow information from Cisco TelePresence System endpoints to a Cisco TelePresence Multipoint Switch, in a multipoint call.

Figure 6-2 Sample Host Level Reporting From a NetFlow Collector Showing TelePresence Endpoints



Note

Figure 6-2 shows a screen capture from the open source ntop NetFlow collector.

The IP addresses of the TelePresence devices have been replaced by a hostname to more easily identify the endpoints. As can be seen, both the actual traffic sent and received, in terms of bytes, as well as the percentage of the overall traffic seen across this particular interface over time are recorded. Such information may be useful from the perspective of determining whether the percentage of bandwidth allocated for TelePresence calls relative to other traffic, across the interfaces of this particular collection point, matches the actual data flows captured over an extended period of time. However, this information must also be used with caution. Flow records are exported by NetFlow based on the following:

- The flow transport has completed; for example, when a FIN or RST is seen in a TCP connection.
- The flow cache has become full. The cache default size is typically 64 K flow cache entries on Cisco IOS platforms. This can typically be changed to between 1024 and 524,288 entries.
- A flow becomes inactive. By default on Cisco IOS platforms, a flow unaltered in the last 15 seconds is classified as inactive. This can typically be set between 10 and 600 seconds.
- An active flow has been monitored for a specified number of minutes. By default on Cisco IOS platforms, active flows are flushed from the cache when they have been monitored for 30 minutes. You can configure the interval for the active timer between 1 and 60 minutes.

- Routing device default timer settings are 15 seconds for the inactive timer and 30 minutes for the active timer. You can configure your own time interval for the inactive timer between 10 and 600 seconds. You can configure the interval for the inactive timer between 10 and 600 seconds.


Long-lived flows such as TelePresence meetings may export flow data while the meeting is still ongoing. Therefore, the amount of data sent and/or received may not reflect the entire flow. In addition, the percentage of overall traffic does not indicate a particular timeframe, but more likely the percentage since collection began on the Netflow collector. The network administrator would benefit more from information that indicated the percentage of traffic during specific time intervals, such as peak times of the work day. Finally, the percentage of overall traffic represents an average over time, not peak usage, which may again be necessary to truly determine whether sufficient bandwidth is provisioned per service class across the medianet infrastructure.

The aggregation of flows based on type of service (ToS) may be useful from a medianet perspective, to characterize the amount or relative percentage of video traffic flows at given points within the network; provided the enterprise has deployed a QoS model that differentiates the various video flows into different service classes. This methodology also assumes a NetFlow collector capable of reporting flows based on ToS markings. NetFlow collectors such as the Cisco NAM Traffic Analyzer provide the ability to monitor and/or generate reports that show traffic flows based on DSCP values. [NAM Analysis of NetFlow Traffic, page 6-15](#) discusses this functionality. If a NetFlow collector that provides aggregation and reporting based on medianet-relevant parameters is not available, it may be necessary in some situations to develop custom applications that show the appropriate level of flow details to provide relevant reporting information from an enterprise medianet perspective.

NetFlow Export of Multicast Traffic Flows

From a medianet perspective, NetFlow version 9 offers the advantage of being able to export flow data from multicast flows. Multicast is often used to efficiently broadcast live video events across the enterprise IP infrastructure, rather than duplicate multiple unicast streams to each endpoint. [Figure 6-3](#) shows an example of multicast flows exported to a generic NetFlow collector.

Figure 6-3 Example of Multicast Flows Captured By a NetFlow Collector




Note

[Figure 6-3](#) shows a screen capture from the open source ntop NetFlow collector.

228415

Besides individual flows, which may be challenging to identify from all the other flow data, some NetFlow collectors can generate aggregate reporting information regarding the total amount of unicast, broadcast, and multicast flows seen at a given point within the network infrastructure. An example is shown in Figure 6-4.

Figure 6-4 Example of Aggregated Flow Data Reported By a NetFlow Collector



Note

Figure 6-4 shows a screen capture from the open source ntop NetFlow collector.

The combination of individual multicast flow information as well as aggregated flow information may be useful in determining whether sufficient bandwidth has been provisioned across a particular point within the medianet infrastructure to support existing multicast flows.

NetFlow Configuration Example

The configuration snippets in Example 6-1 and Example 6-2 show a basic NetFlow configuration on a Cisco Catalyst 6500 Series Switch as well as on a Cisco IOS router platform. Note that this example shows no flow filtering or sampling, which may be necessary to decrease CPU and/or bandwidth utilization for NetFlow collection in production environments.

Example 6-1 NetFlow Configuration on a Cisco Catalyst 6500 Series Switch

```

mls netflow                                ! Enables NetFlow on the PFC
mls flow ip interface-full                 ! Sets the NetFlow flow mask
mls nde sender                             ! Enables NetFlow device export
!
!
~
!
interface TenGigabitEthernet6/1
description CONNECTION TO ME-EASTCORE-1 TEN5/4
ip address 10.16.100.13 255.255.255.252
ip flow ingress                            ! Enables MSFC NetFlow ingress on the interface

```

```

ip multicast netflow ingress           ! Enables multicast NetFlow ingress on the interface
ip pim sparse-mode
no ip route-cache
load-interval 30
wrr-queue bandwidth 5 35 30
priority-queue queue-limit 30
wrr-queue queue-limit 5 35 30
wrr-queue random-detect min-threshold 3 60 70 80 90 100 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 3 70 80 90 100 100 100 100 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 2 1 0
wrr-queue cos-map 3 1 2
wrr-queue cos-map 3 2 3
wrr-queue cos-map 3 3 6
wrr-queue cos-map 3 4 7
priority-queue cos-map 1 4 5
mls qos trust dscp
!
interface TenGigabitEthernet6/2
description CONNECTION TO ME-EASTCORE-2 TEN1/1
ip address 10.16.100.1 255.255.255.252
ip flow ingress                       ! Enables MSFC NetFlow ingress on the interface
ip multicast netflow ingress         ! Enables multicast NetFlow ingress on the interface
ip pim sparse-mode
no ip route-cache
load-interval 30
udld port
wrr-queue bandwidth 5 35 30
priority-queue queue-limit 30
wrr-queue queue-limit 5 35 30
wrr-queue random-detect min-threshold 3 60 70 80 90 100 100 100 100
wrr-queue random-detect max-threshold 1 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 2 100 100 100 100 100 100 100 100
wrr-queue random-detect max-threshold 3 70 80 90 100 100 100 100 100
wrr-queue cos-map 1 1 1
wrr-queue cos-map 2 1 0
wrr-queue cos-map 3 1 2
wrr-queue cos-map 3 2 3
wrr-queue cos-map 3 3 6
wrr-queue cos-map 3 4 7
priority-queue cos-map 1 4 5
mls qos trust dscp
!
!
~
!
!
ip flow-export source Loopback0      ! Sets the source interface of NetFlow export packets
ip flow-export version 9             ! Sets the NetFlow export version to version 9
ip flow-export destination 10.17.99.2 3000 ! Sets the address & port of the NetFlow
                                           collector

```

Example 6-2 NetFlow Configuration on a Cisco IOS Router

```

interface GigabitEthernet2/0
description CONNECTS to BRANCH LAN SWITCH
ip address 10.31.0.1 255.255.255.252
ip flow ingress                       ! Enables NetFlow collection ingress on the interface
!
~
!

```

```

ip flow-export source Loopback0    ! Sets the source interface of NetFlow export packets
ip flow-export version 9           ! Sets the NetFlow export version to version 9
ip flow-export destination 10.16.4.10 2061 ! Sets the address and port of the NetFlow collector

```

For more information regarding the configuration of NetFlow on Cisco IOS routers, see the *Cisco IOS NetFlow Configuration Guide, Release 12.4* at the following URL:

http://www.cisco.com/en/US/docs/ios/netflow/configuration/guide/12_4/nf_12_4_book.html.

For more information regarding the configuration of NetFlow on Cisco Catalyst 6500 Series Switch platforms, see the following documents:

- Configuring NetFlow—
<http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/netflow.html>
- Configuring NDE—
<http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.2SX/configuration/guide/nde.html>

Cisco Network Analysis Module

The Cisco Network Analysis Module (NAM) enables network administrators to understand, manage, and improve how applications and services are delivered over network infrastructures. The NAM offers the following services:

- Flow-based traffic analysis of applications, hosts, and conversations
- Performance-based measurements on application, server, and network latency
- Quality of experience metrics for network-based services such as VoIP
- Problem analysis using packet captures

From an FCAPS management perspective, the NAM is most applicable as a performance management tool within an enterprise medianet, although both the packet capture and the monitoring statistics can also be used for fault management purposes. The current release of NAM software is version 4.1. The NAM software runs on the platforms listed in [Table 6-2](#). Specific hardware configurations and OS versions required for support of NAM modules and/or software can be found in the documentation for each specific platform.

Table 6-2 **NAM Platform Support**

Cisco Product Platform	NAM Model
Cisco Catalyst 6500 Series Switches and Cisco 7600 Series Routers	WS-SVC-NAM-1-250S or WS-SVC-NAM-2-250S
Cisco 3700 Series Routers; Cisco 2811, 2821, and 2851 Series ISRs; Cisco 3800 Series ISRs; Cisco 2911, 2921, and 2951 Series ISR G2s; Cisco 3900 Series ISR G2s	NME-NAM-120S
Cisco WAVE-574 and Cisco WAE-674 with NAM 4.1 Software Running on a Virtual Blade	NAM-WAAS-VB
Standalone NAM Appliance	NAM 2204 or NAM 2220 Appliance

This document discusses only the use of the Cisco Catalyst 6500 Series Network Analysis Module (WS-SVC-NAM-2). Specific testing was performed with a WS-SVC-NAM-2 with a WS-SUP32P-10GE supervisor within a Cisco Catalyst 6506-E chassis. Other platforms may have slightly different functionality. The WS-SVC-NAM-2 can analyze and monitor network traffic in the following ways:

- The NAM can analyze chassis traffic via Remote Network Monitoring (RMON) support provided by the Cisco Catalyst 6500 Series supervisor engine.
- The NAM can analyze traffic from local and remote NetFlow Data Export (NDE).
- The NAM can analyze Ethernet LAN traffic via Switched Port Analyzer (SPAN), remote SPAN (RSPAN), or VLAN ACL (VACL); allowing the NAM to serve as an extension to the basic RMON support provided by the Cisco Catalyst 6500 Series supervisor engine.


This document discusses only certain functionality of the NAM as it relates to gaining visibility into video flows within an enterprise medianet. A comprehensive discussion of the configuration and monitoring functionality of the NAM is outside the scope of this document. For the end-user and configuration guides for the Cisco Network Analysis Module Software, see the following URL: http://www.cisco.com/en/US/products/sw/cscowork/ps5401/tsd_products_support_series_home.html.

NAM Analysis of Chassis Traffic

The WS-SVC-NAM-2 has the ability to collect basic traffic statistics, per interface, from the supervisor line card within the Cisco Catalyst 6500 chassis. These statistics can be viewed as current rates or as cumulative data collected over time. Current rate data includes statistics such as the following:


- Input and output percentage utilization of the interface
- Input and output packets/second
- Input and output bit or byte rates
- Input and output non-unicast (multicast and broadcast) packets/second
- Input and output discards/second
- Input and output errors/second.

Figure 6-5 shows an example of the monitoring output.

Figure 6-5 Example of WS-SVC-NAM-2 Traffic Analyzer Chassis Interface Statistics

From a medianet management perspective, these statistics can be used for high-level troubleshooting of traffic crossing the particular chassis, because even small rates of packet discards or interface errors can result in degraded video quality. Note, however, that the interface statistics alone cannot be used to determine whether video traffic is being discarded, because packet discards can be occurring only within a particular switch port queue that may or may not hold video traffic. However, as is discussed in [Router and Switch Command-Line Interface, page 6-35](#), many router and switch platforms can show drops down to the level of individual queues within the CLI.


If mini-RMON port statistics are enabled, the WS-SVC-NAM-2 provides slightly more information regarding the types of errors encountered per interface, including undersized and oversized packets, Cyclic Redundancy Check (CRC) errors, fragments, jabbers, and collisions. [Figure 6-6](#) provides an example showing port statistics on the uplink ports of a Cisco Catalyst 6500 switch.

Figure 6-6 Example of W-SVC-NAM-2 Traffic Analyzer Chassis Port Statistics

The port statistics can also provide information regarding the amount multicast traffic crossing the interfaces. When viewed as current rates, the NAM port statistics show the number of multicast packets/second seen by the interface. These can be graphed in real time as well, or viewed as cumulative data. The port statistics do not show current rates in terms of bits/second or bytes/seconds for multicast data, which would be useful for determining bandwidth provisioning for multicast traffic. However, the design engineer can still gain some visibility into the amount of multicast traffic crossing a particular interface on the Cisco Catalyst 6500 through the WS-SVC-NAM-2 port statistics.

If the WS-SVC-NAM-2 is installed within a Cisco Catalyst 6500 chassis that contains a Sup-32 PISA supervisor, you have the option of enabling Network-Based Application Recognition (NBAR) analysis of traffic forwarded through the supervisor, on a per-interface basis. NBAR adds the ability to analyze the traffic statistics collected through the supervisor at the protocol level. An example is shown in Figure 6-7.

Figure 6-7 Example of NAM Traffic Analyzer with NBAR Enabled



As before, the data can be viewed as current rates or as cumulative data. Individual protocol rates can also be graphed in real-time. As can be seen in Figure 6-7, NBAR has the ability to identify audio and video media as RTP streams, along with Real-Time Control Protocol (RTCP) control channels. NBAR can also identify signaling protocols, such as SIP. Therefore, NBAR can provide useful information regarding how much video traffic is crossing interfaces of the particular Cisco Catalyst 6500 chassis. This information may be used in determining whether sufficient bandwidth has been provisioned for a particular type of traffic, such as RTP. However, the combination of the NAM with NBAR still does not specifically identify a particular type of RTP flow as possibly being an IP video surveillance flow or a desktop video conferencing flow. Also, because different RTP flows from different video applications can be configured for different service classes, they may be placed into separate egress queues on the Cisco Catalyst 6500 switch ports. Therefore, simply knowing the aggregate bit rate of RTP flows through an interface still does not necessarily provide the level of detail to determine whether sufficient bandwidth is allocated per service class, and therefore per queue, on the particular Cisco Catalyst switch port. As is discussed in the next section, the NetFlow Data Export and SPAN monitoring functionality of the NAM can provide further detailed information to assist in determining whether sufficient bandwidth has been provisioned per service class.

NAM Analysis of NetFlow Traffic

As mentioned in [NetFlow Strategies Within an Enterprise Medianet, page 6-6](#), the NAM Traffic Analyzer can also function as a NetFlow collector. This allows the NAM to analyze traffic flows from remote devices within the enterprise medianet, without having to use the SPAN and RSPAN functionality


of Cisco Catalyst switches. Although NetFlow provides less information than a SPAN or RSPAN of the actual traffic, the overall bandwidth utilization can be significantly less, and NetFlow can therefore be far more scalable as a mechanism to view traffic flow data throughout a medianet. In this type of configuration, NetFlow traffic statistics can be collected from remote switches and routers throughout the enterprise medianet and forwarded to one or more WS-SVC-NAM-2 modules centrally located, perhaps within a Cisco Catalyst 6500 service switch within a campus data center service module. Alternatively NetFlow traffic may be forwarded to a NAM 2200 Series Appliance.

To configure NetFlow collector functionality within the NAM, each remote NetFlow Data Export (NDE) device must be added to the NetFlow Devices screen of the NAM web-based GUI, as shown in [Figure 6-8](#). An optional SNMP v1/2c read-only community string can be configured to allow the NAM to include the configured description next to interface definitions.

**Note**

Note that the NAM does not currently support SNMP v3.


Figure 6-8 Configuration of NetFlow Devices within the NAM



The NAM allows multiple interfaces on a single physical device to be treated as a single NetFlow custom data source. Interfaces from different devices cannot currently be aggregated into a single data source. This means that redundant pairs of switches or routers that load balance traffic (as shown in [Figure 6-1](#)) appear as multiple data sources to the NAM Traffic Analyzer. You may have to manually combine the results from the individual NetFlow data sources to gain an understanding of the total traffic flows through a given set of redundant devices. The exception to this is if VSS is deployed across a pair of redundant Cisco Catalyst 6500 switches. VSS allows a redundant pair of Cisco Catalyst 6500s to appear

as a single device. Therefore, the NetFlow statistics from multiple interfaces on both switches can appear as a single data set. Figure 6-9 shows an example of how multiple interfaces on a single device are aggregated into a single NetFlow custom data source on the NAM.

Figure 6-9 Configuration of Custom Data Sources on the NAM



From a medianet management perspective, one of the attractive features of the NAM as a NetFlow collector is its ability to monitor and generate reports on traffic flows, based on their DSCP values. If the various video applications running over the network are separated into different service classes, gaining visibility into the amount of traffic per service class allows you to gain visibility into the amount of traffic that a particular application is generating across key parts of the medianet infrastructure. To accomplish this, you first need to create a diffserv aggregation profile that maps traffic with different DSCP values into aggregation groups for reporting. An example of an aggregation profile based on the Cisco enterprise 12-class QoS model is shown in Figure 6-10.

Figure 6-10 Diffserv Profile Based on the Cisco Enterprise 12-Class QoS Model

Profile Name	DSCP Value	Group Description
12-Class-QoS	0	Best-Effort-CS0
	1	
	2	
	3	
	4	
	5	
	6	
	7	
	8	Scavenger-CS1
	9	
	10	Bulk-Data-AF11
	11	Bulk-Data-AF12

The screenshot shows the 'DiffServ Profile Setup' page in the NAM Traffic Analyzer. The page has a sidebar with navigation links: Home, Monitor, Reports, Capture, Alarms, Admin. The main content area shows the 'DiffServ Profile Setup' page. The page has a table with 12 rows, each representing a DSCP value and its corresponding group description. The table is titled 'DiffServ Profile Setup' and has columns for 'Profile Name', 'DSCP Value', and 'Group Description'. The 'Profile Name' is '12-Class-QoS'. The 'DSCP Value' column contains values from 0 to 11. The 'Group Description' column contains descriptions for each DSCP value. The table is as follows:

As can be seen, each of the DSCP markings corresponds to one of the 12 QoS classes. Because assured forwarding (AF) traffic may be marked down (for example from AFx1 to AFx2 or AFx3) within service provider Multiprotocol Label Switching (MPLS) networks, these have been added to the diffserv aggregation profile as separate aggregation groups in the example above. This can provide additional value in that you may be able determine whether traffic within the particular assured forwarding traffic classes is being marked down because the traffic rates are outside the contracted rates of the service

provider network. The downside to this approach, however, is that you may have to manually combine the monitoring and reporting statistics from the separate aggregation groups to gain a view of all the traffic within a single assured forwarding (AFx1, AFx2, and AFx3) class. Alternatively, the AFx1, AFx2, and AFx3 traffic can be placed into a single aggregation group (for instance AF41, AF42, and AF42 all placed into a multimedia-conferencing group). This makes it easier to view the overall amount of traffic within a particular AF class, but at the loss of the information regarding if or how much of the traffic was marked down.

After the diffserv aggregation profile has been created, it must be applied to monitor each data source in which it is desired to see traffic statistics, application statistics, and/or host statistics; based on the aggregation groupings defined within in the profile. An example of this is shown in Figure 6-11, in which the 12-Class-QoS diffserv profile has been applied to the NDE source corresponding to a WAN edge router.

Figure 6-11 Application of the Diffserv Aggregation Profile to an NDE Source

NAM Traffic Analyzer

Setup Monitor Reports Capture Alarms Admin

Chassis Parameters Data Sources Monitor Protocol Directory Alarms Preferences

You Are Here: Setup > Monitor > DiffServ > Monitoring

DiffServ Monitor Setup


Data Source: NDE-me-westwan-1 [Filter] [Clear]

DiffServ Profile: 12-Class-QoS

Monitoring Function	Max Entries
☒ Traffic Statistics	Not applicable
☒ Application Statistics	100
☒ IP Host Statistics	100


← Select an item then take an action → [Apply] [Reset]

When applied, the traffic, application, and/or IP host statistics can be viewed as current rates or cumulative data. Figure 6-12 shows an example of the output from the traffic statistics shown as current rates.

Figure 6-12 Traffic Statistics per Service Class from an NDE Source


The example above shows the breakout of traffic flows from a campus core to a WAN edge switch (as shown in Figure 6-1). This level of traffic analysis may be used to assist in determining whether the provisioning of traffic on existing WAN policy maps is appropriate for the actual traffic levels that cross the WAN interfaces. Policy maps on Cisco IOS router platforms are often configured to allow applications to exceed the allocated bandwidth for a particular service class, if available bandwidth exists on the WAN link. Therefore, just because no drops are being seen on a particular service class on a WAN link, does not mean the provisioned bandwidth is sufficient for the traffic within that service class. The service class may be borrowing from other service classes. Visibility into the amount of actual traffic flows per service class can help ensure that you allocate the appropriate amount of bandwidth per service class.

You can drill down further into each service class to identify particular application flows, based on their TCP or UDP port numbers. This is done through the Diffserv Application Statistics screen, as shown in Figure 6-13.

Figure 6-13 Application Statistics per Service Class from an NDE Source

Here, note again what was previously mentioned in [NetFlow Collector Considerations, page 6-7](#). The NAM itself cannot identify the particular application flows per service class as being IP video surveillance flows, TelePresence flows, or VoD flows. However, if different video applications are separated into different service classes, you may be able to determine to which video application the flows belong. For example, in the network used for the example in [Figure 6-13](#), only Cisco TelePresence traffic was placed in the Real-Time Interactive service class. Therefore, you can easily identify that the flows within [Figure 6-13](#) represent TelePresence meetings. By selecting any one of the flows and clicking the Details button, you can see the host IP addresses that generated the flows. Alternatively, you can drill down into each service class to identify particular hosts responsible for the flows, based on their IP addresses. This is done through the Diffserv Application Hosts screen, as shown in [Figure 6-14](#).

228425

Figure 6-14 Host Statistics per Service Class from an NDE Source

Note that if the particular device is an application-specific video device, such as a Cisco TelePresence System endpoint or an IP video surveillance camera, DNS address translation may be useful to provide a meaningful name that indicates the type of video device instead of an IP address.

NAM Analysis of SPAN/RSPAN Traffic

When configured to analyze traffic that has been sent to the WS-SVC-NAM-2 via the Cisco Catalyst 6500 SPAN or RSPAN features, the NAM provides the same ability to monitor and generate reports for traffic based on service class, as was discussed in the previous section. In addition, the NAM can provide more detailed monitoring of RTP streams included within the SPAN or RSPAN traffic flows. An example of the RTP stream traffic is shown in Figure 6-15.

Figure 6-15 NAM RTP Stream Traffic

NAM Traffic Analyzer

Overview | Apps | Voice/Video | Hosts | Conversations | VLAN | DiffServ | Response Time | Chassis | MPLS

You Are Here: Monitor > Voice/Video > RTP Stream Traffic

RTP Stream Traffic

Current Data: as of Mon 07 Dec 2009, 05:34:39 EST

Auto Refresh

Act Pkt Loss /million

Showing 1-15 of 64 records

#	Source Addr : Port	Dest Addr : Port	Payload Type	SSRC	Act Pkt Loss /million	Worst MOS	Adj Pkt Loss (%)	Jitter (ms)	Total SSC	Status	Start Time
1	10.17.1.11 : 20206	10.16.1.20 : 16422	Video	170897367	116.010	-	11.60	0.64	851	Inactive	12-02-09 12:55:47 America/New_York
2	10.17.1.11 : 25076	10.16.1.20 : 16642	Video	150471625	117.556	-	11.76	0.70	90	Inactive	12-03-09 01:41:57 America/New_York
3	10.17.1.11 : 20206	10.17.64.11 : 24458	Video	537447363	49.690	-	4.97	0.42	104	Inactive	11-24-09 07:15:06 America/New_York
4	10.17.1.11 : 20206	10.17.64.11 : 24458	Video	1031274411	49.508	-	4.95	0.23	102	Inactive	11-24-09 07:15:06 America/New_York
5	10.17.1.11 : 20206	10.17.64.11 : 24458	Video	1759311971	49.825	-	4.98	0.25	105	Inactive	11-24-09 07:15:06 America/New_York
6	10.17.1.11 : 26662	10.17.64.11 : 22260	Video	269940599	61.065	-	6.11	0.48	25	Inactive	12-04-09 02:53:01 America/New_York
7	10.17.1.11 : 26662	10.17.64.11 : 22260	Video	938462907	61.303	-	6.13	0.50	25	Inactive	12-04-09 02:53:01 America/New_York
8	10.17.1.11 : 26662	10.17.64.11 : 22260	Video	1094383967	61.189	-	6.12	0.31	25	Inactive	12-04-09 02:53:01 America/New_York
9	10.17.1.11 : 27246	10.26.1.11 : 22100	Video	1786561635	152.820	-	15.28	0.47	2044	Inactive	12-01-09 03:09:16 America/New_York
10	10.17.1.20 : 16514	10.16.1.29 : 10000	Unknown	672430737	47.823	-	4.78	0.00	0	Inactive	12-02-09 10:22:20 America/New_York

As can be seen from Figure 6-15, the NAM has the ability to collect detailed performance data from RTP flows down to individual synchronization sources (SSRCs), including packet loss counts for the session, packet loss percentages for the session, jitter, and whether the RTP session is still active. Further information can be viewed by highlighting and selecting the details regarding each individual flow, as shown in Figure 6-16.

Figure 6-16 RTP Flow Details

Current Data: as of Mon 07 Dec 2009, 06:37:57 EST

RTP Stream Details

Source IP Address / Port: 10.17.1.11 / 20206

Destination IP Address / Port: 10.17.64.11 / 24458

Payload Type: Video

SSRC: 537447363

Total Duration Monitored (seconds): 1639

Stream Lifetime

Worst / Avg / Max MOS: - / - / -

Worst / Avg / Min Jitter (ms): 0.66 / 0.42 / 0.27

Worst / Avg / Min Adjusted Packet Loss (%): 50.05 / 4.97 / 0.00

Worst / Avg / Min Actual Packet Loss (%): 50.05 / 4.97 / 0.00

Worst / Avg / Min Seconds of Concealment: 31.0 / 3.5 / 0.0

Worst / Avg / Min Seconds of Severe Concealment: 31.0 / 3.4 / 0.0

Last N Reports

Report Timestamp	Report Duration (seconds)	Worst MOS	Jitter (ms)	Adjusted Pkt Loss (%)	Actual Pkt Loss (%)	Seconds of Concealment	Seconds of Severe Concealment
11-24-09 09:25:04 America/New_York	14	-	0.62	1.10	1.10	1.0	1.0
11-24-09 09:10:59 America/New_York	19	-	0.38	0.00	0.00	0.0	0.0
11-24-09 09:10:12 America/New_York	48	-	0.46	0.84	0.84	1.0	1.0
11-24-09 06:43:00 America/New_York	21	-	0.38	0.00	0.00	0.0	0.0
11-24-09 06:41:22 America/New_York	35	-	0.66	42.62	42.62	27.0	26.0

Close

Here you can view the flow in increments over time, to see whether the packet loss and high jitter levels were a one-time event during the session, or were continuous throughout the session. This level of detail can be used to assist in identifying performance issues within the network down to the level of individual cameras within a multi-screen (CTS-3000) TelePresence meeting.

Cisco IP Service Level Agreements

Cisco IPSLAs are an active traffic monitoring utility for measuring network performance. IPSLA support is included within most Cisco IOS router platforms, Cisco Catalyst switch platforms (including the Cisco Catalyst 6500, Cisco Catalyst 4500, and Cisco Catalyst 3750E Series), and some Cisco video endpoints such as Cisco TelePresence Systems endpoints. IPSLAs operate in a sender/responder configuration. Typically a Cisco IOS router or switch platform is configured as a source (the IPSLA sender) of packets, otherwise known as IPSLA probes, which are crafted specifically to simulate a particular IP service on the network. These packets are sent to the remote device (the IPSLA responder), which may loop the packets back to the IPSLA Sender. In this manner, enterprise medianet service level parameters such as latency, jitter, and packet loss can be measured. There are a variety of Cisco IPSLA operations, meaning that various types of IP packets can be generated by the IPSLA sender and returned by the IPSLA responder. Depending on the particular platform, these can include the following operations:

- UDP jitter
- ICMP path jitter
- UDP jitter for VoIP
- UDP echo
- ICMP echo
- ICMP path echo
- HTTP
- TCP connect
- FTP
- DHCP
- DNS
- Data Link Switching Plus (DLSW+)
- Frame Relay

For a discussion of each of the different IPSLA operations and how to configure them on Cisco IOS router platforms, see the *Cisco IOS IPSLAs Configuration Guide, Release 12.4* at the following URL: http://www.cisco.com/en/US/docs/ios/12_4/ip_sla/configuration/guide/hsla_c.html.

