Defense Research Engineering Network
Net Centric Digital IPTV Video Transport Bridge

Anthony Sullivan, Christopher Healy, Ken Wernle, Joe Gaucher,
U.S. Army TEDT-WS-SE-D, WSMR, NM, USA 88002
U.S. Navy 7.2.3 Enterprise Architect, Lakehurst, NJ 08733
USAF 46th Test Group 746TSS/XPX, Holloman AFB, NM 88330-7717
Video Furnace Inc, Libertyville, IL 60048

ABSTRACT

Various approaches are described in this document for transporting digital video over the Defense Research Engineering Network (DREN), Ethernet and SONET networks, along with solutions to operational and technical challenges. Video formats include high resolution complementary metal–oxide–semiconductor (CMOS), high bit-depth infrared, and high frame rate parallel digital video. Commercial analog and digital frame grabbers are utilized, as well as software running under Linux, Solaris, Mac, and Microsoft Windows XP. No other specialized hardware is required. A network configuration using independent video virtual local area network (VLAN)'s for video channels provides efficient transport for high bandwidth data. A framework is described for implementing both uncompressed and compressed streaming with standard and non-standard video. The Video Furnace Internet Protocol Television (IPTV) Bridge and Translator Multiplexor solution is presented as a Joint Net Centric distribution solution of video services to all platforms and operating systems with a zero footprint client server appellate. AeroStar Unmanned Air System integrated with the Video Furnace IPTV solution transports live UAV video to Joint Base Dix, McGuire, and Lakehurst in real time with latencies of less than 200 msec achieved.

Keywords: Defense Research Engineering Network, IPTV, Video Compression, Streaming Video, Digital Video Network, Joint Net Centric Distribution, Zero Footprint, Net Enabled Command Capability

1 INTRODUCTION

In the Net Centric Test & Evaluation (NCT&E) community, there is a requirement for a variety of test scenarios to transport and distribute motion imagery data from both low and high speed imaging systems. The utilization of analog and film cameras has satisfied this requirement in the past. However, increased environmental concerns and cost, along with demands for reduced turn around time with increased performance drives the need to replace these legacy systems with fully digital imaging systems. As imaging systems are replaced, the old "Sneaker net" transport system is being replaced by high speed broad band networks. Captured digital motion imagery data can be transported long haul to remote locations for viewing in real-time, processing and review. Both compressed and uncompressed data transport is required because motion imagery data is used for a variety of purposes. There is a need for real-time use of the motion imagery data in the Test & Engineering (T&E) environment, latencies of less than 200 msec "glass-to-glass" is required in many instances. This paper discusses methodologies utilized in the transport of uncompressed Gigabit/second class "Screaming" video and the techniques used to encode and transport traditional compressed streaming video with solutions to operational and technical challenges.

Due to a variety of reasons, transporting analog video via traditional methods is being phased out. National Television System Committee (NTSC) low bandwidth analog video has traditionally been transported via coaxial cable, fiber and microwave radio frequency (RF). Each method has pros (simplicity, well understood) and cons (potential for significant signal loss, high cost via RF). The Test Support Network Internet Protocol (TSNIP) installed at White Sands Missile Range (WSMR) has evolved into a very large fiber based network infrastructure. The TSNIP transports the NTSC video after it has been converted to a packetized digital format. Up to 211 Mbits/second bandwidth allocations, depending on the digitized format, that is required for uncompressed NTSC video. When compression is acceptable, this rate can be reduced to 2-10 Mbits/sec while still satisfying the < 200 ms propagation delay requirement.

Analog cameras can be replaced with digital cameras, with the current improvements in digital camera frame rates, sensitivities, resolution, etc. Film cameras can also be replaced while significantly reducing environmental impacts, reducing costs and improving performance. Multi-wavelength imaging is supported with the introduction of digital cameras. A significant reduction in production time is possible, without the need to process and "sneaker net" film. With this recognition however, comes a new set of challenges including transport, short and long term storage of the digital imagery data. The transport of film replacement data can easily result in data rates of 130 Mbits/sec to over 1 GByte/sec. Although short, tens of meter fiber runs of GB/sec data rates are being performed over interfaces such as Very Short Reach (VSR) parallel optics OC-192/STM-64 interface, optimized for network intra-PoP interconnections. This paper discusses data transfer rates of up to Gigabit Ethernet bandwidths. Transfer efficiencies as high as 93 percent have been achieved over various network topologies. This includes combinations of fast and Gigabit Ethernet networks and OC-12 through OC-192 optical backbones.

The utilization of digital cameras may lead to the belief that transporting data via a network would be more straightforward than via analog systems. However, this is not the case. The lack of standards and industry immaturity in this area requires additional effort. Digital camera output formats vary significantly and include such formats as American

National Standard ANSI/TIA/EIA-422-B (formerly RS-422), Low-voltage differential signaling (LVDS), Camera Link, etc. A method for converting these electrical formats into a packetized format is required. By utilizing custom software, this has been accomplished and is discussed herein along with details of the network configuration utilized for transport.

2 BACKGROUND

The two main commercial approaches for streaming video use hardware and software; this includes the streaming of signals from analog sources such as cameras, VCR, and the streaming of computer desktops. Dedicated hardware solutions are costly, proprietary, inflexible, and do not support computer desktops or non-standard video formats such as Camera Link. Software solutions have a high latency, which is not a concern for many industries. In fact, a delay is often required to allow images (such as half-time show wardrobe malfunction) to be blocked prior to transmission. However, sub-second latency is required for Net Centric T&E ranges, where personnel are making decisions in real-time. Less than 200 ms or even tens of ms latency requirements exist when using the data with tracking and decision-making hardware.

Hardware streaming methods perform compression in custom or programmable application-specific integrated circuit (ASICS). In almost all cases, the systems are based on the Moving Picture Experts Group (MPEG) 2 compression algorithm. When a low-latency requirement is invoked on the hardware vendor, additional costs are typically incurred. Most of the time, the hardware solutions support broadcast formats only and do not provide mechanisms for interfacing to machine vision type interfaces, whether they are standard, such as Camera Link, or non-standard, such as parallel low-voltage differential signaling (LVDS). Traditionally, high-speed A/D converters have been designed with parallel CMOS digital outputs. However, as applications demand faster speeds, higher resolution and smaller form factors, CMOS outputs have become a major constraint on chip designers, as well as system engineers. In order to break through the CMOS output driver limitations and meet customer demands, LVDS drivers are integrated into high-speed converters, signaling a new era in the A/D converter market.

Access to the video data software solutions running on a Personal Computer (PC) must interface to some form of frame grabber or PC interface port such as Universal Serial Bus (USB) or Institute of Electrical & Electronics Engineers (IEEE) 1394. Commercial based streaming solutions typically support only Video for Windows (VFW)/Directshow compliant frame grabbers and as a result only support standard broadcast or web-cam type formats. Windows based support for machine vision interfaces is made possible through a proprietary frame grabber and software interfaces. As a result, it is difficult for streaming vendors to support this industry because of the lack of a standard driver interface and Application Program Interface (API). To provide the interface to the non-standard and typically unsupported video formats the streaming video software developed at WSMR uses the Matrox family of frame grabbers. Commercial, off-the-shelf software encoders do not support the low latency requirements of the Net Centric T&E community. To accomplish the streaming of data from a variety of cameras and formats,

which support low latency, a custom software application had to be written. This is briefly discussed in the paper.

The key issue in accomplishing high bandwidth or low latency streaming solutions is determining who "owns" the pipe. A high bandwidth, uncompressed screaming video demonstration was possible because of dedicated ownership of the backbone. Total bandwidth availability is usually limited when multiple users are on the same pipe. Coordination with numerous network administrators is required when traversing cross-country over the Defense Research Engineering Network (DREN), even when minimal bandwidth is utilized (2-10Mbits).

3 REQUIREMENTS

T&E environment streaming video performance requirements are driven by numerous factors. Compressed and uncompressed video factors are described below.

The T&E community of interest, National Aeronautics and Space Administration (NASA) and others use a variety of optical instruments to track and image a variety of low and high dynamic targets. The instruments rotate in both azimuth and elevation. The instruments have the ability to acquire and track targets using a variety of methods including pointing data from radar, telemetry, Global Positioning System (GPS), and other systems, as well as manual tracking of the object. When a target is temporarily "lost" due to loss in external data or loss of video track, the operator must quickly engage manual track until a re-lock can occur. Operators may be 10's of miles away from the tracking instrument due to safety reasons. The manual tracking method necessitates a man-in-the-loop in the servo control system tracking instrumentation.

Tests were performed at WSMR to determine threshold tracking performance degrades as the latency of the image to the human operator increases. To determine a realistic requirement for a man-in-the-loop tracking latency requirement, a software application was developed that allows an evaluator the ability to "dial-in" a latency for video that is being displayed to an operator that manually tracks representative objects of interest with a T&E tracking instrument. The software utilizes a large ring buffer to allow for setting latencies in increments of video frames. The range of latency tuning is approximately 90 milliseconds up to about one second. During testing, a number of experienced operators were asked to track targets. The evaluator varied the latency without the operator knowing the settings. The evaluator monitored the tracking errors and phase lag at the various latency settings. While the tracking performance is highly dependent on the target dynamics, the final results showed a noticeable degradation at 300 milliseconds, reasonable performance at 250 milliseconds, and an acceptable minimal degradation at 200 milliseconds. As a result, to eliminate the prospect of decreasing the overall performance of man-in-the-loop tracking operations with the switch to digital technologies, the metric of 200 milliseconds was chosen. It needs to be noted that while 200 milliseconds was chosen as a benchmark, a lower latency may be required for certain challenging scenarios.

In addition to the tracking requirement, the operator also requires observation videos when operating the instrument remotely. These videos are used for verifying equipment operation, scanning the area for weather issues, etc. Observation videos require less

stringent latency performance but do require good quality video. A “blocky” streaming video that drops numerous frames is not adequate for this scenario.

Due to a variety of reasons including safety, security, efficiency, and costs, it is desirable to minimize the hardware in the field required to operate a tracking instrument. At WSMR, support vans and tracking instrumentation with very expensive electronics equipment are driven into the field to remote locations in order to image test hardware. This setup significantly adds to operation and maintenance cost as well as to the workload of the operators. Driving to a remote location at 2:00 a.m. while encountering a variety of desert and road hazards can result in operator overload and procedure fatigue failure (PFF), and consequently, misconfigured systems and data loss. To the extent practicable, limiting hardware in the field and “remoting” it to centralized or main campus locations can greatly improve these issues. The current method of optics operation at WSMR is to locate a master control van in a centralized area which provides remote operation capabilities. Redundant recording hardware of analog video signals is located in the master control van. However, the main data product is recorded at the optics instrumentation site including film and high speed digital data as well as analog recordings. The emergence of digital cameras and the Test Support Network (TSN) with high bandwidth capacities and robust commercial telecommunications equipment initiates a framework that will allow for the removal of large amounts of equipment from the field. High speed data can be locally buffered and then transferred long haul via fiber for recording. The term Locally Buffered - Remote Recorded (LB-RR) is a phrase for the future. The high frame rate and high resolution recordings are currently buffered in the camera’s RAM during operation and are later downloaded to a recording device. High-speed RAM and Hard Drive digital recorders temporarily record GB/sec data, which is then downloaded to storage area network (SAN) systems for reduction and archival. The 18-200 MB/sec type data is locally buffered and downloaded in real-time for remote recording. The uncompressed, high bandwidth test described herein is an example of the feasibility of the long-term goal of LB-RR.

