



Office of Information Technologies

IT Architecture Initiative

VIDEO DISTRIBUTION and NETWORK EVOLUTION at CORNELL

Prepared by R. David Vernon

INTRODUCTION

As part of the Information Technology Outreach Initiative, the Office of Information Technology (OIT) is producing a series of papers outlining directions in information technology architecture. The papers are intended to facilitate understanding and dialogue about information technology trends at Cornell, with the ultimate goal of the utilization and interoperability of information technology services through

SYNOPSIS

Support for widespread video distribution is a critical design goal of the design strategy at Cornell. This paper briefly outlines the following issues related to video distribution:

- n Evolution of video support networks at Cornell.
- n Three types of video services required by Cornell.
- n Recommended temporary transition strategies.
- n Final IP-based data network design issues and directions.
- n Practical implications.

EVOLUTION OF VIDEO SUPPORT NETWORKS (1980s- FUTURE)

The video support network at Cornell has undergone an evolution since its original installation in the 1980s. As needs and technology have changed, the network has evolved from the original analog coaxial-based broadband to the current video fiber network, and is currently looking to the next stage, an IP-based data network.

FROM MID-SPLIT BROADBAND TO THE VIDEO FIBER NETWORK (VFN)

In the early 1980s Cornell installed an analog broadband system¹ to provide video services to approximately 35 locations on campus. Aging hardware, the need to recover cost and a limited technology growth path led to the decision to retire this system. At that time Cornell Information Technologies (CIT) started a limited deployment of point fiber-based video services, now called the Video Fiber Network (VFN), using this technology to provide services until a unified data network strategy could be implemented.

FROM VFN TO AN IP-BASED DATA NETWORK

While the VFN provides very high quality video, it is not economical or technically feasible for the long-term. The cons include limited fiber and conduit space, limited supported locations, and the high cost of configuring esoteric hardware. Video services at Cornell will be best served by engineering the data network to be a widespread and cost-effective resource to support Cornell's video service needs.

The transmission of video over the IP data network is technically achievable. However, the transition from the VFN to an IP-based network requires review and understanding of the trade-off's of cost, access, and video quality. The basic differences are:

- n The VFN, based on dedicated, point-to-point fiber links attached to costly decoding software (codecs) provides very high quality video, but at high cost per transmission and to a limited number of campus locations and therefore a limited number of patrons.
- n IP-based video services can be scaled to meet the needs of tens of thousands of Cornell patrons, at lower costs, but initial quality of service may fall short of the demands of select transmission requirements.

At this time it is the belief of the Office of the Vice President of Information Technology (OIT) that the quality and value of video services achievable over an IP-based network meet the vast majority of our patron needs. Furthermore we believe that investment should be minimized and maintained for only as long as its premium video quality is worth the higher cost to the patron base when compared to ever improving quality video services.

CORNELL VIDEO SERVICE REQUIREMENTS

The IP-based video network solution will have to address three types of video services:

- n Video on demand services based on Unicast.
- n One to many video broadcast services based on Multicast.
- n Video conferencing.

¹ A mid- and sub-split 300Mhz coax analog broadband system.

The following three sections further define these needs.

UNICAST STREAMING

For the purpose of this document, unicast streaming is defined as a single-point transmission of encoded video that often leverages client buffering or network caches to improve video display quality.

Caching a percentage of the video on the local client or network before playback is an effective way of working around bandwidth limitations of local, campus, and data networks. In unicast streaming, the video is being sent in only one direction. A delay a buffer or cache creates does not negatively affect the quality of the video.

However, using unicast streaming for "broadcast events" has drawbacks. (A broadcast event is one that is sent to many destinations at the same time, as opposed to a single destination.) The main issues are:

- n Network use: Unicast streaming requires a dedicated flow of information for each session. If there are 1,000 people viewing the same video at the same time, there will need to be one thousand unicasts active. Each stream would consume some percentage of the available IP data network.
- n License cost: Many popular commercial streaming servers are licensed based on the quantity of concurrent unicasts supported. This raises the total cost of the service.

It is also important to emphasize that caching is not an acceptable tool for video streaming because a "conference" requires real-time interactive communications. Caching is not suitable for interaction.

Despite these limitations, unicast streaming, enabled with familiar applications such as QuickTime Player, Windows Media Player and RealNetworks, remains the most widely adopted IP video distribution tool at Cornell and throughout the world.

MULTICAST

Multicast is not a system of video encoding or transport protocol such as RTP. It is a strategy for engineering an IP data network that efficiently distributes encoded video. Multicast enabled networks replicate single video streams to multiple destination clients based on defined router rules. While unicast streaming w

² <http://www.cselt.it/mpeg/>

³ <http://www.itu.int/itudoc/itu-t/rec/h/h261.html>

⁴ <http://www.ietf.org/html.charters/avt-charter.html>

on-one video services, multicast enabled networks can broadcast video to one destination ⁵while minimizing network impact and cost.

