UEN Technical Services

Successful LAN Designs

Understanding and Resolving Local Area Networking Challenges in a Changing Environment

Abstract:
An introductory discussion and description of the benefits of an optimized Local Area Network(LAN). How these benefits impact deliverability to and from the UEN Wide Area Network edge for customer requested services within their campus/local environment. Develop clear expectations about UEN best service practices on the wide area network and how network hand-offs from WAN-to-LAN, that is those networks outside the scope of the UEN managed networks, impact the local area network performance and service delivery.

This document will serve as an introduction, the first in a series, of UEN Technical Service distributed technical white papers. In our ongoing effort to improve the delivery of service to our statewide customers, we have added distribution of technical white papers with our new technical forum discussions and improved communications, management, and maintenance policies. Such a broadly conceived architecture will position UEN to implement and delivery various applications, content and service while maintaining a high level of consistency throughout the network to better serve our customers.

Distribution Date: September 11, 2000
Version: 1.0

Table of Contents

The Changing Landscape

Over the last couple of years, LAN switching has greatly increased network performance by replacing shared media with dedicated bandwidth. Users benefit from direct access to their networks, and the bottlenecks of shared Ethernet or token ring disappear as point-to-point switching is deployed.

But as applications arrive to take advantage of switching's improved throughput, bottlenecks are starting to emerge at a higher level of the network. These new bottlenecks stem from switching's roots as a Layer 2 bridging technology--switched networks are flat domains that must be subnetted just like bridged networks to alleviate broadcast overhead. Without the subnetting performed by Layer 3 routing, LAN and switching infrastructures do not scale. Large, flat-switched networks are subject to the broadcast storms, spanning tree loops and inefficient addressing limitations that brought routers into bridged networks in the first place more than a decade ago.

Why Upgrade Your LAN?

Many schools are discovering that their existing LAN infrastructures are not adequate for today's Internet educational model. Strategic high-bandwidth applications, more powerful desktops and servers, and a growing number of educators and students are driving the need for higher-capacity and more intelligent networks. Some of the reasons schools will be migrating from shared-hub to high-speed switched infrastructures, would be to take advantage of improved LAN performance and access to LAN intensive applications such as IP Telephony, multicast video/data services and video streaming.

This document will address the school and district considerations in preparing to upgrade the LAN environment.

Signs that a LAN Upgrade Is Due

Until the mid 90's, LAN congestion was rarely a problem. LANs could move more data per second than computers could output it. Now, advances in computing and communications technology have shifted the bottleneck from the computer to the LAN. The output from one device can occupy the full bandwidth capacity of a typical LAN. When several such devices share this bandwidth, congestion often results.

Here are some indicators to watch for the signs of network congestion:

User dissatisfaction--The ultimate sign of LAN congestion is user dissatisfaction. If users are happy, then the network is generally functioning at acceptable levels, at least for the moment. Similarly, if users are dissatisfied with the performance of their networked applications, this fact overrides all statistics indicating that performance is within acceptable ranges. Causes of network congestion that may or may not have yet affected users are detailed in the following paragraphs.

Higher network utilization--Existing users may be using the network more often. Schools can measure network utilization--the percentage of time that a communications path is busy--with protocol analyzers or Remote Monitoring (RMON) devices, or get statistical information from UEN about how busy their circuits are and what type of traffic is traversing their school WAN connection link. The higher the traffic load, the higher the network utilization.

For many school environments, the LAN is approaching an excessive load when network utilization reaches the following levels:

20 percent averaged over an eight-hour workday
30 percent averaged over the busiest hour of the day
50 percent averaged over the busiest 15 minutes of the day

These figures are only guidelines, because network performance will vary based upon the number of stations on the LAN, application behavior, and traffic patterns. However, if these thresholds are exceeded, the LAN is at risk for experiencing a degradation in performance, even to the point of crippling the network.

Increased collision counts--If multiple stations on a shared, hub-based network begin sending Ethernet packets at the same time because they all sense a quiet network, a "collision" between packets results. When a collision occurs, each participant waits a random amount of time and tries to send its packet again. An increase in the number of packet collisions on an Ethernet LAN can indicate rising LAN congestion.

Application performance degradation--When the LAN is congested, applications do not perform as well as they do otherwise. File transfers take longer and terminal sessions are sluggish. In extreme circumstances, sessions may time out and disconnect, and applications or operating systems may actually crash, requiring a system restart. A decrease in application performance can be an indication of LAN congestion.

