Internet Multimedia Server—VideoCharger

By Eddie Ho and Sam Juliano

Until recently, the World Wide Web was novel and entertaining, but far from dynamic. Now, thanks to a rapidly growing group of multimedia-enabling tools, such as Java, ActiveX, Shockwave, Quicktime, and video streaming, the excitement of surfing the Web is finally picking up its pace. For many interactive, content-rich environments, the scalable VideoCharger for AIX can deliver end-to-end multimedia content to the desktop using the open Web infrastructure. This article details the emerging technology.

Video and audio are the most natural form of human communications for concepts and ideas. In our society we are constantly surrounded by multiple media types, each with the intent of delivering the strongest messages to capture our mindshare.

Starting at home, television is the information delivery system. TV is a broadband, one-way broadcast system that can deliver high-quality video, but it does not provide interaction. In the business environment, most commercial applications are non-graphic and optimized for fast data transfer in the transactional environment. Today's application designer rarely considers multiple datatypes because of the programming environment and network constraints.

The World Wide Web is the emerging hypermedia information retrieval system with built-in support for multiple datatypes. This infrastructure enables the flow of multimedia content to both business and consumer environments. The Internet frontier has gradually extended to accommodate audio and video.

Internet Multimedia Challenges

Multiple requirements make video delivery across digital networks extremely challenging. The challenge is compounded by many network providers in the Internet industry, each with different levels of capability and support. Figure 1 provides a reality check in today's combined Internet and intranet environments. It shows a typical data-intensive environment where different business units access each other's information.

The availability of multimedia on the Internet poses several challenges:

Network bandwidth. Digital video requires high-speed throughput to produce acceptable quality. Most current LAN architectures were not designed to handle multiple concurrent users in a video environment. This problem is exacerbated when the client PC is attached through a slow modem connection. The bitmap-intensive pages have converted the World Wide Web into the "World Wide Waiting" room for many.

Isochronous media delivery. Digital video is sensitive to timing and cannot tolerate any jitter or latency. Although video components of a stream can tolerate some data error and loss, the audio stream is less tolerant of data irregularities. Dropped packets and jitter can be distracting and render audio totally unintelligible in the worst case.

Filesystem infrastructure. Digital video consumes great quantities of storage space. For example, a digital video file in MPEG-1 format recorded at 1.2 Mbits per second will require 1.2 GB storage for a two-hour movie. Most UNIX® filesystems are not designed to manage and stream many large files.