IPSLAs as a Pre-Assessment Tool

From an FCAPS management perspective, IPSLAs are most applicable as a performance management tool within an enterprise medianet. They can be used to pre-assess the ability of the IP network infrastructure to support a new service, such as the deployment of Cisco TelePresence, between two points within the network. Because most video flows are RTP-based, the UDP jitter IPSLA operation typically has the most relevance from a medianet pre-assessment perspective.

**Note**

Note that many video flows use other transport protocols. For example, both VoD and MJPEG-based IP video surveillance may use HTTP as the transport protocol instead of RTP.

The usefulness of IPSLAs as a pre-assessment tool depends to a large extent on the knowledge of the medianet video flow that is to be simulated by IPSLA traffic, and whether an IPSLA operation can be crafted to accurately replicate the medianet video flow. This can be particularly challenging for high


definition video for several reasons. First, video flows are sent as groups of packets every frame interval. These groups of packets can be bunched-up at the beginning of the frame interval, or spread evenly across the frame interval, depending on how the video application (that is, the transmitting codec) is implemented. Also, each packet within a single frame can vary in size. Operations such as the UDP jitter IPSLA operation transmit fixed-sized packets at regular intervals, similarly to VoIP. Second, high definition video frames often consist of more than ten packets per frame, meaning that the interval between the individual packets sent within a single video frame can vary from less than one millisecond to several milliseconds. Observations on a highly loaded Cisco Catalyst 6500 with a Sup-32 processor have shown that individual UDP jitter IPSLA operations can generate packets with intervals of 4 milliseconds or greater with large packet payload sizes. Smaller platforms such as the Cisco 2800 Series ISR may be capable of generating packets with intervals of only 8–12 milliseconds or greater, depending on the loading of the platform.

**Note**

Although some platforms allow configuration of intervals down to one millisecond, the design engineer may find it necessary to capture a data trace of the IPSLA probes to determine the actual frame rate generated by the IPSLA sender. Partly because the loading on the CPU affects the rate at which IPSLA probes are generated, pre-assessments services of deployments such as Cisco TelePresence are often performed with dedicated ISR router platforms. Likewise, some organizations deploy router platforms at campus and branch locations dedicated for IPSLA functions.

Crafting one or more UDP jitter IPSLA operations that accurately replicate the size of the individual packets sent, the interval between individual packets sent, and the frame-based nature of video can be challenging. These attributes are important to factor in because network parameters such as jitter and packet loss are often largely dependent on the queue depths and buffer sizes of networking gear along the path between the endpoints. Sending a smooth flow of evenly spaced packets, or larger packets less frequently, may result in significantly different results than the actual video flows themselves.


As an example, to accurately pre-assess the ability of the network to handle a flow such as a TelePresence endpoint, you must craft a sequence of packets that accurately simulates the endpoint. [Figure 6-17](#) shows a close-up of the graph from a data capture of the video stream from a Cisco TelePresence CTS-1000 running Cisco TelePresence System version 1.6 software.

Figure 6-17 Detailed Graph of a CTS-1000 Video Stream (Version 1.6)

Approximate Total Bit Rate : 16 Pkts * 1,100 Bytes / Pkt = 17,600 Bytes Per 33 msec Frame
 17,600 Bytes / Frame * 30 Frames / sec = 528 Kbytes / sec
 528 Kbytes / Sec * 8 bits / Byte = 4.224 Mbps

As can be seen from Figure 6-17, TelePresence video packets average slightly under 1100 bytes in size. Each video frame consists of approximately 16 packets, spaced from 1–3 msec apart, spread across the 33 msec frame interval. Based on this analysis, a UDP jitter IP SLA operation consisting of 1060 byte packets with a interval of 2 msec between packets, sent with a ToS value equivalent to CS4 traffic, would simulate the size and packet rate of a TelePresence video stream. The overall data rate would be approximately 1060 bytes/packet * 500 packets/sec * 8 bits/byte = 4.24 Mbps.

Figure 6-18 shows a close-up of the graph from a data capture of a single audio stream from a TelePresence CTS-1000 running Cisco TelePresence System version 1.6 software.


Figure 6-18 Detailed Graph of a CTS-1000 Audio Stream (Version 1.6)

Approximate Total Bit Rate : $225 \text{ Bytes / Packet} * 50 \text{ Packets / Sec}$
 $11,250 \text{ Bytes / Sec} * 8 \text{ Bits / Byte} = 90 \text{ Kbps}$

As shown in [Figure 6-18](#), TelePresence audio packets are approximately 225 bytes in size, sent every 20 msec. Based on this analysis, a UDP jitter IP SLA operation consisting of 225-byte packets with a interval of 20 msec between packets, sent with a ToS value equivalent to CS4 traffic (because Cisco TelePresence sends audio with the same marking as video) would simulate the size and packet rate of a single TelePresence audio stream. The overall data rate would be approximately $225 \text{ bytes/packet} * 50 \text{ packets/sec} * 8 \text{ bits/byte} = 90 \text{ Kbps}$.

As previously mentioned, however, a lightly loaded Cisco Catalyst 6500 with Sup-32 processor was observed to be able to generate packets with a minimum packet interval of only 4 milliseconds. Therefore, one method of simulating the number of packets and their sizes within the TelePresence video stream is to implement two UDP jitter IPSLA operations on the Cisco Catalyst 6500, each with a packet interval of 4 milliseconds, and to schedule them to start simultaneously. A third UDP jitter IPSLA operation can also be run simultaneously to simulate the audio stream. [Figure 6-19](#) shows a close-up of the graph from a data capture of the actual data stream from these UDP jitter IPSLA operations.

Figure 6-19 Detailed Graph of Multiple UDP Jitter IPSLA Operations Simulating a Cisco TelePresence CTS-1000



From Figure 6-19, it appears that UDP jitter IPSLA operations #1 and #2 space their packets each 1 millisecond apart. However, this is just because of the graphing of the data points. The actual data trace reveals that the Cisco Catalyst 6500 switch sends packets from both UDP jitter IPSLA operations roughly back-to-back every four milliseconds. Therefore, the IPSLA-simulated video packets are slightly more clumped together than actual TelePresence video packets, but still considered acceptable from a pre-assessment perspective. The third UDP jitter IPSLA operation generates a simulated audio stream of packets every 20 milliseconds. Note that a CTS-1000 can receive up to three audio streams and an additional auxiliary video stream for presentations. Simulation of these streams are not shown in this example for simplicity. However, the method discussed above can be extended to include those streams as well if needed. Likewise, the method may be used to simulate other video flows simply by capturing a data trace, analyzing the flow, and setting up the appropriate IPSLA operations.

The configuration snippet on Example 6-3 shows the configuration of the UDP jitter IPSLA operations on the Cisco Catalyst 6500 switch that were used to create the simulation from which the data in Figure 6-19 was captured.

Example 6-3 IPSLA Sender Configuration on a Cisco Catalyst 6500 Series Switch

```
ip sla monitor 24
  type jitter dest-ipaddr 10.24.1.11 dest-port 32800 source-ipaddr 10.16.1.1 source-port
  32800 num-packets 16500 interval 2
  request-data-size 1018
  tos 128
ip sla monitor 25
  type jitter dest-ipaddr 10.24.1.11 dest-port 32802 source-ipaddr 10.16.1.1 source-port
  32802 num-packets 16500 interval 2
  request-data-size 1018
  tos 128
ip sla monitor 26
  type jitter dest-ipaddr 10.24.1.11 dest-port 32804 source-ipaddr 10.16.1.1 source-port
  32804 num-packets 3300 interval 20
```

```

request-data-size 183
tos 128
!
ip sla monitor group schedule 1 24,25,26 schedule-period 1 frequency 70 start-time now
life 700
!
~

```

Even though the packet interval has been configured at 2 milliseconds for *ip sla monitor 25* and *ip sla monitor 26*, the real interval between packets was observed to be 4 milliseconds. Sending 16,500 packets spaced at 4 milliseconds apart takes approximately 66 seconds. The configuration of *ip sla monitor group schedule 1* with a schedule period of one second causes the three UDP jitter operations to simultaneously start. The frequency of 70 seconds ensures that the previous operations complete before they begin again. The operation was set to run for approximately 10 intervals, or 700 seconds. Note that the length of time needed to perform a real assessment of a network to support a service such as a CTS-1000 is completely at the discretion of the network administrator. The aggregated output from the IPSLA tests can be displayed via the **show ip sla monitor statistics aggregated details** command on the Cisco Catalyst 6500 switch, as shown in [Example 6-4](#). It shows the packet loss; minimum and maximum jitter; and minimum, maximum and average latency for each of the three UDP jitter IPSLA operations.

Example 6-4 IPSLA Aggregated Statistics on a Cisco Catalyst 6500 Series Switch

```

me-eastdc-1#show ip sla monitor statistics aggregated details
Round trip time (RTT)      Index 24
Start Time Index: .10:53:07.852 EST Mon Nov 16 2009
Type of operation: jitter
Voice Scores:
MinOfICPIF: 0      MaxOfICPIF: 0      MinOfMOS: 0      MaxOfMOS: 0
RTT Values
    Number Of RTT: 94674
    RTT Min/Avg/Max: 1/1/7
Latency one-way time milliseconds
    Number of Latency one-way Samples: 0
    Source to Destination Latency one way Min/Max: 0/0
    Destination to Source Latency one way Min/Max: 0/0
    Source to Destination Latency one way Sum/Sum2: 0/0
    Destination to Source Latency one way Sum/Sum2: 0/0
Jitter time milliseconds
    Number of Jitter Samples: 94664
    Source to Destination Jitter Min/Max: 1/4
    Destination to Source Jitter Min/Max: 1/6
    Source to destination positive jitter Min/Avg/Max: 1/1/4
    Source to destination positive jitter Number/Sum/Sum2: 1781/1849/2011
    Source to destination negative jitter Min/Avg/Max: 1/1/4
    Source to destination negative jitter Number/Sum/Sum2: 1841/1913/2093
    Destination to Source positive jitter Min/Avg/Max: 1/1/6
    Destination to Source positive jitter Number/Sum/Sum2: 3512/3532/3632
    Destination to Source negative jitter Min/Avg/Max: 1/1/6
    Destination to Source negative jitter Number/Sum/Sum2: 3447/3468/3570
    Interarrival jitterout: 0      Interarrival jitterin: 0
Packet Loss Values
    Loss Source to Destination: 0      Loss Destination to Source: 0
    Out Of Sequence: 0      Tail Drop: 0      Packet Late Arrival: 0
Number of successes: 10
Number of failures: 0
Failed Operations due to over threshold: 0
Failed Operations due to Disconnect/Timeout/Busy/No Connection: 0/0/0/0
Failed Operations due to Internal/Sequence/Verify Error: 0/0/0
Distribution Statistics:
Bucket Range: 0-19 ms

```

```

Avg. Latency: 0 ms
Percent of Total Completions for this Range: 100 %
Number of Completions/Sum of Latency: 10/3
Sum of RTT squared low 32 Bits/Sum of RTT squared high 32 Bits: 3/0
Operations completed over thresholds: 0

Round trip time (RTT)   Index 25
Start Time Index: .10:53:07.856 EST Mon Nov 16 2009
Type of operation: jitter
Voice Scores:
MinOfICPIF: 0    MaxOfICPIF: 0    MinOfMOS: 0    MaxOfMOS: 0
RTT Values
    Number Of RTT: 94672
    RTT Min/Avg/Max: 1/1/8
Latency one-way time milliseconds
    Number of Latency one-way Samples: 0
    Source to Destination Latency one way Min/Max: 0/0
    Destination to Source Latency one way Min/Max: 0/0
    Source to Destination Latency one way Sum/Sum2: 0/0
    Destination to Source Latency one way Sum/Sum2: 0/0
Jitter time milliseconds
    Number of Jitter Samples: 94662
    Source to Destination Jitter Min/Max: 1/4
    Destination to Source Jitter Min/Max: 1/7
    Source to destination positive jitter Min/Avg/Max: 1/1/3
    Source to destination positive jitter Number/Sum/Sum2: 2498/2559/2691
    Source to destination negative jitter Min/Avg/Max: 1/1/4
    Source to destination negative jitter Number/Sum/Sum2: 2553/2620/2778
    Destination to Source positive jitter Min/Avg/Max: 1/1/7
    Destination to Source positive jitter Number/Sum/Sum2: 4470/4511/4725
    Destination to Source negative jitter Min/Avg/Max: 1/1/6
    Destination to Source negative jitter Number/Sum/Sum2: 4413/4448/4622
    Interarrival jitterout: 0    Interarrival jitterin: 0
Packet Loss Values
    Loss Source to Destination: 0    Loss Destination to Source: 0
    Out Of Sequence: 0    Tail Drop: 0    Packet Late Arrival: 0
Number of successes: 10
Number of failures: 0
Failed Operations due to over threshold: 0
Failed Operations due to Disconnect/TimeOut/Busy/No Connection: 0/0/0/0
Failed Operations due to Internal/Sequence/Verify Error: 0/0/0
Distribution Statistics:
Bucket Range: 0-19 ms
Avg. Latency: 0 ms
Percent of Total Completions for this Range: 100 %
Number of Completions/Sum of Latency: 10/5
Sum of RTT squared low 32 Bits/Sum of RTT squared high 32 Bits: 5/0
Operations completed over thresholds: 0

Round trip time (RTT)   Index 26
Start Time Index: .10:53:02.892 EST Mon Nov 16 2009
Type of operation: jitter
Voice Scores:
MinOfICPIF: 0    MaxOfICPIF: 0    MinOfMOS: 0    MaxOfMOS: 0
RTT Values
    Number Of RTT: 16500
    RTT Min/Avg/Max: 1/1/8
Latency one-way time milliseconds
    Number of Latency one-way Samples: 0
    Source to Destination Latency one way Min/Max: 0/0
    Destination to Source Latency one way Min/Max: 0/0
    Source to Destination Latency one way Sum/Sum2: 0/0
    Destination to Source Latency one way Sum/Sum2: 0/0
Jitter time milliseconds

```

```

Number of Jitter Samples: 16490
Source to Destination Jitter Min/Max: 1/4
Destination to Source Jitter Min/Max: 1/6
Source to destination positive jitter Min/Avg/Max: 1/1/4
Source to destination positive jitter Number/Sum/Sum2: 440/457/505
Source to destination negative jitter Min/Avg/Max: 1/1/4
Source to destination negative jitter Number/Sum/Sum2: 496/512/558
Destination to Source positive jitter Min/Avg/Max: 1/1/6
Destination to Source positive jitter Number/Sum/Sum2: 571/587/679
Destination to Source negative jitter Min/Avg/Max: 1/1/6
Destination to Source negative jitter Number/Sum/Sum2: 513/529/621
Interarrival jitterout: 0      Interarrival jitterin: 0
Packet Loss Values
  Loss Source to Destination: 0      Loss Destination to Source: 0
  Out Of Sequence: 0      Tail Drop: 0      Packet Late Arrival: 0
Number of successes: 10
Number of failures: 0
Failed Operations due to over threshold: 0
Failed Operations due to Disconnect/Timeout/Busy/No Connection: 0/0/0/0
Failed Operations due to Internal/Sequence/Verify Error: 0/0/0
Distribution Statistics:
Bucket Range: 0-19 ms
Avg. Latency: 1 ms
Percent of Total Completions for this Range: 100 %
Number of Completions/Sum of Latency: 10/10
Sum of RTT squared low 32 Bits/Sum of RTT squared high 32 Bits: 10/0
Operations completed over thresholds: 0

```

For the example above, the IPSLA responder was an actual Cisco TelePresence CTS-1000. Only IPSLA responder operations can be configured on Cisco TelePresence System endpoints; they cannot function as IPSLA sources. Configuration is only via the SSH CLI, as shown in [Example 6-5](#).

Example 6-5 IPSLA Responder Configuration on a CTS-1000

```

admin: utils ipsla responder initiators add net 10.16.1.0/24
admin: utils ipsla responder enable start

```

The configuration above enables the IPSLA responder function for initiators (senders) on the 10.16.1.0/24 subnet. This corresponds to the source of the IPSLA packets from the Cisco Catalyst 6500. By default, the range of ports enabled on the CTS-1000 is from 32770 to 33000. However, the port range can be enabled by including start and end ports within the **utils ipsla responder enable** command. For a discussion of all the commands available via the SSH CLI, including all the IPSLA commands, see the *Cisco TelePresence System Release 1.6 Command-Line Interface Reference Guide* at the following URL: http://www.cisco.com/en/US/docs/telepresence/cts_admin/1_6/CLI/cts1_6cli.html.


The use of IPSLA as a pre-assessment tool can be disruptive to existing traffic on the IP network infrastructure. After all, the objective of the pre-assessment test is to see whether the network infrastructure can support the additional service. For example, if a particular link within the network infrastructure has insufficient bandwidth, or a switch port has insufficient buffering capacity to support existing TelePresence traffic as well as the additional traffic generated from the IPSLA pre-assessment tests, both the existing TelePresence call and the IPSLA operation show degraded quality during the tests. You must therefore balance the possibility of temporarily degrading production services on the network against the value of the information gathered from running an IPSLA pre-assessment test during normal business hours. Running the IPSLA tests after hours may not accurately assess the ability of the network to handle the additional service, because after-hour traffic patterns may vary significantly from traffic patterns during normal business hours. Further, running a successful pre-assessment test after hours may lead to the installation of a production system that then results in degraded quality both for itself and for other production systems during normal business hours.

Finally, when multiple redundant equal-cost paths exist within the medianet infrastructure, Cisco Express Forwarding (formerly known as CEF) load balances the traffic across the equal-cost paths using a hash of the source and destination IP addresses for each session. Each router and switch along the path independently creates its Cisco Express Forwarding table based on IP routing protocols, and load balances sessions across its interfaces that represent equal-cost paths to the next hop along the path to the destination. An IPSLA probe generated by a switch or router has a different IP source address than the actual video device that is being pre-assessed. Therefore, the path taken by the IPSLA probes within a highly redundant network infrastructure may not be exactly the path taken by the actual video traffic from the device. The use of dedicated routers to perform an IPSLA network assessment eases this issue slightly, because the routers can be configured to use the actual IP addresses that the video endpoints will ultimately use. However, any changes to the Cisco Express Forwarding tables, brought about through routing changes or reloading of the switches/routers along the path, may result in a slightly different path established for the traffic when the actual video devices are installed. You should be aware of these limitations of IPSLA within a highly redundant medianet infrastructure.

IPSLA as an Ongoing Performance Monitoring Tool

If configured with careful consideration, IPSLAs can also be used as an ongoing performance monitoring tool. Rather than simulating an actual medianet video flow, IPSLA operations can be used to periodically send small amounts of traffic between two points within the network, per service class, to assess parameters such as packet loss, one-way latency, and jitter. [Figure 6-20](#) shows an example of such a deployment between two branches.

Figure 6-20 Example of IPSLA Used for Ongoing Performance Monitoring



For example, [Figure 6-20](#) shows both TelePresence and desktop video conferencing endpoints. Following the Cisco 12-class QoS model, TelePresence traffic can be marked CS4 and placed within a real-time interactive service class because it traverses both the private WAN links as well as an MPLS service between the branches. Likewise, desktop video conferencing traffic can be marked AF41 and placed within a Multimedia Conferencing service class because it traverses both the private WAN links and MPLS service between the branches. (Note that both traffic types may be remarked as it enters and exits the MPLS network). The configuration snippets in [Example 6-6](#) and [Example 6-7](#) show this type of IPSLA configuration with a pair of Cisco 3845 ISRs, one configured as the IPSLA sender and the other configured as the corresponding IPSLA responder.

Example 6-6 IPSLA Sender Configuration on a Cisco ISR 3845

```
ip sla 10
  udp-jitter 10.31.0.1 32800 source-ip 10.17.255.37 source-port 32800 num-packets 5
  interval 200
  request-data-size 958
  tos 128
  frequency 300
!
ip sla 11
  udp-jitter 10.31.0.1 32802 source-ip 10.17.255.37 source-port 32802 num-packets 5
  interval 200
  request-data-size 958
  tos 136
  frequency 300
!
ip sla reaction-configuration 10 react jitterDSAvg threshold-value 10 1 threshold-type
immediate action-type trapOnly
ip sla reaction-configuration 10 react rtt threshold-value 300 1 threshold-type immediate
action-type trapOnly
ip sla reaction-configuration 10 react jitterSDAvg threshold-value 10 1 threshold-type
immediate action-type trapOnly
ip sla reaction-configuration 10 react packetLossDS threshold-value 1 1 threshold-type
immediate action-type trapOnly
ip sla reaction-configuration 10 react packetLossSD threshold-value 1 1 threshold-type
immediate action-type trapOnly
ip sla reaction-configuration 10 react connectionLoss threshold-type immediate action-type
trapOnly
ip sla reaction-configuration 10 react timeout threshold-type immediate action-type
trapOnly
!
ip sla reaction-configuration 11 react rtt threshold-value 300 1 threshold-type immediate
action-type trapOnly
ip sla reaction-configuration 11 react jitterDSAvg threshold-value 10 1 threshold-type
immediate action-type trapOnly
ip sla reaction-configuration 11 react jitterSDAvg threshold-value 10 1 threshold-type
immediate action-type trapOnly
ip sla reaction-configuration 11 react packetLossDS threshold-value 1 1 threshold-type
immediate action-type trapOnly
ip sla reaction-configuration 11 react packetLossSD threshold-value 1 1 threshold-type
immediate action-type trapOnly
ip sla reaction-configuration 11 react connectionLoss threshold-type immediate action-type
trapOnly
ip sla reaction-configuration 11 react timeout threshold-type immediate action-type
trapOnly
!
ip sla group schedule 1 10-11 schedule-period 5 frequency 300 start-time now life forever
!
~
```

Example 6-7 IPSLA Responder Configuration on a Cisco ISR 3845

```
ip sla monitor responder
ip sla monitor responder type udpEcho ipaddress 10.31.0.1 port 32800
ip sla monitor responder type udpEcho ipaddress 10.31.0.1 port 32802
!
~
```

In the configuration example above, five 1000-byte packets (probes) with a CS4 DSCP marking, each spaced 200 milliseconds apart, are sent every 300 seconds. Likewise, five 1000-byte packets with an AF41 DSCP marking are sent every 300 seconds.

**Note**

The request-data-size parameter within the UDP jitter IPSLA operation specifies only the UDP payload size. The overall packet size on an Ethernet network can be obtained by adding the IP header (20 bytes), UDP header (8 bytes), and Layer 2 Ethernet header (14 bytes).

This is a relatively small amount of traffic that can be used to measure parameters such as jitter, one-way latency, and packet loss per service class on an ongoing basis. As with the Cisco Catalyst 6500 example above, statistics can be viewed via the **show ip sla monitor statistics aggregated details** command on the Cisco 3845 ISR configured as the IPSLA sender. However, in this example, the IPSLA sender has also been configured to send SNMP traps in response to the IPSLA traffic in the following situations:

- When destination-to-source jitter or source-to-destination jitter is outside the range of 1–10 milliseconds
- When the round trip latency is outside the range of 1–300 milliseconds
- When any packet loss occurs
- When the IPSLA operation times out or the IPSLA control session indicates a connection loss

Note that the jitter, packet loss, and round-trip-time latency parameters for the SNMP traps are configurable. The values used here are examples only. The settings chosen on a real implementation depend entirely on the service level targets for the particular traffic service class. For a discussion of each of the various traps and how to configure them on Cisco IOS router platforms, see *Cisco IOS IPSLAs Configuration Guide, Release 12.4* at the following URL:

http://www.cisco.com/en/US/docs/ios/12_4/ip_sla/configuration/guide/hs1a_c.html.

Rather than having to periodically log on to the IPSLA sender to view the statistics, you can simply monitor a central SNMP trap collector to determine whether the jitter, packet loss, and latency targets are being met. The usefulness of this approach depends to a large extent on how often the IPSLA traffic is sent and what the network is experiencing in terms of congestion. If a network is experiencing somewhat continuous congestion, resulting in high jitter (because of queueing) and some packet loss, an IPSLA operation that sends a few packets every few minutes is likely to experience some degradation, and therefore generate an SNMP trap to alert the network administrator. However, even under these circumstances, it may be several IPSLA cycles before one of the IPSLA packets is dropped or experiences high jitter. If the network experiences very transient congestion, resulting in brief moments of high jitter and packet loss, possibly periodic in nature (because of some traffic that sends periodic bursts of packets, such as high definition video frames), it may be many cycles before any of the IPSLA packets experience any packet loss or high jitter. Therefore, you must again balance the amount and frequency of traffic sent via the IPSLA operations against the additional overhead and potential degradation of network performance caused by the IPSLA operation itself. However, if implemented carefully, IPSLA operations can be used to proactively monitor service-level parameters such as jitter, packet loss, and latency per service class, on an ongoing basis. As discussed earlier, you may also choose to implement dedicated routers for IPSLA probes used for ongoing performance monitoring, rather than using the existing routers at branch and campus locations.

Router and Switch Command-Line Interface

The following sections present several Cisco router and switch commands that can be run from the CLI, to gain visibility into traffic flows across an enterprise medianet. As with the functionality discussed in previous sections, a complete listing of all possible CLI commands is outside the scope of this document. Instead, this discussion focuses on commands that assist in determining at a high level whether drops are occurring within router and switch interfaces along the path of a video flow; and more specifically, to determine whether drops are occurring within service classes that are mapped to separate queues within the interfaces on the respective platforms. It is assumed that a QoS model has been implemented in which the various video applications are mapped to different service classes within the medianet infrastructure. The mapping of video applications can be accomplished through classification and marking of the application within the Cisco Catalyst switch port at the ingress edge of the network infrastructure; or by trusting an application-specific device, which is then connected to the Cisco Catalyst switch port, to correctly mark its traffic. [Figure 6-21](#) shows an example of such a QoS model. The reader is encouraged to review [Chapter 4, “Medianet QoS Design Considerations”](#) before proceeding.

Figure 6-21 Cisco RFC-4594 Based 12-Class QoS Model


Application Class	PHB	Admission Control	Queueing and Dropping	Application Examples
VoIP Telephony	EF	Required	Priority Queue (PQ)	Cisco IP Phones
Broadcast Video	CS5	Required	Optional (PQ)	Cisco IP Surveillance, Cisco Enterprise TV
Realtime Interactive	CS4	Required	Optional (PQ)	Cisco TelePresence
Multimedia Conferencing	AF4	Required	BW Queue + DSCP WRED	Cisco Unified Personal Communicator
Multimedia Streaming	AF3	Recommended	BW Queue + DSCP WRED	Cisco Digital Media System (VoD)
Network Control	CS6		BW Queue	EIGRP, OSPF, BGP, HSRP, IKE
Call-Signaling	CS3		BW Queue	SCCP, SIP, H.323
OAM	CS2		BW Queue	SNMP, SSH, Syslog
Transactional Data	AF2		BW Queue + DSCP WRED	Cisco WebEx, Cisco MeetingPlace, ERP Apps
Bulk Data	AF1		BW Queue + DSCP WRED	Email, FTP, Backup Apps, Content Distribution
Best Effort	default		Default Queue + RED	Default Class Traffic
Scavenger	CS1		Min BW Queue (Deferential)	YouTube, iTunes, BitTorrent, Xbox Live

227922

As can be seen from [Figure 6-21](#), IP video surveillance traffic is assigned to the Broadcast Video service class with a CS5 marking; TelePresence traffic is assigned to the Real-Time Interactive service class with a CS4 marking; desktop videoconferencing is assigned to the Multimedia Conferencing service class with an AF4x marking; and VoD/enterprise TV is assigned to the Multimedia Streaming service class with an AF3x marking. After the traffic from the various video applications has been classified and marked, it can then be mapped to specific ingress and egress queues and drop thresholds on Cisco router and switch platforms. Each queue can then be allocated a specific percentage of the overall bandwidth of the interface as well as a percentage of the overall buffer space of the particular interface. This provides a level of protection where one particular video and/or data application mapped to a particular service class cannot use all the available bandwidth, resulting in the degradation of all other video and/or data applications mapped to other service classes. The more granular the mapping of the service classes to separate queues (in other words, the more queues implemented on a platform), the more granular the

control and therefore the protection of service classes. When multiple service classes are mapped to a single queue, separate drop thresholds can be implemented (on platforms that support them) to provide differentiation of service classes within the queue. The implementation of queueing and drop thresholds is viewed as necessary to provide the correct per-hop treatment of the video application traffic to meet the overall desired service levels of latency, jitter, and packet loss across the medianet infrastructure. An example of the mapping of service classes to egress queueing on a Cisco Catalyst 6500 WS-X6704-10GE line card, which has a 1P7Q8T egress queueing structure, as shown in Figure 6-22. Note that the percentage of bandwidth allocated per queue depends on the customer environment; Figure 6-22 shows only an example.