The scene content experienced during T&E operations was extremely varied. Tracking a low, fast target against a high contrast mountainous background resulted in high frequency content changes with an image. Each frame of video was different and compression schemes that rely on multiple frames and buffering techniques were broken down quickly. Motion artifacts quickly destroyed the quality of the delivered video product during these types of scenarios. Success in handling this problem was achieved with single frame compression techniques and advanced full-frame wavelet based codec’s such as MPEG IV. This approach is discussed in Section 5 herein.

Optical video trackers are often engaged, either manually or with an external queue from radar, when an object is tracked. Depending on the scenario, a tight closed loop track is required. Although not practical in the near future, except where dedicated, non-network, fiber exists, the possibility of remoting the video tracker has been discussed. The use of extremely low latencies in the ms range is required. The capability to accomplish this with a streaming source utilizing commercial telecommunications is not currently a requirement, but a future goal.

Many of the videos acquired in the field, must be forwarded to a variety of locations for real-time viewing. This is used at WSMR or at cross-country locations. This requirement necessitates support for both unicast and multicast capabilities. In addition, recast of a stream down an independent network is often required.

To accomplish the requirements discussed above, different technologies are implemented and are further discussed in Sections 4 and 5.

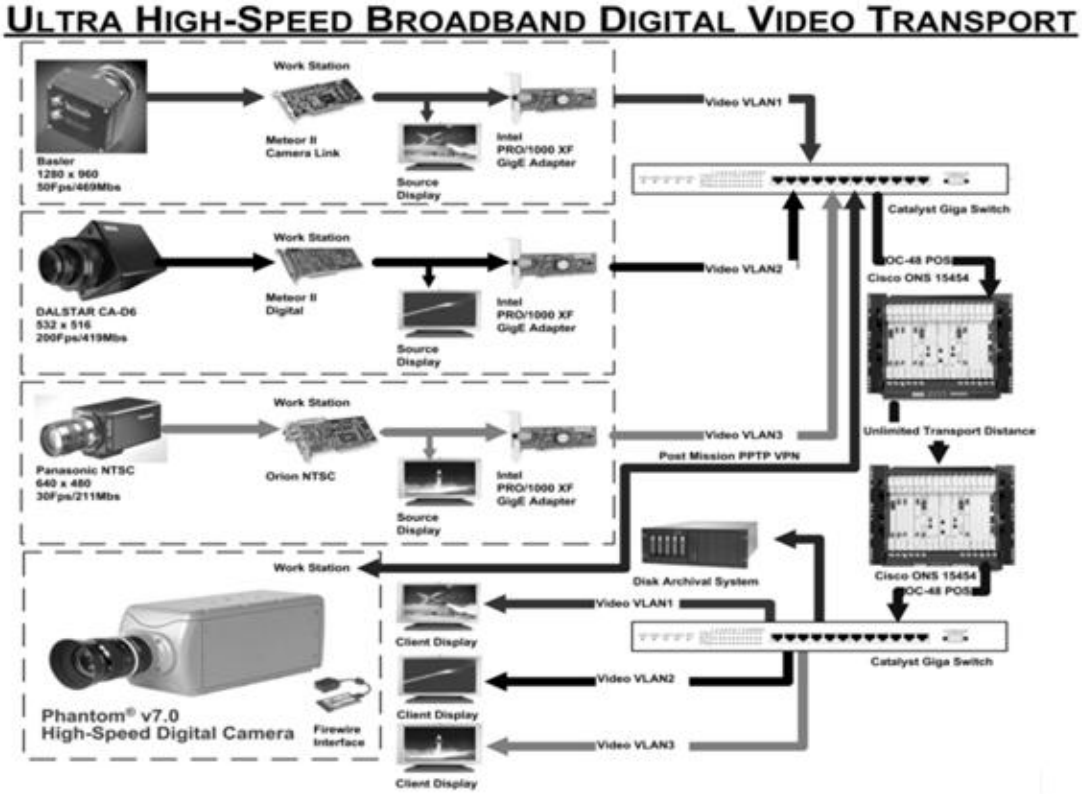

Figure 1: High speed uncompressed video transmission demonstration.

4 NETWORK CONFIGURATION

A critical component of a streaming video solution is obviously the network. Through both insight, and trial and error, several configurations that would consistently support high speed, low latency video transport were successfully utilized. To ensure low latency performance, similarities in the network are required whether uncompressed or compressed video is transported. Figure 1 depicts a setup successfully utilized to transport uncompressed high bandwidth data. The total aggregate data rate in this configuration, which was sustained for continuous nine-hour periods, was approximately 1.2 Gbits/sec. In this configuration, data from three different cameras was captured and transported across a GigE network in an uncompressed format in real-time. The cameras were running from 30 to 200 frames per second. A fourth camera that stored up to 500 frames/second of data in local RAM was also streamed post mission over a Point to Point Tunneling Protocol (PPTP) Virtual Private Network (VPN) configured in the same network. A key component of this network was the establishment of separate “video” VLANS for each camera stream. At the receiving end of the video transport were local streaming client systems that could both capture the data to hard disk and display the streamed video in real-time.

Superior quality streaming video depends on maintaining strict constraints for packet loss, delay, and jitter (contrast this with traditional data network traffic, where the focus is on response time or throughput). We have identified a variety of techniques that were utilized to successfully perform the screaming video and streaming video during both

demonstrations and in support of live missions. The following design features summarize the techniques utilized in the Screaming / Streaming Video data network.

Packet Loss: The packet loss should be kept well below one percent and bursts of consecutive lost packets must be avoided. Packet loss occurs due to congestion or electromagnetic noise. It can also occur when jitter is high and the jitter buffer is too small to compensate. Increased bandwidth and good tuning can often reduce network congestion, which in turn, reduces jitter and packet loss.

Priority Packet II: We utilized Intel Pro 1000 network interface cards for our un-compressed streaming configurations and a variety of Network Interface Card (NICs) for the compressed configurations. When streaming uncompressed, the Intel Priority Packet II utility was utilized to set up priority filters to process high priority network traffic before normal traffic. When using Priority Packet II, filters can be configured to give priority to time-critical traffic. Filters can operate in a global or network component specific manner. Global filters operate on the system as a whole, while network-component-specific filters apply only to chosen network components such as individual adapters, VLANs, or adapter teams. Priority Packet II prioritizes traffic based on priority filters, which contain parameters that are assigned by the user. Priority Packet II provides the following methods for prioritizing traffic: IEEE 802.1p Tagging; IP ToS (Type of Service) Layer 3 Tagging; Support both Differentiated Services Code point (DSCP) and Legacy IP Precedence; and Intel High Priority Queue. When using the Priority Filter Wizard, filters are set up through templates to define the following traffic filter types: Node (MAC) address/Ether type and TCP/IP information. Priority Packet comes with pre-defined priority filter templates, and can be modified at the discretion of the end-user. By prioritizing traffic at the hosts or entry point of the network, network devices can base forwarding decisions on priority information defined in the packet.

IEEE 802.1p Tagging: During operation IEEE 802.1p tagging was implemented within the Cisco Catalyst 10/100/1000 3500 Series Ethernet switches. IEEE 802.1p is an IEEE standard for tagging, adding additional bytes of information to packets with different priority levels. Packets are tagged with four additional bytes, which indicate a priority level. Tags also increase the packet size. When these packets are sent out on the network, the higher priority packets are transferred first by IEEE 802.1p-aware devices. Priority packet tagging (also known as Traffic Class Expediting) allows the adapter to work with other components of the network (switches, routers) to deliver packets based on priority level. Using Priority Packet, 802.1p tagging lets you assign specific priority levels from 0 (low) to 7 (high). Since IEEE 802.1p tagging increases the size of the packets, some hubs and switches will not recognize the larger packets because they exceed the maximum frame size of standard Ethernet packets and will drop them. If these devices do not support 802.1p, High Priority Queuing can be used to prioritize network traffic.

Packet-Loss Concealment (PLC): PLC masks the loss of a packet or two by using information from the last good packet. Packet loss can occur randomly or in bursts. PLC helps with random packet loss. The cost for doing PLC is minimal, since it is usually part of the codec processing.

One Way Delay: When working with private networks or coast-to-coast “leased” networks, it is possible to keep the one way delay well below 100msec for point-to-point connections. One-way delay is equivalent to the sum of the propagation delay, the transport delay, the packetization delay, and the jitter buffer delay. Video quality degrades quickly when the total one-way delay is greater than 150ms. When traffic is required to cover long distances, the network path should be kept as direct as possible. In addition, static routes should be used on the router access control list. Transport delay is the total time spent inside each of the devices in the network, such as switches, routers, gateways, traffic shapers, and firewalls. Some devices add more latency than others. For

example, a software firewall running on a slow PC adds more delay than a dedicated hardware-based firewall. A variety of techniques exist that allow the user to evaluate the number of hops traveled by the video traffic. This information can be used to reduce the number of hops if possible. In addition, excessive latency in individual devices should also receive attention. Packetization delay is the fixed time needed for the codec to do its job. The DIVX codec is regularly used for compressed streaming operations due to its excellent quality and small packetization delay. The Jitter buffer delay is used to dampen variations in packet arrival rates. If the network delay is low and the jitter is high, one can afford to have a larger jitter buffer than in a network where the delay is already high. Several other techniques used in the configuration tuning include increasing the Receive/Transmit and coalesce buffer size and to disable the Receive/Transmit Checksums.

Multicast or Unicast: Both multicast and unicast transmission have been utilized during testing and normal operation. Multicast can be implemented at both the data-link layer and the network layer. For scenarios where multiple clients are viewing the same streams, multicast is a more efficient transport mechanism than unicast. Situations where multicast may not be available across the entire network, the software allows for sending a unicast transmission between two locations, which allows the receiving client to recast the stream as a multicast transmission to other viewers that reside on the same multicast enabled network.

Hardware: The compressed streaming application has been successfully used with a variety of commercial vendor's hardware including Lucent, Calix and Cisco. The screaming, un-compressed demonstration was performed utilizing Cisco 15454 Optical Transport Equipment. We used two, OC-12, 622 Mb Video VLANs. The virtual ports are scalable in size from SONET STS-1 (~50Mbps) or SDH STM-1 (~155Mbps) to STS-24c/VC-4-8c (~1.25 Gbps), or an aggregate of 2.5 Gbps total virtual port bandwidth per card. The payload from a client interface is mapped directly to a virtual port. The virtual ports are then cross-connected to the system's optical interfaces (from 155 Mbps to 10 Gbps) for transport, along with other services, to other network elements.

There are a variety of toolsets and methods utilized by network administrators to verify the quality of the system. Discussion of all the methods and tools utilized for evaluating network performance is beyond the scope of this paper. A few of the toolsets utilized are listed for reader convenience: The Simple Server Redundancy Protocol (SSRP) communicates between the router and Asynchronous Transfer Mode (ATM) switch groups to detect device failures and initiate automatic reroutes within one second of failover time; the Service Management Solution is used for managing and monitoring service-level agreements on the DREN; the VPN/Security Management Solution is used for optimizing VPN performance and security administration; Quality of Service (QoS) features provide value-added functionality such as network-based application recognition (NBAR) for classifying traffic on an application basis; a service assurance agent (SAA) is used for end-to-end QoS measurements, and Resource Reservation Protocol (RSVP) signaling for admission control and reservation of resources; Net Support Management tools are used for both servers and clients along with the advanced LAN Management Solution for Cisco Catalyst multilayer switches.