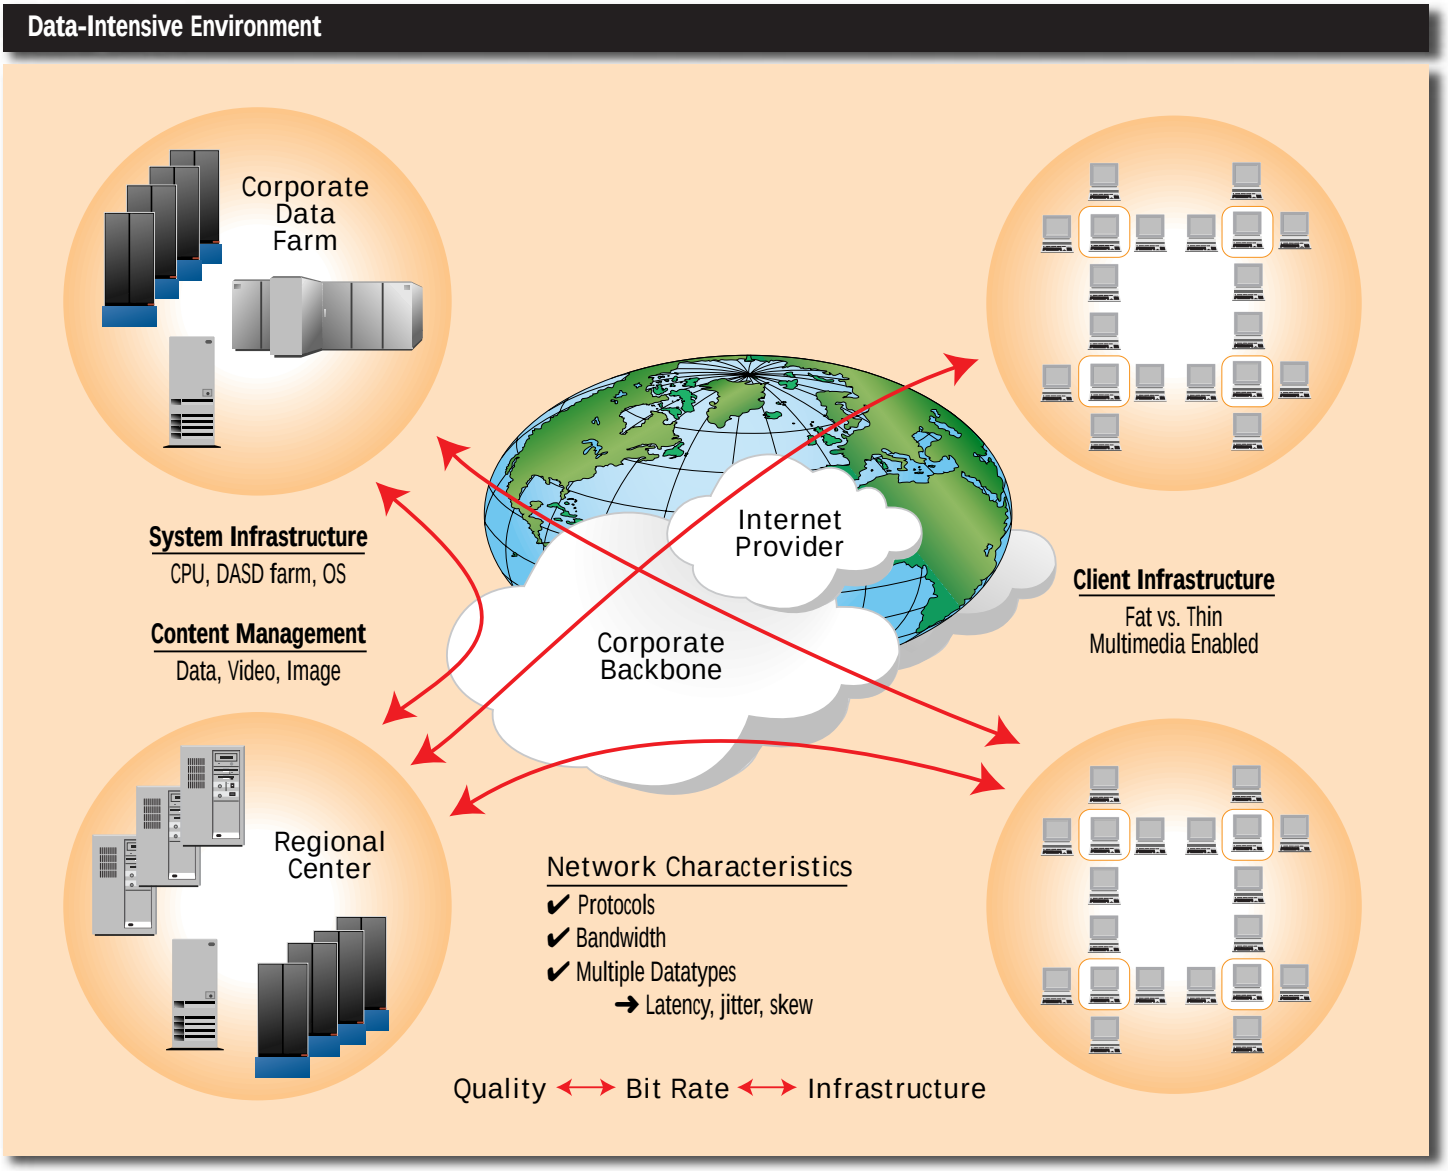

Figure 1. Typical data-intensive environment

Client considerations. MPEG decoder performance plays a key role in video quality because the stream is pushed from the server and the client decoder is required to keep up in a constant environment.