Figure 6-22 Example Mapping of Service Classes to Egress Queueing on a Cisco Catalyst 6500 Line Card with 1P7Q8T Structure



This QoS methodology can also provide enhanced visibility into the amount of traffic from individual video application types crossing points within the medianet infrastructure. The more granular the mapping of individual video applications to service classes that are then mapped to ingress and egress queues, the more granular the visibility into the amount of traffic generated by particular video applications. You can also gain additional visibility into troubleshooting video quality issues caused by drops within individual queues on router and switch platforms.

The following high-level methodology can be useful for troubleshooting video performance issues using the router and switch CLI. As with anything, this methodology is not perfect. Some of the shortcomings of the methodology are discussed in the sections that cover the individual CLI commands. However, it can often be used to quickly identify the point within the network infrastructure where video quality issues are occurring. The steps are as follows:

1. Determine the Layer 3 hop-by-hop path of the particular video application across the medianet infrastructure from end-to-end, starting from the Layer 3 device closest to one end of the video session. The traceroute CLI utility can be used for this function.

2. Determine at a high level whether any drops are being seen by interfaces on each of the Layer 3 devices along the path. The **show interface summary** command can be used to provide this function quickly. If drops are being seen on a Layer 3 device, further information can be gained by observing the specific interfaces in which drops are occurring. The **show interface <interface>** command can be used for this.
3. To determine whether drops are occurring within the specific queue to which the video application is mapped on the platform, various **show** commands that are specific to a particular platform can be used.


The following sections discuss the commands for each of the steps.

Traceroute

Traceroute is a command-line utility within Cisco router and switch products (and also in Unix and Linux systems) that can be used to produce a list of Layer 3 devices between two points within an IP network infrastructure. The Cisco traceroute utility sends a series of UDP packets with incrementing time-to-live (TTL) values from one IP address configured on the Layer 3 router or switch, to the desired destination IP address. Each Layer 3 device along the path either decrements the TTL value and forwards the UDP packet to the next hop in the path; or, if the TTL value is down to 1, the Layer 3 device discards the packet and sends an ICMP Time Exceeded (Type 11) message back to the source IP address. The ICMP Time Exceeded messages received by the source IP address (in other words, the originating router or switch device) are used to create a list of Layer 3 hops between the source and destination addresses. Note that ICMP Time Exceeded messages need to be allowed within the medianet infrastructure for traceroute to work. Also, if a Layer 3 device along the path does not send ICMP Time Exceeded messages, that device is not included in the list of Layer 3 hops between the source and destination addresses. Further, any Layer 2 devices along the path, which may themselves be the cause of the video quality degradation, are not identified in the traceroute output, because traceroute uses the underlying Layer 3 IP routing infrastructure to operate.

The traceroute utility works best when there are no equal-cost routes between network devices within the IP network infrastructure, and the infrastructure consists of all Layer 3 switches and routers, as shown in the sample network in [Figure 6-23](#).

Figure 6-23 Traceroute in an IP Network with a Single Route Between Endpoints



In this network, if the **traceroute** command is run from *me-eastcamp-1* using the VLAN 161 interface as the source interface, the output looks similar to that shown in [Example 6-8](#).

Example 6-8 Example Output from the Traceroute Utility Over a Non-Redundant Path Network

```
me-eastcamp-1#traceroute
Protocol [ip]:
Target IP address: 10.16.5.20
Source address: 10.16.1.1
Numeric display [n]: yes
Timeout in seconds [3]:
Probe count [3]: 4           ! Sets the number of packets generated with each TTL value.
Minimum Time to Live [1]:
Maximum Time to Live [30]: 6      ! Sets the max TTL value of the UDP packets generated.
Port Number [33434]:
Loose, Strict, Record, Timestamp, Verbose[none]: v
Loose, Strict, Record, Timestamp, Verbose[V]:
Type escape sequence to abort.
Tracing the route to 10.16.5.20

 1 10.16.2.2 0 msec 4 msec 0 msec 0 msec
 2 10.16.3.2 0 msec 0 msec 4 msec 0 msec
 3 10.16.4.2 0 msec 0 msec 4 msec 4 msec
 4 10.16.5.20 0 msec 0 msec 0 msec 8 msec
```

Traceroute returns the IP addresses of each Layer 3 hop in the route between *me-eastcamp-1* and *me-eastctms-1*. More specifically, because traceroute traces the hop route in one direction only, it returns the IP address of the interface of each router or switch that is closest to the source IP address. These IP addresses are shown in blue in [Figure 6-23](#). Note that because traceroute is initiated by the *me-eastcamp-1* switch, it does not appear within the traceroute output itself.

If the traceroute command is run from *me-eastcamp-2* using the VLAN 165 interface as the source interface, the output looks similar to that shown in [Example 6-9](#).

Example 6-9 Example Output from the Traceroute Utility From the Other Direction

```
me-eastcamp-1#traceroute
Protocol [ip]: ip
Target IP address: 10.16.1.11
Source address: 10.16.5.1
Numeric display [n]: yes
Timeout in seconds [3]:
Probe count [3]: 4
Minimum Time to Live [1]:
Maximum Time to Live [30]: 6
Port Number [33434]:
Loose, Strict, Record, Timestamp, Verbose[none]: V
Loose, Strict, Record, Timestamp, Verbose[V]:
Type escape sequence to abort.
Tracing the route to 10.16.1.11

 1 10.16.4.1 0 msec 0 msec 0 msec 4 msec
 2 10.16.3.1 0 msec 0 msec 0 msec 0 msec
 3 10.16.2.1 0 msec 0 msec 4 msec 0 msec
 4 * * * *
 5 * * * *
 6 * * * *
Destination not found inside max TTL diameter.    ! Indicates the end device did not return
                                                    ! an ICMP Time Exceeded Pkt.
```

The IP addresses returned from traceroute run in this direction are shown in red in [Figure 6-23](#). Note that because traceroute is initiated by the *me-eastcamp-2* switch, it does not appear within the traceroute output itself. Because a single path exists between the CTS-1000 and the Cisco TelePresence Multipoint Switch, the same Layer 3 hops (routers and switches) are returned regardless of which direction the

traceroute is run, although different IP addresses corresponding to different interfaces are returned, and the list of devices is reversed. Therefore, you need to run the traceroute in only one direction to understand the media flows in both directions. However, note that this may not necessarily be the case in a network with multiple equal-cost paths. Also note that the CTS-1000 does not return ICMP Time Exceeded packets, and therefore the traceroute utility times out. For a TelePresence endpoint, this can be rectified by directing the traceroute to the IP Phone associated with the TelePresence endpoint. However, be aware that some video endpoints may not respond to UDP traceroute packets with ICMP Time Exceeded packets.

Single-path non-redundant IP network infrastructures are somewhat counter to best practices for network designs with high availability in mind. Unfortunately, the use of traceroute within an equal-cost redundant IP network infrastructure can sometimes return unclear results regarding the path of an actual video flow between two endpoints. An example of why this occurs can be seen with the output of two **traceroute** commands run on a Cisco Catalyst 4500 Series switch to a TelePresence Cisco TelePresence Multipoint Switch (IP address 10.17.1.20), as shown in [Example 6-10](#).

Example 6-10 Example Output from the Traceroute Utility

```
me-westcamp-1#traceroute
Protocol [ip]:
Target IP address: 10.17.1.20
Source address: 10.24.1.1
Numeric display [n]: yes
Timeout in seconds [3]:
Probe count [3]: 4
Minimum Time to Live [1]:
Maximum Time to Live [30]: 10
Port Number [33434]:
Loose, Strict, Record, Timestamp, Verbose[none]: v
Loose, Strict, Record, Timestamp, Verbose[V]:
Type escape sequence to abort.
Tracing the route to 10.17.1.20


 1 10.17.100.37 8 msec 0 msec 4 msec 0 msec
 2 10.17.100.17 0 msec 0 msec 4 msec 0 msec
 3 10.17.100.94 0 msec 0 msec 0 msec 0 msec
 4 10.17.101.10 4 msec 0 msec 0 msec 0 msec
 5 10.17.101.13 0 msec 4 msec 0 msec 0 msec
 6 10.17.1.20 0 msec 4 msec 0 msec 0 msec

me-westcamp-1#traceroute
Protocol [ip]:
Target IP address: 10.17.1.20
Source address: 10.26.1.1
Numeric display [n]: yes
Timeout in seconds [3]:
Probe count [3]: 4
Minimum Time to Live [1]:
Maximum Time to Live [30]: 10
Port Number [33434]:
Loose, Strict, Record, Timestamp, Verbose[none]: v
Loose, Strict, Record, Timestamp, Verbose[V]:
Type escape sequence to abort.
Tracing the route to 10.17.1.20

 1 10.17.100.37 0 msec 0 msec 0 msec 0 msec
 2 10.17.100.29 4 msec 0 msec 0 msec 0 msec
 3 10.17.100.89 0 msec 4 msec 0 msec 0 msec
 4 10.17.101.10 0 msec 0 msec 0 msec 4 msec
 5 10.17.101.13 0 msec 0 msec 0 msec 4 msec
 6 10.17.1.20 0 msec 0 msec 0 msec 0 msec
```

The output from this traceroute was obtained from the network shown in [Figure 6-24](#).

Figure 6-24 Test Network Used for Traceroute Example



As can be seen from [Example 6-10](#) and [Figure 6-24](#), the first traceroute is run using the source interface VLAN 241 on switch *me-westcamp-1*, which has IP address 10.24.1.1. The output is Route #1: from *me-westcamp-1* to *me-westdist-3* to *me-westcore-1* to *me-westdc7k-2* (VDC #1) to *me-westdc7k-2* (VDC #2) and finally to *me-westctms-1*. The second traceroute is run using the source interface VLAN 261 on switch *me-westcamp-1*, which has IP address 10.26.1.11. The output is Route #2: from *me-westcamp-1* to *me-westdist-3* to *me-westcore-2* to *me-westdc7k-2* (VDC #1) to *me-westdc7k-2* (VDC #2) and finally to *me-westctms-1*. Note that the devices greyed out in [Figure 6-22](#) do not show up at all within the output of either **traceroute** command. These include any Layer 2 devices as well as some Layer 3 devices, and also the actual Cisco TelePresence System endpoints, because the traceroute is initiated from the switches. The two traceroutes follow different routes through the redundant network infrastructure. This is because Cisco Express Forwarding

switching, which itself is based on IP routing protocols, by default load balances sessions based on a hash of the source and destination IP address, when equal-cost paths exist. Therefore, different source and destination address pairs may yield different routes through an equal-cost redundant path network infrastructure. Cisco Express Forwarding switching can be configured for per-packet load balancing. However, this is not recommended because it can result in out-of-order packets for voice and video media. Therefore, you may not be able to tell from the traceroute utility alone whether the route returned through the network is the actual route taken by the video media, because the source IP address of the video endpoint is different than that used for the traceroute utility on the router or switch. Ideally, if traceroute can be run on the video endpoint itself, the actual route followed by the media through the network infrastructure can more easily be determined. However, most video endpoints such as Cisco TelePresence endpoints, Cisco IP video surveillance cameras, and Cisco digital media players (DMPs) do not currently support the traceroute utility.

On some switch platforms, such as Cisco Catalyst 6500 Series platforms, the **show ip cef exact-route <source ip address> <destination ip address>** command may be used to determine the actual route taken by the media flow of interest. An example of the output using the actual source IP address of a TelePresence CTS-1000, 10.24.1.11, and the destination IP address of the Cisco TelePresence Multipoint Switch, 10.17.1.20, is shown in [Example 6-11](#).

Example 6-11 Example Output from show ip cef exact-route Command

```
me-westdist-3>show ip cef exact-route 10.24.1.11 10.17.1.20
10.24.1.11      -> 10.17.1.20      : GigabitEthernet1/1 (next hop 10.17.100.29)
```

As can be seen, the actual route taken by the video and audio streams from the CTS-1000 follows Route #2 from *me-westdist-3* to *me-westcore-2* within this hop, and not Route #1 from *me-westdist-3* to *me-westcore-1*. The same command can be run on all the switches in the path that support the command to determine the actual route of the video flow in question.

When the initial switch, from which the traceroute utility is run, has equal-cost paths to the first hop along the path to the destination, the output becomes somewhat undeterministic. This is because traceroute packets generated by the switch are CPU-generated, and therefore process-switched packets. These do not follow the Cisco Express Forwarding tables within the switch that generated them. Instead, the switch round-robins the multiple UDP packets generated, each with a given TTL value, out to each next hop with equal cost to the destination. The result is that only some of the hops corresponding to equal-cost paths appear in the traceroute output. However, the list of the actual hops returned by the traceroute depends on the Cisco Express Forwarding tables of the downstream switches and routers. An example of this behavior is shown in [Example 6-12](#) and [Figure 6-25](#).

Example 6-12 Example Output from Traceroute on a Switch with Redundant Paths

```
me-westcamp-1#traceroute
Protocol [ip]:
Target IP address: 10.17.1.20
Source address: 10.24.1.1
Numeric display [n]: yes
Timeout in seconds [3]:
Probe count [3]:
Minimum Time to Live [1]:
Maximum Time to Live [30]:
Port Number [33434]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Type escape sequence to abort.
Tracing the route to 10.17.1.20


 1 10.17.100.37 0 msec
   10.17.100.42 0 msec
```

```

10.17.100.37 0 msec
2 10.17.100.21 4 msec
10.17.100.17 0 msec
10.17.100.21 0 msec
3 10.17.100.94 0 msec 0 msec 0 msec
4 10.17.101.10 0 msec 0 msec 0 msec
5 10.17.101.13 0 msec 4 msec 0 msec
6 10.17.1.20 0 msec 0 msec 4 msec

```

Figure 6-25 Test Network Used for Redundant Switch Traceroute Example



The traceroute shown in [Example 6-12](#) is again run from source interface VLAN 241 with source IP address 10.24.1.11 to the Cisco TelePresence Multipoint Switch with destination IP address 10.20.1.11. Because the switch from which the **traceroute** command is run has equal-cost paths to the first hop in the path, both switches *me-westdist-3* and *me-westdist-4* appear as the first hop in the path. Both paths then converge at the next switch hop, *me-westcore-1*, with *me-westcore-2* not showing up at all in the

traceroute output. However, note that a video traffic session (consisting of a source IP address and a destination IP address) that is Cisco Express Forwarding-switched through the router follows one or the other first hop through *me-westdist-3* or *me-westdist-4*, and not both hops, as indicated within the traceroute output. Again, the use of the **show ip cef exact-route** command on switches along the path may be necessary to determine the exact route of the video flows.

show interface summary and show interface Commands

After you have discovered the path of the actual video stream, possibly from using a combination of **traceroute** and the **show ip cef exact-route** command on switches along the path, a next logical step in troubleshooting a video quality issue is to see at a very high level whether interfaces are dropping packets. The **show interface summary** command can be used on Cisco Catalyst switch and IOS router platforms for this purpose (note that this command is not supported on Cisco Nexus switch platforms). [Example 6-13](#) shows an example output from this command on a Cisco Catalyst 6500 platform.

Example 6-13 Partial Output from the show interface summary Command on a Cisco Catalyst 6500 Switch

```
me-westcore-1#show interface summary
```

```

*: interface is up
IHQ: pkts in input hold queue      IQD: pkts dropped from input queue
OHQ: pkts in output hold queue    OQD: pkts dropped from output queue
RXBS: rx rate (bits/sec)          RXPS: rx rate (pkts/sec)
TXBS: tx rate (bits/sec)          TXPS: tx rate (pkts/sec)
TRTL: throttle count

```

Interface	IHQ	IQD	OHQ	OQD	RXBS	RXPS	TXBS	TXPS	TRTL
Vlan1	0	0	0	0	0	0	0	0	0
* GigabitEthernet1/1	0	0	0	0	1000	1	0	0	0
GigabitEthernet1/2	0	0	0	0	0	0	0	0	0
...									
* TenGigabitEthernet3/1	0	0	0	0	1000	1	2000	1	0
* TenGigabitEthernet3/2	0	0	0	0	1000	1	1000	1	0
TenGigabitEthernet3/3	0	0	0	0	0	0	0	0	0
TenGigabitEthernet3/4	0	0	0	0	0	0	0	0	0
* GigabitEthernet5/1	00	0	0	0	1000	1	2000	3	0
* GigabitEthernet5/2	00	0	0	0	2000	2	0	0	0
* Loopback0	0	0	0	0	0	0	0	0	0

The **show interface summary** command can be used to quickly identify the following:

- Which interfaces are up on the switch or router, as indicated by the asterisk next to the interface
- Whether any interfaces are experiencing any input queue drops (IQD) or output queue drops (OQD)
- The amount of traffic transmitted by the interface in terms of bits/second (TXBS) or packets/second (TXPS)
- The amount of traffic received by the interface in terms of bits/second (RXBS) or packets/second (RXPS)

The **show interface summary** command may need to be run multiple times over a short time interval to determine whether drops are currently occurring, rather than having occurred previously. Alternatively, the **clear counters** command can typically be used to clear all the counters on all the interfaces.

However, simply because an interface is determined to be experiencing drops does not necessarily mean that the interface is relevant to the path of the video flow in question. You may still need to run the **show ip cef exact-route** command, or consult the IP routing tables via the **show ip route** command to determine whether the particular interface experiencing drops is along the path of the video flow. [Example 6-14](#) shows an example output from both of these commands.

Example 6-14 Example Output from the show ip route and show ip cef exact-route Commands

```
me-westdist-3#show ip route 10.17.1.0
Routing entry for 10.17.1.0/24
  Known via "eigrp 111", distance 90, metric 6144, type internal
  Redistributing via eigrp 111
  Last update from 10.17.100.29 on GigabitEthernet1/1, 2w2d ago
  Routing Descriptor Blocks:
  * 10.17.100.17, from 10.17.100.17, 2w2d ago, via GigabitEthernet5/3
    Route metric is 6144, traffic share count is 1
    Total delay is 140 microseconds, minimum bandwidth is 1000000 Kbit
    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 4
  10.17.100.29, from 10.17.100.29, 2w2d ago, via GigabitEthernet1/1
    Route metric is 6144, traffic share count is 1
    Total delay is 140 microseconds, minimum bandwidth is 1000000 Kbit
    Reliability 255/255, minimum MTU 1500 bytes
    Loading 1/255, Hops 4

me-westdist-3#show ip cef exact-route 10.24.1.11 10.17.1.20
10.24.1.11      -> 10.17.1.20      : GigabitEthernet1/1 (next hop 10.17.100.29)
```

[Example 6-14](#) shows that the IP routing tables indicate that there are equal-cost paths to IP subnet 10.17.1.20 through next hops 10.17.100.17 and 10.17.100.29, via interfaces GigabitEthernet5/3 and GigabitEthernet1/1, respectively. The asterisk next to the 10.17.100.17 route indicates that the next session will follow that route. However, the output from the **show ip cef exact-route** command shows that the Cisco Express Forwarding table has already been populated with a session from source IP address 10.24.1.11, corresponding to the CTS-1000, to destination IP address 10.17.1.20, corresponding to the Cisco TelePresence Multipoint Switch, via interface GigabitEthernet1/1. Therefore, when troubleshooting drops along the path for this particular video flow, you should be concerned with drops shown on interface GigabitEthernet1/1.

Having determined which relevant interfaces are currently experiencing drops, you can drill down further into the interface via the **show interface <interface>** command. [Example 6-15](#) shows an example output from this command on a Cisco Catalyst 6500 platform.

Example 6-15 Example Output from the show interface Command

```
me-westdist-3#show interface gigabitethernet1/1
GigabitEthernet1/1 is up, line protocol is up (connected)
  Hardware is C6k 1000Mb 802.3, address is 0018.74e2.7dc0 (bia 0018.74e2.7dc0)
  Description: CONNECTION TO ME-WESTCORE-2 GIG1/25
  Internet address is 10.17.100.30/30
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s
  input flow-control is off, output flow-control is off
  Clock mode is auto
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:04, output 00:00:00, output hang never
  Last clearing of "show interface" counters 00:13:53
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 0! Input & output
```

```

! queue drops.

Queueing strategy: fifo
Output queue: 0/40 (size/max)
30 second input rate 0 bits/sec, 0 packets/sec
30 second output rate 0 bits/sec, 0 packets/sec
L2 Switched: ucast: 117 pkt, 22493 bytes - mcast: 184 pkt, 14316 bytes
L3 in Switched: ucast: 14 pkt, 7159 bytes - mcast: 0 pkt, 0 bytes mcast
L3 out Switched: ucast: 0 pkt, 0 bytes mcast: 0 pkt, 0 bytes
  374 packets input, 53264 bytes, 0 no buffer
  Received 250 broadcasts (183 IP multicasts)
  0 runts, 0 giants, 0 throttles
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored ! May indicate link-level errors
  0 watchdog, 0 multicast, 0 pause input
  0 input packets with dribble condition detected
  282 packets output, 32205 bytes, 0 underruns
  0 output errors, 0 collisions, 0 interface resets ! May indicate link-level errors.
  0 babbles, 0 late collision, 0 deferred
  0 lost carrier, 0 no carrier, 0 PAUSE output
  0 output buffer failures, 0 output buffers swapped out

```

The **show interface** command can provide an instantaneous display of the current depth of the input and output queues, as well as a running total of input and output drops seen by the interface. This can be used to detect possible congestion issues occurring within the switch interface. It also provides additional detail in terms of the type of traffic: unicast versus multicast switched by the interface. More importantly, the **show interface** command provides additional detail regarding potential link level errors, such as CRCs, collisions, and so on. These can be the result of cabling issues or even duplex mismatches between switch interfaces that are difficult to detect, but can be the cause of degraded video quality as well. Note that changing the load interval from the default of 5 minutes to a lower value, such as 60 seconds, can provide increased visibility, so that the statistics are then more up-to-date.

Platform Specific Queue-Level Commands

Because a relevant interface along the path of the video flow in question is experiencing drops does not necessarily mean that the drops are occurring within the queue that holds the particular video application traffic. You may need to run additional platform-specific commands to display drops down to the queue level to determine whether video degradation is occurring on a particular switch or router. The following sections discuss some of these platform-specific commands.

Cisco Catalyst 6500 Series Commands

When QoS is enabled on Cisco Catalyst 6500 Series switches, the **show queueing interface** command allows you to view interface drops per queue on the switch port. [Example 6-16](#) shows the output from a Cisco Catalyst 6500 WS-X6708-10GE line card. Selected areas for discussion have been highlighted in bold.

Example 6-16 Output from Cisco Catalyst 6500 show queueing interface Command

```

me-eastcore-1#show queueing interface tenGigabitEthernet 1/1

Interface TenGigabitEthernet1/1 queueing strategy:  Weighted Round-Robin
  Port QoS is enabled
  Trust boundary disabled

Trust state: trust DSCP
Extend trust state: not trusted [COS = 0]
Default COS is 0
  Queueing Mode In Tx direction: mode-dscp      Transmit queues [type = 1p7q4t]:
  Queue Id      Scheduling  Num of thresholds

```

```

-----
      01          WRR          04
      02          WRR          04
      03          WRR          04
      04          WRR          04
      05          WRR          04
      06          WRR          04
      07          WRR          04
      08          Priority      01

WRR bandwidth ratios:    1[queue 1] 25[queue 2]  4[queue 3] 10[queue 4] 10[queue
5] 10[queue 6] 10[queue 7]
queue-limit ratios:      1[queue 1] 25[queue 2]  4[queue 3] 10[queue 4]
10[queue 5] 10[queue 6] 10[queue 7] 30[Pri Queue]

queue tail-drop-thresholds
-----
1      70[1] 100[2] 100[3] 100[4]
2      70[1] 100[2] 100[3] 100[4]
3      100[1] 100[2] 100[3] 100[4]
4      100[1] 100[2] 100[3] 100[4]
5      100[1] 100[2] 100[3] 100[4]
6      100[1] 100[2] 100[3] 100[4]
7      100[1] 100[2] 100[3] 100[4]

queue random-detect-min-thresholds
-----
1      80[1] 100[2] 100[3] 100[4]
2      80[1] 100[2] 100[3] 100[4]
3      70[1] 80[2] 90[3] 100[4]
4      70[1] 80[2] 90[3] 100[4]
5      70[1] 80[2] 90[3] 100[4]
6      70[1] 80[2] 90[3] 100[4]
7      60[1] 70[2] 80[3] 90[4]

queue random-detect-max-thresholds
-----
1      100[1] 100[2] 100[3] 100[4]
2      100[1] 100[2] 100[3] 100[4]
3      80[1] 90[2] 100[3] 100[4]
4      80[1] 90[2] 100[3] 100[4]
5      80[1] 90[2] 100[3] 100[4]
6      80[1] 90[2] 100[3] 100[4]
7      70[1] 80[2] 90[3] 100[4]

WRED disabled queues:

queue thresh cos-map
-----
1      1      0
1      2      1
1      3
1      4
2      1      2
2      2      3 4
2      3
2      4
3      1      6 7
3      2
3      3
3      4
4      1
4      2
4      3

```

```

4      4
5      1
5      2
5      3
5      4
6      1
6      2
6      3
6      4
7      1
7      2
7      3
7      4
8      1      5

queue thresh dscp-map
-----
1      1      1 2 3 4 5 6 7 8 9 11 13 15 17 19 21 23 25 27 29 31 33 39 41 42 43 44 45
47
1      2
1      3
1      4
2      1      0
2      2
2      3
2      4
3      1      14
3      2      12
3      3      10
3      4
4      1      22
4      2      20
4      3      18
4      4
5      1      30 35 37
5      2      28
5      3      26
5      4
6      1      38 49 50 51 52 53 54 55 57 58 59 60 61 62 63
6      2      36
6      3      34
6      4
7      1      16
7      2      24
7      3      48
7      4      56
8      1      32 40 46

Queueing Mode In Rx direction: mode-dscp
Receive queues [type = 8q4t]:
Queue Id      Scheduling  Num of thresholds
-----
01            WRR            04
02            WRR            04
03            WRR            04
04            WRR            04
05            WRR            04
06            WRR            04
07            WRR            04
08            WRR            04

WRR bandwidth ratios:  10[queue 1]   0[queue 2]   0[queue 3]   0[queue 4]   0[queue
5]   0[queue 6]   0[queue 7]  90[queue 8]

```

```

queue-limit ratios:      80[queue 1]   0[queue 2]   0[queue 3]   0[queue 4]   0[queue
5]  0[queue 6]   0[queue 7]  20[queue 8]

```

```

queue tail-drop-thresholds
-----

```

```

1      70[1] 80[2] 90[3] 100[4]
2      100[1] 100[2] 100[3] 100[4]
3      100[1] 100[2] 100[3] 100[4]
4      100[1] 100[2] 100[3] 100[4]
5      100[1] 100[2] 100[3] 100[4]
6      100[1] 100[2] 100[3] 100[4]
7      100[1] 100[2] 100[3] 100[4]
8      100[1] 100[2] 100[3] 100[4]

```

```

queue random-detect-min-thresholds
-----

```

```

1      40[1] 40[2] 50[3] 50[4]
2      100[1] 100[2] 100[3] 100[4]
3      100[1] 100[2] 100[3] 100[4]
4      100[1] 100[2] 100[3] 100[4]
5      100[1] 100[2] 100[3] 100[4]
6      100[1] 100[2] 100[3] 100[4]
7      100[1] 100[2] 100[3] 100[4]
8      100[1] 100[2] 100[3] 100[4]

```

```

queue random-detect-max-thresholds
-----

```

```

1      70[1] 80[2] 90[3] 100[4]
2      100[1] 100[2] 100[3] 100[4]
3      100[1] 100[2] 100[3] 100[4]
4      100[1] 100[2] 100[3] 100[4]
5      100[1] 100[2] 100[3] 100[4]
6      100[1] 100[2] 100[3] 100[4]
7      100[1] 100[2] 100[3] 100[4]
8      100[1] 100[2] 100[3] 100[4]

```

```

WRED disabled queues:      2 3 4 5 6 7 8

```

```

queue thresh cos-map
-----

```

```

1      1      0 1
1      2      2 3
1      3      4
1      4      6 7
2      1
2      2
2      3
2      4
3      1
3      2
3      3
3      4
4      1
4      2
4      3
4      4
5      1
5      2
5      3
5      4
6      1
6      2
6      3
6      4

```

```

7      1
7      2
7      3
7      4
8      1      5
8      2
8      3
8      4

queue thresh dscp-map
-----
1      1      0 1 2 3 4 5 6 7 8 9 11 13 15 16 17 19 21 23 25 27 29 31 33 39 41 42 43 44
45 47
1      2
1      3
1      4
2      1      14
2      2      12
2      3      10
2      4
3      1      22
3      2      20
3      3      18
3      4
4      1      24 30
4      2      28
4      3      26
4      4
5      1      32 34 35 36 37 38
5      2
5      3
5      4
6      1      48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63
6      2
6      3
6      4
7      1
7      2
7      3
7      4
8      1      40 46
8      2
8      3
8      4

```

Packets dropped on Transmit:
 BPDU packets: 0

```

queue          dropped  [dscp-map]
-----
1              0  [1 2 3 4 5 6 7 8 9 11 13 15 17 19 21 23 25 27 29 31 33 39
41 42 43 44 45 47 ]
2              0  [0 ]
3              0  [14 12 10 ]
4              0  [22 20 18 ]
5              0  [30 35 37 28 26 ]
6              0  [38 49 50 51 52 53 54 55 57 58 59 60 61 62 63 36 34 ]
7              0  [16 24 48 56 ]
8              0  [32 40 46 ]

```

Packets dropped on Receive:
 BPDU packets: 0

```

queue                dropped  [dscp-map]
-----
1                    0      [0 1 2 3 4 5 6 7 8 9 11 13 15 16 17 19 21 23 25 27 29 31
33 39 41 42 43 44 45 47 ]
2                    0      [14 12 10 ]
3                    0      [22 20 18 ]
4                    0      [24 30 28 26 ]
5                    0      [32 34 35 36 37 38 ]
6                    0      [48 49 50 51 52 53 54 55 56 57 58 59 60 61 62 63 ]
8                    0      [40 46 ]

```

The information within the first highlighted section can be used to quickly verify that the queueing and bandwidth ratios have been set correctly for the traffic service class of interest that is crossing the particular interface. As can be seen, the line card has a 1p7q4t egress queueing structure, meaning one priority queue and seven additional queues, each with four different drop thresholds. Egress queueing is configured to use a weighted-round robin (WRR) algorithm. The WRR bandwidth ratios are used by the scheduler to service the queues, which effectively allocates bandwidth across the seven non-priority queues based on the weight ratios. Note that the priority queue is always serviced first, and therefore has no weight. The queue-limit ratios allocate available egress queue space based on the ratios as well. Note that egress queueing space for the priority queue is included.