One final potential failure point relates to the data generating source. In the case of software based streaming systems running on Windows PCs, miscellaneous housekeeping services can affect performance. Extra services should be shut down to ensure maximum performance, especially when transmitting and receiving high bandwidth uncompressed streams. Figure 2 shows the Windows 2000 services that were running during high speed data transfer operations.

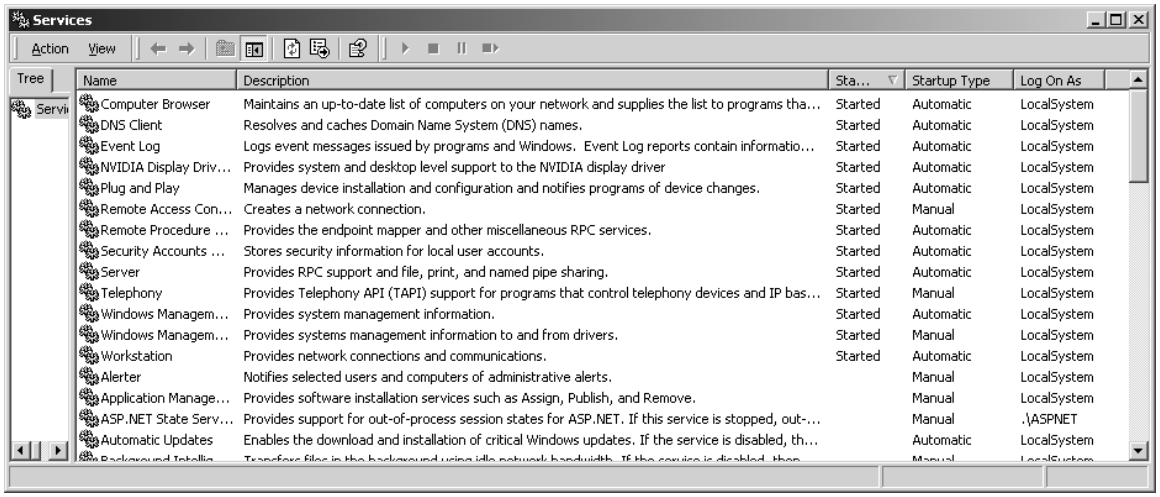

Figure 2: Configuration of Windows services for high speed video transfer tests.

5 WSMR SCREAMING / STREAMING VIDEO

Software applications developed at WSMR, and are used to accomplish the initial streaming design goals include two streaming applications and one client, or viewing application. Of the two streaming applications, one streams live video via an interface to Matrox frame grabbers, while the other streams all or a portion of Windows 2000/XP desktop. The client application is used to view streaming content from either of the two source applications.

5.1 *Hardware and software configuration*

The video streaming server consists of a high-end Windows 2000/XP system and a Matrox frame grabber. Multiple versions of Matrox frame grabbers are supported, including the MeteorII family, the Genesis, the Orion, and the CronosPlus. Support for the Odyssey and Helios systems is currently in process. With this support, a wide variety of video formats are supported, including RS-170, NTSC, RS-422, LVDS, and Camera Link. These interfaces have been tested and confirmed. A typical configuration for a streaming server utilizes a ruggedized rack-mount Windows server. To stream compressed full resolution NTSC video at 30 Hz, a CPU speed of at least 2.4 GHz is required. In general, a dual 3.06 GHz or higher CPU motherboard is utilized to provide as much overhead and flexibility as possible. Memory bandwidth is also a key parameter in determining the performance of the system when using compression. When performing uncompressed “screaming” video transfer, CPU requirements are actually quite reasonable, with a 1 GHz Pentium III system capable of streaming and receiving 500+ Mbit/second streams.

The software architecture of the streaming video applications is based on C/C++ for the underlying functionality and Microsoft Foundation Classes (MFC) for the GUI interface. Development is performed within the Microsoft Visual Studio v6.0. The Matrox Imaging Library (MIL-Lite) is utilized for frame grabber access and the Intel optimizing C/C++ compiler and VTune is utilized for generating efficient executables.

When utilizing compression, the server and client applications both rely on the Microsoft Windows Video for Windows (VFW) programming API. The VFW interface is replaced with the DirectX/DirectShow interface, which allows more flexibility and higher performance that the older VFW interfaces. Utilization of the DirectX/DirectShow interface is currently a planned upgrade to the system. By utilizing a standard Windows multimedia API, access to a wide variety of both standard and proprietary compression codec’s is made available. In the current application, any VFW compliant codec can be utilized if it has been installed on both the video server and client systems. After evaluating a majority of the codec which are common and readily available, the Digital Video Express (DIVX) codec, from DIVX Networks, was chosen as the default compression engine for the streaming servers. Testing in the lab showed the DIVX codec to have a superior combination of both very good image quality and CPU efficiency. These two traits are critical to developing a successful software-based video streaming system.

5.2 *Programming methodology*

The general programming philosophy for WSMR streaming video focuses on the concepts of simplicity, flexibility, and efficiency. First and foremost, the system is designed for "real-time" viewing. As a result, if a packet or frame of video is dropped, it is forgotten. The process of trying to recover a lost packet via retransmission does not work in a low latency streaming environment. Certain commercial systems have employed elaborate buffering schemes in an attempt to guarantee that every frame arrives and is displayed. In a noisy or high traffic network environment, this scheme usually gets bogged down and the user is left staring at a re-buffering message of some type. In a clean network environment, the system performs well, but with long latency and a large amount of unnecessary overhead.

The software architecture is centered on three main C++ classes that perform the functions of video capture, video compression and decompression, and Ethernet transmission and reception. The construction and format of these classes make the process of embedding a streaming capability into existing applications quick and simple. The software has a multithreaded, pipelined architecture with an emphasis on simplicity and efficiency, as shown in Figure 3. When a user initiates a streaming operation, the software configures the frame grabber and a number of ring buffers. Multi-stage ring buffer architecture is used for frame grabbing, reformatting (if necessary), compression (if necessary), packetization and finally transmission and local display. Event signals are used to hand-off buffer pointers from one function to the next. The initial process involves grabbing frames via the frame grabber. This process is handled with a set of callback functions that hook to the frame grabber start-of-frame event and to the end-of-frame event. During the start-of-frame callback, a buffer index is incremented and a command is issued to grab the next frame. In the end-of-frame callback, an event trigger is set to notify a waiting worker thread that a buffer is available. This process continues until a series of packets are finally transmitted over the Ethernet. The process of grabbing frames is continued until a flag is set via user GUI menu selection.

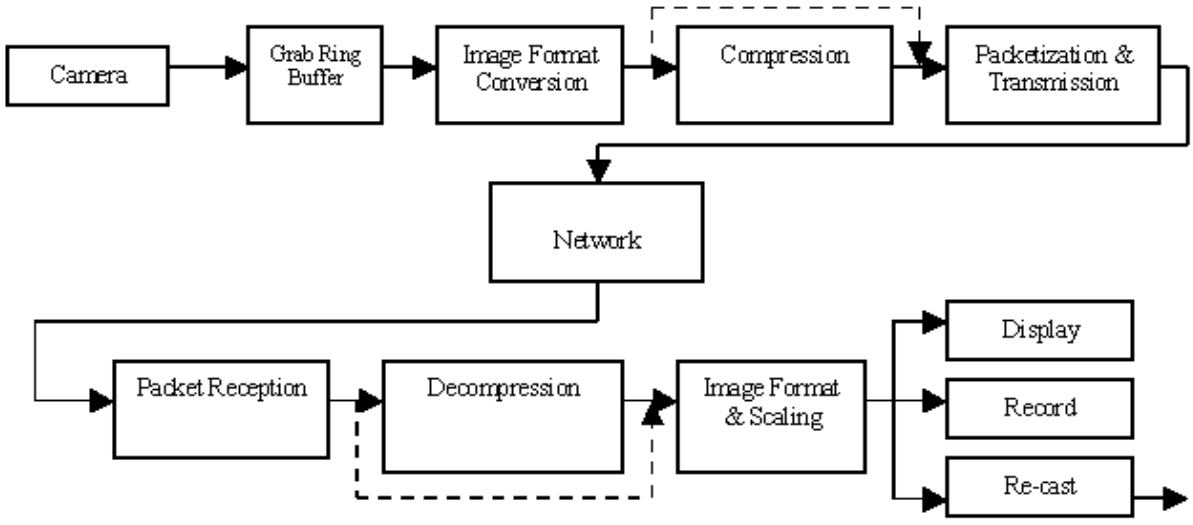

Figure 3: Streaming video data flow.

Due to a requirement for extremely low latency transmission, frames of video are treated independently. When transmitting compressed video, a frame is received, compressed, and transmitted independently of neighboring frames. This approach has a benefit that motion in the video scene will not cause any artifacts in the video stream when decompression is performed. In addition, if a frame or packet is dropped, neighboring frames are not affected. The tradeoff with this approach is that maximum compression efficiency is not achieved. However, with the Motion Picture Experts Group (MPEG)4 based codec, such as DIVX, the stream bandwidths that are achieved are still lower than that of comparable quality MPEG2 streams. In many of the scenarios where the streaming is used at WSMR, motion artifacts have traditionally been a key problem when testing and evaluating MPEG2 based systems. While testing a commercial MPEG2 system using a previously recorded video tape, a small object of interest in the field-of-view actually disappeared.

When non-compressed video is transmitted, a mechanism is used to maintain efficiency that relies on uniform packet sizes. When the streaming configuration is performed in the software, the video format, including the width, height, and pixel format, is analyzed to determine an optimum packet size for the video transmission. The packet size is selected to be uniform, if possible, in that 'n' number of packets equals one frame of video. For instance, with NTSC video the format is 640x480x24 or 7,372,800 bits per frame of video. With User Datagram Protocol (UDP) Ethernet packet transfers, the maximum packet transmission size is 64,494 bytes or 515,952 bits. As a result, the image cannot be sent as one single packet, but must be sent as multiple independent packets. If the image is broken into 15 equally sized packets, the resulting packet size will be 61,440 bytes. This maintains efficient Ethernet transmission, while allowing for an efficient buffering and pointer scheme in the processing and formatting stages of the software. In the event that a particular packet is dropped with an uncompressed transmission, the dropped packet is detected via packet header information and the packet is replaced with the last known good packet from that location in the image frame. Again, the architecture is geared toward achieving the best image quality with a hard low latency requirement.

With respect to compressed video streams, uniform packet sizes are not maintained because of the constant variability in the size of compressed frames. In general, compressed frames from standard video sources fit within one UDP packet. This is not guaranteed, but necessary, if a mechanism is in place which divides the compressed frame into multiple packets. These packets are reassembled on the client side before any decompression operations are performed. This feature is heavily utilized, for instance, when streaming desktops at 1600x1200 resolutions.

When a packet of video is received in the client software, the reverse process of that which occurred in the server software takes place. A thread waits for incoming packets and loads them into a local buffer when they arrive. The packets are then decompressed (if necessary) and loaded directly into a buffer that is made ready for display. This streamlined process is very efficient and supports very low latency operation.