Today limited multicast sessions are used in campus networks throughout the advanced service development and campus to campus multicast transmission via Internet service providers is still under active ⁶testing and development.

VIDEO CONFERENCING

Video conferencing is the combination of audio, video, and digital communication technology that enables live, real-time interaction between two or more local video on demand or video broadcasts, which are usually used for one-way transmission. The primary use for video conferencing is two-way, site-to-site communication.

The nature of a video conference -- fairly static participants requiring moderate bandwidth -- allows for efficient encoding requiring far less bandwidth than most high quality video broadcasts. However, video conferencing places unique quality demands on the network. The perceived worth of an interactive conference can degrade quickly if network bandwidth that which is naturally tolerated by the participants. This level of service is achievable with Cornell's installed data network despite the moderate bandwidth requirements of video conferencing applications.

Currently Cornell's ad hoc packet-based video conferencing strategy is based on the H.323 standard, which is a suite of video/audio codecs and system control protocol for video conferencing. While H.323 does have provisions to support multicasts, current implementations require one unicast for each participant. Multiple participants require additional bridging hardware and software and a Multipoint Control Unit or MCU. Implementation of H.323 video conferencing can be achieved with a growing list of software for desktop computers or via turnkey standalone units, such as those provided by Polycom Corporation.

NETWORK DESIGN ISSUES AND DIRECTIONS

The necessity to retool the data network infrastructure at Cornell will be driven by the need to guarantee assured quality links for video conferencing services and by the demand for "multicasts" of high-resolution real-time video broadcasts. The current network must accommodate the different needs of these two classes of video services. While the issues for unicast, they alone may not dictate a network strategy change because the issues in the current network can be addressed by the use of local host buffers or

⁵ <http://www.ietf.org/internet-drafts/draft-ietf-mboned-mcast-apps-01.txt>

⁶ <http://www.internet2.edu/multicast/>

⁷ <http://www.itu.int/itudoc/itu-t/rec/h/h323.html#dl>

⁸ <http://www.polycom.com/>

caches. Additionally, unicast is primarily focused on streaming a stored video to a recipient, and therefore does not benefit from a multicast enabled network.

However, for all three classes of video service, high quality, full motion video requires greater network bandwidth than low quality static video. The data network at Cornell will not support the total suite of packet video services outlined. The network has three critical areas:

- n Capacity
- n Support for multicast
- n Predictable quality of service

The best approach to retooling the network to support video at Cornell will be to address the following:

- n Available funding
- n Primary limitations of the installed wire plant
- n Good network design

Cornell is currently hobbled by its installed "category 3" building data wiring which has a maximum data rate of 10-MB connections. 10 MB switched Ethernet, if ubiquitously deployed, is arguably adequate for the video services outlined above, but with little additional capacity to support other concurrent network bandwidth.

While pilot programs are exploring the total cost of rewiring the campus, an approach that addresses video service demands within a network supporting a myriad of concurrent applications, is not likely to succeed based on simple over-provisioning of bandwidth. Implementing best practice strategies of engineering data networks to support "quality of service" the bandwidth that Cornell can deliver can be optimized as well. Furthermore, such strategies remain valuable regardless of the base network design and are important tools for future services such as Voice over IP.

MORE ON QUALITY OF SERVICE

Quality of service (QoS) is a networking term that specifies a guaranteed throughput and prioritization of data transmissions. End to End QoS (e2eQoS) is a rubric used to describe services that allow prioritized transmissions across the larger enterprise. It refers not to the video data transmitted, but to a data network design criteria that enable predictable video delivery quality.

⁹ 10-MB data link limitations will be problematic for router to switch, switch to switch, server to server, and server to client connections supporting concurrent video.

QoS enabled network systems can be programmed to allow guaranteed service to applications over a network, which helps assure presentation quality. This is because video conferencing and voice services (even at low bandwidth) are particularly vulnerable to compromised quality due to congested networks. (Congested networks create latency or a lag in the amount of time it takes a packet to travel from destination. This is obviously detrimental to the interactive aspect of a video conference.)

Strategies for implementing QoS might include:

- § Over provisioning. For example, installing hardware that has the ability to handle projected peak demands without degradation of service to any concurrent users.
- § Static hardware configurations that prioritize aggregate data streams based on a classification schema. For example, data is given preference over e-mail or voice. Or, data from a specified list of hardware always get preferential treatment on transmitting hardware.
- § Dynamic QoS negotiation. For example, a process that allows any qualified user to request an application, to reserve a specified level of bandwidth for a specific period of time.