The second highlighted section can be used to quickly verify that a particular traffic service class is mapped to the correct egress queue on the line card. It provides a quick view of the mapping of DSCP values to egress queues and drop thresholds. Further, this can then be used to identify which video applications are mapped to which queues, based on DSCP values. This assumes specific video applications have been mapped to service classes with separate DSCP values. Note that in older Cisco Catalyst 6500 line cards, egress queues may be mapped to internal switch class of service (CoS) values that are then mapped to DSCP values. In such cases, you may need to use the **show mls qos maps dscp-cos** command to display the mapping of DSCP values to internal CoS values within the Cisco Catalyst switch.

Finally, the third highlighted block shows the number of packets dropped by the interface, per transmit queue. This can be used for either performance management, in the case where a particular video application mapped to the queue is experiencing degraded service because of packet loss; or for fault isolation, in the case where a particular video application is dropping the connection because of packet loss.

The same information is also provided for ingress queueing with this particular line card. Note, however, that the various Cisco Catalyst 6500 line cards support different ingress and egress queueing structures, as well as modes of operations. Older Cisco Catalyst 6500 line cards support ingress queueing based on Layer 2 CoS marking only. Ingress queueing may not be used within a routed (non-trunked) infrastructure on Cisco Catalyst 6500 line cards.

Cisco Catalyst 4500/4900 Series Commands

Visibility into traffic flows down at the queue level within a Cisco Catalyst 4500 Series switch depends on the supervisor line card within the switch. For Cisco Catalyst 4500 Series switches with a Supervisor-II-Plus, Supervisor-IV, or Supervisor-V (also referred to as classic supervisors), and for Cisco Catalyst 4900 Series switches, the **show interface counters** command provides a similar ability to view interface drops per queue on the switch port. [Example 6-17](#) shows a partial output from a Cisco Catalyst 4948 switch. For brevity, output from only the first two interfaces and the last interface on the switch are shown. Selected areas for discussion have been highlighted in bold.

Example 6-17 Output from Cisco Catalyst 4948 show interface counters detail Command

```
tp-c2-4948-1#show interface counters detail
```

```

Port InBytesIn UcastPkts InMcastPkts InBcastPkts ! Provides info on ingress multicast packets
Gi1/1 0 0 0 0
Gi1/2 0 0 0 0
...
Gi1/48 500745084 946163 4778144 892284

Port OutBytes OutUcastPkts OutMcastPkts OutBcastPkts ! Provides info on egress multicast
packets
Gi1/1 0 0 0 0
Gi1/2 0 0 0 0
...
Gi1/48 18267775 20009 190696 2

Port InPkts 64 OutPkts 64InPkts 65-127 OutPkts 65-127
Gi1/1 0 0 0 0
Gi1/2 0 0 0 0
...
Gi1/48 5676114 107817 705522 97227

Port InPkts 128-255 OutPkts 128-255 InPkts 256-511 OutPkts 256-511
Gi1/1 0 0 0 0
Gi1/2 0 0 0 0
...
Gi1/48 58703 1700 169614 2283

Port InPkts 512-1023 OutPkts 512-1023
Gi1/1 0 0
Gi1/2 0 0
...
Gi1/48 5859 1461

Port InPkts 1024-1518 OutPkts 1024-1518InPkts 1519-1548 OutPkts 1519-1548
Gi1/1 0 0 0 0
Gi1/2 0 0 0 0
...
Gi1/48 779 219 0 0

Port InPkts 1549-9216OutPkts 1549-9216
Gi1/1 0 0
Gi1/2 0 0
...
Gi1/48 0 0

Port Tx-Bytes-Queue-1Tx-Bytes-Queue-2Tx-Bytes-Queue-3Tx-Bytes-Queue-4 ! Provides
! transmitted byte count per queue
Gi1/1 0 0 0 0
Gi1/2 0 0 0 0
...
Gi1/48 67644 1749266 181312 16271855

Port Tx-Drops-Queue-1Tx-Drops-Queue-2Tx-Drops-Queue-3 Tx-Drops-Queue-4 ! Provides
! packet drop count per queue
Gi1/1 0 0 0 0
Gi1/2 0 0 0 0
...
Gi1/48 0 0 0 0

Port Db1-Drops-Queue-1Db1-Drops-Queue-2Db1-Drops-Queue-3Db1-Drops-Queue-4 ! Provides DBL
! packet drop count per queue
Gi1/1 0 0 0 0
Gi1/2 0 0 0 0
...
Gi1/48 0 0 0 0

```

```

Port      Rx-No-Pkt-Buff  RxPauseFrames  TxPauseFrames  PauseFramesDrop
Gi1/1     0                    0              0              0
Gi1/2     0                    0              0              0
...
Gi1/48    0                    0              0              0

Port      UnsupOpcodePause
Gi1/1     0
Gi1/2     0
...
Gi1/48    0

```

The first two highlighted sections can provide information regarding how many unicast and multicast packets have crossed the interface in the inbound or outbound direction. Multicast traffic is often used to support real-time and VoD broadcasts. The multicast packet count within the switch interface increments from when the switch was reloaded or the counters were manually cleared. Because of this, and because the information does not include the byte count, you cannot use the statistics alone to determine the data rate of multicast traffic across the interface. However, you may be able to gain some useful information regarding the percentage of multicast traffic on the interface based on the ratio of the unicast to multicast packets seen.

The third highlighted section provides two additional pieces of information. First, it indicates the number of queues per interface. Because the output above is from a Cisco Catalyst 4948 switch, four transmit queues per interface are supported. Second, the output indicates the amount of traffic, in bytes, that has been transmitted per queue per interface. Because this is a summation of bytes since the counters were last cleared or the switch reloaded, you must run the command multiple times over a time interval to get a rough estimate of the byte rate over that time period. This can be used to gain an idea of the current data rate of a particular traffic service class across the switch interface.

The final two highlighted sections indicate the number of packets dropped in the egress direction, per transmit queue. You can use this information to assist in troubleshooting a video application performance issue or fault condition caused by packet loss. Note that *Dbl-Drops* are drops that are the result of the dynamic buffer limiting (DBL) algorithm, which attempts to fairly allocate buffer usage per flow through the Cisco Catalyst 4500 switch. You have the option of enabling or disabling DBL per service class on the switch.

To make use of information regarding transmit queues drops shown in [Example 6-17](#), you must understand which traffic classes are assigned to which transmit queues. For Cisco Catalyst 4500 Series switches with classic supervisors as well as Cisco Catalyst 4900 Series switches, the **show qos maps** command can be used to display which DSCP values are mapped to which transmit queues on the switch, as shown in [Example 6-18](#).

Example 6-18 Output from Cisco Catalyst 4948 show qos maps Command

```

tp-c2-4948-1#show qos maps
DSCP-TxQueue Mapping Table (dscp = d1d2) ! Provides mapping of DSCP value to transmit
! queue on the switch
d1 : d2  0  1  2  3  4  5  6  7  8  9
-----
0 :    02 01 01 01 01 01 01 01 01 01
1 :    01 01 01 01 01 01 04 02 04 02
2 :    04 02 04 02 04 02 04 02 04 02
3 :    04 02 03 03 04 03 04 03 04 03
4 :    03 03 03 03 03 03 03 03 04 04
5 :    04 04 04 04 04 04 04 04 04 04
6 :    04 04 04 04

Policed DSCP Mapping Table (dscp = d1d2)
d1 : d2  0  1  2  3  4  5  6  7  8  9

```

```

-----
0 :    00 01 02 03 04 05 06 07 08 09
1 :    10 11 12 13 14 15 16 17 18 19
2 :    20 21 22 23 24 25 26 27 28 29
3 :    30 31 32 33 34 35 36 37 38 39
4 :    40 41 42 43 44 45 46 47 48 49
5 :    50 51 52 53 54 55 56 57 58 59
6 :    60 61 62 63

```

```

DSCP-CoS Mapping Table (dscp = d1d2)
d1 : d2  0  1  2  3  4  5  6  7  8  9
-----
0 :    00 00 00 00 00 00 00 00 01 01
1 :    01 01 01 01 01 01 02 02 02 02
2 :    02 02 02 02 03 03 03 03 03 03
3 :    03 03 04 04 04 04 04 04 04 04
4 :    05 05 05 05 05 05 05 05 06 06
5 :    06 06 06 06 06 06 07 07 07 07
6 :    07 07 07 07

```

```

CoS-DSCP Mapping Table
CoS:   0  1  2  3  4  5  6  7
-----
DSCP:   0  8 16 24 32 46 48 56

```

The highlighted section in the example above shows the mapping of the DSCP values to transmit queues. The vertical column, marked *d1*, represents the first decimal number of the DSCP value, while the horizontal column, marked *d2*, represents the second decimal number of the DSCP value. For example, a *d1* value of 3 and a *d2* value of 2 yields a DSCP decimal value of 32, which corresponds to the CS4 service class. You still need to separately understand the mapping of specific video applications to service classes that are then marked with a particular DSCP value. However, combined with the knowledge of which traffic classes are mapped to which transmit queue, you can use this information to troubleshoot video application performance issues across the Cisco Catalyst 4500/Cisco Catalyst 4900 switch platform.

**Note**

DSCP markings are represented by 6-bit values within the ToS byte of the IP packet. The DSCP values are the upper 6 bits of the ToS byte. Therefore, a DSCP decimal value of 32 represents a binary value of 100000, or the CS4 service class. The full ToS byte would have a value of 10000000 or a hexadecimal value of 0x80.

For the Cisco Catalyst 4500 with a Sup-6E supervisor line card, the mapping of traffic classes to egress queues is accomplished via an egress policy map applied to the interface. The policy map can be viewed through the **show policy-map interface** command. [Example 6-19](#) shows the output from a GigabitEthernet interface. Selected areas for discussion have been highlighted in bold.

Example 6-19 Output from Cisco Catalyst 4500 Sup-6E show policy-map interface Command

```
me-westcamp-1#show policy-map int gig 3/3
GigabitEthernet3/3
```

```
Service-policy output: 1P7Q1T ! Name and direction of the policy map applied to the
                               ! interface.
```

```
Class-map: PRIORITY-QUEUE (match-any)! Packet counters increment across all
22709 packets                               ! interfaces to which the policy map is applied.
```

```
Match: dscp ef (46)
0 packets
```

```

Match: dscp cs5 (40)
      0 packets
Match: dscp cs4 (32)
      22709 packets

police:                ! Byte counters under 'police' line increment per interface.
      cir 300000000 bps, bc 12375000 bytes, be 12375000 bytes
      conformed Packet count - n/a, 10957239 bytes; actions:
        transmit
      exceeded Packet count - n/a, 0 bytes; actions:
        drop
      violated Packet count - n/a, 0 bytes; actions:
        drop
      conformed 2131000 bps, exceed 0 bps, violate 0 bps
priority queue:        ! Byte counters and packet drops under 'priority queue' line
      Transmit: 9877576 Bytes, Queue Full Drops: 0 Packets ! increment per interface.

Class-map: CONTROL-MGMT-QUEUE (match-any)
  17 packets
  Match: dscp cs7 (56)
    0 packets
  Match: dscp cs6 (48)
    8 packets
  Match: dscp cs3 (24)
    9 packets
  Match: dscp cs2 (16)
    0 packets
  bandwidth: 10 (%) ! Byte counters and packet drops under 'bandwidth' line
    Transmit: 1616 Bytes, Queue Full Drops: 0 Packets ! increment per interface.

Class-map: MULTIMEDIA-CONFERENCING-QUEUE (match-all)
  0 packets
  Match: dscp af41 (34) af42 (36) af43 (38)
  bandwidth: 10 (%)
    Transmit: 0 Bytes, Queue Full Drops: 0 Packets

Class-map: MULTIMEDIA-STREAMING-QUEUE (match-all)
  0 packets
  Match: dscp af31 (26) af32 (28) af33 (30)
  bandwidth: 10 (%)
    Transmit: 0 Bytes, Queue Full Drops: 0 Packets

Class-map: TRANSACTIONAL-DATA-QUEUE (match-all)
  0 packets
  Match: dscp af21 (18) af22 (20) af23 (22)
  bandwidth: 10 (%)
    Transmit: 0 Bytes, Queue Full Drops: 0 Packets
  dbl
    Probabilistic Drops: 0 Packets
    Belligerent Flow Drops: 0 Packets

Class-map: BULK-DATA-QUEUE (match-all)
  0 packets
  Match: dscp af11 (10) af12 (12) af13 (14)
  bandwidth: 4 (%)
    Transmit: 0 Bytes, Queue Full Drops: 0 Packets
  dbl
    Probabilistic Drops: 0 Packets
    Belligerent Flow Drops: 0 Packets

Class-map: SCAVENGER-QUEUE (match-all)
  0 packets
  Match: dscp cs1 (8)
  bandwidth: 1 (%)

```

```

Transmit: 0 Bytes, Queue Full Drops: 0 Packets

Class-map: class-default (match-any)
  6 packets
  Match: any
    6 packets
  bandwidth: 25 (%)
    Transmit: 436 Bytes, Queue Full Drops: 0 Packets
  db1
    Probabilistic Drops: 0 Packets
    Belligerent Flow Drops: 0 Packets

```

In [Example 6-19](#), the first highlighted line shows the name of the service policy and direction (outbound or inbound) applied to the interface. The second highlighted section shows the mapping of DSCP markings to each queue defined within the policy map. Directly under that, the number of packets that matched the service class are displayed. Take special note that if a policy map is shared among multiple interfaces, these packet counters increment for all interfaces that have traffic that matches the particular class-map entry. For example, if the policy map named *IP7Q1T* shown in the example above were applied across two uplink interfaces, the packet counters would show the total packets that matched each class-map entry for both interfaces. This can lead to some confusion, as shown in [Example 6-20](#). Selected areas for discussion have been highlighted in bold.

Example 6-20 Second Example Output from Cisco Catalyst 4500 Sup-6E show policy-map interface Command

```

me-westcamp-1#show policy-map int gig 3/1
GigabitEthernet3/1

Service-policy output: 1P7Q1T

Class-map: PRIORITY-QUEUE (match-any)
  15360 packets
  Match: dscp ef (46)
    0 packets
  Match: dscp cs5 (40)
    0 packets
  Match: dscp cs4 (32)
    15360 packets
  police:
    cir 300000000 bps, bc 12375000 bytes, be 12375000 bytes
    conformed 0 packets, 0 bytes; actions:
      transmit
    exceeded 0 packets, 0 bytes; actions:
      drop
    violated 0 packets, 0 bytes; actions:
      drop
    conformed 0 bps, exceed 0 bps, violate 0 bps
  priority queue:
    Transmit: 0 Bytes, Queue Full Drops: 0 Packets

```

Notice in [Example 6-20](#) that interface GigabitEthernet3/1 appears to have seen 15,360 packets that match the PRIORITY-QUEUE class-map entry. Yet, both the policer and the priority queue statistics indicate that no packets that match the PRIORITY-QUEUE class-map entry have been sent by this interface. In this scenario, the 15,360 packets were sent by the other interface, GigabitEthernet3/3, which shared the policy map named *IP7Q1T*. To prevent this type of confusion when viewing statistics from the **show policy-map interface** command on the Cisco Catalyst 4500 with Sup6E, you can simply define a different policy map name for each interface. [Example 6-21](#) shows an example of this type of configuration.

Example 6-21 Partial Configuration Example Showing Separate Policy Map Per Interface

```
class-map match-all MULTIMEDIA-STREAMING-QUEUE
  match dscp af31 af32 af33
class-map match-any CONTROL-MGMT-QUEUE
  match dscp cs7
  match dscp cs6
  match dscp cs3
  match dscp cs2
class-map match-all TRANSACTIONAL-DATA-QUEUE
  match dscp af21 af22 af23
class-map match-all SCAVENGER-QUEUE
  match dscp cs1
class-map match-all MULTIMEDIA-CONFERENCING-QUEUE
  match dscp af41 af42 af43
class-map match-all BULK-DATA-QUEUE
  match dscp af11 af12 af13
class-map match-any PRIORITY-QUEUE
  match dscp ef
  match dscp cs5
  match dscp cs4
!
!
policy-map 1P7Q1T-GIG3/3
  class PRIORITY-QUEUE
    police cir percent 30 bc 33 ms
    conform-action transmit
    exceed-action drop
    violate-action drop
    priority
  class CONTROL-MGMT-QUEUE
    bandwidth percent 10
  class MULTIMEDIA-CONFERENCING-QUEUE
    bandwidth percent 10
  class MULTIMEDIA-STREAMING-QUEUE
    bandwidth percent 10
  class TRANSACTIONAL-DATA-QUEUE
    bandwidth percent 10
    dbl
  class BULK-DATA-QUEUE
    bandwidth percent 4
    dbl
  class SCAVENGER-QUEUE
    bandwidth percent 1
  class class-default
    bandwidth percent 25
    dbl
policy-map 1P7Q1T-GIG3/1
  class PRIORITY-QUEUE
    police cir percent 30 bc 33 ms
    conform-action transmit
    exceed-action drop
    violate-action drop
    priority
  class CONTROL-MGMT-QUEUE
    bandwidth percent 10
  class MULTIMEDIA-CONFERENCING-QUEUE
    bandwidth percent 10
  class MULTIMEDIA-STREAMING-QUEUE
    bandwidth percent 10
  class TRANSACTIONAL-DATA-QUEUE
    bandwidth percent 10
    dbl
  class BULK-DATA-QUEUE
```

```

        bandwidth percent 4
        dbl
    class SCAVENGER-QUEUE
        bandwidth percent 1
    class class-default
        bandwidth percent 25
        dbl
!
~
!
interface GigabitEthernet3/1
    description CONNECTION TO ME-WESTDIST-3 GIG1/13
    no switchport
    ip address 10.17.100.38 255.255.255.252
    ip pim sparse-mode
    load-interval 30
    service-policy output 1P7Q1T-GIG3/1
!
~
!
interface GigabitEthernet3/3
    description CONNECTION TO ME-WESTDIST-4 GIG1/2
    no switchport
    ip address 10.17.100.41 255.255.255.252
    ip pim sparse-mode
    load-interval 30
    service-policy output 1P7Q1T-GIG3/3
!
~

```

Notice that the class-map definitions shown at the top of the configuration example are shared between the policy maps. However, a unique policy map name is applied to each of the GigabitEthernet uplink interfaces.

Referring back to [Example 6-20](#), when a policer is applied to a queue, the bit rates of the data that conform, exceed, and violate the policer committed information rate (CIR) are also displayed within the **show policy-map interface** command. This information can provide a view of how much traffic is currently being handled by a policed queue, and whether sufficient bandwidth has been provisioned on the policer for the service classes handled by the queue. The final two highlighted sections in [Example 6-20](#) provide an aggregate byte count of the packets handled by the particular queue, as well as the number of packets dropped because of insufficient buffer space on the queue. This holds for either the priority queue defined via the priority command, or a class-based weighted fair queueing (CBWFQ) defined via the bandwidth command. You can get an estimate of the overall data rate through a particular queue by running the **show policy-map interface** command several times over fixed time intervals and dividing the difference in byte count by the time interval.

Cisco Catalyst 3750G/3750E Series Commands

When QoS is enabled on the Cisco Catalyst 3750G/3750E Series switches with the **mls qos global** command, egress queueing consists of four queues; one of which can be a priority queue, each with three thresholds (1P3Q3T). The third threshold on each queue is pre-defined for the queue-full state (100 percent). Queue settings such as buffer allocation ratios and drop threshold minimum and maximum settings are defined based on queue-sets applied across a range of interfaces; not defined per interface. The Cisco Catalyst 3750G/3750E Series switches support two queue sets. Ports are mapped to one of the two queue-sets. By default, ports are mapped to queue-set 1. The **show platform port-asic stats drop** command allows you to view interface drops per queue on the switch port. [Example 6-22](#) shows the output from a NME-XD-24ES-1S-P switch module within a Cisco 3845 ISR, which runs the same code base as the Cisco Catalyst 3750G.

Example 6-22 Output from Cisco Catalyst 3750G/3750E show platform port-asic stats drop Command

```
me-eastny-3#show platform port-asic stats drop fast 1/0/1
```

```
Interface Fa1/0/1 TxQueue Drop Statistics
```

```
Queue 0
  Weight 0 Frames 0
  Weight 1 Frames 0
  Weight 2 Frames 0
Queue 1
  Weight 0 Frames 0
  Weight 1 Frames 0
  Weight 2 Frames 0
Queue 2
  Weight 0 Frames 0
  Weight 1 Frames 0
  Weight 2 Frames 0
Queue 3
  Weight 0 Frames 0
  Weight 1 Frames 0
  Weight 2 Frames 0
```

To make use of information regarding transmit queues drops shown in [Example 6-22](#), you must understand which traffic classes are assigned to which transmit queues and which drop thresholds within those queues. For Cisco Catalyst 3750G or 3750E Series switches, the **show mls qos maps dscp-output-q** command can be used to display which DSCP values are mapped to which transmit queues and drop thresholds on the switch, as shown in [Example 6-23](#).

Example 6-23 Output from Cisco Catalyst 3750G or 3750E Series show mls qos maps dscp-output-q Command

```
me-eastny-3#show mls qos maps dscp-output-q
```

```
Dscp-outputq-threshold map:
```

	d1	d2	0	1	2	3	4	5	6	7	8	9
0 :	03-03	02-01	02-01	02-01	02-01	02-01	02-01	02-01	04-01	02-01		
1 :	04-02	02-01	04-02	02-01	04-02	02-01	02-01	03-01	02-01	03-01		
2 :	02-01	03-01	02-01	03-01	02-03	03-01	02-02	03-01	02-02	03-01		
3 :	02-02	03-01	01-03	04-01	02-02	04-01	02-02	04-01	02-02	04-01		
4 :	01-01	01-01	01-01	01-01	01-01	01-01	01-03	01-01	02-03	04-01		
5 :	04-01	04-01	04-01	04-01	04-01	04-01	02-03	04-01	04-01	04-01		
6 :	04-01	04-01	04-01	04-01								

The vertical column, marked *d1*, represents the first decimal number of the DSCP value, while the horizontal column, marked *d2*, represents the second decimal number of the DSCP value. For example, a *d1* value of 3 and a *d2* value of 2 yields a DSCP decimal value of 32, which corresponds to the CS4 service class. This is mapped to queue 1, drop threshold 3 in [Example 6-23](#) (highlighted in bold). Again, you still need to separately understand the mapping of specific video applications to service classes that are then marked with a particular DSCP value. However, combined with the knowledge of which traffic classes and are mapped to which transmit queue and drop threshold, you can use this information to troubleshoot video application performance issues across the Cisco Catalyst 3750G/3750E Series platforms.

To see the particular values of the buffer allocation and drop thresholds, you can issue the **show mls qos queue-set** command. An example of the output is shown in [Example 6-24](#).

Example 6-24 Example Output From Cisco Catalyst 3750G or 3750E Switch Stack show mls qos queue-set Command

```
me-eastny-3#show mls qos queue-set
Queueset: 1
Queue      : 1          2          3          4
-----
buffers   :    30       30       35         5
threshold1 :   100       70      100        40
threshold2 :   100       80      100       100
reserved   :    50      100       50       100
maximum    :   400      100      400       100

Queueset: 2
Queue      : 1          2          3          4
-----
buffers   :    25       25       25         25
threshold1 :   100      200      100        100
threshold2 :   100      200      100        100
reserved   :    50       50       50         50
maximum    :   400      400      400        400
```

In [Example 6-24](#), buffers are allocated according to weight ratios across the four egress queues. *Threshold1* and *threshold2* correspond to the two configurable thresholds per queue, with the third non-configurable threshold being at 100 percent queue depth. The Cisco Catalyst 3750G and 3750E Series switches dynamically share buffer space across an ASIC that may support more than one physical interface. The reserved and maximum settings are used to control the minimum reserved buffer percentage size guaranteed per queue per port, and the maximum buffer percentage size a particular port and queue can dynamically allocate when it needs additional capacity. The combination of drop statistics per queue, mapping of DSCP value to output queue, and the buffer allocations per queue-set, can be used to determine whether sufficient bandwidth has been allocated per service class (and per application if individual video applications are mapped to separate service classes corresponding to different DSCP values) on the Cisco Catalyst 3750G/3750E Series platforms.

When configured in a switch stack, statistics such as those found within the **show platform port-asic stats drop** command are not directly accessible on member switches from the master switch. To determine which switch is the master switch, and which switch you are currently logged into within the switch stack, you can run the **show switch** command. An example of this output is shown in [Example 6-25](#).

