



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

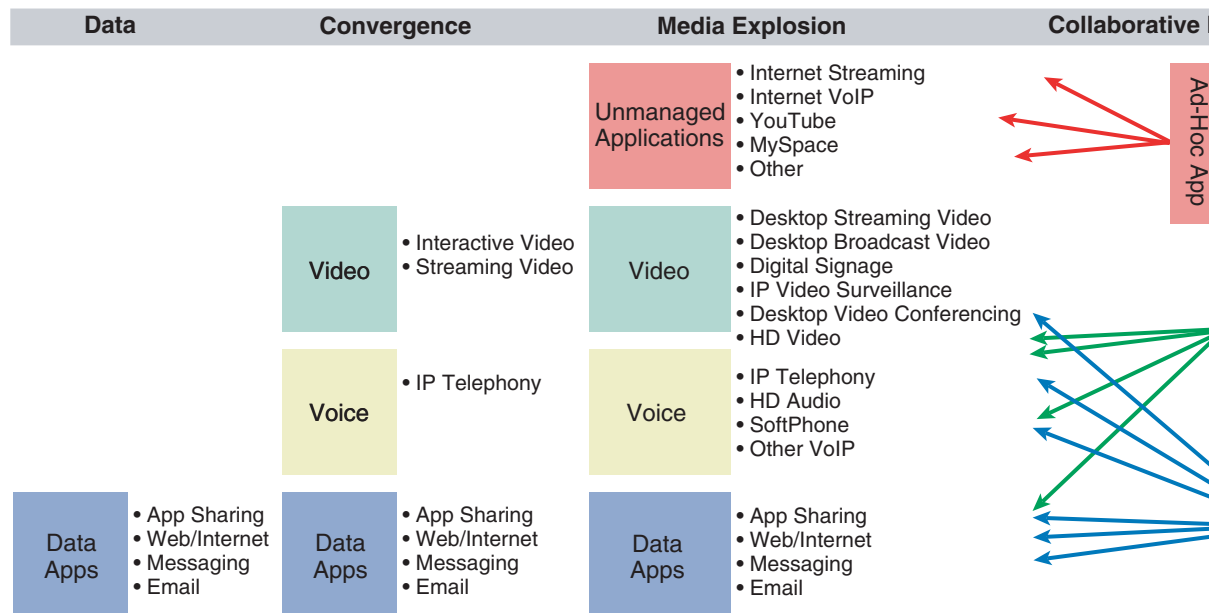
With the first stage of IP convergence, the Cisco Architecture for Voice, Video, and Integrated Data (AVVID) provided the foundation for different applications to effectively and transparently share the same IP network. One of the challenges to overcome with converged networks is to be able to simultaneously 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?