Increased network delay--All LANs have an upper limit on the amount of data they can carry. When temporarily overloaded, the LAN takes longer to distribute the load. Therefore, when the load is heavy, the average delay, also known as service time, increases. This delay makes the network seem slower to users. To directly measure service time, network administrators need special driver software. It is easier to assess congestion using other, more observable criteria, such as the data from network analyzers.

LAN Network Design Principles

Before the discussion of HOW to improve your LAN performance, we must first consider some factors when designing switched LAN networks.

Where does the UEN Network end and the School or District LAN begin?

This document is intended to help clarify and explain where the UEN wide area network demarcation point, or UEN network "edge", exists in your environment. With hundreds of circuits throughout the WAN, it would be difficult to describe how each of these "edge" hand-offs are designed. However, UEN engineers have standardized the edge connections in order to facilitate the rapid growth of school and districts over the past 5 years, and in the following section, we will explore some of these configurations.

The UEN "Fusion" process has developed the groundwork to create a network capable of delivering video and production data traffic via flexible shared connections to the schools/districts/institutions. As multicasting tests are undertaken and completed, this functionality should remain in place into the future. The backbone is currently ATM enabled, which permits configuring of virtual circuits providing Quality of Service(QoS) guarantees for higher priority traffic, such as video. As traffic prioritization is able to be supported farther and farther out on the network, services requiring guaranteed throughput (IP telephony and multicast video) may be delivered to the desktop.

In order to accommodate this network architecture and topology constraints, UEN will work in conjunction with the schools/districts/institutions to provide systems engineering recommendations to achieve successful service delivery at the UEN edge. That is the spirit in which this document has been crafted.

UEN Data Services / UEN Video Services

Formerly UtahLINK (UEN Data Services) and EDNET (UEN Video Services), UEN Technical Services has direct administrative oversight for these services under the direction of UEN Technical Director, Jim Stewart. To demonstrate how successful LAN designs will impact your environment, the following diagrams illustrate the typical UEN "edge" points throughout the state network.

Diagram(s) here on the two

Benefits of LAN Switches (Layer 2 Services)

An individual Layer 2 switch might offer some or all of the following benefits:

Bandwidth---LAN switches provide excellent performance for individual users by allocating dedicated bandwidth to each switch port (for example, each network segment). This technique is known as microsegmenting.

VLANs---LAN switches can group individual ports into logical switched workgroups called VLANs, thereby restricting the broadcast domain to designated VLAN member ports. VLANs are also known as switched domains and autonomous switching domains. Communication between VLANs ***requires*** a router.

Benefits of Routers (Layer 3 Services)

Because routers use Layer 3 addresses, which typically have structure, routers can use techniques (such as address summarization) to build networks that maintain performance and responsiveness as they grow in size. By imposing structure (usually hierarchical) on a network, routers can effectively use redundant paths and determine optimal routes even in a dynamically changing network. Some router functions that are vital in switched LAN designs include: broadcast and multicast control, broadcast segmentation, and media transition.

Campus, or LAN, based routing enabled the new technology in the local networks and intranets -- but the varied information on how Layer 3 differs from other switching and routing technologies can be confusing. This article gives a clear presentation of the key factors to consider when choosing and implementing a Layer 3 switch: http://www.cisco.com/warp/public/784/packet/july98/12.html

How Can LAN Switches Help?

If you presently have a 10-Mbps Ethernet hub and your network is exhibiting signs of congestion, a 10/100 switch can immediately and dramatically improve network performance. Unlike a hub, which forwards packets to all connected ports, a switch forwards packets only to one port: the one connected to the destination of the packet, reducing the overall volume of packets on a network.

In addition, a switch provides higher total capacity than a hub, because it can support multiple simultaneous sessions. For example, in an eight-port 10-Mbps Ethernet switch, four pairs might be communicating simultaneously. The total aggregate bandwidth in use is 40-Mbps--or four times the capacity of a 10-Mbps shared Ethernet hub.

Switched 10/100/1000-Mbps ports can operate at 10-, 100-, or 1000-Mbps, based on the capacity requirements of a given node. 10/100 switches automatically sense the speed of the attached device and configure the port for the proper speed. This feature, called autosensing, simplifies deployment in mixed Ethernet and Fast Ethernet environments.