5.3 *Streaming Server Application*

The WSMR video streaming server application is built as an MFC dialog application, using Visual Studio v6.0. The application auto-detects Matrox frame grabbers and allows the user to select different video formats via cameras that may be attached to the frame grabber. The software supports variable resolutions, frame rates, color formats, and bit-depths. The interface allows the user to select network transmission parameters, such as

time-to-live, port number, and IP address. The user can select unicast or multicast addresses, with the software automatically detecting in which mode to operate. In unicast mode, the transmission is sent to one specific machine, similar to a phone call. In multicast mode, the transmission is made available to any machine on the network, similar to a radio broadcast. To support systems that are CPU challenged, the software supports modes that utilize less CPU for the software compression. These include sub-sampling and binning in both frame rate and resolution. In addition, a smoothing filter is available for reducing the effective bit-rate of the transmitted signal. This is useful in wireless Ethernet environments where the total network bandwidth is roughly equivalent to a single video stream bandwidth. The application is CPU intensive; therefore, a throttle mode is available that steps through the various binning options in an attempt to maintain an overall CPU utilization of less than 80 percent. A CPU monitor, similar to that which is available through the Windows Task Manager, provides the user with a continuous update of the CPU utilization.

5.4 *Streaming Client Application*

The WSMR streaming video client software provides a user with three main functions of viewing, recording, and recasting/transmitting an incoming video stream. The recast feature is very useful when long haul transmission is required with multiple viewers at a remote location. In many cases, all of the network gear from point A to point B may not be multicast enabled. In this scenario, an initial UDP unicast can be performed to transmit the video from one facility to another; a machine at the destination end can then recast the video as a multicast transmission on that local network, allowing multiple simultaneous viewers to see the video. This process of recasting, or daisy chaining, can be carried on as many times as the user requires. The client software also allows for local recording of the video stream to an Audio Video Interleave (AVI) file on a client hard drive. The software also provides a mechanism for sub sampling frames on the incoming video stream. For instance, the software can be configured to decompress and display incoming 30 Hz video at 10 Hz, 5 Hz, 1 Hz, etc. This is useful on underpowered machines or on systems that are running other CPU intensive applications, such as 3D graphics.

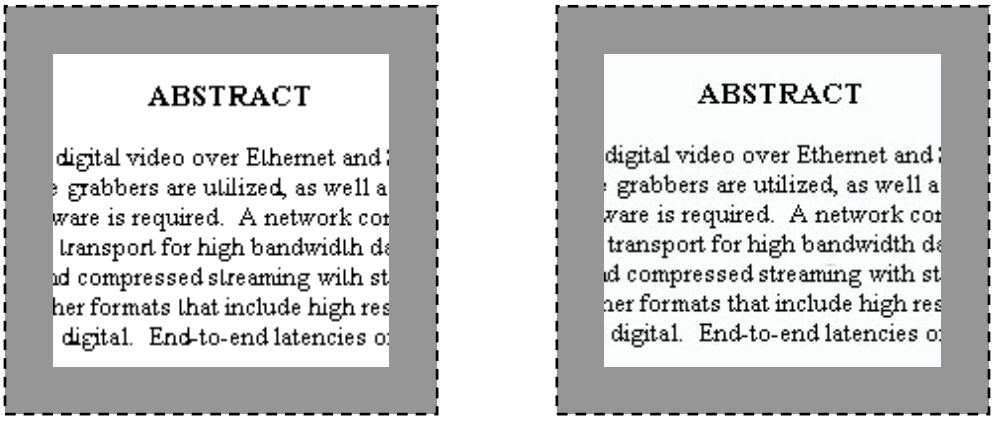

Figure 4: Comparison of source desktop content on left and post DIVX compressed and streamed version of the content on the right.

5.5 *Desktop Streaming*

The final application in the WSMR suite of streaming software is an application that provides the ability to compress and stream all or portions of a Windows desktop. The software captures a snapshot of the desktop graphics in memory, then compresses and transmits the image as if it were from a video source. The video client software is used to view the desktop streams. Any application that runs on the Windows desktop can be streamed in this manner. Resolutions up to 1600x1200 have been streamed over long haul networks. At the 1600x1200 resolution, frame rates of 9 Hz have been achieved on dual 3.06 GHz Pentium P4 systems. The DIVX codec that has been used to date has performed well with a variety of desktop applications. With PowerPoint slides, small text is readable on the client end, as shown with an example labeled Figure 4. Satellite imagery has also been transmitted successfully with minimal compression artifacts.

6 TESTING & MEASUREMENTS

To provide a quantitative measure of the capabilities and performance of the streaming video system, a number of tests were performed using both uncompressed and compressed data. Table 1 summarizes the variety of configurations that have been successfully demonstrated over long haul networks.

To measure the latency of the streaming video, a test suite was built that measures the true glass-to-glass latency of the system. An external measurement system using a series of LED lights driven by a Trak Microwave Model 9000 GPS timing clock and a Dalsa CA-D6-512 high frame rate digital camera allowed for capturing the total latency of the system. A Panasonic WV-CL350 NTSC camera was configured as the video source with a short exposure of less than 1 millisecond. It was pointed at the timing lights, which have a fast decay time and encode timing information from days down to microseconds. The captured video is compressed and transmitted to a second machine, where the resulting video is displayed on a flat panel monitor. This monitor is placed next to the original timing lights. The high speed Dalsa camera is configured such that the flat panel monitors and the source timing lights are in its field of view. The high speed video is then captured to disk using a digital disk recorder. After this data has been captured, individual frames can be analyzed to compare the timing information contained in the source LED lights and that contained in the streaming client display on the flat panel monitor. The accuracy of the test runs is limited by a 5 millisecond exposure setting of the high speed camera. This setting was required to achieve an adequate exposure of the client display on the flat panel monitor. The 5 millisecond limitation was considered adequate for these measurements.

Source	Compression*	Resolution	Bit Depth	Frame Rate	Bit Rate
NTSC	None	640x480	24 bit color	30 Hz	211 Mbits/second
Dalsa CA-D6-512	None	532x516	8 bit mono	220 Hz	461 Mbits/second
Basler 501k	None	1280x960	8 bit mono	66 Hz	619 Mbits/second
Indigo Phoenix IR	None	640x512	14 bit mono	30 Hz	131 Mbits/second
NTSC	DIVX	640x480	24 bit color	30 Hz	< 8 Mbits/second
Basler 501k	DIVX	1280x1024	8 bit mono	10 Hz	< 10 Mbits/second
Indigo Phoenix IR	DIVX	640x512	14 bit mono	30 Hz	< 8 Mbits/second
Windows Desktop	DIVX	1600x1200	24 bit color	5 Hz	< 15 Mbits/second

Table 1: Results of streaming tests

* The uncompressed streaming tests were performed with dual 2.0 GHz Pentium P4 systems running Windows 2000. The compressed streaming tests were performed with dual 3.06 GHz Pentium P4 systems running Windows 2000.

Latency measurements were obtained for compressed NTSC transmissions at the request of an end-user of the system. The design requirement for the system was 200 milliseconds, as stated earlier. The measured latency for the test configuration was approximately 80 milliseconds for local network operations. Figure 5 shows a frame from a round trip latency test from WSMR in New Mexico to Ft. Belvoir in Virginia. The recast feature was used to send the data from Virginia back to the lab in New Mexico. The measured latency in this frame was approximately 190 milliseconds. The third row of lights in the image contained BCD encoded hundreds, tens and single millisecond units. The source lights showed a time of 193 msec in the third row (the fourth light from the right was stuck on) and the same lights in the client display show either 001 or 009 msec (again, the stuck light). The major times in the first two lines are identical.

To test the robustness of the technology in a realistic network environment, systems were configured at WSMR, in New Mexico, Ft. Belvoir, in Virginia, and Ft. Rucker, in Alabama. NTSC and infrared RS-170 video was configured and transmitted from base to base on a 24/7 basis for a period of two straight weeks. No failures or anomalies were reported during the tests. In addition, 1600x1200 desktops containing both Microsoft PowerPoint slides and satellite imagery were successfully transmitted from WSMR to Ft. Belvoir.

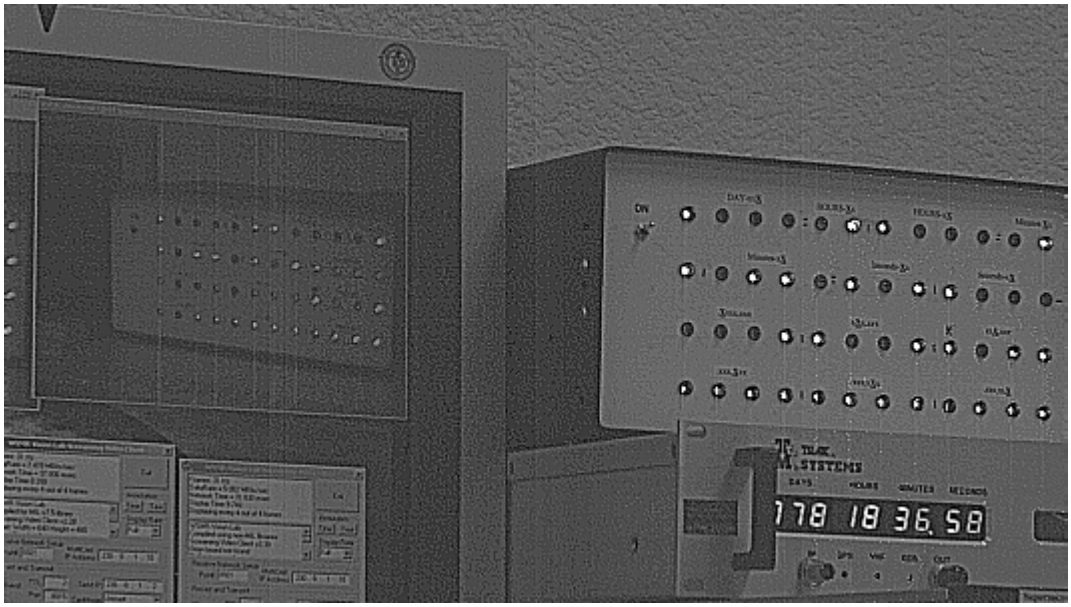

Figure 5: High speed image frame shows ~190 millisecond round trip cross country latency.

Figure 1 shows a block diagram of a configuration that was used to test a full up configuration of an uncompressed high bandwidth scenario. The test involved live streaming of three video sources totaling about 1.2 GBits/second. These sources included an NTSC source at 211 MBits/second, a Dalsa CA-D6-512 running at 419 MBits/second

and a Balser 501k running at 469 MBits/second. Independent VLANS were configured for each channel and multiple Gigabit Ethernet pipes were fed into an OC-192 backbone via Cisco 15454 optical transports. On the receiving end, the data was broken back out onto the Gigabit Ethernet networks and received by independent machines for each channel. This configuration was operated continuously over 9 hour periods and successfully used on a live mission to view and record the data at the client or remote location. Within this same configuration, an additional test was performed to show the utility in remotely downloading and archiving data from a very high speed RAM based camera. In this test 4 GBytes of data from a Vision Research Phantom camera was downloaded after the mission remotely over the same pipe. The download time was improved by a factor of about 4X, when compared to the existing approach, which utilizes a portable notebook computer to download data from the camera. The difference in download times was attributed to the relatively slow notebook hard drive versus the higher end Ultra160 Seagate Cheetah stripe array that was used in the remote download test.