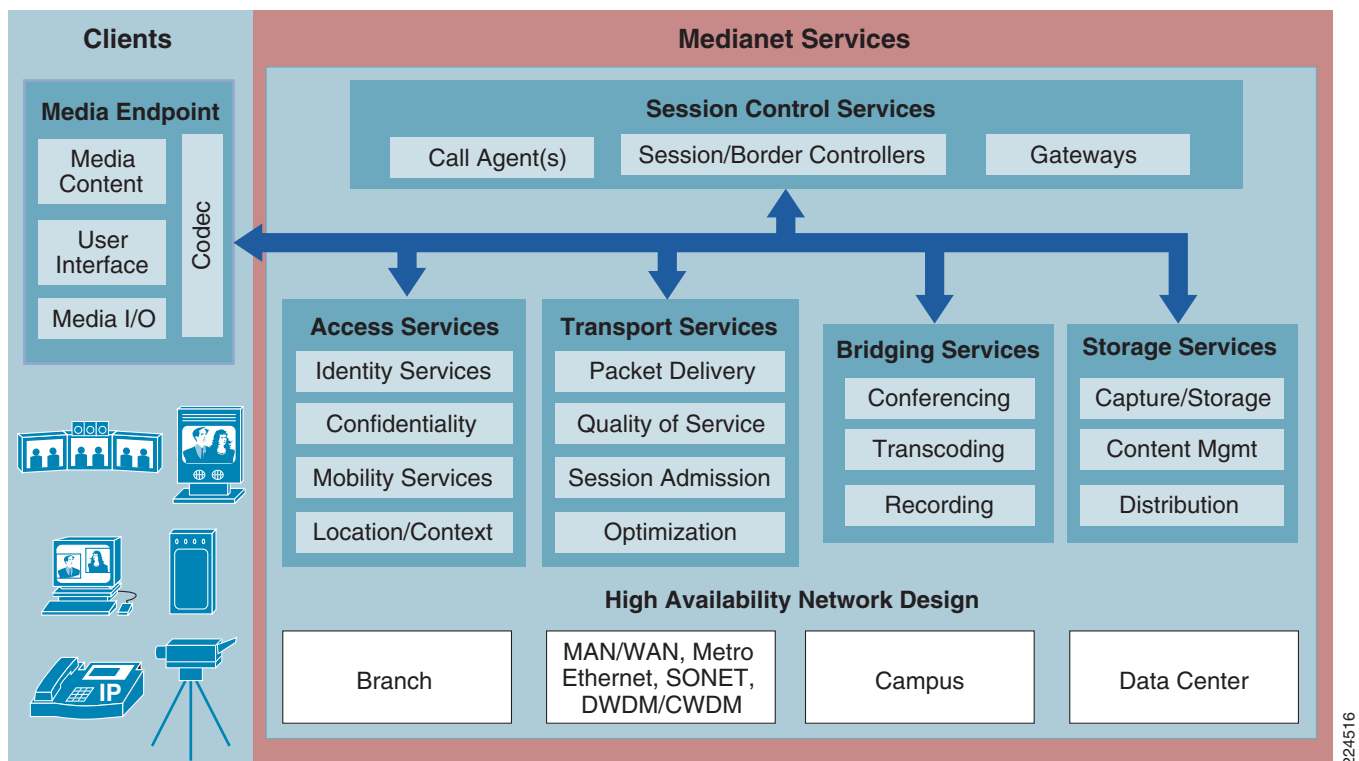
- What will the usage policies be and how will they be enforced?

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.

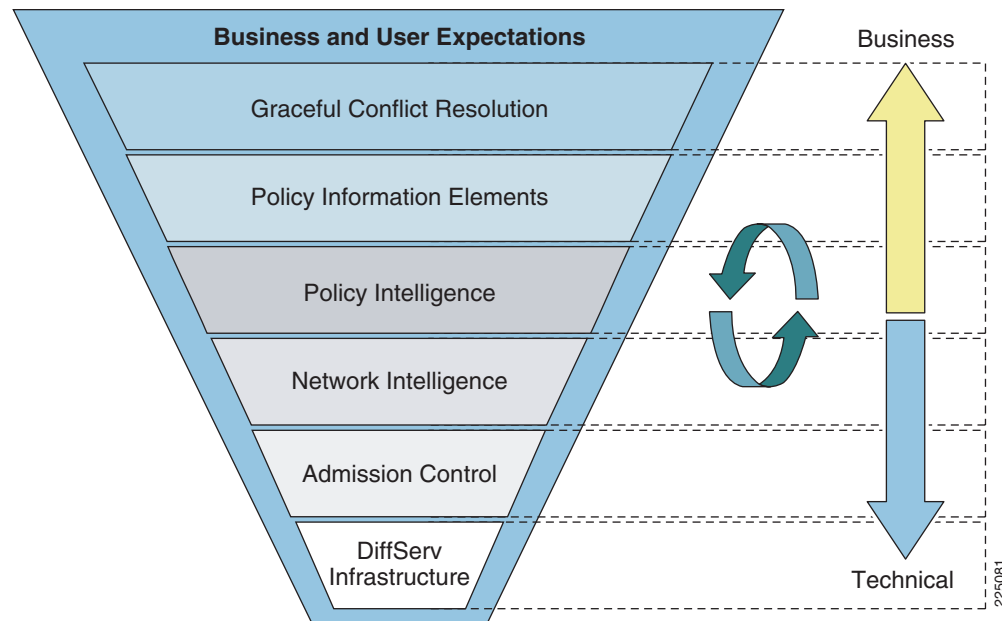
The fourth level of admission control contains policy elements and weights to determine what exactly constitutes a “higher-priority” flow, as defined by the administrative preferences of an organization. Such policy information elements may include-but are not limited to-the following:

