

2 Technology Primer: The Basics of Streaming and Digital Media

Introduction to Streaming and Digital Media Technologies

This chapter introduces streaming and digital media technologies, defines the technical differences between the types of content commonly delivered over the Internet, and highlights the key differences between traditional broadcasting and Internet technologies. Skip this chapter if, when asked by a colleague, customer, or cousin, you can confidently explain *all* of the following:

- The basic technical concepts of digital media delivery, such as streaming, download, loss, latency, jitter, content delivery networks (CDNs), and digital rights management (DRM).
- The difference between the two Internet media delivery approaches—streaming and download—and how those strategies create live, simulated live, and on-demand media playback.
- The cost models to scale traditional broadcasting media and Internet media delivery when building a service to reach a few hundred, a few thousand, or hundreds of thousands of consumers and users.
- The three metrics on which consumers and users evaluate “content quality” for Internet delivery.
- The five strategies currently used to secure content delivery over the Internet to public and private audiences.

Understanding these five areas is essential for building a solid digital media business. If you already understand these points, then jump to Chapter 3, *The Four Keys to a Profitable Streaming or Digital Media Business*. This chapter builds on these concepts and delves into more detail about how businesses can capitalize on the scalability, quality, security, and audience measurement capability of digital media.

By the end of this chapter, you will be able to explain the broad basics of digital media technology to your boss over coffee or your mother over Christmas dinner.

Internet Basics: Built for Reliability, Not Quality

A quick overview of the Internet in general is required before learning about streaming and digital media. Although it may seem that the Internet leapt fully formed from a collective consciousness in the early 1990s, it actually began as a project of the U.S. government and MIT in the early 1970s. Called ARPANET, it was the first attempt to connect computers across the country, and eventually around the world, into a loose community of interrelated processing centers. ARPANET's objective was to build a fully redundant and fully reliable communication system that would operate even in the event of a catastrophic global crisis such as nuclear war or a meteor hitting the Earth. "Networking" is a common word used to describe social gatherings, job hunting, and computer systems, but the concept didn't exist before ARPANET.

In ARPANET and subsequently the Internet, all communications occur in a packet-based system: Data, files, and digital objects are chunked up into subsets of the whole and wrapped in digital packets. These packets are then sent over the network using standardized protocols such as hypertext transfer protocol (HTTP) and file transfer protocol (FTP). These protocols employ a series of query-response and try-retry transactions where information is stored and used.

When you click on a hotlink, your web browser finds where the information is stored using a universal record locator (URL) number and the domain name system (DNS), which connects the URL to the computer location. When a request is made, the protocols send a list of numbers that represent all the packets the requesting computer expects to receive. The content origin, where the content is stored, starts sending numbered packets that correspond to the list. The requester gets them one at a time, in no particular sequence, and reassembles the packets back into the original file.

If a specific packet is missing, the requesting computer sends out another request for the missing data. The requester tries and retries until all the packets are received or until a preset amount of time passes. At this point, the requesting computer displays a “file not found” or “server too busy” notice. When the globe icon spins away in the upper right-hand corner of Microsoft Internet Explorer, the try-retry process is occurring. Packets can be lost during the transfer process for any number of reasons, which will be discussed shortly. In fact, the Internet was designed with the assumption that packets would be lost regularly, as the industry loss average is 1 in 12. When you get a red “X” instead of a picture, the delivery has failed. Clicking the refresh icon will repeat the whole process, sometimes filling in the X. The user can then decide whether to make do with the received information, initiate a new round of requests, or give up all together.

If that doesn’t sound reliable, it’s because it’s not. The Internet isn’t reliable like turning on a light switch, which approaches 100 percent reliability for the 100-year-old technology. The builders of ARPANET and the Internet were not trying to create a system that would deliver every single packet every single time. Instead, they created an architecture that built in the assumption that any individual transaction can fail, and yet the overall network could still function. It places the responsibility on the requesting computer to confirm if everything has been received and to track what has or has not been received. The requesting computer can seek out alternate paths to the origin or even other sources of the same information.

If a particular path goes down, the requester can try a new path. No individual path becomes a single point of failure for the overall network. Theoretically each path is as good as another. With a large enough network that has enough alternate paths, the overall network itself should never fail. During the 9/11 crisis, when the phone system on the entire East Coast crashed due to suddenly oversubscribed circuits and key telecom hubs in lower Manhattan were lost, the Internet slowed down but never stopped. E-mails took longer than usual to go through, but it was a difference of a few minutes versus the usual few seconds. Information requests were rerouted around the world, served from temporary stores or “caches,” and communication continued. The same is not true for telephone service: All carriers were out of operation for a period ranging from hours to days.

In the late 1980s and early 1990s, several major computer industry trends intersected to make the Internet more than just a science experiment or military project:

- Personal computers (PCs) began to be produced in mass quantities, which dropped the average price so that anyone could buy one.

- Following Moore's Law,¹ computing power crossed a critical threshold during this period, which provided powerful processing capabilities in a small, affordable form factor as Intel released the 386 and 486 microprocessors.
- Apple and Microsoft released the Macintosh and Windows 3.0 operating systems, respectively, which significantly simplified PC usage for the average office user or home consumer.
- Desktop productivity software hit the market, which presented a productivity-enhancing reason for businesses and individuals to purchase PCs.
- IBM, Novell, and other companies developed networking software, which allowed multiple PCs to share printers, storage, and other network resources.
- In 1993 the first web browser, Mosaic, was widely adopted. It used hypertext markup language (HTML) to wrap HTTP, FTP, and other protocols into an easy-to-use and visually appealing wrapper around HTTP on Internet protocols.