7 VIDEO Furnace Bridge

There were lessons learned while configuring the Lakehurst/WSMR Bridge from the Lakehurst standpoint. When traversing the Netscreen Firewall and Tipping Point Intrusion Protection System (IPS) we were cautious of straying from our standard rule sets. The firewall rules were configured with very specific source and destination addresses and were limited to only those ports necessary to pass the bridge traffic successfully. We experienced issues when straying from standard IPS rule sets. The nature of the IPS is to protect the internal network from traffic that appears damaging or out of the normal baseline for a particular network. The bridge between Lakehurst and WSMR uses non-standard ports and a fairly large stream of bandwidth. The IPS was not familiar with the traffic initially and identified it as a potential attack thus denying the bridge traffic. Once the offending rules were identified the appropriate adjustments were made allowing the IPS to establish a baseline for the new traffic pattern, thus eliminating the issue.

With packets flowing at optimum size, it became apparent that stream management was an issue for where and how the video streams were to traverse the disparate networks. For instance, a multicast stream was produced at WSMR and needed to be displayed at the Lakehurst facility. Video Furnace was challenged to develop a methodology that could bridge the two facilities so they could enjoin the video streams. With the requirements defined, Video Furnace created a model software subsystem named VFBridge. The objective of this was to take any two points that needed to send/receive the video streams and dynamically bridge the streams to the necessary protocol wrappers for network acceptance. The first iteration was a proof of concept by creating two VFBridges, one for the sending network and one for the receiving network. These modules would “handshake” on a given set of rules for each location. The WMSR multicast stream must translate to a unicast stream for the trip to Lakehurst, then re-instantiated as a multicast stream with the proper addresses for the local network at Lakehurst. Proof of concept model was successful but cumbersome to configure and extremely inefficient with bandwidth as each stream was sent whether it was requested or not. The addresses of streams were static and not dynamically configurable, and a pair of bridges was needed at both locations. However, the bridge subsystem did work.

With the bridge traversal a success, Video Furnace desired to simplify the process and make it more dynamic. It was specified that the ideal bridge would be designed to be less complex and more bandwidth efficient. With these specifications, Video Furnace advanced the VFBridge subsystem to account for a scalable and dynamic environment. Currently, the subsystem only requires one pair of bridges for “n” channels. Data now flows only if the stream is requested, thus reducing bandwidth usage to demand only.

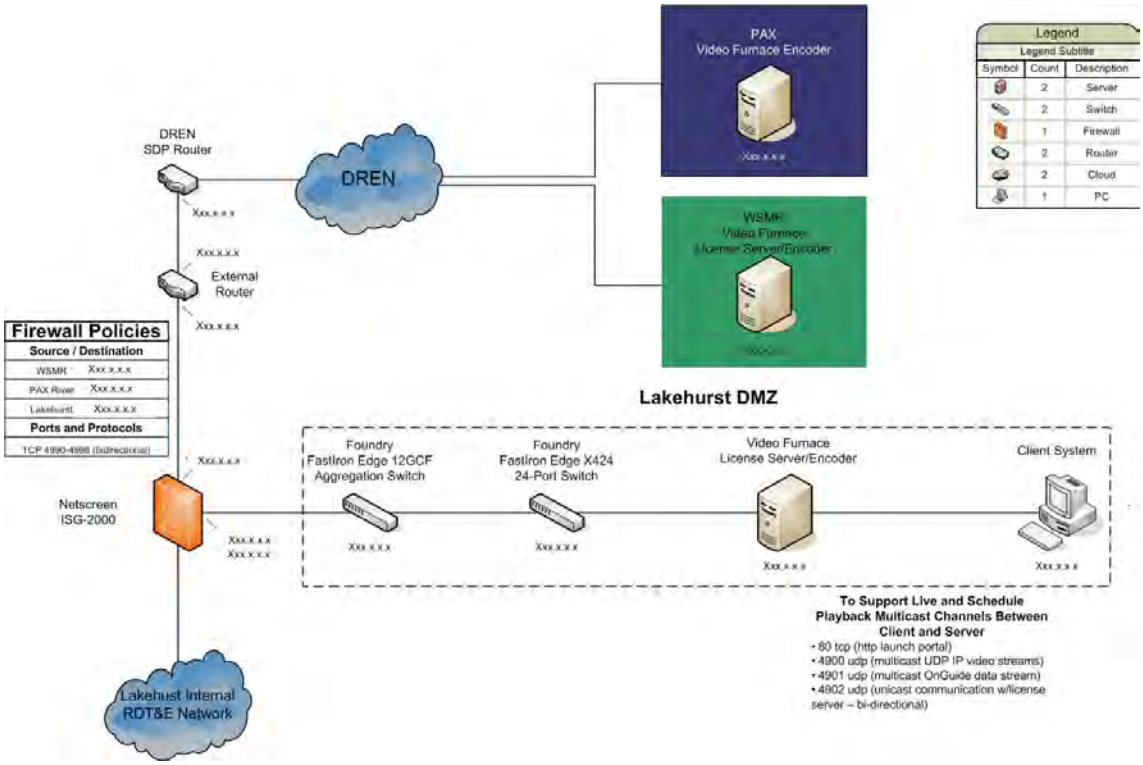

Figure 5: DREN Test Configuration

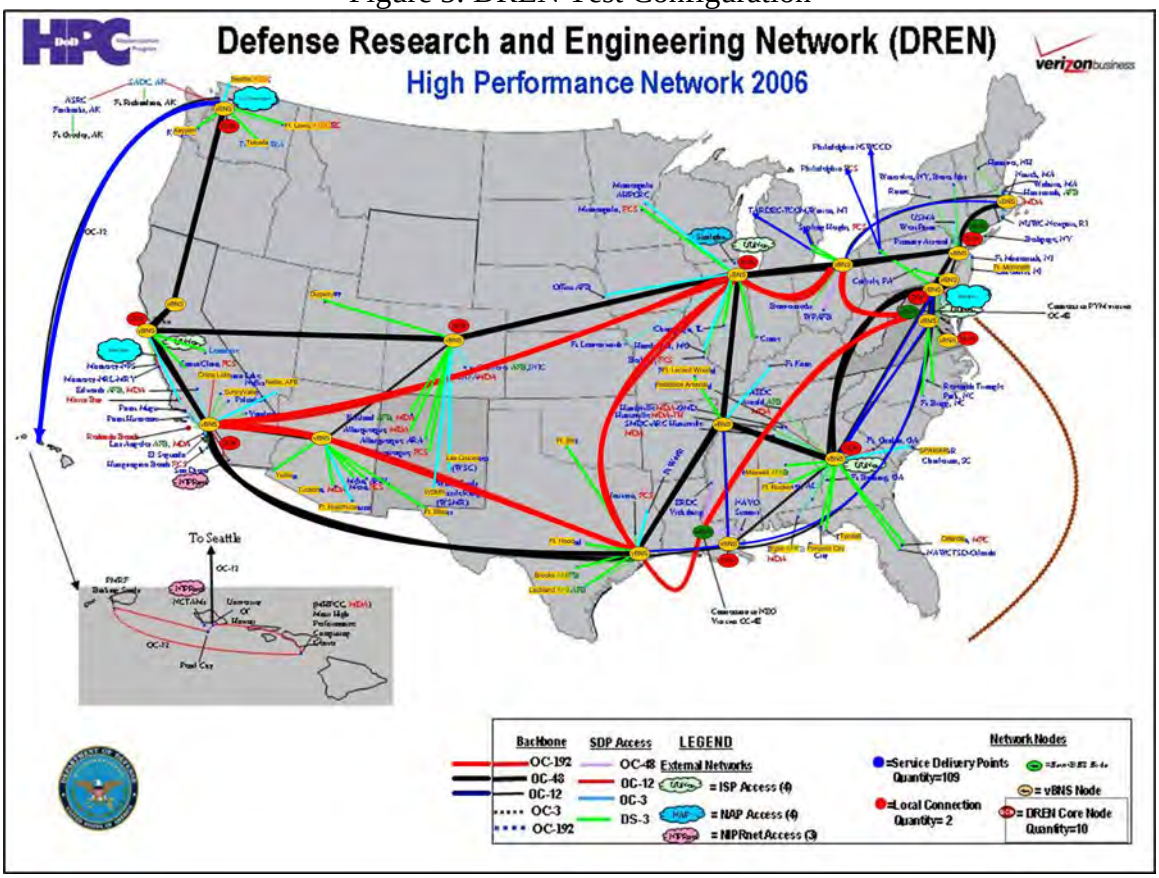

Figure 6: DREN High Performance Network 2006

8 Video Furnace TRANSLATOR/MULTIPLEXOR

WSMR used Video Furnace IP video distribution system because of Video Furnace's compliancy to all computing OS platforms with nothing to install on the receivers. This flexibility was required with the current IP video streams that were produced by FIRECAM at the range. The current system demanded dedicated hardware and software both on the sending and receiving nodes thus making it a complex solution for extending the streams to other nodes. Video Furnace was asked to evaluate this problem and design a software subsystem that could listen to various streams on the network and translate them into their distribution system. Video Furnace uses compliant MPEG Transport Streams which are housed by a global standard; ISO 13818-2 and 14492. The WSMR camera streams were of proprietary format and needed to be translated into ISO standard MPEG streams with proper video and audio standards applied.

Video Furnace developed a subsystem; see Figure below, which allowed WSMR streams to be translated to compliant MPEG streams with proper video and audio to be distributed by the Video Furnace distribution system. An elegant one box solution took twenty camera streams and created an IP line extension for those streams to be seen anywhere on a connected network, by any authorized personnel.

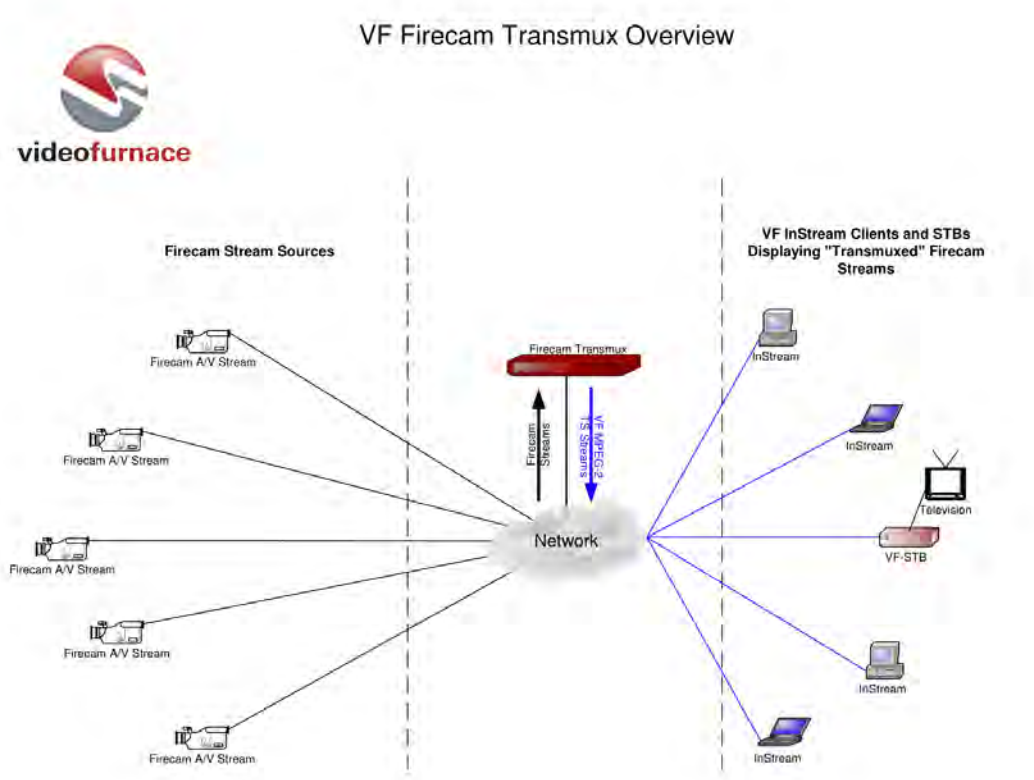

Figure 7: VF FIRECAM Transmux Overview Diagram.