- 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.

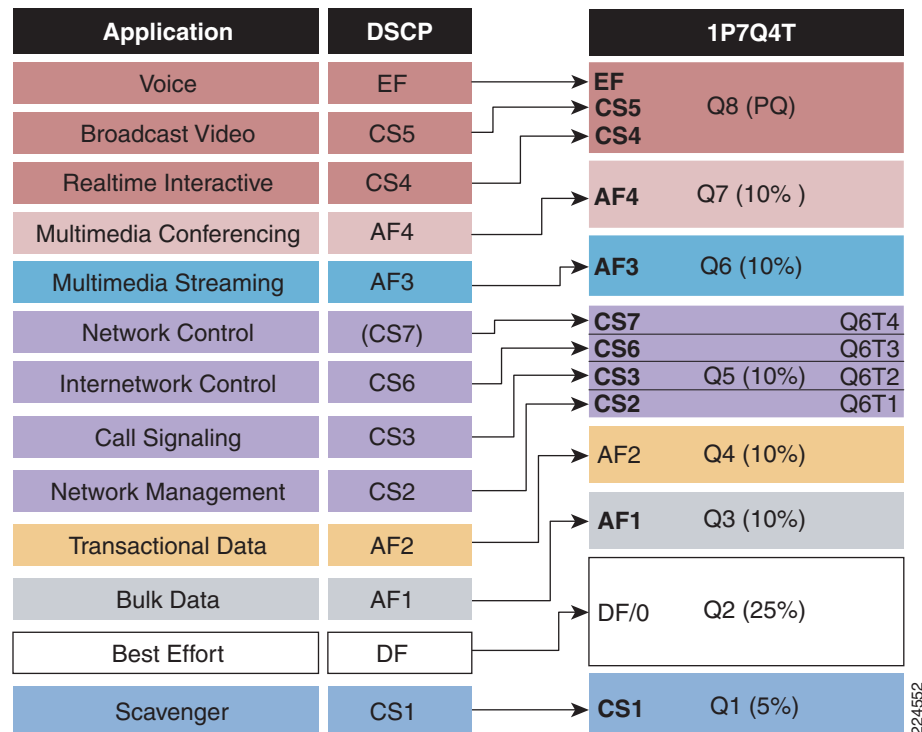
Once traffic has been classified and marked, then queuing policies must be implemented on every node where the possibility of congestion could occur (regardless of how often congestion scenarios actually do occur). This is an absolute requirement to guarantee service levels. In the campus, queuing typically 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