Here is a brief summary of the benefits/advantages of 10/100 switches:

- Improved performance--Switches offer each individual or workgroup 10- or 100-Mbps of bandwidth. Fast Ethernet ports allow high-speed server and backbone connections, alleviating bottlenecks.
- Bandwidth scalability--The more ports on the switch, the more bandwidth that is available. Aggregate bandwidth equals the number of ports times the full-duplex speed of the connection (20-Mbps or 200-Mbps) divided by two.
- Bandwidth management--The network administrator can dedicate 10- or 100-Mbps connections to users who need them most.
- Reduction in collisions--Although user groups sharing a single port will still produce some collisions, there will be considerably fewer collisions than would occur if users across all user groups shared a single 10-Mbps Ethernet connection.
- Improved security--10/100 switches forward packets only to the port associated with the destination address of the packet. Other users don't see the traffic.

Features to Look for in a Switch

Today's switch must provide more than just additional bandwidth. The diversity of applications present on the network makes it necessary for medium-sized networks to employ more sophisticated bandwidth and traffic management capabilities to ensure that all these different applications get the appropriate share of available bandwidth--no more, no less. The right intelligent management features can make the difference between a network that needs frequent overhauling and one that can accommodate significant change and growth, thereby protecting the school's investment and responding more readily to changing educational requirements.

Some of the future-proofing capabilities to look for in a 10/100 switch include the following:

QoS for differentiated services--Not all applications are created equal. For one thing, not all applications have the same level of importance to the school. And not all applications require the same performance from the network. For example, high levels of latency and packet loss can cause some applications to slow down or even time out, hampering the productivity of employees. For multimedia and real-time voice applications, minimizing network latency and jitter (the variance in latency) becomes a critical consideration.

Virtual LANs (VLANs)--In general, there are two different types of networks: physical and logical. Members of a physical network are connected to a common cabling plant and are located near one another geographically. Logical networks, on the other hand, share a common set of requirements and resources, but may support users who are not physically near one another.

Multicast Support—Efficient deployment of video applications throughout the LAN and WAN is clearly highly dependent on multicast versus unicast technology. The biggest challenges are presented in the wiring closet where multicasts may need to fan out to dozens or eventually hundreds of desktops. Simple Layer 2 switches lack the intelligence to prevent flooding of multicast traffic and therefore have limited applicability in video-enabled networks. As the number of active multicast groups and the density of multicast client increase, flooded multicast packets can overload even a robust switched network and can seriously degrade end-system performance due to the burden of extraneous multicast processing. IGMP v2 is a requirement for any site deploying video services on their routed LAN. Cisco's CGMP prevents the flooding of multicasts that provides excellent performance that is not degraded by group membership changes.

Network Broadcasts – why we need them and how they impact our LAN

LAN switches do not filter broadcasts, multicasts, or unknown address frames. The lack of filtering can be a serious problem in modern distributed networks in which broadcast messages are used to resolve addresses and dynamically discover network resources such as file servers. Broadcasts originating from each segment are received by every computer in the switched internetwork. Most devices discard broadcasts because they are irrelevant, which means that large amounts of bandwidth are wasted by the transmission of broadcasts.

However, broadcasts are a function of the network. They allow computers to find resources on the network and service on the network to automatically advertise themselves. Initially, broadcasts are not the problem, it's how they are managed and how they impact your network, which becomes the problem.

In some cases, the circulation of broadcasts can saturate the network so that there is no bandwidth left for application data. In this case, new network connections cannot be established, and existing connections may be dropped (a situation known as a broadcast storm). The probability of broadcast storms increases as the switched internetwork grows. Routers do not forward broadcasts, and, therefore, are not subject to broadcast storms. For more information about the impact of broadcasts, see "Broadcasts in Switched LAN Internetworks." (white paper referral)

What's a VLAN and what benefits do they offer?

A virtual LAN (VLAN) is a switched network that is logically segmented by function, project team, or application, without regard to the physical locations of the users. Any switch port can belong to a VLAN, and unicast, broadcast, and multicast packets are forwarded and flooded only to stations in the VLAN. Each VLAN is considered a logical network, and packets destined for stations that do not belong to the VLAN must be forwarded through a router or bridge.