VideoCharger Server Solutions

The VideoCharger for AIX software is designed to deliver networked video and audio to Internet/intranet clients. This allows businesses to enhance their marketing on the Web by adding video to their Web pages for product demonstrations. The video content is delivered in real-time, thereby eliminating the need to download or save the file before video and audio are played. The delivery is flexible and adjustable to the connection speed of the client.

The combination of Low Bit Rate (LBR), Mid-band, and High Bit Rate video/audio can be streamed across the Internet concurrently to many clients. The ability to regulate video based on connection speed can help to provide a balanced mix of transactional versus multimedia content users. Figure 2 compares the bit rate to frame rate and bandwidth.

Figure 3 shows a typical Internet environment with video-streaming support.

The VideoCharger server, a two-tiered client/server design, runs on AIX 4.2. The components consist of the following:

VideoCharger Server: Enhanced HTTP server that delivers streaming video and audio over the Internet/intranet. The server can efficiently manage, store, and deliver content to the connected clients.

VideoCharger Player for Windows 95: An IBM-provided "helper" application player working together with a browser, such as Microsoft Internet Explorer or Netscape Navigator™. The player allows the user to select, receive, present, and control video streams.

Advantages of VideoCharger

The VideoCharger video delivery system has many innovations, including technologies from IBM Research and the Austin development team. There are several key areas of innovation.

Bit Rate	Frame / sec (fps)	Bandwidth	Resolution
Low Bit Rate	7.5	28.8 Kbits/sec	160 x 120
Mid-band	15	256 Kbits/sec	320 x 240
High Bit Rate	30	1.5 Mbits/sec	320 x 240

Figure 2. Comparison of bit rate, frame rate, and bandwidth

Filesystems

Digital video assets are time-based, continuous data streams of information that must be delivered to the viewer at the proper rate in order to preserve human perception of motion and sound. The proper data rate must be maintained from the server to the client; otherwise the viewer will see "glitches" in the video and hear dropouts in the audio.