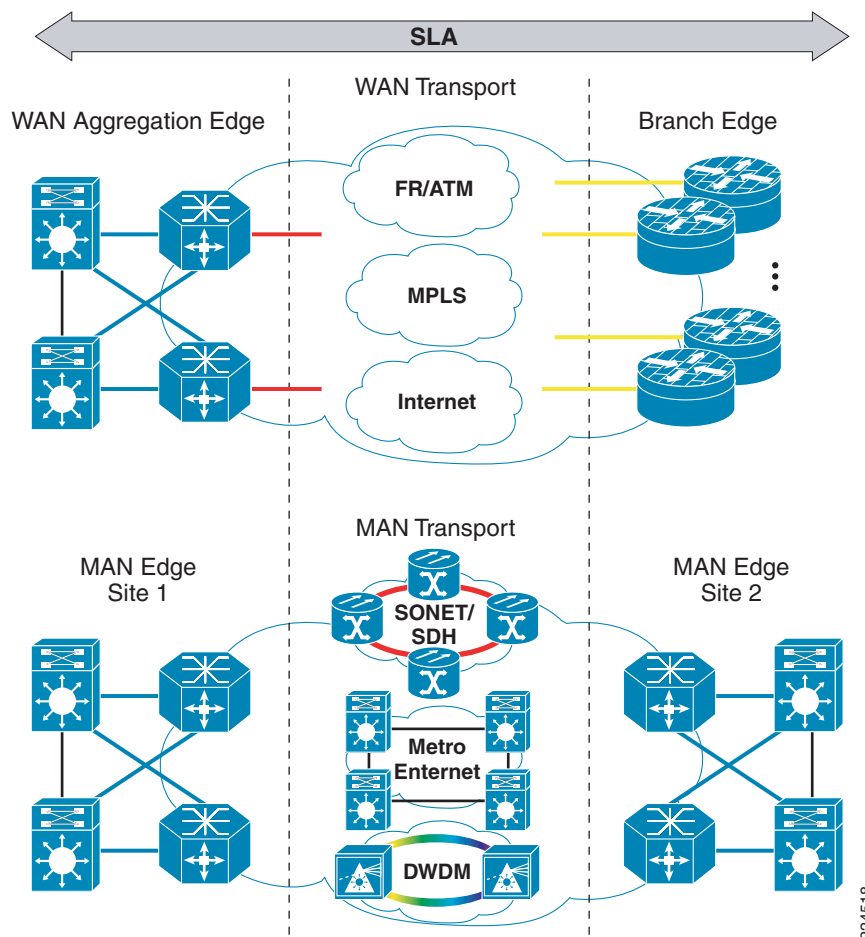



Figure 1-13 Branch Medianet Architecture

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.

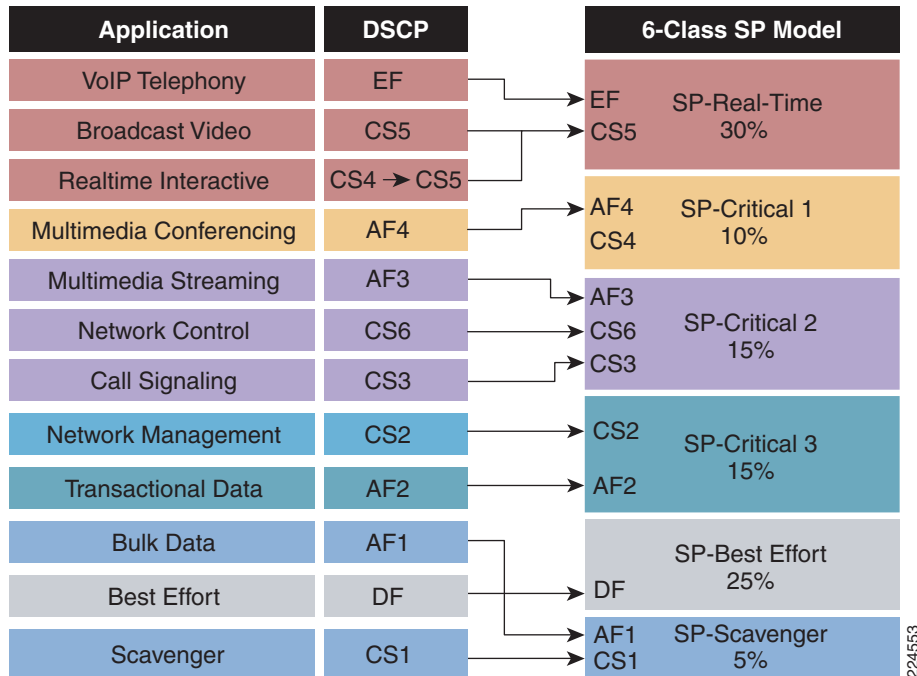
Cisco ISR and ASR product families offer industry-leading QoS implementations, accelerated with low-latency hardware ASICs, that are critical for ensuring the service level for video applications. QoS continues to evolve to include more granular queuing, as well as additional packet identification and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