Due to these advances, the Internet consists of thousands of large and small networks connected together. At one point, there were over 12,000 registered networks contributing resources to the Internet, not counting individual users or unregistered networks. Those networks themselves consist of hundreds or thousands of individual machines. Service providers with thousands of servers and millions of miles of copper wire and fiber optic cables, the “pipes” through which the packets travel, lease usage at generally affordable levels. Universities, corporations, government offices, and individuals all contribute resources to the great Internet cloud.

Static versus Streaming at the Content Level

Until the late 1990s, all content on the Internet was static. Static content, in this context, refers to information that is created and formatted in advance, made available for consumption, and then fully transferred to a consumer or a user's PC before it is used. For example, the typical web page has static content in the form of text, logos, images, and advertisements. Static content also includes individual files such as word processing documents and spreadsheets. It could also refer to a copy

1. Alan Moore, founder of Intel, theorized that the number of capacitors that could fit on a silicon chip, and therefore the computing power of the chip, would double every 18 to 24 months. A corollary to Moore's Law is that the price of silicon chips would be cut in half every 18 months as well. Moore's rough estimate has held remarkably true since the mid-1970s.

of a large database of customer information that is being duplicated from New York to Idaho as a redundant backup.

Static content should generally be delivered as quickly as possible, but once delivered, it doesn't require that an open connection be maintained to the Internet. As discussed above, when a user enters a domain name or clicks on a URL, the web browser sends a request for the web page, information is transferred from the source server to the user's desktop, and the web page appears in the browser. Once the transfer is complete, the connection to the Internet is dormant and unused, or "off," until the next web page is requested.

Some elements are animated within the web browser, but technically they are also static content. For example, a rotating CNN logo is a discrete file that is fully transferred to the user's computer, typically in the form of an animated "GIF" before being activated. The same can be true for assets that are developed in Flash, an application from Adobe. Flash objects can be highly dynamic in appearance, but have been traditionally downloaded before they are used. Adobe has since announced a new platform called Flash Media Server that also supports true live streaming.

Video and audio can also be delivered in this fashion. In a download-then-play model, an audio or video file is completely transferred to the requesting computer before the movie or sound stream is initiated. In this model, digital audio and video are handled like a duplicate customer database by the network resources in the Internet.

There are two potential problems with delivering video and audio files as if they were the same as static content:

1. **File size:** Audio/video files are much larger than most commonly used files. The CNN logo is less than 4 kilobytes (KB) in size, whereas a 30-minute VCR-quality video can exceed 300 megabytes (MB) and a full-length movie over 3 gigabytes (GB). In other words, audio and video can be a hundred to a thousand times larger than the average web page. The transfer time for this file size even over a high-speed cable modem or fast Ethernet network is measured in hours, not seconds. Users must wait for the information to become usable.
2. **File value:** Audio/video files often have a much higher inherent value than commonly used static files. Audio and video content is expensive to create or has an associated artistic value. Because the files are stored on the destination machine, content owners risk theft and illegal sharing by users. The content can then be altered or the monetary value diluted. The importance of content security is directly tied to the value of the content. Movie trailers

are freely available to anyone, and security strategies are designed to prevent content alteration. For copyright-protected or high-value content, security strategies focus on limiting access to ensure that proper fees are collected or that competitors do not gain access to sensitive information.

The security issue is not trivial. In 2001, there were few options to prevent users from duplicating, distributing, or altering the downloaded audio and video once it was completely transferred to a user's machine. The poster child of this problem was Napster, a distributed file-sharing service that allowed anyone to trade digital music files. The end result, after years of litigation, was Napster shut down by court order in March 2002² for infringing on copyright laws. The U.S. Court of Appeals for the 9th Circuit decided that, due to immature technology, Napster was unable to effectively filter out illegally distributed files. The damage was already done: The decision came after tens of millions of songs had been traded without compensating the record industry or the artists.

Enter Streaming Media

Streaming media is a strategy to help jump the hurdles of file size and file value for audio and video content. Rather than transfer information to users' computers (in the streaming industry it is referred to as a "progressive download"), it is more like television: The audio or video data are played as they are received and no data are transferred permanently to the user's computer.

Streaming overcomes the file size problem because, with streaming, users can begin to consume the media as it is received. A 600-MB file starts playing once a sufficient amount of data are transferred into a media player's temporary storage, called a cache or "buffer." The content then looks like any television or radio signal, as images and sound play out on the computer desktop.

Streaming also addresses content security. Content owners can collect fees for content usage by setting up payment gateways at the beginning of the stream. Once access is granted, the stream starts. As a general rule, content owners don't need to worry about consumers or users copying or distributing the files because the users store nothing on their machines to copy or distribute. Although there are a few shareware programs available on the underground market that *can* capture and save a video stream, the tools are not mainstream, nor are they expected to be anytime soon. Only highly sophisticated users have the knowledge or equipment to capture and retransmit the stream.