Example 6-25 Sample Output From Cisco Catalyst 3750G or 3750E Switch Stack show switch Command

```
me-eastny-3#show switch
Switch/Stack Mac Address : 0015.2b6c.1680

Switch#  Role  Mac Address      Priority Version  H/W  Current State
-----
*1       Master  0015.2b6c.1680    15         0      Ready
2       Member  001c.b0ae.bf00    1          0      Ready
```

The output from [Example 6-25](#) shows that Switch 1 is the Master switch, and the asterisk next to Switch 1 indicates that the output was taken from a session off this switch. To access the statistics from the **show platform port-asic stats drop** command on member switches of the stack, you must first establish a session to the member switch via the **session** command. This is shown in [Example 6-26](#).

Example 6-26 Example Output From Member Cisco Catalyst 3750G or 3750E Switch

```
me-eastny-3#session 2
```

```
me-eastny-3-2#show platform port-asic stats drop gig 2/0/24
```

```
Interface Gi2/0/24 TxQueue Drop Statistics
```

```
Queue 0
  Weight 0 Frames 0
  Weight 1 Frames 0
  Weight 2 Frames 0
Queue 1
  Weight 0 Frames 0
  Weight 1 Frames 0
  Weight 2 Frames 0
Queue 2
  Weight 0 Frames 0
  Weight 1 Frames 0
  Weight 2 Frames 0
Queue 3
  Weight 0 Frames 0
  Weight 1 Frames 0
  Weight 2 Frames 0
```

Note that when the **session 2** command is run, the command prompt changed from *me-eastny-3* to *me-eastny-3-2*, indicating that a session to member switch #2 has been established. After the session is established to the remote switch, the **show platform port-asic stats drop** command can be run on an interface, such as GigabitEthernet 2/0/24 shown in the example above, to obtain the drop statistics per queue on the port.

Router Show Policy Map Commands

For Cisco routers, the mapping of traffic classes to egress queues over WAN interfaces is accomplished via an egress policy map applied to the interface, in the same manner as the Cisco Catalyst 4500 with a Sup-6E supervisor. Again, the policy map can be viewed through the **show policy-map interface** command. [Example 6-27](#) shows the output from a Cisco ASR 1000 Series router with a OC-48 packet-over-SONET (POS) interface. Selected areas for discussion have been highlighted in bold.

Example 6-27 Output from Cisco 1000 Series ASR show policy-map interface Command

```
me-westwan-1#show policy-map int pos 1/1/0
POS1/1/0
Service-policy output: OC-48-WAN-EDGE
  queue stats for all priority classes:
    Queueing
    queue limit 512 packets
    (queue depth/total drops/no-buffer drops) 0/0/0
    (pkts output/bytes output) 18577357/16278388540
Class-map: VOIP-TELEPHONY (match-all)
  3347 packets, 682788 bytes
  30 second offered rate 0000 bps, drop rate 0000 bps
  Match: ip dscp ef (46)
  police:
    cir 49760000 bps, bc 1555000 bytes, be 1555000 bytes
    conformed 3347 packets, 682788 bytes; actions:
      transmit
    exceeded 0 packets, 0 bytes; actions:
      drop
    violated 0 packets, 0 bytes; actions:
      drop
    conformed 0000 bps, exceed 0000 bps, violate 0000 bps
  Priority: Strict, b/w exceed drops: 0
```

```

Class-map: REAL-TIME-INTERACTIVE (match-all)
  18574010 packets, 16277705752 bytes
  30 second offered rate 0000 bps, drop rate 0000 bps
  Match: ip dscp cs4 (32)
  police:
    cir 821040000 bps, bc 12315600 bytes, be 12315600 bytes
    conformed 18574010 packets, 16277705752 bytes; actions:
      transmit
    exceeded 0 packets, 0 bytes; actions:
      drop
    violated 0 packets, 0 bytes; actions:
      drop
    conformed 0000 bps, exceed 0000 bps, violate 0000 bps
  Priority: Strict, b/w exceed drops: 0

Class-map: NETWORK-CONTROL (match-all)
  1697395 packets, 449505030 bytes
  30 second offered rate 1000 bps, drop rate 0000 bps
  Match: ip dscp cs6 (48)
  Queueing
  queue limit 173 packets
  (queue depth/total drops/no-buffer drops) 0/0/0
  (pkts output/bytes output) 1644399/446219278
  bandwidth 5% (124400 kbps)

Class-map: CALL-SIGNALING (match-any)
  455516 packets, 157208585 bytes
  30 second offered rate 0000 bps, drop rate 0000 bps
  Match: ip dscp cs3 (24)
  Queueing
  queue limit 173 packets
  (queue depth/total drops/no-buffer drops) 0/0/0
  (pkts output/bytes output) 455516/157208585
  bandwidth 5% (124400 kbps)

Class-map: OAM (match-all)
  0 packets, 0 bytes
  30 second offered rate 0000 bps, drop rate 0000 bps
  Match: ip dscp cs2 (16)
  Queueing
  queue limit 173 packets
  (queue depth/total drops/no-buffer drops) 0/0/0
  (pkts output/bytes output) 0/0
  bandwidth 5% (124400 kbps)

Class-map: MULTIMEDIA-CONFERENCING (match-all)
  0 packets, 0 bytes
  30 second offered rate 0000 bps, drop rate 0000 bps
  Match: ip dscp af41 (34) af42 (36) af43 (38)
  Queueing
  queue limit 347 packets
  (queue depth/total drops/no-buffer drops) 0/0/0
  (pkts output/bytes output) 0/0
  bandwidth 10% (248800 kbps)

Class-map: MULTIMEDIA-STREAMING (match-all)
  0 packets, 0 bytes
  30 second offered rate 0000 bps, drop rate 0000 bps
  Match: ip dscp af31 (26) af32 (28) af33 (30)
  Queueing
  queue limit 173 packets
  (queue depth/total drops/no-buffer drops) 0/0/0
  (pkts output/bytes output) 0/0
  bandwidth 5% (124400 kbps)
  Exp-weight-constant: 4 (1/16)
  Mean queue depth: 0 packets
  class          Transmitted          Random drop  Tail drop  Minimum  Maximum  Mark
                  pkts/bytes pkts/bytes  pkts/bytes thresh  thresh  prob

```

```

Class-map: BROADCAST-VIDEO (match-all)
  771327514 packets, 1039749488872 bytes
  30 second offered rate 0000 bps, drop rate 0000 bps
  Match: ip dscp cs5 (40)
  Queueing
    queue limit 173 packets
    (queue depth/total drops/no-buffer drops) 0/0/0
    (pkts output/bytes output) 771327514/1039749488872
    bandwidth 5% (124400 kbps)
Class-map: TRANSACTIONAL-DATA (match-all)
  0 packets, 0 bytes
  30 second offered rate 0000 bps, drop rate 0000 bps
  Match: ip dscp af21 (18) af22 (20) af23 (22)
  Queueing
    queue limit 173 packets
    (queue depth/total drops/no-buffer drops) 0/0/0
    (pkts output/bytes output) 0/0
    bandwidth 5% (124400 kbps)
Class-map: BULK-DATA (match-all)
  0 packets, 0 bytes
  30 second offered rate 0000 bps, drop rate 0000 bps
  Match: ip dscp af11 (10) af12 (12) af13 (14)
  Queueing
    queue limit 139 packets
    (queue depth/total drops/no-buffer drops) 0/0/0
    (pkts output/bytes output) 0/0
    bandwidth 4% (99520 kbps)
Class-map: SCAVENGER (match-all)
  79 packets, 6880 bytes
  30 second offered rate 0000 bps, drop rate 0000 bps
  Match: ip dscp cs1 (8)
  Queueing
    queue limit 64 packets
    (queue depth/total drops/no-buffer drops) 0/0/0
    (pkts output/bytes output) 79/6880
    bandwidth 1% (24880 kbps)
Class-map: class-default (match-any)
  3209439 packets, 908940688 bytes
  30 second offered rate 1000 bps, drop rate 0000 bps
  Match: any
  Queueing
    queue limit 695 packets
    (queue depth/total drops/no-buffer drops) 0/0/0
    (pkts output/bytes output) 3052981/905185696
    bandwidth 20% (497600 kbps)
    Exp-weight-constant: 4 (1/16)
    Mean queue depth: 1 packets

```

	class	Transmitted pkts/bytes	Random drop pkts/bytes	Tail drop pkts/bytes	Minimum thresh	Maximum thresh	Mark prob
0	3052981/905185696	0/0	0/0	173	347	1/10	
1	0/0	0/0	0/0	194	347	1/10	
2	0/0	0/0	0/0	216	347	1/10	
3	0/0	0/0	0/0	237	347	1/10	
4	0/0	0/0	0/0	259	347	1/10	
5	0/0	0/0	0/0	281	347	1/10	
6	0/0	0/0	0/0	302	347	1/10	
7	0/0	0/0	0/0	324	347	1/10	

The main difference between the router and the Cisco Catalyst 4500 switch with Sup6E is that the router implements queues in software. It is therefore not limited to eight egress queues as is the Cisco Catalyst 4500 with Sup6E. [Example 6-27](#) shows the 12-class QoS model implemented with 12 separate egress queues over the OC-48 POS interface. Each class-map entry highlighted in bold corresponds to a queue. With this model, traffic from multiple service classes do not have to share a

single queue. This can provide a higher level of granularity into the visibility of various video applications, if separate applications are mapped to separate service classes. The traffic rate and drop rate, as well as counts of total packets and bytes outbound, and also counts of total drops for each queue can be seen from the **show policy-map interface** command when such a policy map is applied to the interface.

Simple Network Management Protocol

The Simple Network Management Protocol (SNMP) refers both to a specific protocol used to collect information and configure devices over an IP network, as well as an overall Internet-standard network management framework. The SNMP network management framework consists of the following components:

- Network management stations (NMSs)—Typically a server that runs network management applications, which in turn uses the SNMP protocol to monitor and control network elements.
- Network elements—The actual managed devices (routers, switches, TelePresence codecs, and so on) on the IP network.
- Agents—Software components running within network elements that collect and store management information.
- Managed objects—Specific characteristics of network elements that can be managed. Objects can be single entities or entire tables. Specific instances of managed objects are often referred to as variables.
- Management information bases (MIBs)—Collections of related management objects. MIBs define the structure of the management data through a hierarchical namespace using object identifiers (OIDs). Each OID describes a particular variable that can either be read from a managed object or set on a managed object. MIBs can be standards-based or proprietary. Because SNMP management information uses a hierarchical namespace, individual vendors can extend the management capabilities of their products through proprietary MIBs, which are typically published.

Currently, three versions of SNMP are commonly deployed:

- SNMPv1—The initial version introduced in the late 1980s. The security model used by SNMPv1 consists of authentication only, using community strings (read-only and read/write) that are sent in clear text within SNMP messages. Because of this, SNMPv1 is considered inherently insecure, and read/write capability should be used with caution, even over private networks.
- SNMPv2c—Proposed in the mid 1990s. The “c” in SNMPv2c indicates a simplified version of SNMPv2 that also uses a security model based on community strings. SNMPv2 improved the performance of SNMPv1 by introducing features such as the get-bulk-request protocol data unit (PDU) and notifications, both listed in [Table 6-3](#). However, because SNMPv2c still uses the same security model as SNMPv1, read/write capability should be used with caution.
- SNMPv3—Introduced in the early 2000s, and is currently defined primarily under IETF RFCs 3411-3418. A primary benefit of SNMPv3 is its security model, which eliminates the community strings of SNMPv1 and SNMPv2. SNMPv3 supports message integrity, authentication, and encryption of messages; allowing both read and read/write operation over both public and private networks.

As mentioned above, the SNMP protocol defines a number of PDUs, some of which are shown in [Table 6-3](#), along with the particular version of SNMP that supports them. These PDUs are essentially the commands for managing objects through SNMP.

Table 6-3 *SNMP Versions and PDUs*

Version	PDU	Description
SNMPv1	get-request	Command/response mechanism by which an NMS queries a network element for a particular variable
SNMPv1	response	Command/response mechanism by which an NMS receives information about a particular variable from a network element, based on a previously issued SNMP request message
SNMPv1	get-next-request	Command/response mechanism that can be used iteratively by an NMS to retrieve sequences of variables from a network element
SNMPv1	set-request	Issued by an NMS to change the value of a variable on a network element, or to initialize SNMP traps or notifications to be sent from a network element
SNMPv1	trap	Asynchronous mechanism by which a network elements issues alerts or information about an event to an NMS
SNMPv2	get-bulk-request	Improved command/response mechanism that can be used by an NMS to retrieve sequences of variables from a network element with a single command
SNMPv2	inform-request	Provides similar functionality as the trap PDU, but the receiver acknowledges the receipt with a response PDU

SNMP traps and/or informs (generically referred to a notifications) can be used to send critical fault management information, such as cold start events, link up or down events, and so on, from a medianet infrastructure device back to an NMS. This may be helpful in troubleshooting issues in which a video session has failed. SNMP GET commands can be used to pull statistics medianet infrastructure devices, which may then be used for assessing performance.

[Example 6-28](#) shows basic configuration commands for enabling SNMP on a Cisco Catalyst 6500 Switch.

Example 6-28 *Sample SNMP Configuration on a Cisco Catalyst 6500 Switch*

```
me-westcore-1(config)#snmp-server group group1 v3 priv access 10
me-westcore-1(config)#snmp-server user trapuser group1 v3 auth sha trappassword priv des
privacypassword
me-westcore-1(config)#snmp-server trap-source Loopback0
me-westcore-1(config)#snmp-server ip dscp 16
```

```
me-westcore-1(config)#snmp-server host 10.17.2.10 version 3 priv trapuser
me-westcore-1(config)#snmp-server enable traps
me-westcore-1(config)#access-list 10 permit 10.17.2.10
```

This configuration creates an SNMP group called *group1* that uses SNMPv3 and *access-list 10* to limit access to only the NMS workstation at IP address 10.17.2.10. A userid called *trapuser* is associated with the SNMP group. The userid uses Secure Hash Algorithm (SHA) for authentication with password *trappassword*, and DES for encryption with password *privacypassword*.

The commands **snmp-server enable traps** and **snmp-server host 10.17.2.10 version 3 priv trapuser** cause the switch to send SNMP traps to the NMS workstation. Note that this enables all traps available on the Cisco Catalyst switch to be enabled. The network administrator may desire to pare this down to traps applicable to the configuration of the Cisco Catalyst switch. Finally, the switch is configured to send traps using the Loopback0 interface with the DSCP marking of CS2 (note that not all platforms support the ability to set the DSCP marking of SNMP data).

The SNMP group information can be displayed with the **show snmp group** command shown in [Example 6-29](#).

Example 6-29 Sample Output From show snmp group Command on a Cisco Catalyst 6500 Switch

```
me-westcore-1#show snmp group

groupname: group1                      security model:v3 priv
readview : vldefault                  writeview: <no writeview specified>

notifyview: *tv.FFFFFFFFF.FFFFFFFFF.FFFFFFFFF.F
row status: active      access-list: 10
```

Similarly, the SNMP user information can be displayed with the **show snmp user** command shown in [Example 6-30](#).

Example 6-30 Sample Output From show snmp user Command on a Cisco Catalyst 6500 Switch

```
me-westcore-1#show snmp user

User name: trapuser
Engine ID: 800000090300001874E18540
storage-type: nonvolatile      active
Authentication Protocol: SHA
Privacy Protocol: DES
Group-name: group1
```

Note that the specific management objects that can be accessed via SNMP depend on the platform and software version of the platform. The Cisco MIB Locator, at the following URL, can be helpful in determining supported MIBS: <http://tools.cisco.com/ITDIT/MIBS/servlet/index>.

Application-Specific Management Functionality

The following sections summarize the components that provide application-specific management functionality for each of the four major video application solutions that co-exist over a converged medianet infrastructure: Cisco TelePresence, Cisco Digital Media Suite, Cisco IP Video Surveillance, and Cisco Desktop Video Collaboration.

Cisco TelePresence

Within Cisco TelePresence, application-specific management functionality is distributed among the following four major components of the deployment:

- Cisco TelePresence System Manager
- Cisco TelePresence Multipoint Switch
- Cisco Unified Communications Manager
- Cisco TelePresence System endpoints

Figure 6-26 provides a high-level summary of the main management roles of each of the components of a TelePresence deployment, each of which is discussed in the following sections.

Figure 6-26 *Summary of the Management Roles of the Components of a TelePresence Deployment*




Table 6-4 highlights the application-specific management functionality of each component.

Table 6-4 *Cisco TelePresence Application-Specific Management Functionality*

Management Product /Tool	Management Functionality	Description
Cisco TelePresence Manager	Fault management	- The Cisco TelePresence Manager web-based GUI provides a centralized view of the status of Cisco TelePresence Multipoint Switch devices and Cisco TelePresence System endpoints; including the status of the connectivity between Cisco TelePresence System endpoints and the Cisco Unified Communications Manager, the status of connectivity between Cisco TelePresence System endpoints and the Cisco TelePresence System Manager, and the synchronization of Cisco TelePresence System rooms with the e-mail/calendaring system used for scheduling meetings. - The Cisco TelePresence Manager web-based GUI also provides a centralized view of scheduled meetings, including those that have error conditions.
	Configuration management	- The Cisco TelePresence Manager web-based GUI provides device/element management capabilities, in that the configuration of the Cisco TelePresence Manager itself is accomplished through the GUI. Limited configuration support of the Cisco TelePresence Manager is available via a Secure Shell (SSH) command-line interface (CLI) as well. - The Cisco TelePresence Manager web-based GUI also provides a centralized view of the configuration capabilities of individual Cisco TelePresence System endpoints; including features such as high-speed auxiliary codec support, document camera support, interoperability support, and so on.
	Accounting management	- The Cisco TelePresence Manager interoperates with an e-mail/calendaring system to retrieve information for meetings scheduled by end users, and update individual Cisco TelePresence System endpoints regarding upcoming meetings. - The Cisco TelePresence Manager interoperates with one or more Cisco TelePresence Multipoint Switch devices to allocate segment resources for multipoint meetings scheduled by end users. - The Cisco TelePresence Manager web-based GUI provides a centralized view of ongoing and scheduled meetings for the entire TelePresence deployment, and per individual Cisco TelePresence System endpoint.
	Security management	- The Cisco TelePresence Manager web-based GUI provides a centralized view of the web services security settings of each Cisco TelePresence System endpoint, as well as a centralized view of the security settings of scheduled and ongoing meetings. - The Cisco TelePresence Manager currently provides administrative access via the local user database only.

Table 6-4 *Cisco TelePresence Application-Specific Management Functionality (continued)*

Cisco Unified Communications Manager	Fault management	The Cisco Unified Communications Manager provides limited fault management capability for Cisco TelePresence deployments. The Session Initiation Protocol (SIP) registration status of the Cisco TelePresence System endpoints to the Cisco Unified Communications Manager can be centrally viewed from the Cisco Unified Communications Manager Administration web-based GUI.
	Configuration management	- The Cisco Unified Communications Manager centrally controls the configuration of Cisco TelePresence System endpoints via the Cisco Unified Communications Manager Administration web-based GUI. - The Cisco Unified Communications Manager centrally controls the provisioning (that is, downloading of system load and device configuration) for Cisco TelePresence System endpoints via TFTP/HTTP server functionality.
	Accounting management	Call detail records (CDRs) captured by the Cisco Unified Communications Manager can be used to determine start and stop times for Cisco TelePresence meetings. These may be used to bill back individual departments based on TelePresence room resource usage.
	Performance management	The Cisco Unified Communications Manager Administration web-based GUI provides the ability to statically limit the amount of network bandwidth resources used for audio and video per TelePresence meeting and per overall location.  Note Note that Cisco Unified Communications Manager location-based admission control has no knowledge of network topology.
	Security management	- The Cisco Unified Communications Manager centrally controls the security configuration of Cisco TelePresence System endpoints via the Cisco Unified Communications Manager Administration web-based GUI. - In combination with the Certificate Authority Proxy Function (CAPF) and Certificate Trust List (CTL) Provider functionality, Cisco Unified Communications Manager provides the framework for enabling secure communications (media) and signaling (call signaling, and web services) for TelePresence deployments.

Table 6-4 Cisco TelePresence Application-Specific Management Functionality (continued)

Cisco TelePresence Multipoint Switch	Fault management	The Cisco TelePresence Multipoint Switch provides limited fault management capabilities. The web-based GUI interface can display errors and warnings for scheduled and non-scheduled meetings, as well as system errors.
	Configuration management	- The Cisco TelePresence Multipoint Switch web-based GUI provides device/element management capabilities, in that the configuration of the Cisco TelePresence Multipoint Switch itself is accomplished through the GUI. Limited configuration support of the Cisco TelePresence Multipoint Switch is available via an SSH CLI as well. - The Cisco TelePresence Multipoint Switch web-based GUI also provides the interface for administrators and meeting schedulers to configure static and ad hoc TelePresence meetings.
	Performance management	- The Cisco TelePresence Multipoint Switch web-based GUI provides centralized call statistics for multipoint calls, including SLA parameters such as bit rates, latency, drops, jitter, and so on, per Cisco TelePresence System endpoint. - The Cisco TelePresence Multipoint Switch web-based GUI also provides historical statistics for Cisco TelePresence Multipoint Switch resources including CPU utilization, traffic load per interface, packet discards, TCP connections, memory, and disk usage.
	Security management	- The Cisco TelePresence Multipoint Switch web-based GUI provides the interface for configuration of the security requirements for static and ad hoc TelePresence meetings. - Access control to the Cisco TelePresence Multipoint Switch is via the local database with three roles: administrator, meeting scheduler, or diagnostic technician.
Cisco TelePresence System Endpoint	Fault management	- The Cisco TelePresence System web-based GUI and SSH interfaces both provide device/element management capabilities, including a view of the system status, as well as diagnostics that can be used to troubleshoot the camera, microphone, and display components of the Cisco TelePresence System endpoint. - SIP Message log files accessed through the Cisco TelePresence System web-based GUI can be used to troubleshoot SIP signaling between the Cisco TelePresence System endpoint and Cisco Unified Communications Manager. - Additional Cisco TelePresence System log files can be collected and downloaded via the web-based GUI to provide system-level troubleshooting capabilities. - Status of peripheral devices (cameras, displays, microphones, and so on) can be accessed centrally via SNMP through the CISCO-TELEPRESENCE-MIB.
	Configuration management	The Cisco TelePresence System web-based GUI and SSH interfaces both provide information regarding current hardware and software versions and current configuration of the Cisco TelePresence System endpoint. Limited configuration is done on the Cisco TelePresence System endpoint itself. Most of the configuration is done via the Cisco Unified Communications Manager Administrator web-based GUI.

Table 6-4 *Cisco TelePresence Application-Specific Management Functionality (continued)*


	Accounting management	The Cisco TelePresence System web-based GUI provides access to statistics for ongoing calls, or the previous call if the Cisco TelePresence System endpoint is currently not in a call. Accounting management statistics include the call start time, duration of the call, remote number, bit rate, and the number of packets and bytes transmitted and received during the call. These statistics are also available via an SSH CLI as well as through SNMP.
	Performance management	- The Cisco TelePresence System web-based GUI provides access to statistics for ongoing calls, or the previous call if the Cisco TelePresence System endpoint is currently not in a call. Performance management statistics include parameters such as packet loss, latency, jitter, and out-of-order packets for audio and video media streams. These can be used to assess the performance of the network infrastructure in meeting service level agreements. These statistics are also available via SNMP through the CISCO-TELEPRESENCE-CALL-MIB. - An IP service level agreements (IPSLA) responder within the Cisco TelePresence System endpoint can be enabled, allowing the Cisco TelePresence System endpoint to respond to packets sent by an IPSLA initiator. IPSLA can be used to pre-assess network performance before commissioning the Cisco TelePresence System endpoint onto a production network, or used to assess ongoing network performance or when troubleshooting.
	Security management	- Access control to the individual Cisco TelePresence System endpoints is currently handled via a local database, although the userid and password used for access control are centrally managed via the configuration within the Cisco Unified Communications Manager Administration web-based GUI. - SNMP notifications can be set on the Cisco TelePresence System endpoint to alert after failed access control attempts.

**Note**

Both static location-based admission control and RSVP are considered part of performance management within this document, because the scheduling of resources is not done per end user, but to ensure that necessary resources are allocated to meet service level requirements.

Cisco TelePresence Manager

From a management perspective, the primary functions of Cisco TelePresence Manager are resource allocation, which is part of accounting management; and fault detection, which is part of fault management. Cisco TelePresence Manager allocates Cisco TelePresence System endpoints (meeting rooms) and Cisco TelePresence Multipoint Switch segment resources based on meetings scheduled by end users through an e-mail/calendaring system such as Microsoft Exchange or IBM Lotus Domino. Note that the Cisco TelePresence Manager has no knowledge of the underlying IP network infrastructure, and therefore has no ability to schedule any network resources or provide Call Admission Control (CAC) to ensure that the TelePresence call goes through during the scheduled time. [Figure 6-27](#) shows an example of the resource scheduling functionality of Cisco TelePresence Manager.

Figure 6-27 Cisco TelePresence Manager Resource Scheduling

Cisco TelePresence Manager periodically queries the e-mail/calendaring system to determine whether an end user has scheduled TelePresence rooms for an upcoming meeting. Having previously synchronized the TelePresence rooms defined within the Cisco Unified Communications Manager database with the TelePresence rooms defined within the e-mail/calendaring system database, the Cisco TelePresence Manager then pushes the meeting schedule the IP Phone associated with each TelePresence room. If a multipoint meeting has been scheduled by the end user, the Cisco TelePresence Manager selects an appropriate Cisco TelePresence Multipoint Switch for the meeting, and schedules the necessary resources for the meeting. The Cisco TelePresence Multipoint Switch then updates the end user via an e-mail confirmation.

**Note**

In [Figure 6-27](#) and throughout this chapter, the CTS codec and associated IP phone together are considered the CTS endpoint. The IP phone shares the same dial extension as the CTS codec, is directly connected to it, and is used to control TelePresence meetings.

The other primary management function of the Cisco TelePresence Manager is fault detection. Cisco TelePresence Manager includes the ability to centrally view error conditions in the various components of a Cisco TelePresence deployment. It also allows you to view error conditions that resulted in the failure of scheduled meetings. [Figure 6-28](#) shows an example of the Cisco TelePresence Manager screen used to view the status of Cisco TelePresence System endpoints.

Figure 6-28 Fault Detection Using the Cisco TelePresence Manager

The screenshot shows the Cisco TelePresence Manager web interface. The top navigation bar includes the Cisco logo, the title 'Cisco TelePresence Manager', and user links for 'admin', 'Logout', 'Preferences', 'Help', and 'About'. The left sidebar contains a tree view with categories like 'System Information', 'Support', and 'System Configuration'. The main content area is titled 'Support > Rooms' and has tabs for 'Summary', 'Status', and 'Capability'. The 'Status' tab is active, showing a table of rooms. Above the table are filters for 'Status' (set to 'All'), 'Room Name', and 'Unified CM'. The table shows 8 records, with columns for 'Status', 'Room Name', 'Connectivity' (with sub-columns 'Unified CM/CTS' and 'CTS Man/CTS'), 'CTS' (with sub-column 'CTS Error'), 'Unified CM' (with sub-column 'Profile'), 'Microsoft Exchange' (with sub-columns 'Email ID', 'Subscription', and 'Sync'). A 'Today's Meetings' summary is located in the bottom left corner.

Status	Room Name	Connectivity		CTS	Unified CM	Microsoft Exchange		
		Unified CM/CTS	CTS Man/CTS	CTS Error	Profile	Email ID	Subscription	Sync
OK	room-a0@tp.com	✓	✓	✓	✓	✓	✓	✗
OK	room-a2@tp.com	✓	✓	✓	✓	✓	✓	✗
OK	room-a13@tp.com	✓	✓	✗	✓	✓	✓	✗
OK	c1-room@tp.com	✓	✓	✗	✓	✓	✓	✗
OK	room-a1@tp.com	✓	✓	✓	✓	✓	✓	✗
OK	room-a14@tp.com	✗	✗	✗	✓	✓	✓	✗
OK	room-a3@tp.com	✓	✓	✓	✓	✓	✓	✗
OK	room-a12@tp.com	✓	✓	✓	✓	✓	✓	✗

Today's Meetings:
 With Error: 0
 In Progress: 0
 Scheduled: 0
 Other Errors: 10

As shown in Figure 6-28, a red X indicates some type of error condition that may need to be further investigated. These error conditions can include communication problems between the Cisco TelePresence System endpoint and the Cisco Unified Communications Manager, or between the Cisco TelePresence System endpoint and the Cisco TelePresence Manager itself. Other error conditions include problems within the Cisco TelePresence System endpoint itself, such as an issue with one of the peripherals (cameras, displays, and so on). Still other error conditions include synchronizing the TelePresence room defined within the Cisco Unified Communications Manager with the room definition within the e-mail/calendaring system. The System Status panel in the lower left corner of Figure 6-28 provides information regarding whether any meetings scheduled for the current day had errors. By clicking on the icons within the panel, you can gain additional details about the scheduled meetings and error conditions.