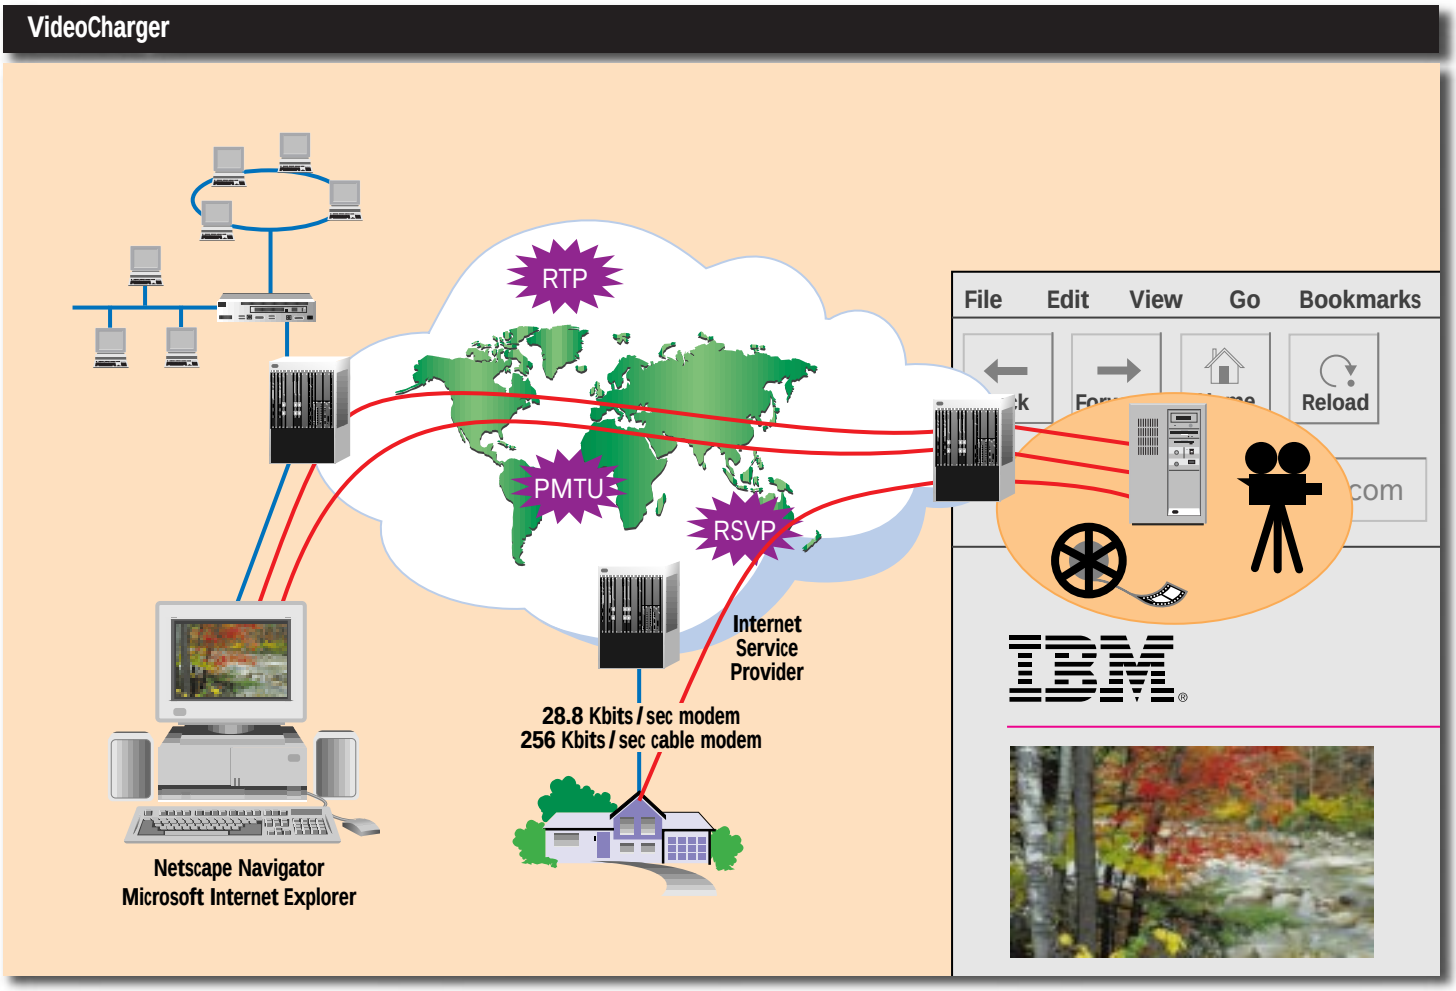

Figure 3. VideoCharger

The VideoCharger Server uses a custom filesystem known as the Multimedia Filesystem. Unlike standard UNIX filesystems designed for storing and retrieving relatively small files, the Multimedia Filesystem was designed for the efficient delivery of digital video assets, providing guaranteed bandwidth for data delivery. The file infrastructure is optimized for the storage, management, and distribution of large digital video files. This design is critical to the balance of the end-to-end delivery system. Some characteristics include the following:

Data Striping. Striping balances the workload evenly across all disks, which increases the aggregate bandwidth of the filesystem. It also permits concurrent access for multiple clients.

Large files and filesystems. Each filesystem can conceptually support up to 64,000 disks (up to 256 terabytes).

Large block reads. Reading large blocks of data (256 KB and up) reduces the number of read requests the server must handle.

Data replication. Replication spreads multiple copies of digital video files across filesystem disks to protect against unexpected disk failure. Replication factors can be assigned to the entire filesystem or to individual files.