The benefits of the Video Furnace system are numerous. Through any authorized portal the viewing technology is delivered with authentication to the receiver's memory space (the system has been tied to the CAC/PKI) and only content that is authorized for that session is viewable. All the stream data is FIPS-197 compliant with 256 bit encryption on

the data and 64 bits on the block ciphers. All key exchanges are done under secure protocol. When the viewing session is done, the viewer technology leaves the memory space of the node and nothing is left behind on the hard drive. The system leaves no footprint whatsoever and requires only a Java compliant browser, thus being compatible with any browser the Army “Gold” chooses.

With this client/server viewing technology the system gives complete control over manifestation and refresh. Thus, when new features are available, a single copy of the viewer is issued to the license server and the next session is completely benefited by any of the technology refreshes. Thus, when new video standards are released later this year, such as H.264 and High Definition, Video Furnace users will benefit immediately after a single server update.

9 DELIVERING SECURE IP VIDEO

9.1 *Overview*

As video and streaming video make greater inroads into network computing environments, a clear and present danger has emerged: security vulnerabilities. While no wide-spread exploitation of these IP video vulnerabilities have been reported, the breaches are serious and wide-spread enough to warrant extra vigilance. This paper outlines the chief areas of weak security in the majority of IP video platforms and the steps that Video Furnace has taken to ensure its products deliver secure IP video.

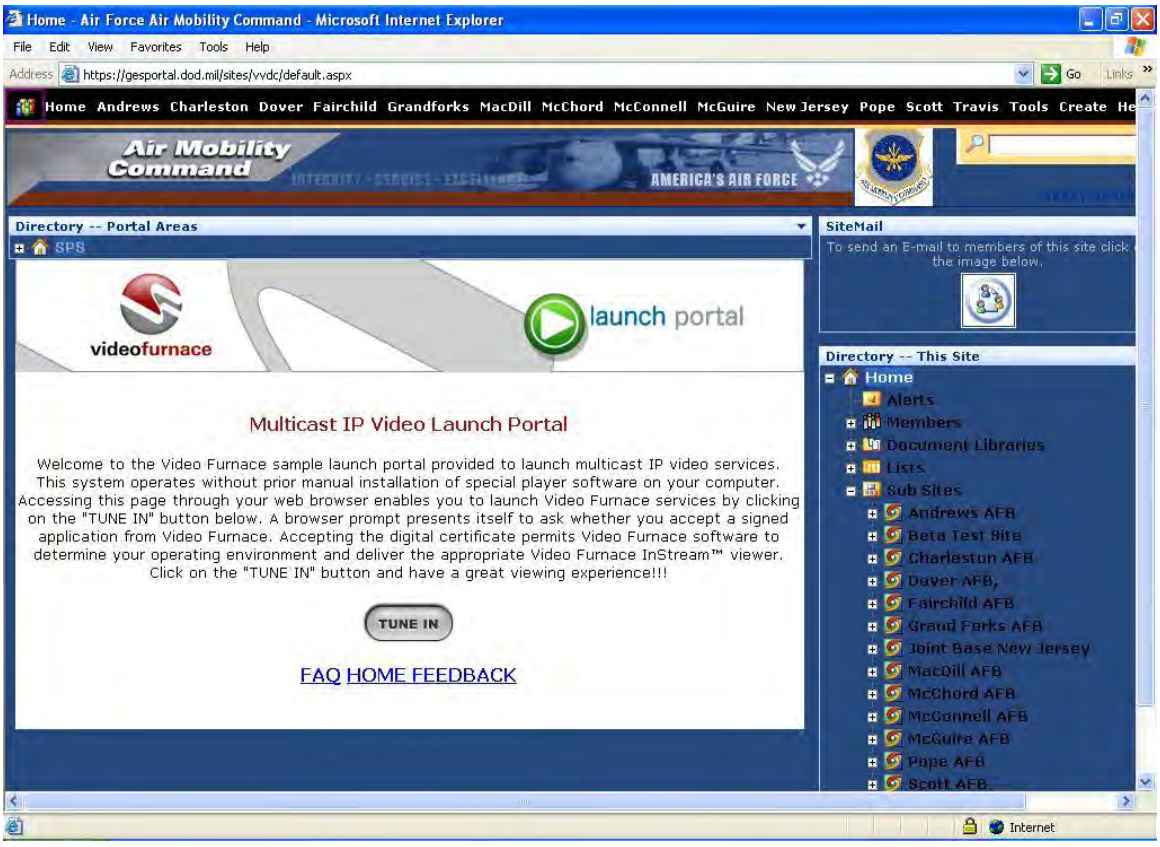

Figure 8: GIG CAC/PKI Enabled Portal for the USAF Air Mobility Command

9.2 *Recent IP Video Security Threats*

While much has been documented regarding browser and operating system vulnerabilities over the years, security and research analysts are only now coming to grips with the threats posed by network usage of IP video and audio. Vulnerabilities have been reported for Windows Media™ Player, Apple® QuickTime®, Firefox® media streaming and a number of unnamed open source media players and IP video streaming packages. Over the last three months, the following has been reported in *Computerworld, Network World, PC World, TechNewsWorld* and elsewhere:

A Quick Time vulnerability allows scripting to run with full user rights without the user's knowledge. Because of the background nature of this security hole, a hacker could insert code to take over system resources, run full-blown hacker applications and collect or maliciously destroy data on the infected machine. In addition, security experts point out that iTunes® software, using similar full-user rights authority, could also become a security breech for unsuspecting users.

Windows Media Player is vulnerable to attack by hackers through either the introduction of HTML code into files ran by the less-restrictive media player or by exploiting user rights breeches in the player itself. These security holes allow hackers to gain access to other Windows resources and to phish for user credentials by easily faking the Windows login/logout sequences.

A flaw in Windows security makes Firefox and Opera users open to attack through the operating system. In other words, even if users are running Firefox or Opera, streaming video or audio in these browsers may still allow the download and launch of malicious code.

Just recently, Adobe Corporation released a vulnerability issue with the current Flash player that could expose a user's machine to an attack. As you know, this is the most widespread media player that is utilized by just about every website in the world. The concerns here are that just a visit to a website that has flash logos, etc and your machine is immediately vulnerable with no actions on the user's behalf.

In addition, in an address to the Black Hat USA 2007 hacker conference, a senior security consultant with iSEC Partners declared the firm had found significant faults in both commercial and open source media streaming programs. Flaws or a particularly disturbing allowed the automatic launching of potentially malicious code. He declined to name the products because he is still disclosing the exploits to the software developers so they can develop and implement patches to fix them.

Figure 9: WSMR MOSS 2007 CAC/PKI enabled portal

9.3 *Key Areas of IP Video Vulnerability*

Examinations of the recent rash of video and audio IP security flaws show the majority of the security concerns center on three areas: media player, browser and data-at-rest vulnerabilities.

9.3.1 Media Player

Media players are often allowed access to operating-system level resources, and they are continually exposed to threats from outside sources. In addition, to provide certain digital rights management and interface controls, the players must frequently run Java, HTML and other scripted code strings. While patches for evolving threats can be issued, each new revision of the player can cause new security holes to be introduced. As long as the media players are intimately tied to the operating system or running in a less secure segment of the operating system, they are susceptible to security flaws.

9.3.2 Browser-based Media Streaming

Similar to the media player, browsers may be too closely tied to the operating system to provide full insulation from malicious code. In addition, unlike media players, these browser-based media streams may be automatically launched from any web site using Java or HTML embedded in the page. This provides virtually no protection from the launch of hacked or embedded malicious code.

9.3.3 Data-at-rest

Un-encoded and unencrypted multimedia files stored on servers or local drives are susceptible to modification from outside rogue software including worms and viruses. In

addition, these “open” files can become infected during transmission or prior to play, because the receiving players and media streamers are only expecting common media file types for processing.

9.4 *Video Furnace Solution*

Throughout the engineering and deployment of Video Furnace solutions, the company has made security and integrity of its delivery system top priority.

9.4.1 Encode and Encrypt

Throughout their lifecycle, media files are vulnerable to malicious attack—from inception, storage, delivery and playback. First and foremost, Video Furnace employs the strongest encryption available in the creation, storage and delivery of media streams: Advanced Encryption Standard (AES). AES’s strong encryption method was approved by the U.S. Department of Defense in 1997, and since its development in 1995 there has been no crack of the AES encryption method despite years of attempts. The AES system uses both variable length bit lengths for the encoding, and variable length keys for encryption and decryption, it is especially resistant to brute-force security attacks. Video Furnace system utilizes AES throughout the entire lifecycle of its multimedia delivery stream; therefore, the media streams themselves are highly resistant to tampering of any kind.

9.4.2 Eliminate Resident Players

Commercial and open source players operate in an open and unsecured Internet environment. Many third-party providers develop player plug-ins for everything from visualizations to compression codec, and are especially susceptible to hacking and malicious code.

Video Furnace has developed micro-client viewing technology that automatically provisions the host (utilizing digital certification technology) upon request of the service. The request from the host is passed through a license manager before issuance. The license manager is open design which allows user specific validations to occur before provisioning the host.

By removing the requirement to manually install a viewer at the host, the provider controls the process of placing the viewer in the host environment, changing it as needed, and removing it from the host when the session is terminated. This design completely removes any possibility of the media player being hacked by outsiders. The Video Furnace InStream™ player is not a browser plug-in, nor is it a Java applet. It is a standalone, runtime executable file that is separate from the browser. This unique design ensures that the stability of the viewing experience is independent of the stability of the Web browser. The Web browser is used to enable the video selection and subsequent automatic player delivery. Once the InStream™ player has been streamed, the user is free to close or minimize the browser, or surf to another site altogether, while continuing to enjoy the video experience.

9.4.3 Deliver Secure Streams

The Video Furnace stream is secured using AES encryption directly to the user’s desktop, minimizing brute force hijacking of the stream. The system adds an additional layer of security as validation by invoking a “watermark” option on streams. With this exclusive system, the stream is watermarked with the end user’s IP or MAC address. Watermarking

combined with AES encrypted streams provide not only security but media streaming tracking as well.

9.5 *Command and Control*

VFCommand and Control (VFC2)™ enables administrators of the system to remotely control any and all of the active Video Furnace clients (InStream viewers or VF-STBs). Administrators can remotely control features such as: changing the tuned channel on a VF-STB in a kiosk, placing kiosk or conference room STBs in or out of standby, sending other commands or messages tailored to one or more clients, or even terminating certain viewing sessions. VFC2™ provides a very flexible method to address messages and commands to all clients, subsets of clients, or a single client tuned to Video Furnace services on the network.