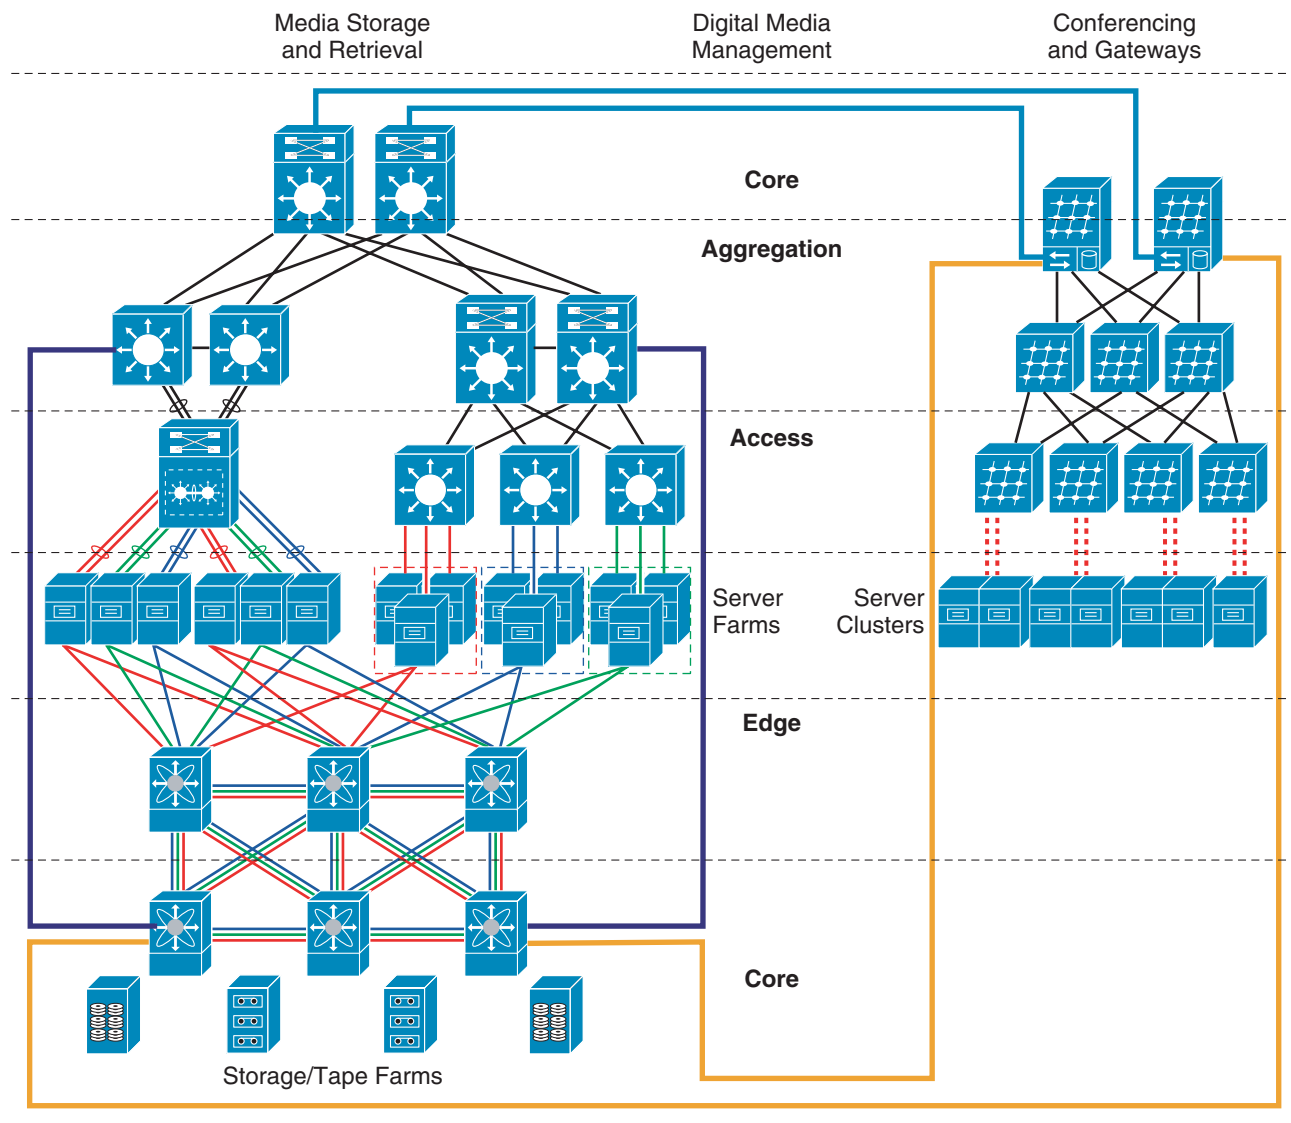
Scalability of WANs with encryption-enabled can suffer from multicast traffic due to the requirements to encrypt the same packet numerous times, once for each branch office connection. The Cisco Group Encrypted Transport VPN (GETVPN) offers a solution, allowing many branch office connections to share 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.