2. Napster Copyright Litigation, U.S. Court of Appeals for the 9th Circuit, March 25, 2002.

Streaming Across the Internet

While streaming media does increase security and enable delivery of large files, it also imposes a significant delivery cost to the network. As shown in Figure 2.1, the Internet infrastructure of the content delivery value chain can be segmented into four separate regions. The First Mile includes all the equipment where the information is stored, such as databases, servers, data center infrastructure, and access connections. The Internet Cloud, or Middle Mile, refers to the vast interconnect networks that make up the Internet, such as thousands of service provider transport pipes, routers, and switches, and other invisible peering points of the Internet. The Internet Edge is actually not a real “edge,” but refers to the point at which a home consumer or business user connects to the Internet at large. The Last Mile denotes the local access connection between the user and the Internet Edge.

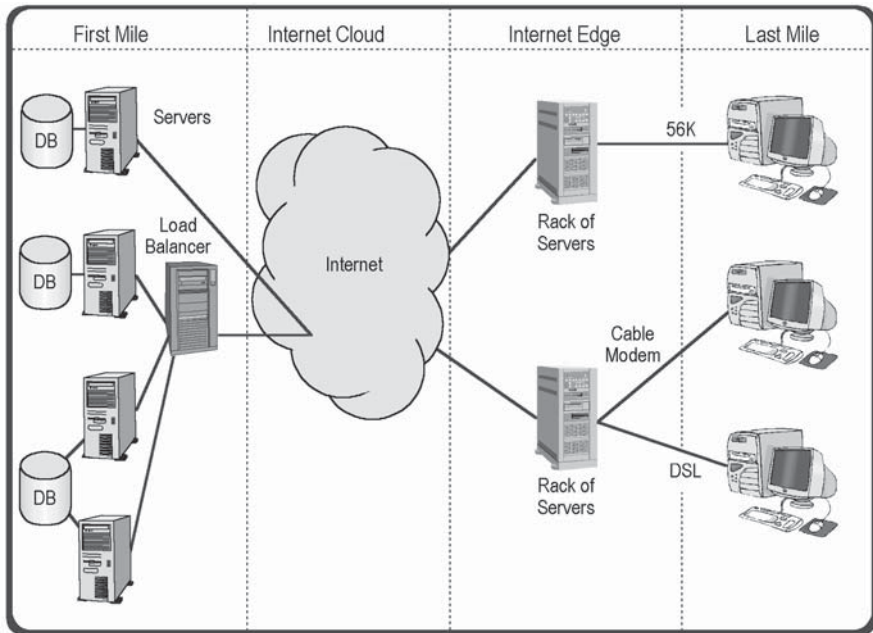


FIGURE 2.1 *The content delivery value chain.*

Behind the scenes, streaming media requires a continuous connection between the viewer and the content source. Unlike static content, the normal state of the network connection between end user and content source is “on”: A link is engaged, and network resources and bandwidth are used until the stream ends or is disconnected. To smooth playback, most streaming media players temporarily store a portion of the content in the buffer, caching anywhere from a few seconds to a few minutes depending on the technology.

The “always on” connection between the user and the content source means that a constant, steady connection must exist between the source site and the end user’s player. Because this is computer technology, the network rarely (if ever) works perfectly. Therefore, there are a host of problems that continually occur in the Internet to prevent a constant, steady connect. Network outages, congestion, or delays at any link in the chain can interrupt playback. The types of interference are generally classified into three areas: latency, loss, and re-buffering.

1. **Latency:** Due to standard Internet protocols, there is an inherent delay, or latency, when a piece of information is requested and when it is delivered. The latency mainly affects the amount of time a user must wait to see the information. In the case of streaming media, latency impacts when a consumer or user can start watching the video. The more network connections, or “hops,” that must be traveled between the requester and the origin, and the larger the file requested, the longer the start-up latency. Some authorization and security systems in the First Mile may also increase latency.
2. **Loss:** Routers, switches, caches, and other network equipment are designed to handle a specific number of simultaneous transactions. If the equipment exceeds its limit, information packets are dropped entirely from the network. The assumption is that the user’s machine will use a try-retry approach of transmission control protocol (TCP)/IP to repeat the request for a missing packet. This approach does not affect static content delivery, except to slow it down. For streaming, though, packet loss translates into skipped frames, missing audio, or stopped playback.
3. **Re-buffering:** Consistency is key in delivering streaming media. The media player’s buffer compensates for some variability by providing a local cache of information. Measuring the rate at which packets are received, the player calculates how much information to buffer before starting playback. A slow but constant connection causes a long buffering period before start-up but would not affect playback quality once it begins. Players can’t compensate if the rate changes once playback starts, and, once the cache is empty, playback stops. Any number of issues can cause the transfer rate to change during playback. Re-buffering refers to a changing transfer rate, which then results in jumpy or stopped playback.

When faced with missing data due to loss, latency, or re-buffering, the media player either stops playback or skips over that segment, presenting the end user with an erratic audio and poor quality video experience.

Internet Distance Does Not Equal Geographic Distance

With digital media, the Internet distance doesn't equal geographic distance. Whether on the public Internet, a private intranet, or a local area network (LAN), the distance between two physical points is rarely a network straight line. A user in Peoria, Illinois, accessing amazon.com may get the information from Seattle, Washington, where Amazon is based. However, to get from Peoria to Seattle, the Internet border gateway protocol (BGP) may route the user through Atlanta, Georgia; Dallas, Texas; *and* Denver, Colorado.

Thus, geographical mileage is only one component of distance, but often not the most important one. In 2006, the average Internet distance between a requesting user's location and the content's stored location was between 5 and 10 network connection hops. International connections can take significantly more. At any network hop, traffic could be slowed down, held up, or stopped altogether. With static content, the variability is not much of a problem. Assuming the information is available either at the primary location or elsewhere, Internet protocols will re-request the missing information until it is fulfilled or until a certain amount of time elapses. There may be a delay, but the content gets through.