Deadline scheduling. This ensures that digital video data delivery takes priority over the OS services.

Error recovery and online reconfiguration. Disks can be added and removed while the system is online with no data loss.

Designed for scalability. A minimal configuration can grow without a major interruption, such as doing a back-up, reformatting a disk, and recopying assets.

Multiple Data Rates

The VideoCharger Server supports clients attached to a variety of network types, such as LANs, WANs, and switching internetworking like ATM. Remote users who access the networks using modems through an Internet Service Provider (ISP) require a minimum of 28.8 Kbits/sec or higher. The server supports bit-rate sensitive content to achieve a balanced quality based on the connected line speed. This encoding is based on the industry-standard H.263 video and 6.3 mode of G723 audio from the videoconferencing industry. This technology provides clients with a 28.8 Kbits/sec modem and acceptable quality of 7.5 frames per second and an 8 kHz audio channel.

Mid-band video quality can be used by Internet/intranet clients attached through ISDN or the emerging cable modem technology. A 256 Kbits/sec connection provides a good quality video resolution of 320 x 240 pixels at about 15 frames per second.

High-quality MPEG-1 video streams recorded at 1.5 Mbits/sec provide a 352 x 240 pixel resolution at 30 frames per second, equivalent to VHS videotape in quality.

The flexible support for multiple data rates allows different clients to attach to the server at various speeds based on its supporting network capabilities. This minimizes the amount of bandwidth used and maximizes the number of concurrent clients that can be attached to the server.

New Internet Protocol Extensions

The VideoCharger Server delivers video content to the PC client using some of the emerging Internet extensions discussed below.

Real-time Transport Protocol (RTP). This new protocol extension was designed for internetworked audio and video conferences that require content synchronization, demultiplexing, media identification, and active-party identification. RTP delivers digital video with a minimum network overhead.

ReSerVation Protocol (RSVP). This protocol reserves network bandwidth in an Internet environment. It allows the request for a specific Quality of Service (QOS) class throughout multiple networks to secure a consistent bandwidth for the delivery of video content. The VideoCharger will take advantage of this extension if the network router supports this function and can better ensure end-to-end quality delivery.

IP Multicasts. This enables the VideoCharger Server to transmit IP datagrams to zero or multiple hosts with minimum streams. The host can, in turn, multicast the video asset to clients within its subnet.

Path MTU Discovery. This support can optimize the Maximum Transmission Unit (MTU) size during the initial network setup. Today's internetworking topology can have at least 15 networks. Each network has different MTU sizes, which can deteriorate the performance during transmission.

VideoCharger Server Component Architecture

The VideoCharger Server is an HTTP server enhanced with the capability to stream digital video data to the IBM-compatible PC clients. The Server can coexist with other Internet servers