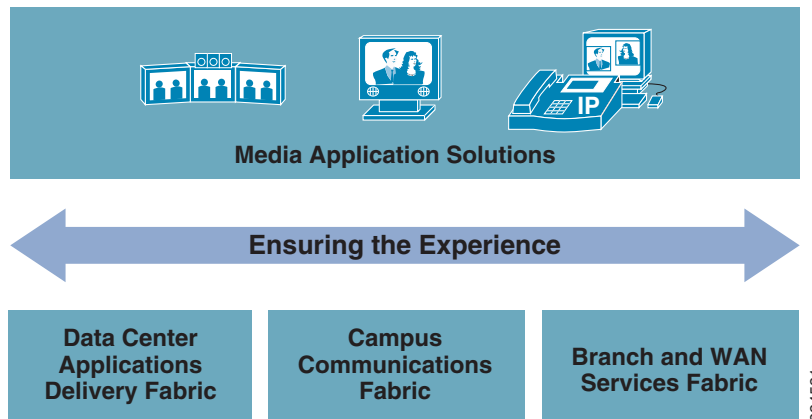
Another consideration is how to manage the increasing volume of video media that contain proprietary, confidential, or corporate intellectual property. Policies and regulatory compliance planning must be in place to manage video content as a company would manage any of its sensitive financial or customer 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/netsol/ns483/c649/ccmigration_09186a008094
- Branch Office Architecture Overview
http://www.cisco.com/application/pdf/en/us/guest/netsol/ns171/c649/ccmigration_09186a00807593b7.pdf
- Data Center Infrastructure Design Guide
http://www.cisco.com/application/pdf/en/us/guest/netsol/ns107/c649/ccmigration_09186a008073377d.pdf
- End-to-End Quality of Service (QoS) Design Guide
http://www.cisco.com/application/pdf/en/us/guest/netsol/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/netsol/ns477/c649/ccmigration_09186a008081c7d5.pdf
- Network Virtualization Path Isolation Design Guide
http://www.cisco.com/application/pdf/en/us/guest/netsol/ns171/c649/ccmigration_09186a0080851cc6.pdf

Websites

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