The Cisco TelePresence Manager also plays a minor role in both configuration management and security management. Cisco TelePresence Manager allows central viewing of specific configured features supported by a particular Cisco TelePresence System endpoint, such as a projector, document camera, or high-speed auxiliary codec support. It also allows you to centrally view the web service security settings for particular Cisco TelePresence System endpoints. Both of these functions are illustrated in Figure 6-29.

Figure 6-29 Cisco TelePresence Manager View of Configuration and Security Options for Cisco TelePresence System Endpoints

Cisco TelePresence Manager

admin

Logout

Preferences

Help

About

Host: tp-c1-ctm...

Support > Rooms

System Information

Support

Dashboard

Scheduled Meetings

Rooms

MCU Devices

Unified CM

System Configuration

Security Settings

Database

Microsoft Exchange

LDAP Server

Discovery Service

MCU Devices

Live Desks

Access Management

System Status

Today's Meetings:

With Error: 0

In Progress: 0

Scheduled: 0

Other Errors: 10

Summary

Status

Capability

Rooms

Status: All Room Name: Unified CM: Filter

Showing 1 - 8 of 8 records

Status	Room Name	CTS Version	Multipoint Conference	Projector	Document Camera	Conference Termination	Interop	HD Interop	Satellite Room	3D FPS	Studio Mode Recording	Web Services Security
	room-a0@tp.com	CTS 1.6.0 (3883)	✓	✗	✗	✓	✓	✓	✗	✗	✓	
	room-a2@tp.com	CTS 1.6.0 (3883)	✓	✗	✗	✓	✓	✓	✗	✗	✓	
	room-a13@tp.com	CTS 1.6.0 (3883)	✓	✗	✗	✓	✓	✓	✗	✗	✓	
	c1-room@tp.com	CTS 1.6.0 (3883)	✓	✓	✗	✓	✓	✓	✗	✗	✓	
	room-a1@tp.com	CTS 1.6.0 (3883)	✓	✗	✗	✓	✓	✓	✗	✗	✓	
	room-a14@tp.com	CTS 1.5.1 (2082)	✓	✗	✗	✓	✓	✗	✗	✗	✗	
	room-a3@tp.com	CTS 1.6.0 (3883)	✓	✗	✗	✓	✓	✓	✗	✗	✓	

84106

A red X indicates that the particular feature is either not configured or currently unavailable on the Cisco TelePresence System endpoint. A locked padlock indicates that web services communications from the Cisco TelePresence System endpoint are secured via Transport Layer Security (TLS). An open padlock indicates that web services communications from the Cisco TelePresence System endpoint are in clear text. Note that this functionality allows the viewing of only certain capabilities configured on the Cisco TelePresence System endpoint. All changes to the Cisco TelePresence System endpoint configuration are handled through the Cisco Unified Communications Manager, which is discussed next.

Cisco Unified Communications Manager

From an overall TelePresence deployment perspective, the primary function of the Cisco Unified Communications Manager is a SIP back-to-back user agent for session signaling. However, the Cisco Unified Communications Manager also plays a central management role for TelePresence deployments. From an FCAPS perspective, the primary roles of the Cisco Unified Communications Manager are in configuration management and security management. The device configuration and software image version for each of the Cisco TelePresence System endpoints is centrally managed through the Cisco Unified Communications Manager Administration web-based GUI, and downloaded to each Cisco TelePresence System endpoint when it boots up. The Cisco Unified Communications Manager therefore plays a central role in the initial provisioning of Cisco TelePresence System endpoints onto the network infrastructure, as well as any ongoing changes to the configuration of the Cisco TelePresence System endpoints. Figure 6-30 provides an example of the Cisco Unified Communications Manager Administrator web page, showing the TelePresence endpoints configured for the particular Cisco Unified Communications Manager.

Figure 6-30 Centralized Configuration Management via the Cisco Unified Communications Manager

The screenshot shows the Cisco Unified CM Administration web interface. The top navigation bar includes links for System, Call Routing, Media Resources, Voice Mail, Device, Application, User Management, Bulk Administration, and Help. The main content area is titled 'Find and List Phones' and shows a list of 8 records found. The table below lists the details for each device.

Device Name(Line)	Description	Device Type	Device Protocol	Status	IP Address	Copy	Super Copy
SEP0019AA044FE3	A13	Cisco TelePresence 1000	SIP	Registered with tp-c1-cm	10.32.1.11		
SEP0019AA04501C	A14	Cisco TelePresence 1000	SIP	Unknown	Unknown		
SEP0019AA044FDA	A12	Cisco TelePresence 1000	SIP	Registered with tp-c1-cm	10.31.1.11		
SEP0019AA044FB6	A3	Cisco TelePresence 1000	SIP	Registered with tp-c1-cm	10.22.1.11		
SEP0019AA044FCE	A1	Cisco TelePresence 1000	SIP	Registered with tp-c1-cm	10.20.1.11		
SEP0019AA044F50	A2	Cisco TelePresence 1000	SIP	Registered with tp-c1-cm	10.21.1.11		
SEP0019AA044FEC	A0	Cisco TelePresence 1000	SIP	Registered with tp-c1-cm	10.19.1.11		
SEP0019AA043D23	B0	Cisco TelePresence 3000	SIP	Registered with tp-c1-cm	10.35.1.11		

The detailed configuration for each Cisco TelePresence System endpoint can be viewed and modified by clicking on each device listed under the Device Name column shown in Figure 6-30. Included within the configuration of each Cisco TelePresence System endpoint is the security configuration. TelePresence security includes the use of Secure Real-time Transport Protocol (SRTP) for confidentiality and data authentication of the audio and video media streams; as well as TLS for confidentiality and data authentication of the SIP signaling and web services signaling between the various TelePresence components. For a thorough discussion of Cisco TelePresence security, see *Cisco TelePresence Secure Communications and Signaling* at the following URL:

<http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/telepresence.html>.


Cisco Unified Communications Manager also plays a role in accounting management, in that call detail records (CDRs) can be captured and used to bill back end users for TelePresence room usage. Cisco Unified Communications Manager can also play a role in performance management, in terms of bandwidth allocation, using static location-based CAC, although it is not in widespread use today for TelePresence deployments. The amount of bandwidth used for the audio and video components of an individual TelePresence call can be centrally controlled per zone via Cisco Unified Communications Manager. Also the total amount of bandwidth allocated for aggregate audio and video traffic to and from a location can be centrally controlled, via Cisco Unified Communications Manager. When a new TelePresence call requested via SIP signaling results in the amount of bandwidth allocated either for the individual call or aggregated for the entire location exceeding the configured zone or location bandwidth, the new call does not proceed. This helps maintain the overall quality of ongoing TelePresence calls. Because static location-based CAC has no knowledge of the underlying network infrastructure, it is typically effective only in hub-and-spoke network designs. Cisco offers location-based CAC integrated with Resource Reservation Protocol (RSVP), using an RSVP agent device, for VoIP and Cisco Unified Communications Manager-based Desktop Video Conferencing. However, this is currently not supported for Cisco TelePresence deployments.

Finally, the Cisco Unified Communications Manager plays a minor role in fault management. The SIP registration state of Cisco TelePresence System endpoints can be centrally viewed, and faults detected, from the Cisco Unified Communications Manager Administration web-based GUI interface, as shown in the Status column of [Figure 6-30](#).


Cisco TelePresence Multipoint Switch

From an overall TelePresence deployment perspective, the primary function of the Cisco TelePresence Multipoint Switch is to provide switching of the video and audio media for multipoint TelePresence calls. However, as with the Cisco Unified Communications Manager, the Cisco TelePresence Multipoint Switch also plays a management role for TelePresence deployments. From an FCAPS perspective, the primary function of Cisco TelePresence Multipoint Switch is in performance management. The Cisco TelePresence Multipoint Switch can collect performance data regarding the Cisco TelePresence System endpoints in an ongoing multipoint meeting. [Figure 6-31](#) shows an example of the call statistics collected by the Cisco TelePresence Multipoint Switch for one of the Cisco TelePresence System endpoints within a three-party multipoint call.


Figure 6-31 Cisco TelePresence Multipoint Switch Performance Statistics for Ongoing Meetings



Call statistics include the maximum jitter seen for the last period (ten seconds), the maximum jitter seen for the duration of the call, latency, and lost packets in both the transmit and receive directions. These statistics are collected by the Cisco TelePresence Multipoint Switch for both the audio and video channels for each of the endpoints. Cisco TelePresence Multipoint Switch call statistics can be used to quickly view whether any leg of a multipoint call is outside the required service level agreement (SLA) parameters of jitter, packet loss, and latency. Statistics regarding the overall status of the Cisco TelePresence Multipoint Switch are also collected, as shown in [Figure 6-32](#). These statistics include CPU loading of the Cisco TelePresence Multipoint Switch, traffic loading for the FastEthernet interfaces, Cisco TelePresence Multipoint Switch memory and disk utilization, open TCP connections, and Cisco TelePresence Multipoint Switch packet discards.

Figure 6-32 Cisco TelePresence Multipoint Switch Statistics for Overall Status

Each of the categories shown in [Figure 6-32](#) can be expanded by clicking on it. For example, the Active CPU Load Average Value * 100 statistics can be expanded, as shown in [Figure 6-33](#). This provides detail regarding CPU utilization on a daily, weekly, monthly, and yearly basis.

Figure 6-33 Expanded Statistics for Active CPU Load Average Value * 100

The statistics collected by the Cisco TelePresence Multipoint Switch can be used to perform long-term trend analysis, allowing you to plan the deployment of additional Cisco TelePresence Multipoint Switch resources before capacity limits are reached and service is degraded.

The Cisco TelePresence Multipoint Switch also plays a role in both configuration management and security management. Static and ad hoc meetings, as well as the security requirements for those meetings, are configured directly on the Cisco TelePresence Multipoint Switch by network administrators or meeting schedulers. Meetings can be configured as non-secured, secured, or best effort. Best effort means that if all endpoints support encryption, the call goes through as secured. However, if any endpoint does not support encryption, the call falls back to an unencrypted or non-secured call. Access control to the Cisco TelePresence Multipoint Switch is controlled through its local database, with the capability of defining three roles: administrators, who have full access to the system; meeting schedulers, who can only schedule static or ad hoc meetings; and diagnostic technicians, who can perform diagnostics on the Cisco TelePresence Multipoint Switch.

Finally, the Cisco TelePresence Multipoint Switch plays a minor role in fault management. The Cisco TelePresence Multipoint Switch logs system errors as well as error or warning conditions regarding meetings. For example, a error message might indicate that a Cisco TelePresence System endpoint cannot join a secure multipoint meeting because it is not configured to support encryption. The error messages can be viewed via the web-based GUI interface of the Cisco TelePresence Multipoint Switch.

Cisco TelePresence System Endpoint

From an overall TelePresence deployment perspective, the primary function of the Cisco TelePresence System endpoint is to transmit and receive the audio and video media for TelePresence calls. However, from an FCAPS management perspective, the Cisco TelePresence System endpoint also plays a role in performance management. The Cisco TelePresence System endpoint collects statistics regarding ongoing TelePresence meetings, or the previous meeting if the device is not in a current call. These can be viewed through the Cisco TelePresence System web-based GUI interface, as shown in the example in Figure 6-34.

Figure 6-34 Cisco TelePresence System Endpoint Call Statistics



228411

As with the Cisco TelePresence Multipoint Switch, statistics are collected for both the audio and video channels for each of the endpoints. The statistics include SLA parameters such as average latency for the period (ten seconds) and for the call; average jitter for the period and the call; percentage of lost packets for the period and the call; as well as total lost packets, out of order packets, late packets, or duplicate packets. They also include some accounting management information such as the call start time, call duration, and the remote phone number; as well as the bandwidth of the call, and the number of bytes or packets sent and received. These statistics can also be collected and stored centrally via SNMP through the CISCO-TELEPRESENCE-CALL-MIB supported by Cisco TelePresence System endpoints running Cisco TelePresence System version 1.5 or higher software. These statistics can then be used for performance analysis and/or billing purposes. A more limited set of call statistics, primarily the accounting management statistics, is available through the SSH CLI, as shown in [Example 6-31](#).

Example 6-31 Call Statistics Available via the SSH Command-Line Interface

```
admin:show call statistics all

Call Statistics
Registered to Cisco Unified Communications Manager : Yes
Call Connected: Yes

Call type       : Audio/Video Call Call Start Time: Oct 27 11:48:29 2009
Duration (sec)  : 2119                      Direction: Outgoing
Local Number    : 9193921003                Remote Number: 9193926001
State           : Answered                  Bit Rate: 4000000 bps,1080p
Security Level  : Non-Secure

-- Audio --
IP Addr Src: 10.22.1.11:25202      Dst : 10.16.1.20:16444
Latency Avg: 1                    Period: 1

Statistics      Left      Center      Right      Aux
Tx Media Type   N/A       AAC-LD      N/A         AAC-LD
Tx Bytes        0         17690311   0           0
Tx Packets      0         105930     0           0
Rx Media Type   AAC-LD     AAC-LD      AAC-LD      AAC-LD
Rx Bytes        0         0          0           0
Rx Packets      0         0          0           0
Rx Packets Lost 0         0          0           0

-- Video --
IP Addr Src: 10.22.1.11:20722      Dst : 10.16.1.20:16446
Latency Avg: 1                    Period: 1

Statistics      Center      Aux
Tx Media Type   H.264      H.264
Tx Bytes        1068119107 0
Tx Packets      1087322    0
Rx Media Type   H.264      H.264
Rx Bytes        1067246669 0
Rx Packets      1055453    0
Rx Packets Lost 1876      0

-- Audio Add-in --
IP Addr Src: 10.22.1.11:0      Dst : 0.0.0.0:0
Latency Avg: N/A              Period: N/A

Statistics      Center
Tx Media Type   N/A
Tx Bytes        0
Tx Packets      0
Rx Media Type   N/A
```

```

Rx Bytes          0
Rx Packets        0
Rx Packets Lost   0

```

In addition to passive collection of statistics during calls, Cisco TelePresence System endpoints can also function as IPSLA responders, as of Cisco TelePresence System version 1.4 or higher. IPSLA can be used to pre-assess network performance before commissioning the Cisco TelePresence System endpoint onto a production network. Optionally, IPSLA can be used to assess network performance when troubleshooting a performance issue of a production device. See [Network-Embedded Management Functionality, page 6-2](#) for more information regarding the use of IPSLA for performance management.

The Cisco TelePresence System endpoint also supports extensive fault management capabilities through diagnostics that can be used to troubleshoot the camera, microphone, and display components of the Cisco TelePresence System endpoint. These diagnostics can be accessed through either the web-based GUI interface of the Cisco TelePresence System endpoint, or through the SSH CLI. Additionally, SIP log files stored within the Cisco TelePresence System endpoint can be accessed through the web-based GUI to troubleshoot call signaling between the Cisco TelePresence System endpoint and the Cisco Unified Communications Manager. Finally, the status of each component (displays, microphones, speakers, and so on) of the Cisco TelePresence System endpoint can be accessed centrally via SNMP through the CISCO-TELEPRESENCE-MIB. This management information base (MIB) is supported on Cisco TelePresence System endpoints running software version 1.5 and higher.

The Cisco TelePresence System endpoint itself also plays a minor role in configuration management and security management. In terms of configuration management, the configuration of the Cisco TelePresence System endpoint, including specific hardware and software levels of each component (displays, microphones, speakers, and so on), can be viewed through the web-based GUI interface, or accessed through the SSH CLI. However, modifications to the configuration of the Cisco TelePresence System endpoint is primarily controlled centrally by the Cisco Unified Communications Manager. In terms of security management, access to the Cisco TelePresence System endpoint is via its local database. However, the userid and passwords are configured centrally within the Cisco Unified Communications Manager and downloaded to the Cisco TelePresence System endpoint. Cisco TelePresence System 1.6 introduces password aging for the SSH and web-based GUI interface of the Cisco TelePresence System endpoints. The security settings of the Cisco TelePresence System endpoint are controlled via the Cisco Unified Communications Manager centrally, as discussed previously. Finally, the Cisco TelePresence System endpoint also supports the ability to generate SNMP traps for authentication failures when attempting to access the system. This can be used to monitor the Cisco TelePresence System endpoints against brute-force password attacks.

Cisco TelePresence SNMP Support

As of this writing (CTS version 1.6), CTS, CTMS, and CTS Manager support the MIBs listed in [Table 6-5](#). Future versions of Cisco TelePresence may add additional SNMP MIB support.

Table 6-5 *MIB Support in TelePresence Endpoints (CTS, CTMS, and CTS-MAN)*

MIB Name	Description
CISCO-SYSLOG-MIB	Provides an SNMP interface into syslog messages
CISCO-CDP-MIB	Provides Ethernet neighbor information, such as the attached IP phone and upstream switch
HOST-RESOURCES-MIB	Provides system operating system information such as system CPU, memory, disk, clock, and individual process information

Table 6-5 MIB Support in TelePresence Endpoints (CTS, CTMS, and CTS-MAN) (continued)

RFC-1213-MIB	Provides basic MIB2 structure/information such as system uptime, system description, SNMP location, and SNMP contact
IF-MIB	Provides Ethernet interface statistics, such as bytes and packets transmitted and received, as well as interface errors
UDP-MIB	Provides the number of inbound and outbound UDP packets, as well as drops
TCP-MIB	Provides the number of inbound and outbound TCP packets, connections, and number of TCP retransmissions
CISCO-TELEPRESENCE-MIB	Provides notification on peripheral and user authentication failures; also allows for the remote restart of the CTS device
CISCO-TELEPRESENCE-CALL-MIB	Provides detailed call statistics for TelePresence meetings
CISCO-ENVMON-MIB	Provides system temperature
SNMP protocol-specific MIBs: • SNMP-FRAMEWORK-MIB • SNMP-MPD-MIB • SNMP-NOTIFICATION-MIB • SNMP-TARGET-MIB • SNMP-USM-MIB • SNMP-VACM-MIB	Provides information relating to the SNMP daemon configuration and current state

IP Video Surveillance

For information regarding the medianet management functionality of the Cisco IP Video Surveillance solution, see the *Cisco IP Video Surveillance Design Guide* at the following URL:
http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/IPVS/IPVS_DG/IPVS_DG.pdf.

Digital Media Systems

For information regarding the medianet management functionality of the Cisco Digital Media Systems solution, see the *Cisco Digital Media System 5.1 Design Guide for Enterprise Medianet* at the following URL: http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/DMS_DG/DMS_DG.html.

Desktop Video Collaboration

Future revisions of this document will include discussion regarding medianet management functionality for Cisco Desktop Video Collaboration solutions.

Summary

This design chapter has focused on functionality that can be used to provide increased visibility and management of video flows within an enterprise medianet. From a high-level perspective, the functionality can be separated into two broad categories: application-specific management functionality and network-embedded management functionality. Application-specific management refers to functionality within the components of a particular video solution: Cisco TelePresence, Cisco IP Video Surveillance, Cisco Digital Media Systems, and Cisco Desktop Video Collaboration.

Network-embedded management refers to functionality embedded within the medianet infrastructure itself, which allows both visibility and management of video flows. These include specific embedded software features such as NetFlow and IPSLA, the Cisco router and Cisco Catalyst switch CLI itself, and also hardware modules such as the Cisco NAM embedded within Cisco Catalyst 6500 Series Switches. By implementing a QoS model that separates the various video applications into different service classes, which are then mapped to separate queues and drop thresholds within Cisco router and switch platforms, you can gain additional visibility into the video applications themselves by collecting flow information based on DSCP aggregation, as well as monitoring the router and switch queues. Typically, the more granular the QoS model (that is, up to 12 service classes) and the more queues and drop thresholds deployed throughout medianet infrastructure devices, the greater the visibility and ability to manage the flows.




CHAPTER 7

Medianet Auto Configuration

Medianet auto configuration is designed to ease the administrative burden on the network administrator by allowing the network infrastructure to automatically detect a medianet device attached to a Cisco Catalyst switch via the Cisco Medianet Service Interface (MSI) and configure the switch port to support that particular device. [Figure 7-1](#) shows an example with a Cisco digital media player (DMP) and a Cisco IP Video Surveillance (IPVS) camera connected to a Cisco Catalyst switch.

Figure 7-1 Example of Auto Configuration



From an FCAPS perspective, auto configuration is part of configuration management. The current medianet auto configuration functionality includes two features:

- [Auto Smartports](#)
- [Location Services](#)

Auto Smartports

Auto Smartports (ASP) macros are an extension to Cisco Static Smartports macros. With Static Smartports, either built-in or user-defined macros can be applied manually to an interface by a network administrator. Macros contain multiple interface-level switch commands bundled together under the macro name. For repetitive tasks, such as multiple interfaces which require the same configuration,

Static Smartports can reduce both switch configuration errors and the administrative time required for such configuration. ASP macros extend this concept by allowing the macro to be automatically applied to the interface based upon built-in or user-defined trigger events. The mechanisms for detecting trigger events include the use of Cisco Discovery Protocol (CDP) packets, Link-Level Discovery Protocol (LLDP) packets, packets which include specific MAC addresses or Organizational Unique Identifiers (OUIs), and attribute-value (AV) pairs within a RADIUS response when utilizing ASP macros along with 802.1x/MAB.

**Note**

Triggering an ASP macro by passing a RADIUS AV pair to the Catalyst switch has not been validated at the time this document was written.

Platform Support

Table 7-1 shows Cisco Catalyst switch platforms and IOS software revisions which currently support ASP macros.

Table 7-1 Platform and IOS Revision for Auto Smartports Support

Platform	ASP IOS Revisions	Enhanced ASP IOS Revisions
Catalyst 3750-X Series Switches	12.2(53)SE2	12.2(55)SE
Catalyst 3750, 3560, 3750-E, and 3560-E Series Switches	12.2(50)SE, 12.2(52)SE	12.2(55)SE
Cisco ISR EtherSwitch Modules ¹	12.2(50)SE, 12.2(52)SE	12.2(55)SE
Catalyst 4500 Series Switches	IOS 12.2(54)SG	Future Release
Catalyst 2975 Series Switches	12.2(52)SE	12.2(55)SE
Catalyst 2960-S and 2960 Series Switches	12.2(50)SE, 12.2(52)SE, 12.2(53)SE1	12.2(55)SE

1. This applies to ISR EtherSwitch Modules which run the same code base as Catalyst 3700 Series switches.

There are basically two versions of ASP macros, which are referred to as ASP macros and Enhanced ASP macros within this document. This is due to differences in functionality between ASP macros running on older IOS software revisions and ASP macros running on the latest IOS software revisions. Table 7-2 highlights some of these differences.

Table 7-2 Partial List of Feature Differences Between ASP Macros and Enhanced ASP Macros

Feature	ASP Macros	Enhanced ASP Macros
Macro-of-last-resort	No	Yes
Custom macro	No	Yes
Ability to enable/disable individual device macros	No	Yes
Ability to enable/disable individual detection mechanisms	No	Yes
Built-in ip-camera macro	Yes, without AutoQoS	Yes, with AutoQoS

Table 7-2 *Partial List of Feature Differences Between ASP Macros and Enhanced ASP Macros*

Feature	ASP Macros	Enhanced ASP Macros
Built-in media-player macro	Yes, with MAC-address/OUI trigger	Yes, with CDP trigger or MAC-address/OUI trigger
Built-in phone macro	Yes	Yes
Built-in lightweight access-point macro	Yes	Yes
Built-in access-point macro	Yes	Yes
Built-in router macro	Yes	Yes
Built-in switch macro	Yes	Yes
Built-in detection mechanisms	CDP, LLDP, mac-address, and RADIUS AV pair	CDP, LLDP, mac-address, and RADIUS AV pair

Throughout this document, the term “ASP Macros” is generally used to refer to both the non-enhanced and enhanced Auto Smartports macro functionality. The term “Enhanced ASP Macros” is only used when specific features which are supported by the enhanced Auto Smartports functionality are discussed.

As mentioned above, from a medianet perspective the primary benefit of ASP macros is to ease the administrative burden of provisioning medianet devices onto the IP network infrastructure. [Table 7-3](#) lists the medianet devices currently supported by built-in ASP macros.

Table 7-3 *Medianet Devices with Built-in ASP Macros*

Device	Models	Software Revisions and Comments
Cisco IPVS Cameras	CIVS-IPC-2400 Series, CIVS-IPC-2500 Series, CIVS-IPC-4300, CIVS-IPC-4500 ¹	Revision 1.0.7. CDP detection mechanism only.
Cisco DMPs	Cisco DMP 4305G, Cisco DMP 4400G	Revision 5.2.1. OUI detection mechanism only.
Cisco DMPs	Cisco DMP 4310G	Revision 5.2.2. CDP or OUI detection mechanisms

1. Cisco 5000 Series IPVS cameras currently do not support CDP.

Auto Smartports also has built-in macros for devices which are not specific to a medianet. These devices include routers, switches, access-points, and lightweight (CAPWAP/LWAP enabled) access-points.

Switch Configuration

Auto Smartports macro processing is enabled globally on supported Catalyst switches with the command:

```
macro auto global processing
```

This command also automatically enables ASP macro processing on all switchports. This could lead to unwanted consequences when first enabling ASP macros on a Catalyst switch. For example, the network administrator may not want Auto Smartports to automatically change the configuration of existing

uplink ports connected to infrastructure devices such as switches and routers. Such changes could result in an unintentional service outage when first enabling ASP macros. The network administrator should first disable ASP macro processing on interfaces where it is not desired. The following examples show how to disable ASP macro processing at the interface-level for a single interface and a range of interfaces.

Single Interface	Range of Interfaces
<code>interface GigabitEthernet1/0/1</code> <code>no macro auto processing</code>	<code>interface range GigabitEthernet1/0/1 - 48</code> <code>no macro auto processing</code>

**Note**

The **no macro auto processing** interface-level command will currently not appear within the switch configuration—even if it has been typed in—until the **macro auto global processing** global command is entered into the switch configuration. Therefore, the network administrator must manually keep track of which interfaces they have disabled for ASP macro processing before enabling the **macro auto global processing** global command.

The **macro auto global processing** command has one or two optional forms as shown below, depending upon the Catalyst switch platform.

Catalyst Access Switches	Catalyst 4500 Series Switches
<code>macro auto global processing fallback cdp</code>	<code>macro auto global processing fallback cdp</code>
	Or: <code>macro auto global processing fallback llbp</code>

These forms of the command may be used when the network administrator has deployed 802.1x or MAB and wishes either CDP packets or LLDP packets to be used for ASP macro trigger detection—after 802.1x/MAB authentication is successful. This functionality may also be enabled per interface with the following interface-level command:

```
macro auto processing fallback <fallback method>
```

The fallback method can either be CDP or LLDP, depending upon the platform, as discussed above. [Security Considerations](#) further describes the use of MAB with CDP fallback.

**Note**

Since none of the medianet devices currently support an 802.1x supplicant, all testing was performed utilizing MAB with CDP fallback only.

By default, all built-in ASP device macros (also referred to as ASP scripts) are enabled when ASP macro processing is enabled on a Catalyst Switch. [Table 7-4](#) shows the built-in device ASP macros, any configurable parameters which can be passed into the macros when they execute, and the default values of those parameters. These can be displayed through the **show macro auto device** command on the Catalyst switch.