9.6 *IP Video's Time Has Come*

Affordable and far less complex to implement than in the past, video over IP is becoming a must-have for organizations, schools, financial institutions, government, and more. Leading businesses and institutions are finding that IP video and its myriad uses can vastly enhance company performance, greatly improve employee productivity and significantly cut operating costs. Companies leveraging IP video not only enjoy a competitive advantage today but are poised to benefit from new applications far ahead of their competition.

9.6.1 IP Video: Saves Time, Money and Offers Competitive Edge

The use of IP-based video in the enterprise is growing dramatically as companies, schools, government and other organizations increasingly discover that their IP networks have the potential to cost-effectively deliver high-quality live and on-demand video applications for event broadcasts, intra-company addresses, digital signage, security monitoring, public/cable TV programming distribution, mission-critical communications, corporate training, and more. All of this is possible without additional IT resources or end-user training of any kind. Businesses around the world are utilizing IP video over networks to help employees be more productive, to keep personnel better informed and to streamline training initiatives.

In studies done over recent years, researchers conclude that retention of key information is improved by 42 percent when video is used. (1) Retention, of course, is the key to effective communication and training. While retention is important, video also makes an impact like no other media. CEOs and corporate managers are increasingly relying on video to make a lasting impression. IP video is the cost-effective means to do so. IP video also gives organizations an affordable and manageable solution to make video assets available anytime, anywhere. In essence, think of IP video as a vast, easily searchable library of corporate video assets available from any authorized computer.

9.6.2 Corporate Video Goes Mainstream

Far from being a "bleeding edge" or niche application, IP video is moving quickly into the mainstream of America's leading companies, such as British Petroleum, Ford, Hearst Publishing, Knight-Ridder, Monsanto, the U.S. Department of Defense and Wal-Mart. These companies realize that video, like no other network application, has the power to motivate, inform and empower their employees, while delivering greater productivity and higher performance.

Today's industry leaders utilize IP video for unlimited reasons:

a) Businesses around the world are utilizing IP video over networks to help employees be more productive, to keep personnel better informed and to streamline training initiatives.
b) Sharing customer testimonials provides more impact and credibility than any written case study.
c) Delivering multiple live financial feeds on every financial analyst's computer.
d) Providing new product instruction, quality assurance training and new employee orientation, all of which can be accomplished in a fraction of the time.
e) Making clips of video information available on demand, which turns Video On Demand into Information On Demand, a powerful, searchable database of video information.
f) Recruiting new employees, which is far more effective using high-quality video tours, information clips and video-taped corporate highlight reels.
g) Allows automatic capturing and cataloging surveillance and reconnaissance videos for later analysis on virtually any authorized computer (captured clips can be processed by sophisticated video analysis software).

IP video is limited only by the corporation's imagination. With production and implementation costs now at an affordable level, major corporations are rolling out custom video applications at an exponential rate. Some analysts estimate that a significant portion of current conferences, training initiatives and sales training will shift to IP-based video in a matter of five to six years. Companies that deploy IP video solutions now reap even greater rewards in the not-so-distant future.

It is no longer a question of "if" corporations will utilize IP video, but a matter of "when." Corporations and organizations that deploy the technology today will certainly be the leaders in IP video productivity tomorrow.

9.6.3 Breaking News and Executive Communications Using IP Video Broadcasts

In a 24 by 7 business environment, information is the key to staying competitive. When employees need crucial, timely information, video hits home like no e-mail, voice mail or instant message could. IP video can be routed to computers live. Corporate information can be broadcast over live or stored video to multitudes of employees in a fraction of the time it takes to orchestrate a communication program. When information needs to get out fast, when time is crucial, video communicates with the most impact at the lowest total cost. And with IP video broadcasts, you alone control the message, the timing and the distribution. Video, like no other network application, has the power to motivate, inform and empower employees, delivering greater productivity and higher performance.

In addition, executives are finding that they can infuse their corporate communications with a personal touch by directly addressing employees with video. These videos can be stored for later use and rebroadcast to investors, board members or employees who may have missed it the first time. Content and delivery decisions rest solely with the organization. Corporations, universities and the military utilize IP video today to increase efficiency, improve productivity and to gain a competitive edge:

Brokerage firms feed multiple live financial broadcasts to trader's computers, keeping them informed of critical market information throughout the day. Reporters at news organizations all over the world monitor live, multiple video feeds on computers to keep apprised of breaking news and current events all day long CEOs at multinational

corporations broadcast key messages and announcements to employees in real time. Enterprises display informational content on televisions or monitors. Campus and corporate security broadcast alerts and instructions to students or staff during crisis situations. Schools deliver live, pre-scheduled and VOD programming to students, faculty and staff. The military delivers classified and unclassified video to personnel who need it.

When fast-breaking news hits, where do you turn? Think of IP video as your organization's own cable news network.

9.6.4 Information On-Demand

Just as we now expect the Internet, intranet and corporate network to provide immediate access to a variety of knowledge sources, video information can be cataloged and retrieved as effortlessly as other information sources.
Unlike videotapes and DVDs, IP video materials can be easily stored and retrieved with the click of a mouse. Live video feeds can be captured and saved automatically, allowing organizations to offer seminars, new-hire orientations, routine broadcasts, digital displays, focus groups, usability testing sessions, and more, at any time, at a fraction of the cost of less efficient paper or media-based materials. (2) Media production houses charge thousands of dollars for the same capabilities, while IP video has capture and search capabilities already built in.

9.6.5 IP Video Training

Broadcasting video training to computers is an efficient, cost-effective platform for high-retention training. In fact, according to a recent report by Global Industry Analysts, Inc. the market for e-learning is rising rapidly. Revenues for e-learning are expected to top $17.5B in 2007 in the U.S. alone. (3) A growing percentage and driver of this network-based training is video.

The bottom line is clear: major corporations and organizations are using video to train and teach at an increasing rate. Why are many of the world's most competitive companies investing part of their IT budgets in IP video training? They are investing part of their IT budgets in IP video training because it is extremely effective.

9.6.6 IP Video Cuts Costs and Saves Time

First and foremost, IP video training is cost-effective. The alternative is live training which is extremely expensive. On-site training, where trainers visit individual locations, is the least expensive alternative to video training but is still an enormous cost to many organizations. Live trainers must be present, materials must be developed and sessions must then be attended by numerous employees. Schedule conflicts, missed sessions, lost work hours and travel expenses all contribute to the total cost. In the most established alternative, many corporations send employees to centrally located training centers, further increasing costs and loss of productivity. Video training delivered over existing IP networks, on the other hand, requires no travel and fewer hours lost off the job. Costs drop even more when you consider that these assets can be reused without paying trainers for logging additional hours. What if your employees could have valuable training at any time, without the cost of trainers or travel?

9.6.7 IP Video Is Responsive

Traditional training materials are costly, have lengthy production times and are not easily distributed. IP video, on the other hand, can be developed and deployed in a fraction of

the time and distributed to numerous computers. When timely information is essential to keeping an edge in the marketplace, shortened training cycles are imperative.
When new products, services or equipment make retraining essential, you want the benefits as fast as possible.

Corporations, organizations and schools are using IP video in a wide range of training situations:

At one Fortune 500 company, new hires are learning about their jobs at their computers. Rather than fly employees across the country to learn new systems, one corporation has system training available 24 by 7 on its network.

Students at many major universities access missed lectures and archived lessons on their computers from their campus dorm rooms.

9.6.8 Potential Barriers to IP Video in Business

Many people erroneously equate video over the network with YouTube, music videos and other entertainment products, so it is no wonder some businesses are reluctant to open their critical networks to IP video applications. But business video and multimedia solutions are a breed apart from Internet diversions, and companies that exclude IP video from their mix of productivity tools run the risk of falling behind the competition. Business video cuts costs, provides an archive of easily retrievable content and gives companies a cutting-edge platform for communication. With the economic and productivity case for video well documented in business environments, CEOs and CTOs have to ask:

How can we deploy IP video without opening our network floodgates to unwanted video from the Internet?

Another potential concern for deploying IP video is the associated costs and complexity of early IP video systems. In some of the initial solutions, nearly all the organization's switches had to be upgraded to handle the issues of broadcasting quality video without degrading the signal. If deployed incorrectly, the solution could affect the performance of the core network. Some even required specialized and costly hardware to distribute the video. Further, nearly all of the early-stage solutions still require media players to be installed on every computer that will view video. Version control and software maintenance then become an additional expense and support issue. So network administrators have to ask: How can we implement IP video with the least impact and cost to our existing network?
Those desktop media players may, in fact, be a network security risk. As reported in Computerworld, eWeek, TechNews and other IT industry journals, most commercial media players – including Microsoft Media Player and Apple QuickTime – have security holes that malicious files and viruses can exploit. In addition, these players make it problematic to control what media can be downloaded, again, perhaps, opening the network to any media available on the Internet. Network security experts must ask themselves: How can we implement IP video securely and provide only the content we want our users to access?

IP Video in the Enterprise: Hearst Tower, New York

When Hearst Corporation created a convergent network of digital voice, video and data in its new Manhattan building, Hearst Tower, it chose Video Furnace to deliver more

than 50 channels of live broadcasting and VOD services to keep editorial staff throughout the building apprised of current events and breaking news.

> *Charles Montplaisir, vice president of information technology for Hearst Magazines in New York, said in the Chicago Sun-Times that sending video over Internet Protocol (IP) has proven cheaper and more flexible than setting up a cable-TV system inside Hearst's headquarters. "Every employee has access on the computer to video content distributed by Video Furnace," Montplaisir said. According to the Sun-Times story, most employees see only in-house channels, though top executives and conference rooms have access to a wide array of programming. Overall, the move to video over IP saved Hearst hundreds of thousands of dollars. Companies that exclude IP video from their mix of productivity tools run the risk of falling behind the competition.*

9.7 *Not All IP Video is Created Equal*

Video Furnace, a leading developer of enterprise-class IP video solutions, has addressed all of the above concerns associated with deploying IP video over business networks. First and foremost, the Video Furnace solution does not require desktop or laptop software to be installed, simplifying not only support issues but also eliminating security vulnerabilities as well.

With Video Furnace, organizations can control which people view particular content, when they can view it and even how often they can have access to materials. The Video Furnace solution is an organization's own video broadcasting service, complete with permission-only channels and access. In addition to secure delivery to computers without agent software installations, Video Furnace supports IP-enabled set-top boxes, allowing the hook-up of televisions and monitors without the need for a desktop or laptop computer. Video Furnace is primarily a software product; upgrades, updates and even the initial installation have less short- or long-term impact on an existing network.

9.8 *A Future-Proof Picture*

Once security and performance issues are addressed, IP video has a multitude of further applications in organizations. The importance of software, though, in the converged network is paramount. Solutions that deliver IP video using primarily software are more readily adaptable and upgradable at a lower cost. In addition, software-centric solutions are a stable platform for building additional functions into the video network.
We will, in a few short years, find that high-quality computer video is as common as spreadsheet software and Web browsers. Video truly will be simply another application on the network.