In a flat, bridged network all broadcast packets generated by any node in the network are sent to and received by all other network nodes. The ambient level of broadcasts generated by the higher layer protocols in the network---known as broadcast radiation---will typically restrict the total number of nodes that the network can support. In extreme cases, the effects of broadcast radiation can be so severe that an end station spends all of its CPU power on processing broadcasts.

Broadcast radiation---Broadcast radiation can become fatal---that is, 100 percent of host CPU cycles can be consumed by processing broadcast and multicast packets. Because of delays inherent in carrier sense multiple access collision detect (CMSA/CD) technologies, such as Ethernet, any more than a small amount of broadcast traffic will adversely affect the operation of devices attached to a switch. Although VLANs reduce the effect of broadcast radiation on all LANs, there is still a scaling issue as to how many hosts should reside on a given VLAN. A router allows for larger network designs because a VLAN can be subsegmented depending on traffic patterns. However, in a nonoptimal network design, a single router can be burdened with large amounts of traffic.

VLANs have been designed to address the following problems inherent in a flat, bridged network:

Scalability issues of a flat network topology
Simplification of network management by facilitating network reconfigurations

VLANs solve some of the scalability problems of large flat networks by breaking a single bridged domain into several smaller bridged domains, each of which is a virtual LAN. Note that each virtual LAN is itself constrained by the scalability issues described in "Broadcasts in Switched LAN Internetworks." It is insufficient to solve the broadcast problems inherent to a flat-switched network by superimposing VLANs and reducing broadcast domains. VLANs without routers do not scale to large campus environments. Routing is instrumental in the building of scalable VLANs and is the only way to impose hierarchy on the switched VLAN internetwork.

VLANs offer the following features:

- Broadcast control---Just as switches isolate collision domains for attached hosts and only forward appropriate traffic out a particular port, VLANs refine this concept further and provide complete isolation between VLANs. A VLAN is a bridging domain, and all broadcast and multicast traffic is contained within it.
- Security---VLANs provide added security by separating administrative systems with students, for example.
- Performance---The logical grouping of users allows, for example, an educators making intensive use of a networked station or using a multicast application to be assigned to a VLAN that contains just that educator and the servers he or she needs. The educator’s work does not affect the rest of the school, which results in improved performance for the educator (by being on a dedicated LAN) and improved performance for the rest of the school (whose communications are not slowed down by the educator's use of the network).
- Network management---The logical grouping of users, divorced from their physical or geographic locations, allows easier network management.

Step-by-Step Upgrade Planning Scenario

This section provides step-by-step guidelines for planning a hub-to-switch upgrade. This section also helps you select switch features that will be useful for your current and future school/district LAN environment.

Step 1: Forecast bandwidth requirements

Measure traffic on your existing network and project your bandwidth requirements for the next three years. Keep in mind that applications are consuming increasing volumes of bandwidth and that intranet-based business applications will further increase the network load. When you arrive at an estimate, it's a good rule of thumb to double or triple the number; this is the amount of total bandwidth you should purchase.

To calculate the bandwidth capacity that a switch can deliver, multiply the number of ports by the bandwidth of the ports (100-Mbps in half-duplex operation, 200-Mbps in full-duplex operation), and then divide by 2.

For example, a switch with 12 full-duplex 10/100-Mbps ports and 2 Gigabit Ethernet ports delivers 3200-Mbps of total bandwidth:

[(12 x 200-Mbps) ÷ 2) + (2 x 2000-Mbps) ÷ 2] = 3200-Mbps

Step 2: Segment the network

When you install a switch, each segment is its own collision domain. Therefore, to minimize unnecessary traffic, you should group network devices into separate segments. The devices in a segment do not need to be physically collocated; if the switch has a VLAN feature, you can create "logical" segments that span different floors or buildings. For example, you might create one segment each for finance, graphics design, and engineering. Creating logical segments not only reduces network traffic, it also increases security by restricting confidential traffic to one segment of the network.

If the congestion problem persists after you segment the network, you might consider replicating resources--for example, adding servers--so that users can access them locally without having to cross the backbone.

Step 3: Choose a LAN backbone

Use the calculation above to determine if you need to use Fast Ethernet or Gigabit Ethernet in your network backbone.

