White Paper

Using UDP/IP to Transport Video Streams

Edwin Hoffman CDO

RAPTOR NETWORKS TECHNOLOGY INC

OVERVIEW

Video/IP is often a major requirement for today's data networking. Many IT departments are finding that video is rapidly becoming a priority in their network plans because there are now many applications that need Video/IP. The list includes:

IPTV	Delivers normal TV channels over IP-based networks directly to PC or MAC workstations. Several trading companies and many news services now provide these facilities to some or all of their employees.
Video conferencing	Performs two-way or multiway conferences between remote sites using both video and audio facilities. Quite often, a collaborative computing application is attached to the video conference that allows all conference participants to work jointly on the documents and diagrams.
Video monitoring	Monitors, stores, and recalls video and audio sequences acquired from various sources. Often these may be fixed video sources that are monitoring hazardous industrial processes or checking other processes for quality of output. These applications often use complex video analysis software to recognize certain changes in the video field.
Video surveillance	Monitors and often stores video for security purposes, usually requiring long-term storage. Audio sequences are sometimes included as well. Manual control is sometimes required, but systems that automatically move video cameras to preset positions and automatically detect movement within the video field are becoming more common.
Video distribution	Moves high-resolution and high-quality video from process to process and site to site without resorting to video tape or DVD production for the transport. This is slowly becoming a major requirement since video tape and DVD formats can easily be stolen or illegally copied while being moved between sites.

Most of these applications will require that an existing IP network start carrying the new video data. However, like many networking applications, Video/IP does benefit from a dedicated IP network design. The resulting performance improvements often justify the associated increased costs. Where these video applications have to be delivered over existing IP networks, the network manager must take care that the new video data does not adversely affect the present network and that the video data is delivered at the quality level that the application requires. With these requirements in mind, this document will discuss the technical requirements of the various applications and the situations that usually cause problems for network managers. Finally, it will show how the Ether-Raptor family of layer 2/3/4 switches is uniquely designed for video IP transport and meets the needs of these kinds of applications.

INTERNET PROTOCOL TV (IPTV)

IPTV is typically produced by passing analog video from an existing TV source (cable, RF, or satellite) through a digital video encoder. These devices can convert the analog video into any of several types of encoded data streams such as:

•MPEG (MPEG-1 and MPEG-2 are mostly used in this area)
•H323 or H261 (used often for video surveillance)
•MJPEG (Motion JPEG).

Sometimes several TV channels are available to the network via these encoders, and these channels are delivered over the IP network using IP multicast protocol. This technique minimizes the amount of bandwidth required to deliver multiple channels to what can be tens or hundreds of workstations. In this situation, the network manager has to care for two areas of concern:

- All of these Video/IP systems use User Datagram Protocol/Internet Protocol (UDP/IP) to deliver the video data. | Many of the sources produce large video frames, and in turn, since Ethernet limits each packet to a length of 1518 bytes, create fragmented IP packets to transport the video data. However, most layer 2/3 switches handle UDP/IP traffic poorly and focus instead on high performance for Transmission Control Protocol/Internet Protocol (TCP/IP). Fragmented packet series also create problems in many layer 2/3 switches because they have to rebuild the entire series (up to 64 kilobytes) in memory and refragment it back out, adding latency to the series. And, when many video streams run, the switch can exhaust memory, which causes all data transport to slow or fail.

- Many layer 2/3 switches perform poorly with IP multicasting because they implement the multicast protocols in software. Many layer 2/3 switches as well as many layer 2 switches consequently perform IP multicasting at speeds as low as 50,000 PPS (packets per second), which is nevertheless adequate when delivering a single, multicast stream over a single 10/100 Mbps Ethernet port.

These switches fail when transporting multiple IP multicast streams, such as over a Gigabit Ethernet switch uplink. Here the packet-per-second rate limitations of the switch can be overrun very quickly, so the switch starts to drop frames and the quality of the video becomes very poor. In these cases, latency is the main enemy. As the traffic congestion increases and frames are dropped more frequently, the video stream can become unusable very quickly. Consequently, network switches for IPTV applications should have excellent performance characteristics with UDP/IP, fragmented packets, and IP multicast traffic.