10 UAV SYSTEMS & OPERATIONS VALIDATION PROGRAM (USOVP) USAF 46th Test Group

The Unmanned Aerial Systems (UAS) Technical Analysis and Applications Center (TAAC) were established in 1999 with the mission to promote safe integration of UAS in the National Airspace System (NAS). In order for UAS to be flown in the USA outside of special use airspace, a Certificate of Authorization (COA) or an experimental

airworthiness certificate must be issued by the FAA. A COA has been obtained by New Mexico State University to operate the Aerostar and Orbiter UAS in the NAS and the TAAC conducts flight operations within restricted airspace when appropriate. The TAAC has been involved in the testing of various UAS platforms and has produced a certification roadmap for the UAS Flight Test Center in Las Cruces, NM. The FAA has authorized the TAAC to operate the UAS Flight Test Center and provide a vehicle for UAS certification while gaining reliability data for future systems.

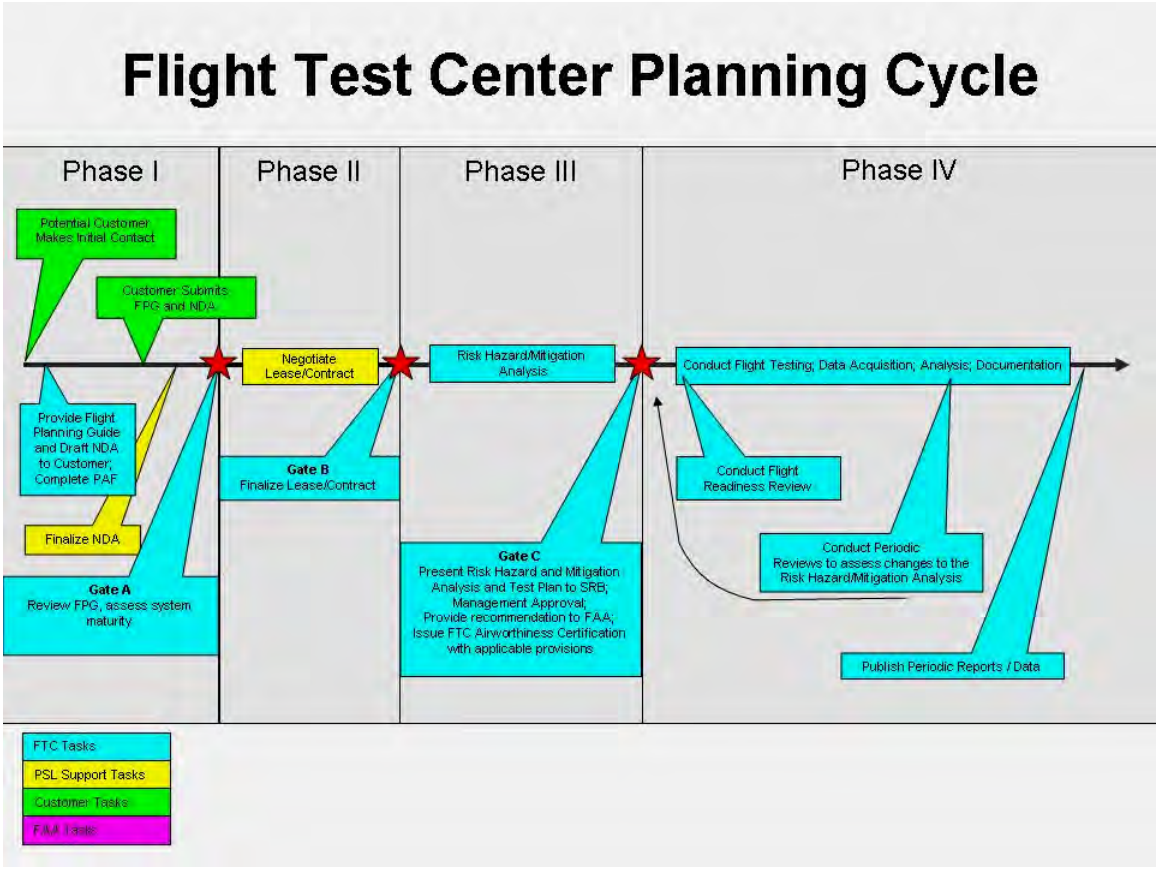

Figure 10: Flight Test Center Planning Cycle.

TAAC has secured funding to implement the UAV Systems and Operations Validation Program (USOVP). USOVP is a unique and dedicated program for operations and performance evaluation of UAS that has been in place since 2003. USOVP is the funding vehicle created to establish operation flight capability for UAS. In addition, PSL, the 46th Test Group at Holloman Air Force Base, and White Sands Missile Range developed a joint regional UAV Test and Evaluation Center (UTEC). The UTEC was formed to develop "file and fly" capabilities to facilitate the use of UAS in the NAS.

TAAC has supported various exercise, demonstrations and Test and Evaluation (T & E) efforts and has established the capability for routine flights supporting critical technologies testing and validation of processes. The following list describes some of the events and elements that have been achieved to execute them.

Figure 11: Aerostar Optical Sensors

10.1 ***Sense and Avoid (SAA) system development***

i) Integrate an optical sensor array on a UAV to detect, evaluate, and execute an avoidance maneuver if applicable. Flights in restricted airspace with intruder aircraft were used to develop this system.
ii) Develop an optical sensor array for use on a UAV to detect and evaluate aircraft in the vicinity. This prototype was tested on a manned surrogate UAV aircraft.
iii) I installed an acoustical sensor array on a UAV. The system collects performance data during routine flights in NAS with chase aircraft under provisions of COA.

10.2 ***Exercise support***

iv) Falcon Virgo – flew in support of Air Defense Artillery training exercise in restricted airspace in conjunction with manned aircraft.
v) Angel Thunder – proved ISR support of USAF CSAR exercise under provision of COA from Playas, NM.

10.3 ***Test and Evaluation***

vi) Evaluated performance of GPS unit developed for Army ground units in a flight environment. Performed in the vicinity of Las Cruces, NM with transitions in and out of restricted airspace in accordance with the provisions of the COA and coordination with ZAB and WSMR.
vii) Evaluated performance of anti-IED device and validated integration in a UAV. Evaluation flights were performed in restricted airspace.
viii) More information may be obtained http://www.psl.nmsu.edu/uav/.

Figure 12: Aerostar Unmanned Air System

11 RESULTS AND CONCLUSIONS

Flexible software based streaming video solutions are developed by the Army at White Sands Missile Range, tested in a Net Centric environment and ported to the DREN for proof of concept demonstrations. The system is tailored to meet T&E requirements but performs positively under any general scenario. Both uncompressed and compressed transmission is supported and is successfully demonstrated via long haul Synchronous Optical Network (SONET) backbone transmissions. A key requirement of end-to-end latencies of less than 200 milliseconds is met with a measured value of approximately 80 milliseconds. An additional capability critical to the T&E community is the ability to interface with many non-standard flavors of specialized video, such as high resolution CMOS, high bit-depth IR, and high frame rate parallel digital. Need to add additional information in regards to “lesson learned” with the DREN and Intrusion detection devices along with firewalls, priority queuing of packets on a unicast tunnel.

The main drawback to the system is the significant amount of CPU operations required when compressing video during transmission. To address this issue, a few modifications to the system will be carried out. First, buffer copies and format conversions are not performed in an optimized manner. These operations are not vectorized by the Intel C/C++ compiler, so utilization of MMX, SSE2, or Intel Intrinsic holds promise for significantly decreasing the amount of CPU operations required to perform the buffer copy and format conversion operations. In addition, porting from the legacy video for Windows library to the DirectShow library holds promise for performance improvements.

The Video Furnace IPTV solution provides the most secure distribution of IP video streams over the existing network. The system allows for capture of analog or IP video streams, applying FIPS 197 compliant AES encryption and Forward Error Correction (FEC) to create a seamless viewing experience by any node on the network. This is all accomplished by the InStream viewer technology which is a patent-pending client server application that supports Windows, Mac OS/X, Linux, Solaris 10 and Set Top Box and requires nothing to be installed on the end node. The viewer is placed in the memory of the node at initiation through net worthy authentication and leaves memory when it is terminated. This allows for a system that is completely zero footprint. This closed system allows for the most secure environment for IP video viewing and reporting.

Media streaming security vulnerabilities will persist and proliferate if media player, browser and data-at-rest issues are not specifically addressed. Video Furnace, the leading provider of enterprise-class, mission-critical, secure video distribution, has employed specific and significant design criteria to ensure the safe delivery and deployment of IP video in network computing environments.

12 ACKNOWLEDGEMENTS

The authors would like to thank Mr. Gilbert Harding of the Department of the Army at White Sands Missile Range for his efforts in supporting this development and for providing the resources necessary to accomplish innovative work within the DoD and WSMR. Special thanks to Mr. Howard Weinzimmer, President and CEO of Video Furnace and Mr. Joe Gaucher, CTO/Founder of Video Furnace, for providing critical company specific information and product information which was essential for the flow of this technical document. Mr. Weinzimmer provided beta test equipment and expert technical support during the development of the Video Furnace IPTV Bridge and Translator Multiplexors. The authors would also like to show appreciation to Mr. Peter Butt from NAVAIR, for his support and dedication to the advancement of DREN capabilities and services. Additionally, we are obliged to Mr. Dave Telles, for the use of his timing lights during latency testing. Mr. Christopher Healy, Navy Enterprise Architect, is recognized for the indubitable can-do attitude of his technical team, especially Mr. Scott Brownrigg. In addition, the authors would also like to thank Amanda Thomas, Thomas Archuletta and their team at Cisco Systems for the outstanding support that they provided during this project.

13 REFERENCES

1. Peter D. Symes, “Video Compression,” McGraw-Hill Inc., 1998.
2. K. Castleman, “Digital Image Processing,” Prentice-Hall Inc., 1996.
3. A. Sadka, “Compressed Video Communications”, John Wiley & Sons, 2002.
4. I. Richardson, “H.264 & MPEG-4 Video Compression”, John Wiley & Sons, 2003.
5. Matrox Imaging Library (MIL-Lite) v7.5 User Guide, 2003.
6. Cisco ONS 15454 system documentation, 2002.
7. Jack Germain, TechNewsWorld, “Media Player Exploits: New Vectors, New Threats,” September 26, 2007.
8. Greg Keizer, Computerworld, “Security Researcher Finds Flaw in Windows Media Player,” September 19, 2007.
9. Jordan Robertson, Associated Press, “Media Players Have Significant Flaws,” August 2, 2007

14 Glossary of Common Terms

Converged Network – A single network capable of carrying voice, data and video.
Information on Demand – Because IP video segments can be tagged with key words before storage on the corporate network, clips of video information can be called up on demand. Tape-based video sessions must be manually labeled and cataloged.
IP Multicast – An efficient, one-to-many transmission of data over an IP network.
IP Network – A network in which transmission of information is done using IP protocol, which is part of the TCP/IP protocol suite. IP has become the global standard for networking.
IP Video (Internet Protocol Video) – Video signal carried over a standard Ethernet IP network.
IPTV (Internet Protocol Television) – The process of providing television services through the use of IP networks.
Set-Top Box (STB) – An electronic device that allows the reception of video signals on a television or monitor.
Usability Testing – The live testing of products by actual customers, often tape recorded for later analysis. These tapes and tests are used to determine how easy or difficult a given product is to use. Current tape-based usability testing is difficult to track and excerpt when compared to computer or IP-based session recordings.
VOD (Video on Demand) – A service that enables end-users to interactively request and receive stored video.