Streaming media is not as tolerant of network delays as static content, primarily because the overriding qualification for streaming delivery is not speed of download. A 30-minute video cannot be viewed in twenty minutes. Instead, streaming media requires constant bandwidth availability and consistent connection reliability during transmission. As shown in Figure 2.2, there are five areas where providers of technologies and services have improved the capability for delivering streaming media:

1. **Higher content quality:** Content owners must balance playback quality with delivery costs and performance constraints. New encoding/decoding formats, or "codecs," are allowing content owners to do more with less, but there is still a long way to go before the bit rate, frame size, and other aspects of playback are comparable to television-quality video.
2. **Better servers in the First Mile:** Internet service providers (ISPs) and content providers increase the serving rate and simultaneous user capacity of their servers or supplement their networks with streaming-capable caches to meet the volume of incoming streaming requests.
3. **Greater backbone bandwidth in the Internet Cloud:** Larger network pipes have been installed to the point where some analysts predict that the fiber will rot in the ground before it is ever fully used. In addition, bandwidth allocation technologies allow service providers and enterprises to cordon

off segments of any specific bandwidth link for a specific application, such as streaming media. Traffic shapers can't control bandwidth usage over the entire connection between the viewer and the content source, only the links to which they are attached. However, they do help the backbone itself meet the stringent streaming requirements.

4. **Faster infrastructure connections in the First Mile, Internet Cloud, and Internet Edge:** A large bandwidth pipe was only half the battle. The Internet junctions of routers, switches, etc., have also been upgraded to be capable of consistent 500 Kbps or greater per user data transmission.
5. **Faster Last Mile connections:** This is probably the largest inhibitor technically. In response, users have gradually adopted digital subscriber lines (DSL) and high-speed cable modem services for connections faster than dial-up access can provide. With a theoretical maximum of 56.6 Kbps, dial-up services are too slow to deliver even low-quality streaming media. In 2006, half of U.S. households subscribe to broadband access, and a shift to much higher data rates in the range of 24 Mbits/sec to 100 Mbits/sec have begun. By 2010, U.S. broadband penetration of 75 percent is likely, and 10 percent to 20 percent of U.S. households will subscribe to very high-speed broadband.

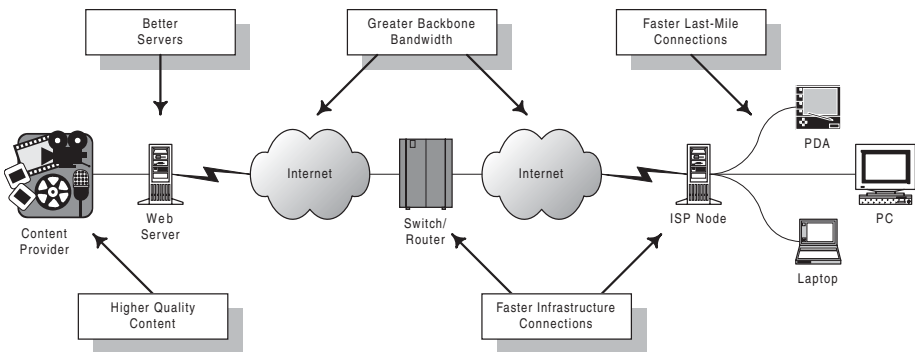


FIGURE 2.2 Internet bottlenecks for content delivery. (From Michael Hoch, *Understanding the Technology of Digital Media: A “Streaming” Primer*, June 2003, Aberdeen Group.)

Security is also a critical feature. It is so important that it will be addressed in greater detail later in this chapter and in subsequent chapters. At a high level, digital rights management (DRM) technologies have improved the speed of authorization at a streaming gateway, and have also made download-then-play a viable option. Encryption technology prevents use of files unless the user acquires a unique, registered key or digital certificate, which would then bind the media with

a specific physical machine. It is extremely difficult to decrypt the file without a key or to use an encrypted file on a different computer without first receiving a key from the content owner. Most DRM applications use 128-bit encryption or higher. The first 64-bit decryption was completed in September 2002, but only after four years of 24-hour-a-day processing by over 330,000 computers.³

Technical Variations of Internet Media from Traditional Broadcasting

It is important to understand the differences between traditional technology and the technology that powers the Internet, streaming, and downloaded content delivery. “Traditional” in this case includes nearly every distribution channel available to move audio and video content from where it is created and/or stored to where it is consumed. The most obvious channel is traditional broadcasting, where content is delivered via transmission towers and cable operators over wavelengths and frequencies. Broadcasting includes such public applications as radio and television programming. Traditional models also include the private broadcast of audio and video, such as closed-circuit TV, corporate television channels, remote training facilities, etc. Finally, traditional models can also include distribution of audio and video on physical media, such as compact discs, cassette tapes, videotapes, and digital video discs.

In this context, those who work behind the Internet intend to replace anything traditional with a network-based technology or service. There are several strong and proven reasons for wanting to replace traditional delivery models with Internet delivery. These include:

- Speed benefits, such as faster time to market and shorter production cycle times.
- Cost benefits, such as lower distribution costs, lower marketing costs, and lower customer service costs.
- Improved customer quality through direct-to-consumer marketing, increased reach to new audiences and geographic regions, and the ability to gather granular customer usage data.