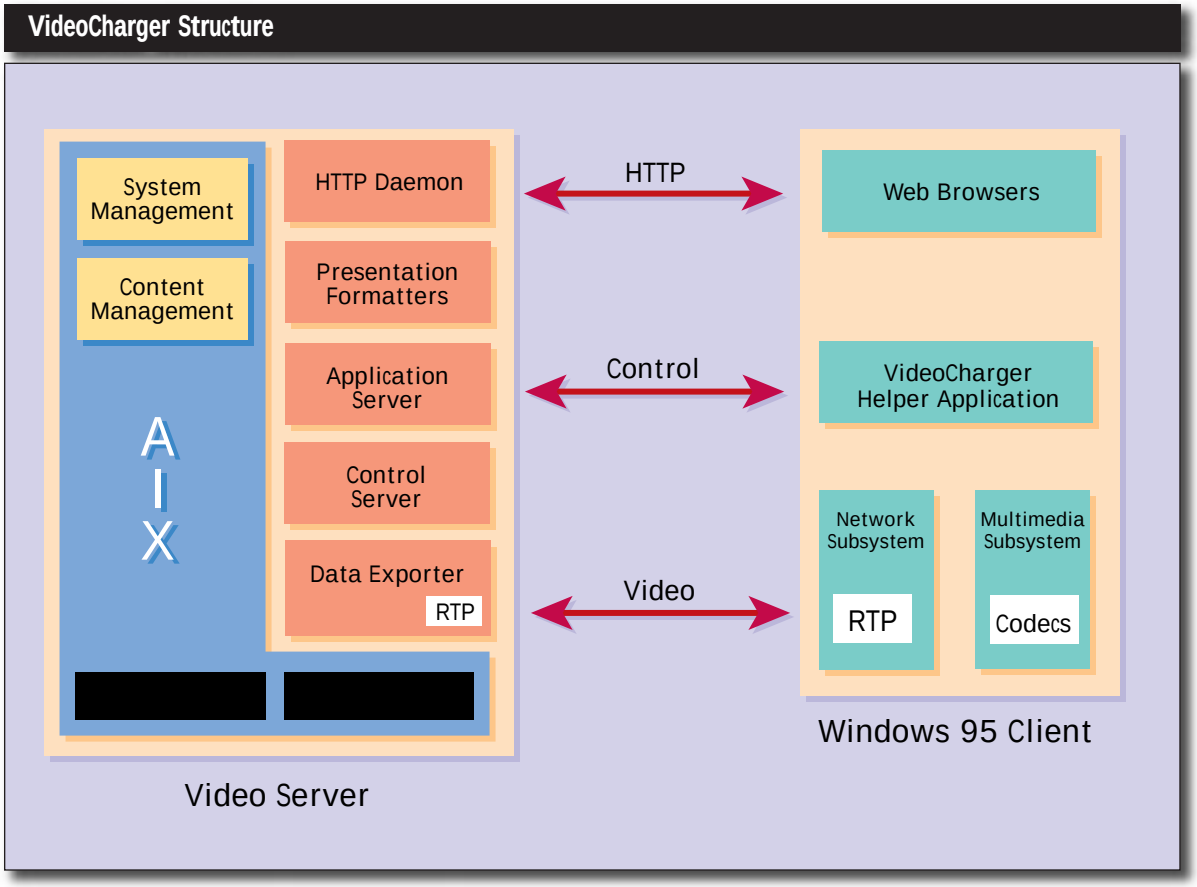

Figure 4. Components of the VideoCharger Server

within the complex using Uniform Resource Locators (URLs) in HTML documents (Web pages). The component infrastructure is scalable from a single system to multiple systems supporting hundreds of concurrent video streams. Figure 4 shows the components of the VideoCharger Server.

Each component of the VideoCharger is described below.

Data Exporter. This interface module controls the rate at which the data is pushed over the Internet. The server can support multiple data exporter modules. Each module runs on a separate server—the key component for scalable data delivery.

Control Server. This supervisory module coordinates the operation of all components and regulates the number of video streams to ensure sufficient resources.

Application Server. The module interface to the client provides video stream control functions, such as video pause, stop, start, and manipulation of the stream.

Multimedia Filesystem. This 64-bit filesystem provides real-time support for playing and recording continuous media.

The VideoCharger allows real-time control of the video stream between the server and the client. It provides support in two ways:

- Real-time feedback from the client to the server at the network level
- User interface commands (VCR buttons) sent from the client to the application server for stop, pause, pace, and position. This support is provided by three HTTP Common Gateway Interface (CGI-bin) programs called Presentation Formatters (PF):
  - — Video Selection formatter for video selection from the list
  - — Video on Demand formatter for query and search of video item
  - — Broadcast Scheduler formatter for video scheduling to be multicast

Figure 5. VideoCharger Player for Windows 95

VideoCharger Client Functions

The plug-in for the VideoCharger Player, provided with the Server product, runs on Windows 95. You can also download it free from http://www.rs6000.ibm.com/solutions/video-servers. The client can access the video content using a Web browser such as Netscape Navigator or Microsoft Internet Explorer. The Player consists of the following:

- Player helper applications
- Networking subsystem, including streaming network interface, application server interface, RTP
- Multimedia subsystem, including ActiveMovie subsystem, streaming source filter, Low Bit Rate (LBR) codec

Figure 5 shows the VideoCharger Player pop-up. Figure 6 shows the interaction between the client and various components of the servers.