VIDEO CONFERENCING

Video conferencing rarely uses multicast protocols, even when multiple sites are conferences and separate UDP/IP streams (often based on H323 protocol) are created between sites. Video conferences often take place over slower lines (ISDN, T1, or DSL), necessitating a high level of video compression to reduce the bandwidth requirements.

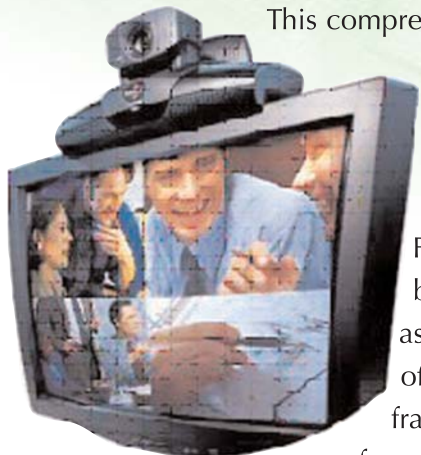

This compression actually reduces the performance in typical layer 2/3 switches because latency is now the major problem. Large latency variations are common in most switches. When these changes occur between video frames, the video quality of the conference can quickly degenerate to an unusable level.

Fragmented IP transport of video also causes problems in many layer 2/3 switches because they are flow-based. Thus, when they encounter a large video frame sent as a fragmented IP stream, they cannot forward the second and subsequent packets of the stream correctly as a single data flow. Instead, they rebuild the entire video frame inside their own memory, apply the appropriate forwarding rules, and then refragment it to send it out the egress port to the final destination. Relying on such software-based processes to rebuild and refragment each video frame creates large and variable latencies in typical switches. Video conferencing is often troublesome with these types of switches, since variable latency is the single, biggest impediment to high-quality video. Therefore, networks that must support video conferences should use layer 2/3 switches that are known to handle fragmented IP traffic well.

VIDEO MONITORING

Video monitoring, including process monitoring, requires clear and precise video data images. This is true whether someone is viewing the images as they occur, or whether the data is simply being stored or delivered to automatic checking software. For this reason, video monitoring has resisted the move to digital video. An example of video monitoring that requires extremely high-quality images is in the field of nuclear energy. Video monitoring of areas hazardous to human beings obviously has to be high of quality, and the operator cannot afford to miss a major event in a nuclear reactor. Movement detector software is often used to “see” movement in a video field that does not normally have movement. Using poor quality network transports that can drop frames or create jitter causes the software to detect incorrect changes all the time.

Handling the UDP/IP traffic is a major problem for many layer 2/3 switches. Most of these switches forward UDP via software, enabling them to focus on hardware-based support in their ASIC devices for the more popular TCP/IP protocol. This makes their UDP/IP performance very poor, with throughput as low as 10,000 PPS in some products. Most switch vendors will not even publish their UDP/IP throughput, reporting only the TCP/IP forwarding rate. Congestion problems abound in switches that have this behavior.

A simple two-stream video application will often cause frame loss, since the loss of even one packet in a fragmented IP stream causes the entire video frame to be lost. Serious jitter is also typical, caused by latency variations that are due to the buildup of congestion-related latency. It is essential in video monitoring to choose network switches that excel in handling UDP/IP and fragmented IP traffic while maintaining low-latency variations.

VIDEO SURVEILLANCE

Video surveillance has been carried out in an analog domain for many years and is just beginning to move to the digital world. The main driving force is the prohibitively high cost of analog video switching. Also, using digital video gives the operator greater flexibility by allowing the video streams to be viewed by any video-capable workstations, like PCs or MACs.

Video surveillance and video monitoring tend to use the same technology, although surveillance applications until very recently tolerated poor quality video. Several important legal cases and the more frequent use of surveillance video in criminal trials have created a need for higher resolution video sources.