3. Distributed Team Collaborates to Solve Secret Key Challenge, *RSA Security*, September 26, 2002.

- New application benefits, such as small-audience programming for entertainment or corporate usage and other types of art, entertainment, or applications that were not previously economically feasible or technically possible.

The speed benefits are readily understandable. Because any location on the Internet is accessible from almost any other location, Internet distribution avoids physical inhibitors such as the time it takes to duplicate, package, ship, unload, stock, select, and purchase the content. Content can be delivered from the origin to the user as quickly as the Internet can carry it. For private distribution, the expectation is that video over networks will become as ubiquitous as voice over a network, and that a video conversation will be as easy as a phone call.

The Fundamental Difference in Cost Structures Between the Old and New

The costs benefits and the quality differences between the old and the new models are not as easily understood. Setting aside the quality discussion, both models require upfront costs to develop and format the content. These costs may change in the Internet world, but in the current environment, most content development processes are not changing. Studio space, production houses, creative teams, and artists, as well as other costs involved in creating movie, music, advertising, corporate, and personal content are still in place. A host of technological tools are rapidly developing to help these areas, but that discussion is beyond the boundaries of this book.

The distribution, delivery, and customer/user support costs are drastically different between traditional models and the Internet world. Both have fixed and variable costs, but the relative weight of fixed versus variable is significant. In a traditional model, content owners are heavily weighted toward fixed distribution costs. These costs include physical broadcasting facilities, such as capturing and encoding facilities; licensing and transmission costs for wavelengths and channels; and maintenance and storage costs.

Traditional Broadcasting Cost Structures Are Heavy on Fixed, Low on Variable

The fixed costs to run a television or radio station for a specific region or for a closed-circuit television system for internal corporate use increases by specific steps. Initial start-up costs include reaching a specific geographic audience, either

due to the transmission radius of a broadcasting tower or the licensing agreements with a given cable television provider. As an example, WBCN in Boston is an alternative radio station broadcasting to 4.5 million listeners within the major Boston metropolitan area. The station has two towers in three locations around the region, which gives WBCN a broadcast radius of about twenty miles. Nielsen/Arbitron statistics show that WBCN has a 7.2 rating, which indicates that, out of the 4.5 million potential listeners, the station has an average of 300,000 listeners.

It is difficult to calculate the expense to set up a similar station. Because WBCN has built its broadcast network gradually over the past 25 years, the costs today would differ significantly from what WBCN actually spent. Those costs would not be transferable to other regions either, because frequency licenses are priced by metro region and real estate in Boston is much more expensive than real estate in Nebraska.

However, assume that the fixed start-up costs are sunk: They are over and done with, currently appearing in the company's financial statements as a debt liability against which payments are made on a predictable schedule. The main variable cost is the per-user cost whereby the station pays royalties to content owners based on the number of users within the geographic region. The fees vary based on the content's popularity, the station's listener base (size, demographics, etc.), the number of content plays, and the negotiation strength of the station with the content owner. WBCN's primary income, advertising revenue, is also per-user dependent, which is driven by the station's listener base and its negotiation strength with advertisers. As long as the incoming ad revenue covers three things—the outgoing royalties, the existing debt that represents the sunk costs, and the fixed costs of running the station (including sales, general, and administrative costs)—WBCN is profitable. Even as part of a large, multichannel operator such as Clear Channel, each station within the overall operation is accountable for profitability along these lines. Television broadcasting has a similar business model, with the added level of cable operator distribution costs that supplement or replace the broadcast tower.

What happens if WBCN wants to increase revenue? The most direct way is to increase its listener base. One strategy is to capture more of the 4.5 million potential listeners. The station can do this by changing formats to something more popular than alternative rock, increasing marketing, promotions, and other awareness strategies, or purchasing a competitor. In these cases, the costs would be calculated as sales, general, and administrative costs (SG&A). A project budget would be created and justified based on the overall success of the specific program. There would be no ongoing operational costs for these projects. The additional costs per additional listener would be nearly insignificant, and the additional ad revenue from reaching a broader base would (hopefully) offset that charge.

However, to add another five miles to its broadcast radius would require a significant additional capital investment. Fixed costs to build a new tower, hire operating personnel, secure real estate for the transmitter, etc., would be hundreds of thousands, if not millions, of dollars. Once the tower was up, the per-user variable cost would be roughly equivalent to the fully or partially depreciated region at a fraction of a cent per listener. The cost model then resembles a large initial step, relatively flat for a range of listeners, then moderate-to-large steps to increase geographic reach, and therefore listener base.

Streaming Media Cost Structures Are Low on Fixed, Heavy on Variable

The Internet model is nearly the opposite of the digital or streaming media model. The upfront costs can be as low as \$5,000, which include mid-range Intel-based server hardware, a copy of Windows.NET server software with Windows Media Server, and a connection to the Internet. This is a simple, straightforward, and cheap model. However, the per-user variable costs for digital or streaming media can be extremely high, primarily because delivering digital media over the Internet incurs a delivery charge for each and every person who requests it.

There are three elements to Internet delivery costs that correlate to the segments of the content delivery value chain. The First Mile costs are the servers, both for the servers themselves and for surrounding services to keep the servers going, such as personnel, data center space, electricity, and network connectivity gear (routers, switches, firewalls, etc.). Each server can handle a given number of simultaneous connections, as many as 10,000+ or as few as 100, depending on the output capacity.⁴ IT departments are cautious with their server capacity and generally add a new server when the steady load uses up 40 to 50 percent of the connections available. The objective is to limit the risk that a sudden surge in users, or “flash traffic,” will overwhelm the servers themselves. Server load balancers spread out incoming requests over the available servers.