Step 4: Decide if you need Layer 2 or Layer 3 switching in the backbone

If you plan to deploy IP telephony, have a network of more than 250 users, or plan to run multicast applications, you should support Layer 3 switching in the backbone. Layer 3 switches are essentially wire-speed routers. These devices are ideally suited to delivering high-performance routing between VLANs or LAN segments. Layer 3 switches also provide advanced security services, such as access control lists (ACLs). With a Layer 3 switch handling all inter-LAN routing, traditional routers can have far more capacity to handle WAN traffic, including Internet access.

Step 5: Plan for remote intranet and extranet connections

If you want to provide intranet access to your schools, or enable teachers, parents or students to access your extranet, you will need a router. Choose a switch that supports the features of your router, such as IP multicasts, security, and management interface.

Step 6: Check wiring and network interface cards

Fast Ethernet requires the use of Category 5 cabling. If you want to connect the switch to devices more than 200 meters away--for example, between buildings--you will need a fiber connection. In this case, look for a 10/100 switch that offers options for 100BaseFX or Gigabit Ethernet uplinks.

Step 7: Install the new 10/100 switches

Installing switches requires simply disconnecting the old hubs and connecting the new switch.

Step 8: Install the new router

Depending on your choice in Step 4, you now should install your campus router for distribution and management of your networked environment. You can create VLANs, specify packet filtering, and select other management options using the switch/router software. However, connecting and configuring the router properly can require administrative overhead that has historically been supported on the UEN side of the edge. UEN will continue to support this function, but will impact service deployment to the LAN for advanced video and data services.

IP Telephony

To prepare for integrated IP telephony and Internet business applications-- may not perform well, or may not perform at all, without a more robust LAN infrastructure. For example, some applications—such as IP telephony--are sensitive to network delays and packet loss. Many schools may be interested in deploying IP telephony because it enables innovative new integrated data/voice applications such as unified messaging. It also saves money on intradistrict phone calls by enabling companies to piggyback IP-encapsulated voice traffic onto the excess capacity of an existing WAN and bypass traditional telephone network tolls. However, to perform at a level acceptable to users, IP telephony has its own special set of network requirements for latency and jitter, meaning that QoS mechanisms must be present on the network.

In addition to addressing the delay considerations for voice, QoS enables schools to give preferential treatment to certain data applications that are more important to the school than others. Without network prioritization or other QoS capabilities, mission-critical applications would receive the same treatment as less strategic applications, such as Web surfing.

Multicast Service

Planning for IP Multicast in LAN Environment

Today, most Internet and intranet applications operate between one sender and one receiver. This includes many of the application UEN now supports. In many emerging UEN WAN applications, one sender will transmit to a group of receivers simultaneously. Examples include transmitting a message to employees/staff, video and audio broadcasting, interactive video distance learning, transmitting data from a centralized data warehouse to multiple educational branches, collaborative computing, and many others.

With IP multicast, applications send one copy of information to a group address, reaching all recipients who want to receive it. Without multicasting, there are two options: The same information may be carried over the network multiple times (one time for each recipient), consuming bandwidth unnecessarily. Or the information can be broadcast to everyone on the network, incurring unnecessary overhead on uninterested end stations.

Support for IP Multicast requires multicast-enabling server and clients systems and at least a portion of the network infrastructure of routers and the switches that interconnect them. Multicast lends itself readily to a staged implementation beginning in the UEN Video Services Facilities and then expanding to encompass localized network distribution. In order to properly design and upgrade your LAN environment for the use of multicast over the UEN WAN, we recommend the following requirements be considered to ensure stable multicast performance and LAN success:

1. Server and client hosts must have an IP protocol stack supporting multicast as specified in Internet RFC 1112 and recommend by the UEN Advanced Technical Engineers.

2. Depending on the network infrastructure, receiving hosts may also benefit from having intelligent NIC cards that can filter out multicasts to unwanted groups, preventing the host CPU from unnecessary interruption.

3. A high-performance ROUTED backbone with a switched connection from the backbone to both the sender and receiver hosts provides a highly scalable LAN infrastructure for multicast. Shared media networks, that is SWITCHED/LAYER 2 only, will be inadequate for even low-bandwidth audio and video applications or limited pilot projects on the LAN.