Table 7-4 ASP Built-in Device Macros

Macro Name	Cisco Device	Configurable Parameters	Defaults
access-point	Autonomous Access Point	NATIVE_VLAN	VLAN1
ip-camera	Video Surveillance Camera	ACCESS_VLAN	VLAN1
lightweight-ap	CAPWAP / LWAP Access Point	ACCESS_VLAN	VLAN1
media-player	Digital Media Player	Access_VLAN	VLAN1
phone	IP Phone	ACCESS_VLAN, VOICE_VLAN	VLAN1, VLAN2
router	Router	NATIVE_VLAN	VLAN1
switch	Catalyst Switch	NATIVE_VLAN	VLAN1

As listed in [Table 7-2](#), one of the benefits of implementing Enhanced ASP macros is the ability to enable/disable individual built-in device macros. This can be accomplished through the following global switch command:

```
macro auto global control device <list of devices separated by spaces>
```

The list of devices includes one or more of the macro names listed in [Table 7-4](#). For example, in order to enable only the built-in ip-camera and media-player ASP macros, the network administrator would configure the following command on a switch platform which supports Enhanced ASP macros:

```
macro auto global control device ip-camera media-player
```

Built-in device macros can also be enabled/disabled per interface with the following interface-level command:

```
macro auto control device <list of devices separated by spaces>
```

The list of devices includes one or more of the macro names listed in [Table 7-4](#). [Security Considerations](#) discusses some potential security reasons why the network administrator may choose to restrict which macros are enabled on a particular switch platform.

With regular ASP macro support, the only way the network administrator can “disable” a built-in macro is to override the macro in such a manner that does nothing. [Overriding Built-in Macros](#) discusses this further.

For the most part, the only parameters which can be passed into the built-in ASP macros are VLAN parameters, as shown in [Table 7-4](#). These can be passed using the following global switch configuration command:

```
macro auto device <device> <line>
```

The device is one of the macro names listed in [Table 7-4](#) and line is one of the following forms:

ACCESS_VLAN=<vlan>	Used for ip-camera, lightweight-ap, and media-player macros
NATIVE_VLAN=<vlan>	Used for access-point, router, and switch macros
ACCESS_VLAN=<vlan> VOICE_VLAN=<vlan>	Used for the phone macro

For example, in order to set the access VLAN to VLAN302 for IPVS cameras which use ASP macros, the network administrator would configure the following global switch command:

```
macro auto device ip-camera ACCESS_VLAN=VLAN302
```

From a network design perspective, the ability to set the VLAN for medianet devices is important for two reasons. First, the default macro parameters typically set the access VLAN to VLAN1. Cisco SAFE security best practices have long recommended that network administrators utilize a VLAN other than VLAN1 for devices. Second, the ability to set the VLAN allows different medianet devices to be placed on separate VLANs. This may be beneficial from a traffic isolation perspective, either for QoS or for security purposes. For example, a network administrator may wish to separate all IPVS cameras on a particular Catalyst switch to a VLAN which is separate from normal PC data traffic. The downside of this is that all devices of a particular type are placed into the same VLAN by Auto Smartports. For example, currently there is no ability to place certain DMPs into one VLAN and other DMPs into another VLAN. This may be desirable if two departments within an organization each control their own sets of DMPs and the content to be displayed.

By default, three mechanisms for detecting ASP trigger events are enabled automatically when ASP macro processing is enabled on a Catalyst Switch. These detection mechanisms are shown in [Table 7-5](#).

Table 7-5 *ASP Detection Mechanisms*

Detection Mechanism Name	Description
cdp	Instructs the switch to look for ASP triggers within CDP packets.
lldp	Instructs the switch to look for ASP triggers within LLDP packets.
mac-address	Instructs the switch to look for either full MAC addresses or the OUI portion of MAC addresses which match in list contained within either a built-in or user-defined MAC-address trigger.



Note

The list above does not include the use of an RADIUS AV pair to return a trigger name, which can be used when 802.1x/MAB authentication is enabled as well as ASP macros.

[ASP Macro Details](#) details how ASP macros are triggered. With Enhanced ASP macros, the network administrator can disable any of the detection mechanisms via the following global switch configuration command:

```
macro auto global control detection <list of detection mechanism names>
```

The list of detection mechanism names corresponds to one or more of the detection mechanism names in [Table 7-5](#). For example, in order enable only CDP and MAC address detection mechanisms on a given Catalyst switch, the network administrator can configure the following global switch configuration command:

```
macro auto global control detection cdp mac-address
```

Detection mechanisms can also be enabled/disabled per interface with the following interface-level command:

```
macro auto control detection <list of detection mechanism names>
```

From a network design perspective, it may be beneficial to disable unused detection mechanisms if the network administrator knows that there are no devices which will utilize a particular mechanism. This can prevent unexpected switchport configuration changes due to accidental triggering of an ASP macro. For instance, medianet specific devices such as Cisco DMPs and IPVS cameras do not currently support the LLDP protocol. Therefore a network administrator who is interested in using Enhanced ASP macros

to ease the administrative burden of configuring these devices across the network infrastructure may decide to enable only CDP and MAC address detection mechanisms. Finally, note that for regular ASP macros, there is no method of disabling a particular ASP detection mechanism.

One additional command is worth noting. Normally ASP macros are applied to an interface upon detecting a trigger event after link-up and removed upon a link-down event by an anti-macro. Since it is recommended that the interface begin with a default interface configuration (with exceptions when using Location Services, the custom macro, or 802.1x/MAB), the link-down returns the interface to its initial default configuration. The **macro auto sticky global** configuration command causes the macro which is applied upon link-up to remain applied upon link-down. The **macro auto port sticky** interface-level configuration command has the same effect on a port-by-port basis.

The benefit of the **macro auto sticky** and **macro auto port sticky** commands is that the macro is only run once when the medianet device is first seen on the interface, versus every time the interface is transitioned from down to up. The running configuration of the switch always shows the applied macro as well, regardless of whether the device is currently up or down. This may be beneficial from a troubleshooting perspective. The downside is that ASP macros which include the **switchport port-security** command may cause the interface to go into an error-disabled state should another device with a different MAC address be placed onto the switchport.

This document is primarily concerned with the built-in ip-camera and media-player ASP macros, since they relate directly to medianet devices. The built-in access-point and lightweight-ap ASP macros were not evaluated for this document. Future revisions of the Medianet Reference Design may include design guidance regarding wireless connectivity and video. The built-in phone macro was evaluated only from the perspective of its effect on medianet devices such as Cisco TelePresence (CTS) endpoints and desktop video conferencing units which consists of a PC running software daisy-chained to an IP phone.

ASP Macro Details

An understanding of the implementation of ASP will assist in general troubleshooting, customization, and security considerations. The macros are fairly transparent and supported by several useful **show** commands and debugging tools. The logical flow is organized into three distinct functions: detection, policy, and function. Detection is used to determine that an actionable event has occurred and selects an appropriate method to classify the device. Three detection methods are available. These are neighbor discovery using either LLDP or CDP, Mac Address, or 802.1x identity. They can be seen with the IOS **exec** command:

```
sh macro auto event manager detector all
```

No.	Name	Version	Node	Type
1	identity	01.00	node0/0	RP
2	neighbor-discovery	01.00	node0/0	RP
3	mat	01.00	node0/0	RP

A detail of each detector is also available that lists more information concerning events and variables that are passed back and forth between the IOS event manager and the ASP detector. The details are not explained here. It is only important to know that ASP starts when IOS calls one or more of these detectors and passes information such as interface, mac-address, and link events into the detector. The detectors are associated to a policy that can examine the variables that are passed and make a decision. The policy generates an event trigger. These policy shell scripts do the major work within ASP. The link between detector and policy can be seen with the **show** command:

```
sh macro auto event manager policy registered
```

Typically six policies are registered with the three detectors. Output from the **show** command is summarized in [Table 7-6](#).

Table 7-6 Policies Associated With ASP Detectors

	Detector	Policy	Event
1	neighbor-discovery	Mandatory.link.sh	link-event down
2	neighbor-discovery	Mandatory.link2.sh	link-event admindown
3	neighbor-discovery	Mandatory.lldp.sh	lldp update
4	mat	Mandatory.mat.sh	use-mat-db yes hold-down 65.000
5	neighbor-discovery	Mandatory.cdp.sh	cdp update
6	identity	Mandatory.identity.sh	aaa-attribute {auto-smart-port}

As an example, when a link-event down occurs, neighbor discovery will run the script Mandatory.link.sh. Details of the script can be seen with the command:

```
sh macro auto event manager policy registered detailed <policy script>
```

The scripts can be read with a little background in programming. It is possible to register user-generated scripts, although the details of that procedure are not included in this document. There are significant differences in the system scripts packaged in Auto Smartports and those found in Enhanced Auto Smartports. Each script fetches information from the router configuration, such as the current macro description. Based on the calling event, passed variables, and interface configuration, the policy script generates a trigger. Triggers are mapped to shell functions. This mapping can be seen with the command:


```
sh shell trigger
```

This displays all of the mapped triggers. However ASP is only relevant to those triggers that map to a function that contains AUTO_SMARTPORT in the function name. Arguments are passed into the shell function from the trigger. Common arguments include \$LINKUP, \$INTERFACE, \$TRIGGER, and \$ACCESS_VLAN. With this information, the function applies the appropriate configuration to the interface of the switch. The functions can be customized. The shell function details can be seen with the command:

```
show shell function
```

As an example, consider the case where a CDP packet is received on an interface that was previously configured with the appropriate ASP configuration. Neighbor-discovery calls the script Mandatory.cdp.sh. The script first checks to see if CDP detection is available; if so, then the CDP capabilities are checked. If the host bit is set, then the CDP platform type is compared against known types. The previous trigger is noted by pulling the macro description from the interface configuration. Another check is made to see if discovery is enabled for that particular type of device. If so, then the script continues to check the other capabilities bits for Phone, Access Point, Router, or Switch. If the host bit is set in conjunction with the phone bit, then the phone trigger takes precedence. Finally a trigger is generated and mapped to a shell function. Different policies can generate the same trigger. For example, both Mandatory.link.sh and Mandatory.cdp.sh can generate a CISCO_DMP_EVENT trigger, but pass the variable LINKUP with a different value into the shell function. The event policy has the logic to handle various situations, such as the case where the new trigger is the same as the previous trigger that was last used to configure the interface. The event policy also checks to see if the interface is configured with a sticky macro. These are not removed when the link is down. As discussed previously, this could result in an err_disabled state if a different device is attached to a sticky interface with port security. Sticky configurations should not be used if the intent is to dynamically configure the interface based on device discovery when devices move from port to port.

The relationship between the various components is shown in [Figure 7-2](#). The example flow shows the result of a CDP event.

Figure 7-2 Auto Smartports Event Flow

229924

Medianet Devices with Built-in ASP Macros

The following devices are currently supported by built-in ASP macros.

Cisco IPVS Cameras

Cisco IPVS cameras support CDP as the detection mechanism for executing the built-in ip-camera ASP macro. There are slight differences in the built-in ip-camera macro applied depending upon the platform (Catalyst access switch or Catalyst 4500) and upon whether the platform supports Enhanced ASP macros or regular ASP macros. The example in [Table 7-7](#) shows the switchport configuration applied after a link-up event for a Catalyst access switch, both with regular ASP Macros and Enhanced ASP Macros. The configuration assumes the initial switchport configuration was a default configuration (meaning no configuration on the interface).

Table 7-7 Configuration Example 1—Switchport Configuration Resulting from the Built-in IP-Camera Macro

Regular ASP Macro	Enhanced ASP Macro
<pre>! interface GigabitEthernet1/0/40 switchport access vlan 302¹ switchport mode access switchport block unicast switchport port-security mls qos trust dscp macro description CISCO_IPVSC_EVENT spanning-tree portfast spanning-tree bpduguard enable !</pre>	<pre>! interface GigabitEthernet1/0/40 switchport access vlan 302 switchport mode access switchport block unicast switchport port-security srr-queue bandwidth share 1 30 35 5 queue-set 2 priority-queue out mls qos trust device ip-camera mls qos trust dscp macro description CISCO_IPVSC_EVENT auto qos video ip-camera spanning-tree portfast spanning-tree bpduguard enable !</pre>

1. Access VLAN set by **macro auto device ip-camera ACCESS_VLAN=VLAN302** global configuration command.

Brief explanations of the commands are shown in [Table 7-8](#).

Table 7-8 Summary of ASP Commands

Command	Description
switchport access vlan 302	Configures the switchport as a static access port using the access VLAN specified through the following manually configured global command: macro auto device ip-camera ACCESS_VLAN=302
switchport mode access	The port is set to access unconditionally and operates as a nontrunking, single VLAN interface that sends and receives nonencapsulated (non-tagged) frames.
switchport block unicast	By default, all traffic with unknown MAC addresses is sent to all ports. This command blocks unicast packets with unknown MAC addresses received by this port from being sent to other ports on the switch. This feature is designed to address the cam table overflow vulnerability, in which the cam table overflows and packets are sent out all ports.
switchport port-security	Enables port security on the interface. Defaults to one secure MAC address. Defaults to set the port in error-disable state upon a security violation. SNMP Trap and Syslog message are also sent.
auto qos video ip-camera	Automatically configures QoS on the port to support a Cisco IPVS camera. Causes the following interface level commands to be added: <pre>srr-queue bandwidth share 1 30 35 5 queue-set 2 priority-queue out mls qos trust device ip-camera mls qos trust dscp</pre> Causes global configuration changes to the switch configuration to occur as well.

Table 7-8 **Summary of ASP Commands**

srr-queue bandwidth share 1 30 35 5	Sets ratio by which the shaped round robin (SRR) scheduler services each of the four egress queues (Q1 through Q4 respectively) of the interface. Bandwidth is shared, meaning that if sufficient bandwidth exists, each queue can exceed its allocated ratio. Note that the priority-queue out command overrides the bandwidth ratio for Q1.
queue-set 2	Maps the port to the 2nd queue set within the switch. Catalyst 3560, 3750, and 2960 Series switches support two queue sets.
priority-queue out	Enables egress priority queuing. Automatically nullifies the srr-queue bandwidth share ratio for queue 1 since the priority queue is always serviced first (unlimited bandwidth).
mls qos trust device ip-camera	Enables the QoS trust boundary if CDP packets are detected indicating the connection of a IP surveillance camera to the interface.
mls qos trust dscp	Classify an ingress packet by using the packet's DSCP value.
macro description CISCO_IPVSC_EVENT	Description indicating which built-in macro has been applied to the interface, in this case the built in ip-camera macro.
spanning-tree portfast	When the Port Fast feature is enabled, the interface changes directly from a blocking state to a forwarding state without making the intermediate spanning-tree state changes.
spanning-tree bpduguard enable	Puts the interface in the error-disabled state when it receives a bridge protocol data unit (BPDU). This should not occur on a port configured for access mode.

The main difference between the Enhanced ASP macro and the regular ASP macro is that the Enhanced ASP macro includes the **auto qos video ip-camera** interface-level command. AutoQoS has been extended in IOS version 12.2(55)SE on the Catalyst access switches to support video devices as well as VoIP. Among other things, the **auto qos video ip-camera** command causes DSCP markings from the device to be trusted when the switchport detects CDP from the attached Cisco IPVS camera. On Catalyst access switches, the **auto qos video ip-camera** command also causes changes to the queue-sets, which globally affect the switch configuration. These global changes—which result from the AutoQoS commands within ASP macros—are not reversed when the anti-macro is run, returning the interface to its default configuration. Instead the global configuration changes remain within the running configuration of the switch. The network administrator may need to manually access the switch in order to save these changes in the running configuration into the startup configuration. Note also that minor disruptions to switch processing may occur the first time the queue-sets are modified. However, this occurs only when the first switchport configured for Enhanced ASP macros detects an IPVS camera. Subsequent switchports which detect an IPVS camera do not cause further changes to the queue-sets, since they have already been modified. For further discussion of the effects of AutoQoS video, see the Medianet Campus QoS Design 4.0 document at:
http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoS_Campus_40.html

**Note**

Cisco recommends a DSCP setting of CS5 for IPVS cameras. However this is not currently the default value which ships in the firmware. The network administrator may have to manually change the DSCP value to CS5 within the IPVS cameras.

Cisco Digital Media Players (DMPs)

Cisco 4310G DMPs running revision 5.2.2 are the only DMPs which support CDP as the detection mechanism for executing the built-in media-player ASP macro. The CDP detection mechanism for DMPs only works for Enhanced ASP macros as well. However, the MAC address detection mechanism automatically works for Cisco 4305G, 4400G, and 4310G DMPs for both Enhanced ASP macros and regular ASP macros. Catalyst switches which support ASP macros have a built-in MAC address trigger which matches on the OUI values of 00-0F-44 or 00-23-AC, corresponding to Cisco DMPs.

The built-in media-player ASP macro is the same regardless of whether the platform supports Enhanced ASP macros or regular ASP macros. The example in [Table 7-9](#) shows the switchport configuration applied after a link-up event for a Catalyst access switch. The configuration assumes the initial switchport configuration was a default configuration (meaning no configuration on the interface).

Table 7-9 Configuration Example 2—Switchport Configuration Resulting from the Built-in Media-Player Macro

Regular and/or Enhanced ASP Macro
<pre>! interface GigabitEthernet2/0/8 switchport access vlan 282¹ switchport mode access switchport block unicast switchport port-security priority-queue out mls qos trust dscp macro description CISCO_DMP_EVENT spanning-tree portfast spanning-tree bpduguard enable !</pre>

1. Access VLAN set by **macro auto device media-player ACCESS_VLAN=VLAN282** global configuration command.

Brief explanations of the commands are shown in [Table 7-10](#).

Table 7-10 Summary of ASP Commands

Command	Description
<code>switchport access vlan 282</code>	Configures the switchport as a static access port using the access VLAN specified through the following manually configured global command: macro auto device media-player ACCESS_VLAN=282
<code>switchport mode access</code>	The port is set to access unconditionally and operates as a nontrunking, single VLAN interface that sends and receives nonencapsulated (non-tagged) frames.
<code>switchport block unicast</code>	By default, all traffic with unknown MAC addresses is sent to all ports. This command blocks unicast packets with unknown MAC addresses received by this port from being sent to other ports on the switch. This feature is designed to address the cam table overflow vulnerability, in which the cam table overflows and packets are sent out all ports.

Table 7-10 Summary of ASP Commands

<code>switchport port-security</code>	Enables port security on the interface. Defaults to one secure MAC address. Defaults to set the port in error-disable state upon a security violation. SNMP Trap and Syslog message are also sent.
<code>priority-queue out</code>	Enables egress priority queuing. Automatically nullifies the srr-queue bandwidth share ratio for queue 1, since the priority queue is always serviced first (unlimited bandwidth).
<code>mls qos trust dscp</code>	Classify an ingress packet by using the packet's DSCP value.
<code>macro description CISCO_DMP_EVENT</code>	Description indicating which built-in macro has been applied to the interface, in this case the built-in media-player macro.
<code>spanning-tree portfast</code>	When the Port Fast feature is enabled, the interface changes directly from a blocking state to a forwarding state without making the intermediate spanning-tree state changes.
<code>spanning-tree bpduguard enable</code>	Puts the interface in the error-disabled state when it receives a bridge protocol data unit (BPDU). This should not occur on a port configured for access mode.

The network administrator should note that MAC-address triggers are executed only after a timeout of either CDP or LLDP triggers. The timeout value is roughly 65 seconds. In other words, when deploying DMPs which do not support CDP, or deploying DMPs on Catalyst switch platforms which do not support Enhanced ASP macros, the Catalyst switch listens for CDP or LLDP triggers for approximately one minute. After the timeout, the switch executes the built-in MAC-address trigger corresponding to the DMP.

It is also important for the network administrator to understand the order in which certain services start when devices such as DMPs boot up. When using dynamic IP addressing, CDP should be sent before any DHCP packets are sent. This is because the access VLAN is often passed into the ASP macro. A device which acquires an IP address before the ASP macro has run will acquire an IP address corresponding to the default VLAN (VLAN 1). When the ASP macro subsequently runs, the device is moved onto a different access VLAN. Therefore, the device will need to release the existing IP address and acquire a new IP address. Typically this is done when the device sees the line-protocol transitioned when the VLAN is changed on the switch port. The built-in macros do not transition the link upon VLAN reassignment. Failure to release and renew the IP address results in an unreachable device, since its IP address corresponds to the wrong VLAN. This issue also exists when using the built-in MAC address trigger to execute the built-in media-player ASP macro for DMPs.

Medianet Devices without Built-in ASP Macros

The following devices are not currently supported by built-in ASP macros.

Cisco TelePresence (CTS) Endpoints

Currently there are no built-in ASP macros for Cisco TelePresence (CTS) endpoints within the Catalyst switch software. CTS endpoints consist of one or more codecs and an associated IP phone. As of CTS software version 1.6(5), both the codec and the phone send CDP packets to the Catalyst switch with the phone bit enabled within the capabilities field of CDP packets. Catalyst switchports currently apply the built-in phone ASP macro for attached CTS endpoints, based on the CDP trigger from the combination

of the IP phone and codec, assuming the phone macro is enabled globally on the Catalyst switch. For customers who have both Cisco IP phones and CTS endpoints attached to the same Catalyst switch and who wish to use ASP macros, this is a likely scenario.

The application of the built-in phone ASP macro does not cause CTS endpoints to stop working, provided the network administrator has deployed the TelePresence endpoint to share the voice VLAN with IP phones. However, the configuration is not optimal or recommended for CTS endpoints. The application of the built-in phone ASP macro includes the interface-level **auto qos voip cisco-phone** command. This applies AutoQoS VoIP to both the global configuration of the Catalyst switch as well as the interface. The current AutoQoS VoIP configuration only identifies, marks, and polices EF and CS3 traffic from an IP phone accordingly. Since TelePresence is recommended to be configured to send traffic with a CS4 DSCP marking, the AutoQoS VoIP configuration does not address TelePresence traffic at all. However, the traffic from the TelePresence codec is still trusted at the ingress port. Therefore the TelePresence traffic still crosses the network with a CS4 marking.

A recommended work-around for this situation is to disable ASP macros via the **no macro auto processing** interface-level command for Catalyst switchports which support Cisco TelePresence endpoints. Either manually configure the switchports or use Static Smartports with the recommended configuration to support a CTS endpoint.


Note

As of IOS version 12.2(55)SE, Catalyst access switches support AutoQoS for CTS endpoints with the **auto qos video cts** command. For more information, see the *Medianet Campus QoS Design 4.0* guide: http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoS_Campus_40.html.

Other Video Conferencing Equipment

Cisco desktop video conferencing software which consists of a PC daisy chained off of a Cisco IP phone will exhibit similar characteristics as Cisco CTS endpoints when implementing ASP macros. The attached Cisco IP phone will result in the built-in phone ASP macro being executed. The resulting configuration may not be optimal for desktop video conferencing.

No built-in ASP macros currently exist for Cisco Tandberg video conferencing equipment at the time this document was written. Therefore, it is recommended to either manually configure the switchports or use Static Smartports with the recommended configuration to support Cisco Tandberg video conferencing equipment.

Overriding Built-in Macros

Generally the built-in ASP macros support the requirements of most customers while easing the deployment of medianet devices onto the network infrastructure. However, sometimes there may be reasons why a network administrator may wish to change the functionality of a built-in ASP macro. For example, with regular ASP macros, there is no ability to disable an individual built-in macro, such as the switch or router macros. Since these macros automatically configure the port as a trunk allowing all VLANs, there may be potential security issues with allowing them to run. The network administrator may desire to override the existing macro in such a manner that it is effectively disabled.

Alternatively, the network administrator may wish to only slightly modify the function of an existing built-in ASP macro. For example, as previously mentioned, the deployment of sticky macros in a dynamic environment causes Auto Smartports to be less effective due to the **switchport port-security**

interface-level command within both the built-in ip-camera and media-player macros. An overridden macro may be configured in order to modify or remove port-security for these devices if the network administrator desires to use sticky macros.

Built-in macros can be overridden by creating a new macro with the same name as an existing built-in macro. These overridden macros can be located in one of three places:

- Embedded within the switch configuration
- A standalone macro file within the switch flash
- A standalone macro file accessed remotely by the switch

The partial configuration example in [Table 7-11](#) shows an overridden switch macro embedded within the configuration of a Catalyst access switch.

Table 7-11 Configuration Example 3—Overridden ASP Macro within the Switch Configuration

```

!
macro auto execute CISCO_SWITCH_EVENT {
    if [[ $LINKUP -eq YES ]]; then
        conf t
            interface $INTERFACE
                macro description $TRIGGER
                description ROGUE SWITCH
DETECTED - PORT ENABLED
                switchport mode access
                shutdown
            exit
        end
    else
        conf t
            interface $INTERFACE
                no macro description
                description ROGUE SWITCH
DETECTED - PORT DISABLED
                no switchport mode access
            exit
        end
    fi
}
!

```

The overridden switch macro example above simply causes the interface to be put into a shutdown state when the switchport detects the presence of another switch via the CDP triggering mechanism.

The benefit of embedding an overridden macro directly within the switch configuration is the ability to view the macro directly from the configuration. The downside is that the network administrator may need to duplicate the same overridden macro on every switch which requires it. This can be both time consuming and error prone in large deployments, limiting the overall ability to scale Auto Smartports deployments.

The second method of overriding a built-in macro is to put the overridden macro in a file located within the flash memory of the switch. In order to override a built-in macro from a flash file, the network administrator needs to include the **macro auto execute <trigger name> remote <remote file location>** command within the global configuration of the switch. The example in [Table 7-12](#) shows the command line added to a Catalyst access switch to override the built-in media-player ASP macro and the file contents of the overridden macro itself.

Table 7-12 Configuration Example 4—Overridden Macro within a Flash File on the Switch**Global Configuration Command**

```
!
macro auto execute CISCO_DMP_EVENT remote flash:DMP_macro.txt
!
```

Contents of the Flash File Overriding the Built-in Macro

```
me-w-austin-3#more DMP_macro.txt
  if [[ $LINKUP -eq YES ]]; then
    conf t
      interface $INTERFACE
        macro description $TRIGGER
        switchport access vlan $ACCESS_VLAN
        switchport mode access
        switchport block unicast
        mls qos trust dscp
        spanning-tree portfast
        spanning-tree bpduguard enable
        priority-queue out
      exit
    end
  fi
  if [[ $LINKUP -eq NO ]]; then
    conf t
      interface $INTERFACE
        no macro description
        no switchport access vlan $ACCESS_VLAN
        no switchport block unicast
        no mls qos trust dscp
        no spanning-tree portfast
        no spanning-tree bpduguard enable
        no priority-queue out
        if [[ $AUTH_ENABLED -eq NO ]]; then
          no switchport mode access
        fi
      exit
    end
  fi
```

The benefit of this method is that a single overridden macro file can be created centrally—perhaps on a management server—and copied to each switch which needs to override the built-in ASP macro. This can help reduce the administrative burden and potential for errors, increasing the scalability of Auto Smartports deployments.

The downside is that there is no method to validate that the overridden macro actually functions correctly when it is typed in a separate text file and subsequently downloaded to the Catalyst switch. It is recommended that the network administrator test any overridden macros—perhaps using a non-production lab or backup switch—before deploying them in order to avoid such errors. Errors in overridden macros will cause macro processing to immediately exit. This can result in un-deterministic results in the configuration of an interface, based upon where in the macro the error occurred. The network administrator should also note that currently no error or warning will be generated to the switch console or syslog when the macro exits due to an error.

A second downside is that there is no method to validate that the overridden macro is correct for the particular model of switch to which it is being downloaded. There are slight command differences between Catalyst access switches and Catalyst 4500 Series switches which could cause a macro written for the wrong switch model to execute incorrectly. This can again result in un-deterministic results in

the configuration of an interface, based upon where the command differences occurred. In order to avoid this potential issue, the network administrator may choose to include the Catalyst switch model within the file name of the overridden macro. This gives the network administrator a quick visual indication if the file being downloaded is correct for the switch model.

The third method of overriding a built-in ASP macro is to put the overridden macro in a file on a remote server and configure the switch to access the file when it needs to run the macro. In order to override a built-in macro from a file on a remote server, the network administrator needs to include the **macro auto execute** <trigger name> **remote** <remote file location> command again within the global configuration of the switch. However, this time the remote file location includes the protocol, network address or hostname, userid and password, and path to the file on the remote server. The example in [Table 7-13](#) shows the command line added to a Catalyst access switch to override the built-in media-player ASP macro. The contents of the overridden macro itself are the same as that shown in [Table 7-12](#).