Server expense is analogous to that for broadcasting transmitters, but on a vastly smaller scale. Once the server availability hits a specified level, a new server is added. Much like building a new tower extends the radio station’s reach to an additional 500,000 listeners, adding a new server increases the concurrent users by 100–10,000. The difference is that while increasing broadcast range can cost

4. The actual number of simultaneous connections depends on the encoding bit rate, length of file, and a series of other criteria. This is discussed in more detail in Chapter 5.

millions of dollars, increasing the number of available simultaneous connections is usually less than \$20,000, and can be less than \$10,000.

Few companies actually host all of their servers themselves. Service providers with digital media expertise handle highly specialized IT functions like media serving more effectively and efficiently than a company whose main focus is something non-IT oriented, such as producing original content or making automobiles. The two most common approaches are either (1) to outsource all hosting/serving operations, or (2) to design a hybrid model where average, daily operations are handled in-house, but service providers are engaged on an as-needed basis to increase capacity for big events or as a protection against flash crowds.

The Middle Mile costs for digital media delivery are captured in one word: bandwidth. Generally measured in megabits per second (Mbps), bandwidth costs are a function of many elements, such as the file size; distance delivered; time of day; whether it is live stream, on-demand, or download-then-play; and aggregate volume of content delivered. File size can also vary depending on the encoding rate, playback duration, and compression strategies employed. To give a rough estimate for delivery costs, prime time delivery of an Internet radio broadcast at 18 Kbps to 10,000 users can cost \$0.05 to \$0.10 per listener. Using peer-to-peer technologies and CDNs may get that down to \$0.02 per listener or lower. To the best of the author's knowledge, no Internet radio station has reached a high enough penetration rate, and therefore enough bandwidth volume, to negotiate bandwidth costs down to sub-cent per-listener costs.

Delivering radio audio streams is much less expensive than delivering video, due to the large file size. Supplemental delivery strategies, such as multicasting, CDNs, managed delivery, or peer-to-peer technologies are essential to reduce the per-user cost of delivery, and they are discussed in greater length in Chapter 5. These strategies don't have the same effect in all situations. CDNs and peer-to-peer require an audience of some scale before they become cost-effective, generally in excess of 50,000 users or more. Managed delivery solutions can aid in delivering content that doesn't need to be used instantly, but don't work in real-time streaming or webcasting.

The Last Mile delivery costs are not as immediately tangible as the First and Middle Mile costs. These costs are associated with how a company measures usage of the content, ensures playback quality at the desktop, and collects any fees that may be associated with content usage. The costs are also associated with DRM systems to secure content, as well as the damages inflicted by theft or breaches of security. The closest analogy in the entertainment world is the Nielsen ratings system, which collects a "statistically significant" sampling in order to determine overall user behavior.

Until recently, executives at most media and entertainment companies believed these Last Mile costs were too high, partly because the magnitude of the costs were unknown (How do we measure usage? How much content is being stolen?) and partly because of the lingering impact of early failures in the industry (Napster, the “free” Madonna concert of 2001). Risk of theft is playing a smaller role as DRM systems mature. A rough estimate of the Last Mile costs per user ranges from \$0.02 to \$0.10. This estimate depends upon the size of the operation, which can include DRM, monitoring, payment reconciliation systems, and customer support. Again, service providers offer packages for these systems to either supplement or replace internal systems.

Most companies interviewed by the author indicated that they use a per-user metric when determining the costs of their Internet ventures. The deciding question is: How much will it cost to build and operate a system for our user base? From there, a consumer-oriented company can determine the per-user charge for their service or the per-impression revenue they need from advertisements. For corporations and private usage, the total system costs often erroneously consider the entire system as a fixed cost. Consequently, when the company increases or expands the usage of their system, they are not prepared for the additional per-user First, Middle, and Last Mile costs involved. These companies should take their cue from consumer companies and determine where the break-even number of customers is, not just the total system costs.

In summary, this discussion makes clear that, although traditional broadcasting is concerned with covering fixed costs, the Internet world must focus on covering variable costs. The advantage of the lower start-up costs on the Internet is that businesses can be built to serve smaller audiences, both public and private. The disadvantage is that, on the Internet, the per-user costs to scale operations have driven countless companies out of business.

Why Broadcast Quality is Different From Internet Quality

Broadcast quality is as different from Internet quality as a spotlight is from a flashlight: Both do the same thing, but to radically different degrees and for widely different purposes. Broadcast quality has matured over the past fifty years to become almost as reliable as a light switch. It is assumed by any user of a television or radio that, if the signal quality is poor, it is most probably the receiver or the cable provider that is at fault. Rarely does a consumer consider calling NBC if *The Nightly News with*

Tom Brokaw is fuzzy or cuts out mid-broadcast. Instead, the television is smacked, and the cable operator's customer service lines are flooded with complaints.

On the Internet, quality responsibility is a complex mix for most consumers. The service provider is certainly the first stop, as in the television model. But what if one site works and another doesn't? The blame for poor quality quickly devolves to the site owner. If CNN.com is available and NBC.com is not, then NBC is blamed.