4. All switches are NOT equally well suited for multicast. The most appropriate switches have a switch architecture that allows multicast traffic to be forwarded to a large number of attached group members without unduly loading the switch fabric. This allows the switch to provide support for the growing number of new multicast applications without impacting other traffic. Layer 2 switches also need some degree of multicast-awareness to avoid flooding multicasts to all switch ports. Multicast control in Layer 2 switches can be accomplished in a few ways:

    a. VLANs can de defined to correspond to the boundaries of the multicast group.

    b. Layer 2 switches can snoop IGMP queries and reports to learn the port mappings of multicast group members. This allows the switch to dynamically track group membership. But snooping every multicast data and control packet consumes a lot of switch processing capacity and therefore can degrade forwarding performance and increase latency.

    c. Taking advantage of the Generic Attribute Registration Protocol (IEEE 802.1p) will allow the end system to communicate directly with the switch to a join a 802.1p group corresponding to a multicast group. This shifts much of the responsibility for multicast group configuration from Layer 3 to Layer 2, which may be most appropriate in large flat-switched networks.

    d. The traditional role of the router as a control point in the network can be maintained by defining a multicast router-to-switch protocol, such as the Cisco Group Multicast Protocol (CGMP), that allows the router to configure the switch's multicast forwarding table to correspond to the current group membership.

The purpose of the Cisco Group Management Protocol (CGMP) and Internet Group Management Protocol (IGMP) Snooping is to restrain multicast traffic in a switched network. By default, a LAN switch floods multicast traffic within the broadcast domain, and this can consume a lot of bandwidth if many multicast servers are sending streams to the segment.

Multicast traffic is flooded because a switch usually learns MAC addresses by looking into the source address field of all the frames it receives. But, since a multicast MAC address is never used as source address for a packet and since they do not appear in the MAC address table, the switch has no method for learning them.

Prepare today for multicast on the LAN. Procurement and design decisions should take careful consideration in deployment of inexpensive, yet ineffective LAN switching solutions. Please refer to the UEN Advanced Technical White Paper on Multicast, release date 10/11/00, for more detailed information on UEN IP Multicast issues.

Conclusion/Recommendations

UEN is bringing “network fusion” (multiservice networking) a standards-based, open-systems architecture for converged networking. UEN Technical Services is continually engaged in the continuing evolution of a plan for enterprise multiservice networking that will successfully deliver the framework for an open architecture. The architecture comprises three distinct building blocks:

- local area infrastructure
- wide area infrastructure
- connectivity from the UEN edge to LAN campus environment

Out-of-the-box integration is of inestimable value to today's dynamic enterprise. By leveraging a common underlying infrastructure, an integrated solution maximizes functional interoperability and minimizes costs. Moreover, functional solutions reduce their local footprint making them appropriate for all heterogeneous organizations and systems, regardless of size. A solution selected for an school's current topology should also anticipate change and potential growth without overloading the current system.

….more here

REFERENCE:

LAN Technologies
http://www.cisco.com/warp/public/473/

LAN Switching
http://itc.broadwing.com/expert_advice/white_papers/lan.asp
http://www.cisco.com/cpress/cc/td/cpress/ccie/ndcs798/nd2023.htm

Convergence of Layer 2 and Layer 3 in today's LAN
http://www.foundrynet.com//WP3.pdf

Layer 3 Switching Demystified
http://www.cisco.com/warp/public/cc/so/neso/lnso/cpso/l3c85_wp.htm

Broadcasts in Switched LAN Internetworks
http://www.cisco.com/cpress/cc/td/cpress/ccie/ndcs798/nd20e.htm

Campus Network Services
http://www.cisco.com/warp/public/cc/pd/si/casi/ca6000/tech/cns_wp.htm

Requirements for the Wiring Closet
http://www.cisco.com/warp/public/cc/so/neso/lnso/wire/wclos_wp.htm

Converged Networks
http://www.3com.com/technology/tech_net/white_papers/500671.html

Optimizing Your Network Design
http://www.cisco.com/cpress/cc/td/cpress/design/topdown/td0512.htm

Overview of IP Multicast
http://www.cisco.com/warp/public/cc/techno/protocol/ipmu/tech/ipmc_wp.htm

Configuring Multicast
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat5000/rel_5_2/config/multi.htm

Future of Telephony
http://www.sonusnet.com/sonus/Backend/prod/whitepapers/telephony.pdf

Preparing Your Campus for AVVID Deployment
http://www.cisco.com/warp/public/cc/so/neso/vvda/avvid/prepc_wp.htm

Technical Considerations for Converging Data, Voice, and Video Networks
http://www.cisco.com/warp/public/cc/so/neso/vvda/avvid/tecon_wp.htm