Server Scalability

The VideoCharger is designed for scalability; all components of the server can be integrated in a single system for small environments, or can be scaled up using multiple systems with distributed components for a very large environment. Users can grow their configuration by adding more data exporter systems to achieve a higher level

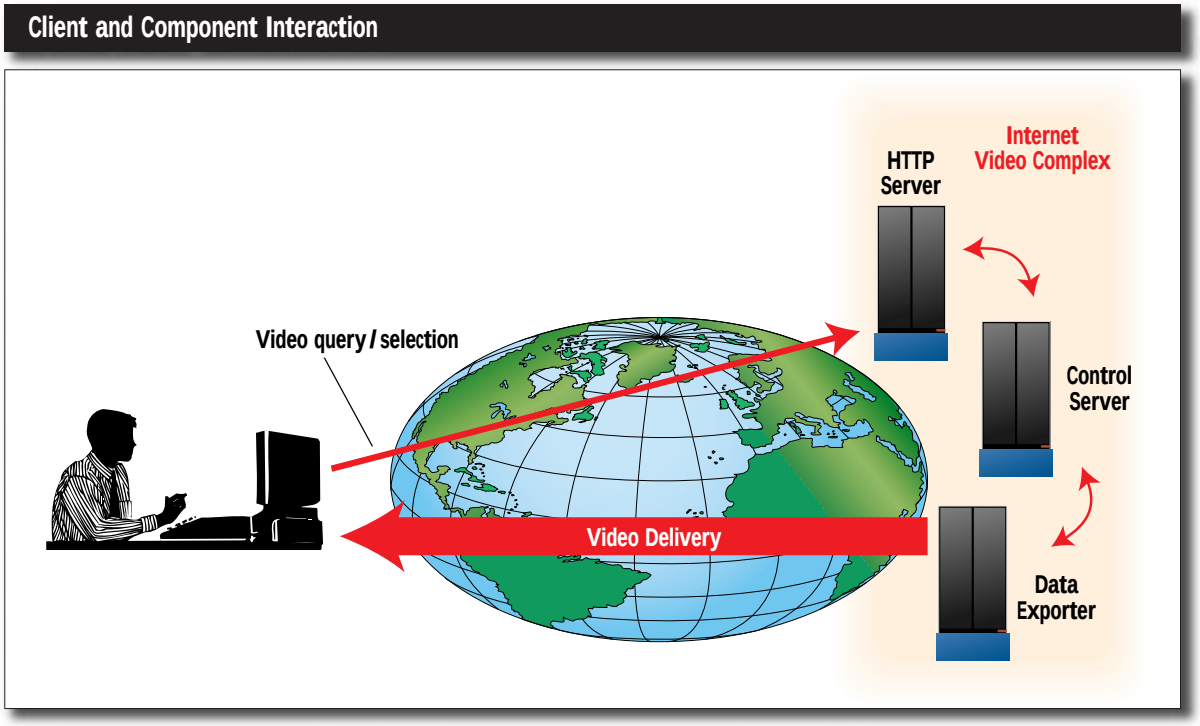

Figure 6. Interaction between client and components

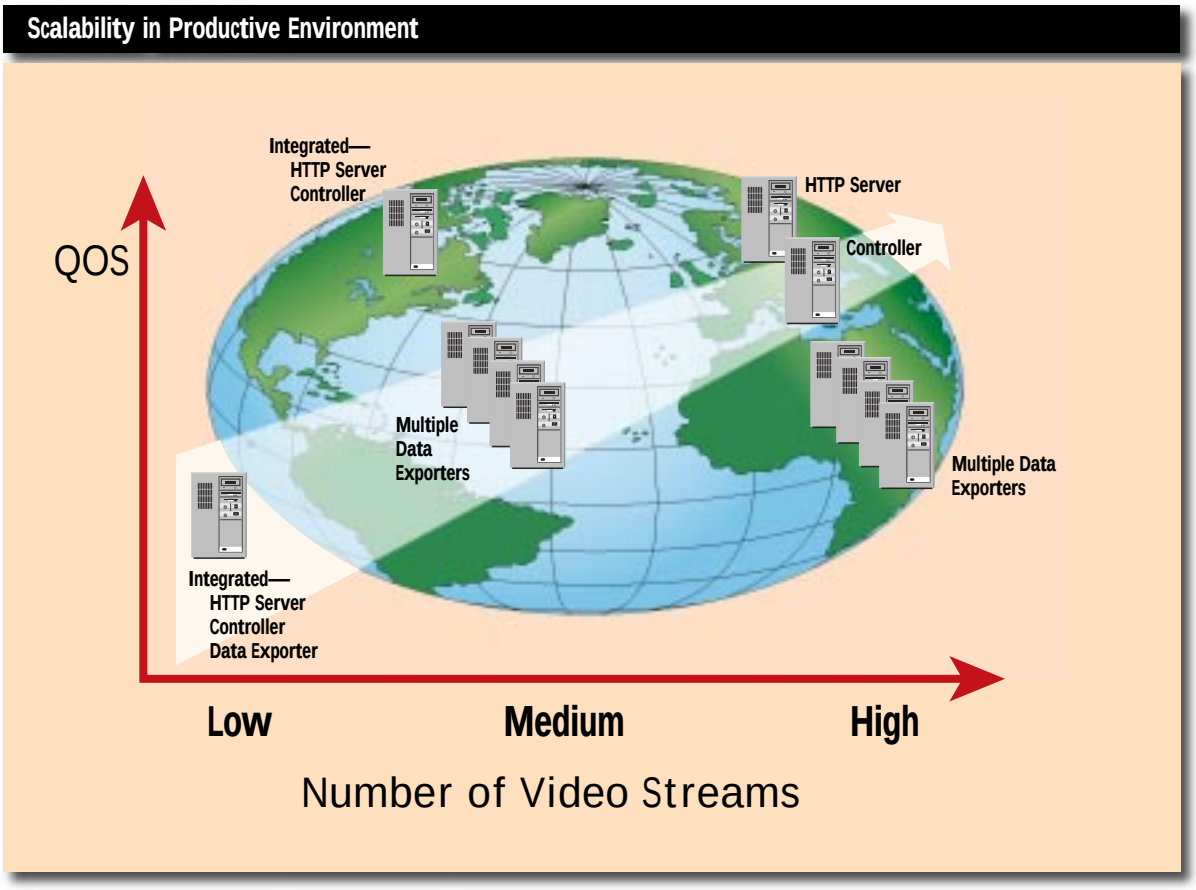

Figure 7. Scalability in productive environments

of performance with a larger number of multimedia streams. Figure 7 highlights the flexibility and scalability in production environments.

Conclusions

Multimedia is changing the way we work. It allows us to be more productive by bringing all data types to the desktop. The spiral effect of faster server systems and higher speed information networks can bring end users from anywhere closer to the real-time environment. The VideoCharger Server technology provides building blocks to launch the second generation wide area Internet-based system with dynamic content delivery from anywhere.

Eddie Ho, *IBM Corporation, 11400 Burnet Road, Austin, TX 78758. Mr. Ho is a programming consultant in the RS/6000 Executive Briefing Center. He has a BS in Computer Science from the University of Wisconsin and an MS in Computer Science from North Dakota State University.*

Sam Juliano, *IBM Corporation, 11400 Burnet Road, Austin, TX 78758. Mr. Juliano is a staff programmer in Internet Multimedia Development. He has a BS in Electrical Engineering from the University of Texas in Austin.*