Data priority is a major issue when using a legacy IP network to deliver IP-based video to end stations. In some cases, video traffic can displace important data for other applications; in other cases, data transmission interferes with video. A layer 2/3 switch must be able to handle priority properly for successful video applications.

DiffServ priority marking is often used to ensure that an IP video stream is given the bandwidth it requires. However, layer 2/3 switches rarely reach directly to the edge of the network where the end station resides. Instead, layer 2 edge switches have the job of delivering the IP video stream correctly.

This situation requires a layer 2/3 switch that can map DiffServ IP priority to 802.1p Ethernet packet priority and send Ethernet priority-tagged packets to the end stations. Not all layer 2/3 switches can do this mapping, while most layer 3 routers do not do this.

VIDEO DISTRIBUTION

The IP video distribution market is in its infancy. Video distribution was traditionally accomplished using video film, and more recently using DVDs sent via overnight transport. Even these straightforward methods have problems - for example, Steven Spielberg's film *ET* was copied illegally while being transported across the U.S. by train. *Terminator 3* was available on DVD in Europe before it was released in theaters. Now digital video delivery is becoming a major force. The new Harry Potter film was sent to post-production by digital delivery, and video conferences via digital links took place throughout the shooting of the film. Digital video distribution uses high-resolution and high-quality video links (often MPEG-2 or MPEG-4) that need up to 45 Mbps of bandwidth and some times even B1 video is used, which can raise the bandwidth bar as high as 611 Mbps. Video distribution uses UDP/IP transport, so the 10,000–50,000 PPS throughput limitation in many layer 2/3 switches often appears as a problem along the distribution path. Fragmented IP is also a major problem because of latency issues. Very high-quality video is often transported over Asynchronous Transfer Mode (ATM) links, but this solution has two major problems—high price and insufficient bandwidth.

The price of service for Gigabit Ethernet is dropping very fast relative to ATM OC-12. Consequently, many companies are moving to Gigabit Ethernet-based IP networks that have plenty of bandwidth compared to the OC-12 bandwidth of 622 Mbps. The bandwidth problem arises because an encompassed B1 video stream with 18 sound channels needs around 611 Mbps to arrive perfectly. ATM OC-12 is barely able to accommodate this and often prevents any other data from flowing over the same link. A Gigabit Ethernet link, using a layer 2/3/4 switch that can forward UDP/IP and multicast traffic at wire speed, can handle such a video stream while still leaving over 300 Mbps for other kinds of traffic.

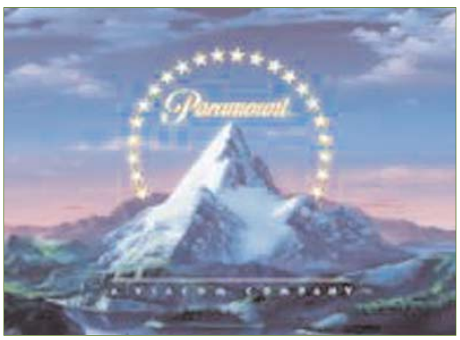

Another bandwidth consideration is the number of video channels that must be sent over an IP link. A film or entertainment company may want to stream high-quality video to a number of sites while delivering more than one stream at a time. Such applications normally need multiple OC-12 ATM links, which would be ruinously expensive, or multiple Gigabit Ethernet links, which would require multiple fibers and thus be quite expensive as well.

The Ether-Raptor system from Raptor Networks Technology, Inc. provides a unique alternative by transporting Ethernet over 10 Gigabit. Taking advantage of this bandwidth, the Ether-Raptor system can simultaneously carry 12 channels of B1 video up to 40 km over a single, directly-connected fiber. This means, for example, that theaters in the Orange County area of Southern California could have direct feeds from Hollywood film lots and show high-quality video (up to IMAX standard) in real time. Using the Ether-Raptor system and its 10-Gigabit Ethernet transport, would not have frame loss or jitter, and most likely none of the annoying scratches that appear in ordinary film theaters.