While over provisioning and static hardware configuration strategies are relatively simple in concept and deployment, dynamic provisioning of QoS services requires a rather complex coordination of broader resources. Dynamic resource reservation implies the exchange and auditing of relevant policy, security, user, hardware and billing information, via a suite of developing protocols to all active components of network infrastructure. Although dynamic provisioning of assured quality links is an efficient use of installed hardware, the complexity of this system implies increased management overhead. In turn, the implementation of QoS at Cornell will not be a simple exercise and will likely evolve over time.

An initial easy to implement strategy combining simple over-provisioning of bandwidth with aggregate IP stream prioritization may progress over time to a complex system guaranteeing end-to-end service levels via dynamic QoS negotiation protocols developed with an eye on national trends. The value at any point in time of a simple vs. complex QoS network architecture. As there is sure to be a growing demand driven by new QoS aware applications, long-term success will require the monitoring and demand to allow ample lead-time for network re-engineering to extend or enable sophisticated tools.

Possibly more challenging than building initial QoS enabled networks will be the policy implications implied. Dynamic brokering of "who" is allowed to use what percentage of the network infrastructure, in a fashion that campus and national "trust," will require a common and well understood strategy for bandwidth allocation.

¹⁰ Among other information resources see: <http://www.ietf.org/html.charters/rsvp-charter.html>, <http://www.ietf.org/html.charters/rap-charter.html>, <http://www.ietf.org/html.charters/intserv-charter.html>

¹¹ Among other information resources see: <http://www.ietf.org/html.charters/diffserv-charter.html>

recovery. In addition, a QoS enabled network must be seen as a primary "client" central directory services, as¹² campus LDAP will be the likely archive of policy information.

As the diligent implementation of QoS and multicast strategies evolve into functions that the Cornell campus will leverage for video and other data distribution exploration of enhanced data caching strategies must also be explored. The massive amounts of high-quality video at guaranteed quality levels might be achieved through preemptive distribution, or caching of video for unicast applications network proximity to the viewing client. Distributed File Systems, or cache network hardware may afford equivalent, or superior unicast service quality dependent only on QoS and multicast. The economies of network caching approaches distribution are well demonstrated by Internet caching providers¹³ such as Akamai¹⁴

PRACTICAL IMPLICATIONS

Departments and individuals with the need to transmit video should be cognizant of evolving Cornell data network potential. Until the data network is retooled, transmission quality may not be suitable for specific applications and /or a department is encouraged to work with CIT and local video support resources to experiment with video delivery via the data network in order to gain an understanding of the limit against the alternative, though expensive, VFN video transmission service.

Clearly, as the data network evolves a greater percentage of Cornell campus-based video will be met cost effectively. However, it is important for departments and individuals to realize that Cornell based QoS and Multicast strategies do not imply equivalent service to the larger Internet. Transmissions of video to off campus locations will allow the architecture of the larger Internet community - where the quality of the service is as good as the worst network link.

There are a number of initiatives, most notably those being driven by Internet Service Providers on the issue of e2eQoS across the wide-area Internet. As wide-area network e2eQoS requirements become clear, Cornell's QoS strategy must evolve and be able to meet the most likely outcome will be a fairly complex QoS architecture based on real-time protocols used to dynamically coordinate resources and allocate costs.

Today departments may find reasonable video transmission for videoconferencing and quality video broadcast to locations connected to I2 even without participating pilots, as I2 is currently uncongested and uniformly deployed. However, departments are limited in their ability to deliver quality video over Cornell's commodity Internet and will need to explore dedicated ground connections or satellite services for video transmission.

¹² <http://www.openldap.org/>

¹³ <http://www.akamai.com/>

¹⁴ <http://www.ibeam.com/>

CLOSING THOUGHTS

As it is the intent of OIT to enable video services outlined in this paper, traditional support for VFN services while Cornell targets and coordinates n investments over an extended period with the following network design goals:

- n Reengineering the campus data network to support multicast.
- n Reengineering the campus data network to support QoS.
- n Policy dialogue.
- n Advanced Network Video caching tools (when appropriate).

The transmission of video over the IP data network is technically achievable the transition from old models of limited very high quality video transmissi supporting thousands of video transmissions at a quality level acceptable to service requires prudent investment and studied review. Paramount among the forge a common understanding of cost benefit of initial compromise in the qu transmission vs. the maintenance of high cost esoteric systems such as the V evolution of video and network services has many challenges and will require proactive communication, and world class technical expertise to succeed.

Over time the evolving ability to distribute video over the data network to a part of a growing suite of services that embody the larger video service s This strategy will also encompass video development for broadcast across the or the world and maintaining massive repositories for video on demand assets support and the development of new and enabling resources are both required grand vision of ubiquitous, high quality, digital video delivery at Cornell.