Quality variations as users move from page to page within a site certainly cause consternation and a wounded brand. With entertainment and media, quality is paramount. Major movie studios still will not actively use Internet distribution until the experience for the consumer is on par with other media. Poor performance can hurt a valuable brand that has been built over time.

Corporate digital media also has quality concerns, but not to the same degree. In most instances, business uses of digital media do not need to equal television or radio quality. This is because, in a business context, users have no choice in how they consume the media. There is no comparison made in the user's mind between the viewing quality of a sales presentation and a made-for-TV movie. Instead, the important threshold for corporate digital media is "business quality access"—quality sufficient to convey the information required—so that presentation does not detract or distract from the message conveyed. For example, with a "talking head" lecture series the audio needs to come through clearly and a representation of the speaker is important, but the video portion doesn't need to show every wrinkle or blink of the speaker.

Start-up Time

Users of digital media evaluate quality based on three parameters: start-up time, playback consistency, and playback size (for video). In a culture where consumers are conditioned for immediate gratification, patience is not a virtue but a sign of weakness. Users expect the same results from all of their digital media, whether it is playing from a CD in the disc drive or from a hosted server half a world away. Studies performed by various research groups in the late 1990s came up with the eight-second rule: If a user has to wait longer than eight seconds between when the content is requested and when it is received, the user will either hit the refresh button or move on to a more responsive site. For e-commerce, that means a lost sales opportunity. For business environments, that means a failed communication cutting down the productivity of their employees.

Start-up time should be measured as the time between when the request is made and when it is fulfilled. For small objects it may only take milliseconds.

As more audio and video is used, the time stretches into seconds, minutes, or even hours. Sometimes the delay is intentional: Not everything on the Internet is designed to be delivered instantly. Video delivery services such as MovieLink use download manager clients to transfer content over a period of time ranging from thirty minutes to several hours. Although this model does not fit current customer expectations, one advantage is that bandwidth costs can be lowered by a factor of ten. There are other business problems with the download model, such as poor availability of content and consumer preferences to watch motion picture entertainment sitting on the couch with a wide-screen television and home theater surround sound. However, long start-up delays certainly factor heavily into preventing wider adoption of entertainment to the desktop.

In a business environment, start-up delay isn't much of a factor. Although users lose productivity, one effective alternative is to load training materials onto CD-ROMs and videotapes and then express mail them around the country or across the globe. Corporate users need predictable delivery, not instantaneous delivery. If the video is delivered by 9 a.m. Tuesday morning, and the user knows that is the delivery time, then it can take whatever time necessary to be delivered as long as it arrives by the scheduled time. The start-up latency would then be calculated from how long after 9 a.m. the user had to wait for their delivery to arrive.

Playback Consistency

Playback consistency refers to the steadiness at which the digital or streaming media plays out after it has been initiated. Nothing starts phones ringing off the hook at call centers like interruptions, skips, frozen frames, dropped video, or other inconsistent playback. Consistency is the key. Users set their expectation of what playback should be during the first five seconds of an audio or video track. If the experience is poor, the user will not continue with the playback. If the quality is acceptable during the initial period, that quality becomes the base level against which the rest of the media experience is judged. Variations from that base cause the user or consumer to believe that there is a problem.

As in start-up time, the consumer audience is much more demanding than their business user counterparts. Whether or not a fee is involved, consumers judge playback consistency against other entertainment media. Television shows do not stop and start. Audio quality may vary from one compact disc to another, but each compact disc has the same audio quality from beginning to end.

Playback Size (for Video)

Size does matter, particularly during video playback. IP-delivered video the size of a postage stamp is no longer considered acceptable in most situations, especially not for broadband consumers or corporate users. At the same time, the video does not need to cover a full 17-inch screen. Consumers want a choice of either a higher bit rate, and therefore larger format video, or a lower bit rate and smaller format video. Advances in compression technologies from Microsoft, RealNetworks, Apple, Adobe and the MPEG-LA Group allow VCR-quality playback in a near full-screen size of 300 Kbps.

Contrary to the initial expectation, business users didn't need a screen size similar to consumers, as corporate streaming and digital media is just as effective—if not more so—in a smaller format. The reason for the smaller format is that video in a business setting is generally one part of a larger application. Video and audio add value to the overall application, but are not the primary value themselves. That is not true for most consumer uses of digital and streaming media, where the value is inherent to the content itself. Therefore, in playback size, consumers demand to see what they are paying for, either through direct fees or through viewing the neighboring advertisements. In a business environment, audio and video help convey the message, but they are not the entire message, so the playback size can be small and still be effective.

Typical Quality Measurements for Streaming and Digital Media

Content owners need to know who is watching their movies and listening to their audio streams, regardless of whether they are addressing a consumer or user audience. In the consumer world, audience monitoring and measurement is critical to charging advertising fees. Without size and demographic information, the cost per impression quickly becomes miniscule, which is why revenue from Internet banner and pop-ads dropped through the floor during the tech wreck of 2001 and 2002. About 2,000 ads were served during June 2001, but there was only a 0.02 percent click-through rate that took the viewer to the ad sponsor. Why so low? It is impossible to tell. Perhaps the ad was not visible or vibrant enough to capture the consumer's attention. More likely, the ad was shown to any random person associated with the site instead of potential customers of the advertiser.