TRANSPORT CARRIERS

Frequently the weakest link in a digital video system is not the video source or the source network, but the delivery network that relies on transport carriers. These carriers often use the same Wide Area Network (WAN) systems for video delivery that they use for delivering ordinary data traffic, which typically have the same kinds of problems handling video's UDP/IP and multicast IP traffic encountered in LAN switches and routers.
Keeping in mind that carriers often multiplex various data sources over a single link, the customer needs to make sure that:

- Bandwidth for the application is available.
- Quality of Service for the video stream is maintained throughout the transport network.
- Carrier systems are capable of handling high-speed UDP/IP and multicast IP streams.

VIDEO MONITORING EXAMPLE

Now consider an example of video surveillance and monitoring in a factory that is combining process control monitoring and surveillance to reduce costs, or highway monitoring where many cameras are under control. When data from many cameras or other video sources converges in a single area, this data is typically fed to a local network switch. If each video source is connected to a 10/100 Mbps interface, it will not stress the capabilities of that interface. If this switch has a typical forwarding capability of only 10,000 PPS for UDP/IP and 50,000 PPS for multicast IP traffic, however, its uplink will start to have problems as the number of video sources increases and the packet rate rises.

If the switch in this example does not support priority or a priority system is not implemented, the video streams greatly interfere with other data traffic on these links. If this is a process control application, then other control data is probably flowing over these uplinks. In a highway scenario, other data streams may be controlling signals and roadside signs and sending telemetry data from road sensors. It would be inadvisable to prevent these data streams from arriving at the control center because video was choking the uplink.

Redundant communication links can also provide a networking challenge. Twin links are often used to assure reliable data delivery. Such links are often diversely routed to avoid the situation where a backhoe or some other construction equipment inadvertently destroys all communications. The network device needs to be able to handle fast switchover to the backup link, and needs to be able to quickly change its routing tables (in the case of layer 2/3 switches) to reestablish communications.

ETHER-RAPTOR BENEFITS

The Ether-Raptor family of layer 2/3/4 switches is optimized for video transport and avoids all the potential problems outlined above because:

- Fragmented IP traffic is handled correctly in the Ether-Raptor switch, allowing video data streams to flow naturally and at wire speed. The Ether-Raptor switch does not reassemble or refragment each frame, instead using hardware integrated directly into the switch hardware to forward each fragmented packet individually at wire speed.

- UDP/IP forwarding is performed at wire speed in hardware via Ether-Raptor switches, ensuring that the switch will not lose video frames.

- Multicast IP is handled directly in hardware at wire speed, enabling hundreds of Multicast IP streams to travel over Ether-Raptor Gigabit Ethernet or 10-Gigabit links without jitter or frame loss. The Ether-Raptor switch can handle up to 10,000 H261 Multicast IP streams over a single 10-Gigabit link.

- The Ether-Raptor switches can prioritize video and other traffic correctly using both the DiffServ standard at layer 3 and the 802.1p standard at layer 2. Furthermore, it correctly maps DiffServ priority to 802.1p priority for data delivery to end stations via subsequent layer 2 switches. The Ether-Raptor switch has specific facilities for handling IP multicast priority that allow it to prioritize one multicast relative to another multicast. This capability does not exist in other layer 2/3/4 switches.

- Ether-Raptor switches allow users to provide layer 2/3/4 control right to the edge of their networks with cost-effective, edge-type switching systems.

SUMMARY

In any serious application that requires Video/IP, it is essential to consider a network switch that can correctly control UDP/IP traffic, fragmented IP traffic, and multicast IP traffic. The layer 2/3/4 switches in the Ether-Raptor family, with their wire-speed support for all aspects of video traffic, are uniquely suited for these kinds of video applications.

Corporate Headquarters: 1241 E. Dyer Road, Suite 150 Santa Ana, California 92705
Phone: 949-623-9300 / Fax: 949-623-9400 / Web: www.raptor-networks.com / E-mail: info@raptor-networks.com

CD-WP300 10/27/04