Table 7-13 **Configuration Example 5—Example Configuration for Overriding a Macro via a File on a Remote Server**

Global Configuration

```
!
macro auto execute CISCO_DMP_EVENT remote ftp://admin:cisco@10.16.133.2/DMP_macro.txt
!
```

The switch is capable of using the following protocols for download of remote macros: FTP, HTTP, HTTPS, RCP, SCP, and TFTP. In production networks, it is recommended to implement a secure protocol, such as SCP or HTTPS.

The benefit to this approach is that the overridden ASP macro files can again be managed centrally on a management server. This further eases the administrative burden of not having to manually copy the macro file to each switch which requires it. This is particularly useful when changing the behavior of an overridden macro that is already deployed on switches throughout the network infrastructure.

The network administrator should note that the overridden ASP macro file is downloaded to the Catalyst switch every time a link-up or link down event occurs. The switch does not cache the macro file; it simply requests the file every time there is a link-up or link-down event on the port. Testing did not investigate any potential scalability implications for processing on the switch, since multiple ports on the same switch may simultaneously request a file for download in order to apply the macro, particularly in scenarios where the switch has just been reloaded. This method also has potential scalability implications for the remote server, since it may have to process multiple simultaneous downloads from multiple ports on a single switch and from multiple switches throughout the network.

A downside to this method is that if the remote server is unavailable, the overridden ASP macro will not be run and the device will end up with a default configuration on the Catalyst switch. In many cases, the device will not function since it may be on the wrong VLAN. If the interface is already up and configured via the overridden ASP macro when the medianet device is removed, the configuration will remain on the Catalyst switchport if the remote server is unavailable. This is because the anti-macro will not be run to clean-up the switchport configuration. If another device is subsequently connected to the switchport, the resulting switchport configuration could be somewhat un-deterministic. This situation should be avoided. The remote server could be unavailable either due to a network error or a server error. Therefore, it is recommended that the network administrator implement both network-level redundancy as well as server-level redundancy in order to ensure the availability of the remote ASP macro when utilizing this method. Obviously the built-in router Auto Smartport macro should not be used to configure the interface that would route to the FTP server. Extra care will also be needed if the account password is to be changed on a reoccurring basis due to a security policy.

Finally, as with the previous method, there is no mechanism to validate the overridden ASP macro has no errors or is the correct macro for the model of switch to which it will be automatically downloaded. It is again recommended that the network administrator test any overridden macros—perhaps using a non-production lab or backup switch—before making them available for automatic download in order to avoid such errors.

**Note**

CiscoWorks LMS is targeted to add support for managing Auto Smartports macros in an upcoming release. Future updates to this document may include details about LMS as it relates to ASP macros.

Macro-of-Last-Resort

As highlighted in [Table 7-2](#), Enhanced ASP macros support a feature known as the macro-of-last-resort (also referred to as the LAST_RESORT macro). The macro-of-last-resort is a built-in ASP macro which is run if no other trigger event is seen and therefore no other ASP macro (built-in or user-defined) is run. Without the use of the macro-of-last-resort, devices such as end-user PCs—which typically will not trigger any ASP macros—may end up with a default switchport configuration, depending on whether the custom macro has been overridden. This may not be the desired switchport configuration, particularly if the network administrator uses a VLAN other than VLAN1 for the normal data VLAN. The custom macro is discussed in [Custom Macro](#).

**Note**

The use of a VLAN other than the default (VLAN 1) for the data VLAN is consistent with Cisco SAFE security guidelines.

The built-in macro-of-last-resort is enabled on Catalyst switches which support Enhanced ASP macros via the following global configuration command:

```
macro auto global control trigger last-resort
```

The macro-of-last-resort can also be enabled per interface with the following interface-level command:

```
macro auto control trigger last-resort
```

The only parameter which can be passed into the built-in ASP macro-of-last-resort is the access VLAN. This can be passed using the following global switch configuration command:

```
macro auto execute CISCO_LAST_RESORT_EVENT built-in CISCO_LAST_RESORT_SMARTPORT  
ACCESS_VLAN=<vlan>
```

For example, in order to set the access VLAN to VLAN100 for the macro-of-last-resort, the network administrator would configure the following global switch command:

```
macro auto execute CISCO_LAST_RESORT_EVENT built-in CISCO_LAST_RESORT_SMARTPORT  
ACCESS_VLAN=100
```

The example in [Table 7-14](#) shows the switchport macro-of-last-resort configuration applied after a link-up event for a Catalyst access switch. The configuration assumes the initial switchport configuration was a default configuration (meaning no configuration on the interface).

Table 7-14 Configuration Example 6—Switchport Configuration Resulting from the Built-in Macro-of-Last-Resort

```

!
interface GigabitEthernet1/0/7
 switchport access vlan 1001
 switchport mode access
 load-interval 60
 macro description CISCO_LAST_RESORT_EVENT
 spanning-tree portfast
 spanning-tree bpduguard enable
!

```

1. Access VLAN set by **macro auto execute CISCO_LAST_RESORT_EVENT built-in CISCO_LAST_RESORT_SMARTPORT ACCESS_VLAN=100** global configuration command.

Brief explanations of the commands are shown in [Table 7-15](#).

Table 7-15 Summary of ASP Commands

Command	Description
switchport access vlan 100	Configures the switchport as a static access port using the access VLAN specified through the following manually configured global command: macro auto execute CISCO_LAST_RESORT_EVENT built-in CISCO_LAST_RESORT_SMARTPORT ACCESS_VLAN=100
switchport mode access	The port is set to access unconditionally and operates as a nontrunking, single VLAN interface that sends and receives nonencapsulated (non-tagged) frames.
load-interval 60	Sets the interval over which interface statistics are collected to average over 60 seconds.
macro description CISCO_LAST_RESORT_EVENT	Description indicating which built-in macro has been applied to the interface, in this case the built in last-resort macro.
spanning-tree portfast	When the Port Fast feature is enabled, the interface changes directly from a blocking state to a forwarding state without making the intermediate spanning-tree state changes.
spanning-tree bpduguard enable	Puts the interface in the error-disabled state when it receives a bridge protocol data unit (BPDU). This should not occur on a port configured for access mode.

When the device is removed from the interface, the anti-macro will return the switchport to a default interface configuration. The macro-of-last-resort can also be overridden. This allows the network administrator to implement a completely custom default switchport configuration for devices which do not match any built-in or user-defined ASP macros.

Since the macro-of-last resort executes if no other triggering events are seen, including MAC-address trigger events, there could be a delay of over one minute between the time the switchport interface becomes active and the execution of the macro-of-last-resort. During this time period, the device will be active on the default VLAN (VLAN 1)—unless it was left on a different VLAN by the custom macro.

An end-user PC which uses DHCP could obtain an IP address before the switch moves the VLAN configuration to that specified by the macro-of-last-resort if the default VLAN contains a DHCP server. When the switchport is subsequently moved to the new VLAN, the end-user PC should release and renew the DHCP lease based on the line protocol transition of the switch. The network administrator may wish to test the PC hardware and operating systems deployed within his/her network to ensure they function properly before deploying Enhanced Auto Smartports. An alternative is to simply not provision a DHCP server on the default VLAN. Typically most DHCP clients will try to DISCOVER a server for more than the macro-of-last-resort timeout, however it is possible the end-user PC will timeout when attempting to obtain a DHCP address. In this case, the end-user may need to manually re-activate DHCP again after the macro-of-last-resort has moved the PC to the correct VLAN in order to obtain an IP address corresponding to the correct VLAN.

**Note**

Testing with the macro-of-last resort did not include the use of 802.1x/MAB on the end-user PC. Therefore, no design guidance around the interaction of 802.1x/MAB and the macro-of-last-resort is provided in this document at this time.

Custom Macro

The custom macro is a built-in Enhanced ASP macro which is automatically executed upon an interface link down event. The following example output from the **show shell function** **CISCO_CUSTOM_AUTOSMARTPORT** exec-level command shows the built-in custom macro.

```
me-w-austin-3>show shell function CISCO_CUSTOM_AUTOSMARTPORT
function CISCO_CUSTOM_AUTOSMARTPORT () {
    if [[ $LINKUP -eq YES ]]; then
        conf t
            interface $INTERFACE
        exit
    end
fi
if [[ $LINKUP -eq NO ]]; then
    conf t
        interface $INTERFACE
    exit
end
fi
}
```

By default, the custom macro does nothing at all unless it is overridden by the network administrator. The network administrator may choose to override the custom macro to provide functionality, such as a VLAN configuration other than the default VLAN (VLAN 1) to a port when there is no device connected to it. The following two examples illustrate possible uses of the custom macro.

Example Scenario #1

The network administrator has pre-configured all unused ports to be on the data VLAN, instead of the default VLAN. If a DMP device is connected to a switchport configured for Enhanced ASP macros, it will be recognized as a DMP and moved into the VLAN specified by the network administrator through the built-media-player Enhanced ASP macro. For example, the port may be moved to a DMP VLAN. If the DMP device is then removed, the DMP anti-macro executes, removing the switchport from the DMP VLAN (which places it into the default VLAN). The custom macro will then execute, moving the switchport into the VLAN specified within the overridden custom macro. This may correspond to the data VLAN again. If a normal PC device is subsequently placed onto the same switchport, the PC will immediately come up within the data VLAN. It will remain there since it will not trigger any other

built-in Enhanced ASP macros. This scenario assumes the macro-of-last-resort has not been enabled. Therefore, in this example, the custom macro provides the network administrator an alternative method of placing devices which do not trigger any built-in Enhanced ASP macros (such as normal PCs) onto a VLAN other than the default VLAN.

The advantage of the custom macro in this scenario is that the device does not have to wait until the macro-of-last resort is executed to be moved into the correct VLAN. This may help minimize issues regarding PCs acquiring an incorrect DHCP addresses because they were moved to another VLAN by the macro-of-last-resort. However, the network administrator should be careful of medianet devices such as DMPs and IPVS cameras accidentally getting the wrong IP addresses, since they initially come up within the data VLAN as well. Finally, the network administrator may have to manually pre-configure all unused switchports be within the data VLAN initially. The custom macro will not be run until another macro has been run on the port and the device has subsequently been removed.

Example Scenario #2

The network administrator has pre-configured all unused ports to be on an unused or isolated VLAN, instead of the default VLAN. If a DMP device is connected to a switchport configured for Enhanced ASP macros, it will be recognized as a DMP and moved into the VLAN specified by the network administrator through the built-media-player Enhanced ASP macro. For example, the port may be moved to a DMP VLAN. If the DMP device is then removed, the DMP anti-macro executes, removing the switchport from the DMP VLAN (which places it into the default VLAN). The custom macro will then execute, moving the switchport into the VLAN specified within the overridden custom macro. This may correspond to the unused or isolated VLAN in this scenario. If a normal PC device is subsequently placed onto the same switchport, the PC will immediately come up within the unused or isolated VLAN. If the macro-of-last-resort has been enabled, it will trigger, moving the device into another VLAN, such as the normal data VLAN. If the PC is then removed from the switchport, its anti-macro will execute, removing switchport from the data VLAN (which places it into the default VLAN). Then the custom macro will again execute, moving the switchport back into the unused or isolated VLAN.

In this scenario, the custom macro provides the network administrator a method of placing unused ports into an unused or isolated VLAN—which is more consistent with Cisco SAFE guidelines. If the unused or isolated VLAN has no DHCP server, then devices will not accidentally get the wrong IP address before they are subsequently moved into their correct VLANs by the Enhanced ASP macros. However, PCs may have to wait longer until the macro-of-last-resort executes in order to become active on the network. Finally, the network administrator may have to manually pre-configure all unused switchports to be within the unused or isolated VLAN initially. The custom macro will not be run until another macro has been run on the port and the device has subsequently been removed.

Overridden Custom Macro

Table 7-16 shows an example of an overridden custom macro.

Table 7-16 Configuration Example 7—Overridden Custom Macro Within the Switch Configuration

```

!
macro auto execute CISCO_CUSTOM_EVENT
ACCESS_VLAN=402 {
    if [[ $LINKUP -eq YES ]]; then
        conf t
            interface $INTERFACE
        exit
    end
fi
    if [[ $LINKUP -eq NO ]]; then
        conf t
            interface $INTERFACE
            switchport access vlan
$ACCESS_VLAN
        exit
    end
fi
}
!

```

The overridden macro example simply places the switchport into VLAN 402 when it goes into a link down state. Note that the VLAN can either be hardcoded into the overridden macro or passed in via a variable declaration as shown in this example.

Security Considerations

CDP and LLDP are not considered to be secure protocols. They do not authenticate neighbors, nor make any attempt to conceal information via encryption. The only difficulty in crafting a CDP packet is that the checksum is calculated with a non-standard algorithm. Even this has been reversed engineered and published in public forums. As a result, CDP and LLDP offer an attractive vulnerability to users with mal-intent. For example, by simply sending in a CDP packet with the “S” bit (otherwise referred to as the switch bit) set in the capabilities TLV, the switch can be tricked into configuring a trunk port that will pass all VLANs and accept 802.1d BPDUs from the hacker. This could be used in man-in-the-middle (MIM) attacks on any VLAN in the switch. Below is an example of a CDP spoofing device that has set the switch bit. Notice that the platform is set to a DMP. An obvious give-away in this example is the host name, CDP Tool1, which was chosen to be obvious. Normally the host name would have been selected to make the device appear to be a legitimate DMP device.

```

me-w-austin-3#sh cdp neigh g1/0/39
Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater, P - Phone,
                  D - Remote, C - CVTA, M - Two-port Mac Relay

Device ID         Local Intrfce   Holdtme    Capability  Platform  Port ID
CDP Tool1         Gig 1/0/39      30         S           DMP 4305G  eth0

```

Because the switch policy ignores the platform, this field can be used to make the entry appear to be legitimate while still tricking the switch to configure a trunk, as shown below.

```

!
interface GigabitEthernet1/0/39
location civic-location-id 1 port-location
floor 2
room Broken_Spoke

```

```

switchport mode trunk
srr-queue bandwidth share 1 30 35 5
queue-set 2
priority-queue out
mls qos trust cos
macro description CISCO_SWITCH_EVENT
macro auto port sticky
auto qos trust
end

```

Fortunately with Enhanced ASP macros, the user is allowed to disable specific scripts. The recommendation is to only enable host type macros. Switches, routers, and access-points are rarely attached to a network in a dynamic fashion. Therefore, ASP macros corresponding to these devices should be disabled where possible.

As discussed previously, ASP allows the use of remote servers to provide macros. Secure sessions such as HTTPS should be used. If the MIM above is used in conjunction with an unsecured remote configuration, the network administrator has released full control of the device to the hacker.

Authenticating Medianet Devices

Device authentication has been an ongoing concern since the early days of wireless access. The topic is the subject of several books. A common approach is to enable 802.1x. This authentication method employs a supplicant located on the client device. If a device does not have a supplicant, as is the case with many printers, then the device can be allowed to bypass authentication based on its MAC address. This is known as MAC-Authentication-Bypass or MAB. As the name implies, this is not authentication, but a controlled way to bypass that requirement. Currently all ASP medianet devices must use MAB if device authentication is in use, since these devices do not support an 802.1x supplicant. With MAB the client's MAC address is passed to a RADIUS server. The server authenticates the devices based solely on its MAC address and can pass back policy information to the switch. Administrators should recognize that MAC addresses are not privileged information. A user can assign a locally-administered MAC address. Another key point is that MAB and ASP can happen independently of one another. A device may use MAB to get through authentication and then use CDP to trigger an ASP event. Security policy must also consider each independently. A user could hijack the MAC address from a clientless IP phone, then spoof CDP to trigger the SWITCH_EVENT macro. The risk is greatly reduced by following the recommendation to turn off ASP support for static devices such as switches and routers.

MAB with ASP can be configured as shown in the example in [Table 7-17](#).

Table 7-17 Configuration Example 7—MAB with ASP

```

!
interface GigabitEthernet1/0/7
description me-austin-cl040 (new_home)
switchport mode access
authentication event fail action authorize vlan 301
authentication event no-response action authorize vlan 301
authentication host-mode multi-domain
authentication order dot1x mab
authentication priority dot1x mab
authentication port-control auto
mab eap
end
!

```

If the built-in macros have been overridden by the user, care should be taken to ensure they do not interfere with the MAB configuration. This includes the anti-macro section that removes the applied macro.

CDP Fallback

This feature is used to provide an alternate trigger method when RADIUS does not return an event trigger. If this feature is not used, then an authenticated device that does not include a RADIUS trigger will execute the LAST_RESORT Macro if enabled. The network administrator may want to disable CDP fallback to prevent CDP spoofing tools from hijacking a MAC-Address known to be authenticated by MAB. This does not prevent the device from being authenticated, but it does prevent the device from assuming capabilities beyond those of the true MAC address. While there is an incremental security gain from this approach, there are service availability concerns if the RADIUS server does not provide a recognized trigger event. As noted previously, this has not been fully validated at the time of this writing.

Guest VLANs and LAST_RESORT Macro

With MAB enabled, the MAC address is sent to a RADIUS server for authentication. If the MAC address is unknown, MAB may direct the interface to join a Guest VLAN if the switch is configured to do so. This is independent of any action invoked via ASP. As a result, there could be inconsistencies in VLAN assignment between MAB and ASP. In this case, the MAB result takes precedence, as shown in [Table 7-18](#).

Table 7-18 *Precedence Between ASP and MAB*

ASP recognized device	MAB Authenticated	Result
NO	NO	GUEST VLAN
NO	YES	LAST RESORT VLAN
YES	NO	GUEST VLAN
YES	YES	ASP ASSIGNED VLAN

The LAST RESORT VLAN corresponds to the access VLAN configured for the macro-of-last-resort, assuming the network administrator has enabled its use. The final VLAN assignment may not be the initial VLAN that was configured on the interface when line protocol initially came up. The timing is important. If the client's DHCP stack successfully obtains an IP address prior to the final VLAN assignment, the client may become unreachable. In this case, the client should be reconfigured to use static addressing. In most situations, MAB and ASP will complete the VLAN assignment prior to DHCP completion. One area of concern arises when CDP packets are not sent by the client. In this case, a MAC-address-based ASP will wait 65 second prior to executing a trigger. The client may have completed DHCP and will not be aware that a VLAN change has occurred. If MAB was also enabled, an unknown client will be placed in the GUEST_VLAN. VLAN reassignments as a result of ASP are transparent to the client's stack. This is also the case if a VLAN is manually changed on an enabled interface. Manual VLAN changes are accompanied by shutting and no shutting the interface. ASP does not do this for the built-in system macros.

Verifying the VLAN Assignment on an Interface

The best method to determine if an ASP has executed correctly is to validate the interface configuration. The macro description can be used to determine which macro has executed. The administrator should also review the configuration settings put in place by the macro. However, when MAB and ASP are running concurrently, the configuration cannot be used to determine the state of the interface. Instead the **show interface switchport** command may be used. The following example shows that the interface has executed the LAST_RESORT macro and therefore could be in VLAN 100 or VLAN 301, depending on the authentication result.

```
!
interface GigabitEthernet1/0/39
description Overridden Macro-of-Last-Resort (Port Active)
switchport access vlan 100
switchport mode access
authentication event fail action authorize vlan 301
authentication event no-response action authorize vlan 301
authentication port-control auto
mab eap
macro description CISCO_LAST_RESORT_EVENT
end
```

The **show** command below indicates that the device was not authenticated and it currently is in VLAN 301:

```
me-w-austin-3#sh int g1/0/39 swi
Name: Gi1/0/39
Switchport: Enabled
Administrative Mode: static access
Operational Mode: static access
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: native
Negotiation of Trunking: Off
Access Mode VLAN: 301 (VLAN0301)
! <additional lines omitted >
!
```

ASP with Multiple Attached CDP Devices

In some situations, there may be two CDP devices on a single interface. A common case is seen with Cisco TelePresence. In this situation both the CTS codec and IP phone appear as CDP neighbors on a single port of the switch. There are other situations that could also arise, such as a downstream hub with multiple LLDP or CDP devices, although in a practical sense this is quite uncommon. Another case may be a CDP spoofing tool. In any case, the script will make an initial determination based on the first trigger selected. Once the macro has configured the interface with a macro description, no further configuration changes will be made. If a user incorrectly removes the macro description, the interface will be reconfigured on the next trigger event. Because only the first trigger is significant, there may be some concern as to which script will run when multiple devices are present. In the case of the CTS, the phone script will be triggered regardless of whether the codec or phone presents its CDP packet first. This is because the phone bit is set in both the CTS codec and its associated IP phone in the capabilities TLV and the script will override any host trigger with a phone trigger. Even if the codec presents a CDP packet first, the phone trigger will execute.

If a hub is attached to an ASP port, several built-in macro scripts include port security that would likely err_disable the switch interface. In the academic situation where two different classes of CDP or LLDP devices may be attached to a hub, where port security is not being used and where each different type is a known ASP class device, then the first CDP packet seen would set the port configuration. Subsequent

CDP packets from adjacent devices will not cause the interface to change configurations. Hubs are rarely seen in today's networks. Even small four port devices are typically switching. Medianet devices would not typically be found attached via a hub, therefore the LAST_RESORT macro would likely be applied to any switchport supporting Enhanced ASP.

Deployment Considerations

When deploying Auto Smartports, the network administrator does not necessarily have to enable ASP macros across the entire switch. Instead the network administrator may wish to consider initially enabling ASP macros only on a range of interfaces. This method of incremental deployment may facilitate a smoother transition from the paradigm of manual configuration to that of auto configuration. For example, if the network administrator is only beginning the transition toward a medianet by deploying digital signage and IP video surveillance cameras over a converged IP infrastructure, he/she may choose to set aside the first several ports on access switches for DMPs and/or IP cameras. ASP macro processing would only need to be enabled for these "reserved" switchports. All end-user PCs and uplinks to other switches, routers, or access points would still be done via either Static Smartports or manual configuration. This methodology works best if the medianet devices (DMPs and IPVS cameras) are placed on a separate VLAN or VLANs from the data VLAN. The macro-of-last resort can be used to simply "quarantine" the medianet device to an unused VLAN if the built-in ASP macro failed to trigger. With this method, the network administrator can still gain the administrative advantages of auto configuration for what may be hundreds or thousands of medianet specific devices, such as IP cameras and DMPs across the network infrastructure. Normal change control mechanisms can be maintained for uplink ports and infrastructure devices such as routers, switches, and access points, since they do not utilize ASP macros in the initial phased rollout of auto configuration. The network administrator can disable the unused built-in ASP macros for these devices, as well as unused detection mechanisms. As the network administrator becomes more familiar with the use of ASP macros, the deployment can then be extended to include other devices as well as infrastructure connections if desired.

Location Services

Location Services is another feature of the Medianet Service Interface (MSI) that provides the ability for the Catalyst switch to send location information to a device via CDP or LLDP-MED. Future benefits of a medianet device learning its location from the network infrastructure may be the ability to customize the configuration of the device based upon its location or the ability to automatically display content based on its learned location.

Catalyst access switches support the ability to pass either civic location information or emergency location information (ELIN) to devices via CDP or LLDP-MED in IOS revision 12.2(55)SE. Catalyst 4500 Series switches support the ability to pass either civic location information or ELIN to devices via LLDP-MED only in IOS revision 12.2(54)SG. This document will only address civic location information.

Civic location is discussed under various IETF proposed standards, including RFCs 4119 and 5139. Civic location information can be configured on a global basis (for location elements which pertain to the entire switch) and on an interface-level basis (for location elements which pertain to the specific switchport). The configuration example in [Table 7-19](#) shows an example of civic location information configured both globally and on a switchport.

Table 7-19 Configuration Example 8—Example Civic Location Configuration

```

!
location civic-location identifier 1
  building 2
  city Austin
  country US
  postal-code 33301
  primary-road-name Research_Blvd
  state Texas
  number 12515
!
!
interface GigabitEthernet1/0/39
  location civic-location-id 1 port-location
    floor 2
    room Broken_Spoke
!

```

The location of the switch—identified via the **location civic-location identifier 1** global command—corresponds to the following hypothetical address: 12515 Research_Blvd, building 2, Austin, Texas, US, 33301. The location of the switchport extends the civic location via the **location civic-location identifier 1 port-location** interface-level command to identify the device as being in the Broken_Spoke room on floor two. The use of civic location in this manner does require the network administrator to manually keep accurate records as to which switchports are wired to which rooms within the facility.

**Note**

There are limitations regarding the total size of the location information which can be sent via CDP and LLDP. The network administrator should keep the location information size under 255 bytes.

The network administrator can enable or disable the sending of location information via CDP on all ports for the entire switch with the **cdp tlv location** or **no cdp tlv location** global commands. For individual switchports, the network administrator can enable or disable the sending of location information via CDP with the **cdp tlv location** or **no cdp tlv location** interface-level commands. The network administrator can enable or disable the sending of location information via LLDP-MED for individual switchports with the **lldp-med-tlv-select location** or **no lldp-med-tlv-select location** interface-level commands.

Currently the only medianet specific device which supports Location Services is the Cisco 4310G DMP running revision 5.2.2 software. [Figure 7-3](#) shows the configuration of a Cisco 4310G DMP with location information which has been passed to the DMP via CDP from an attached Catalyst 2960-S switch.

Figure 7-3 GUI Interface of a Cisco 4310G DMP Showing Location Passed via CDP

The screenshot displays the 'DIGITAL MEDIA PLAYER' configuration interface. The left sidebar contains a navigation menu with categories: Settings (Show IP), Startup, Display Attributes, Network, Flash, Video, Remote Mappings, NTP, Internal Storage, Syslog, Display Actions (Flash Playback, Media Playback, Serial Interface), Administration (DMP Service Account, DMP Management, Manage WAAS Share, Restore Default Settings, Upgrade Firmware, Save and Restart DMP), and About (Hardware and Firmware Versions). The main configuration area is divided into three sections: 'Wired Network Configuration', 'HTTP Proxy', and 'Medianet Services'. The 'Wired Network Configuration' section includes fields for DMP MAC Address (00:0f:44:01:64:ee), Dynamic IP Addressing (DHCP) (Enabled), IP Address (10.28.2.99), Subnet Mask (255.255.255.0), Default Gateway (10.28.2.1), Primary DNS Server (10.17.1.9), Using NAT (No), NAT IP Address (0.0.0.0), and Verify Link (Enabled). The 'HTTP Proxy' section includes Use HTTP Proxy (Disabled), Proxy Server IP Address, Port, and No Proxy List. The 'Medianet Services' section includes MediaNet Enabled (On), Timeout (ms) (30000), Switch IP Address (172.26.135.162), Switch Name (me-v-austin-3.me.com), Switch Port (GigabitEthernet1/0/12), VLAN (282), Location ID, and Location URL (34=Research_Bldg28=Broken_Spoke27=2825=2824=33301819=1251583=Austin81=Texas). The Location URL field is highlighted with a red rectangle. An 'Apply' button is at the bottom right. A vertical timestamp '22/9/25' is on the far right.

**Note**

CiscoWorks LMS is targeted to add support for Location Services in an upcoming release. Future updates to this document may include details around LMS as it relates to Location Services.

Summary

Auto configuration can help facilitate the transition of the network infrastructure towards a medianet by easing the administrative burden of having to manually configure multiple switchports for devices such as digital media players (DMPs) and IP video surveillance (IPVS) cameras. The Auto Smartports (ASP) feature allows the network infrastructure to automatically detect a medianet device attached to a Cisco Catalyst switch via the Cisco Medianet Service Interface (MSI) and configure the switchport to support that particular device. Additionally, Location Services allow the switchport to send civic location information to the medianet device. Such location information may be used in the future for functionality such as customizing the configuration of the device based upon its location or automatically displaying content based upon the learned location of the medianet device.

References

- Medianet Campus QoS Design 4.0:
http://www.cisco.com/en/US/docs/solutions/Enterprise/WAN_and_MAN/QoS_SRND_40/QoS_Campus_40.html
- Auto Smartports Configuration Guide, Release 12.2(55)SE:
http://www.cisco.com/en/US/docs/switches/lan/auto_smartports/12.2_55_se/configuration/guide/aspcg.html
- Configuring LLDP, LLDP-MED, and Wired Location Service:
http://www.cisco.com/en/US/docs/switches/lan/catalyst3750x_3560x/software/release/12.2_55_se/configuration/guide/swlldp.html