Why Nielsen and Arbitron Have Had Trouble Measuring a Digital Media Audience

Audience monitoring and measurement for broadcast entertainment has been handed over to the Nielsen rating system and Arbitron. Using a specific statistical methodology that spreads measuring devices among a “representative population” of television viewers, these companies then collect, standardize, and aggregate the data, so that viewer’s behavior is only represented to content providers and advertisers in aggregate. It is an imperfect yet perfectly acceptable solution, particularly when there is no “feedback loop” between a television set or radio receiver and the broadcaster to tell the content provider what each user is doing.

With IP-based systems, there is a feedback loop. The TCP/IP protocols require a handshake between requester and receiver for each connection made. There is a unique IP address for every computer, network element, or server in the world. Authorization and authentication systems commonly use IP addresses and user name/password registered users to control access to content. Combining these technologies, it is possible to determine what every user is doing at any given time on the Internet.

The problem with this type of monitoring is twofold: privacy and volume. Most nations have consumer privacy acts that prevent companies from manipulating user data in a way that would target an individual user for a specific product or service. In western nations, there is constant tension between personal privacy and public responsibility. Some people express it like this: My right to swing my fist ends when it hits your face. The ability to identify any individual’s actions on the Internet raises the hackles of personal privacy advocates. Certainly the attacks of 9/11 and other terrorism threats have imposed limitations on individual privacy, because the threat of a fist hitting someone’s face feels real and present. In addition, given the differences in laws between different nations and the interconnectedness of the Internet, anti-pornography laws and distribution of other illegal content poses a legal quandary. Considering these complicated issues, the pendulum will continue to swing from side to side for some time.

The second half of the problem is the sheer volume of information that is moved around the Internet. The average web page takes 14–20 individual transactions to transfer to an end user’s machine. The latest count of web pages in the world was around nearly 5 trillion. Tracking all those transfers would lead to a tremendous amount of data collected at every second of every day. Which data are the most important? How do you cut through the noise to find the relevant information?

Most digital and streaming media content providers are satisfied with information that is very basic: the number of unique visitors, the number of content downloads, and the number of simultaneous connections. While important for capacity and operations planning by the IT departments and service providers, these numbers quickly lose their relevance in a business context. Advertisers pay to target market segments with particular demographic or behavioral criteria, not to reach 2 million faceless visitors per month.

The current trend in website viewers is to create “communities of interest,” in which individuals agree to give up a little anonymity in return for privileged access or priority service. Individual sites can then begin to understand the geographic, demographic, and behavioral makeup of their constituencies.

The sampling services of Nielsen and Arbitron are nearly ineffective in the Internet environment. Most communities of interest are too small to register as a statistically significant target group. For those groups that are large enough, the geographic spread makes targeting that group through alternative forms of media difficult if not impossible. In the end, the value of companies like Nielsen and Arbitron will be to aggregate the mass of data and extract meaningful user groups, cross-tab correlations, and other business-relevant information. Thus far, privacy and volume have thwarted most of these efforts. Again, there is a distinct difference between monitoring and measuring the consumer-oriented audience and the business environment as legal systems are still working out what is “private” within a corporate network.

Aside from legal/illegal activities, many enterprises require some sort of user monitoring and measurement to prove regulatory compliance to the relevant regulatory board, such as the Occupational Health and Safety Administration (OSHA) or the Securities and Exchange Commission (SEC). Digital and streaming media can be used for training and education in order to provide employee safety training, public disclosure of financial announcements, and other actions that comply with the particular regulatory board. Audience monitoring and management tools become essential for these applications.

Just as the privacy issue is not as important, neither is the volume problem at present. Validation must occur for perhaps hundreds or maybe thousands of employees. Rarely does the volume reach the level of even a “small” mass consumer audience of a few hundred thousand.

Content Security: Digital Media Means a User Can Record an Exact Copy

This section on technical background concerns content security. Unlike other mediums, digital audio and video content enables near-perfect duplication at a low cost. Therefore, copies of media are virtually the same as the original. Also, the original does not degrade over time or from the duplication process, so a single original can be the source for as many copies as a user wishes to make. This is why the profit margins of Microsoft and other software companies are 25–75 percent.

This is also why owners of digital content are paranoid about theft: If the copy is as good as the original, why should someone purchase the original? Below are four strategies currently employed to secure content.

Digital Rights Management and Content Encryption

Encryption scrambles the piece of content itself, so that it cannot be viewed without the decryption key. Encryption is separate from encoding: It puts the file in a virtual safe, while encoding formats the file for a particular media platform. Many DRM solutions use asynchronous systems, i.e., the key is not tied to the safe but kept somewhere else on the Internet. The user requests and receives the file but needs to perform a separate transaction in order to secure the appropriate decryption key. The danger with this type of delivery is that smart thieves crack open the safe and start spreading the file in an unencrypted format. Anarchist hackers fighting for “free information everywhere” believe it their mission in life to crack every safe and share the contents. Once a piece of content starts spreading unencrypted, the value is lost to the content owner.

DRM is undergoing a rapid evolution. As each new security system is hacked, a stronger and better one is built to replace it. Microsoft, for example, is developing a counterintuitive form of DRM for music files. This includes placing the key on the physical medium itself, so that a CD cannot be played unless the consumer has the physical CD. If the file is copied to a computer or another disc and is no longer on the same physical CD as the key, it becomes permanently encrypted. Ultimately, this system, too, will be broken, and another one will take its place. The challenge is that each successive generation of solutions gets progressively more expensive. It is conceivable that we may end up with an unbreakable system that costs far too much to use on a widespread basis.

Walled Gardens

Walled garden content security places a virtual barrier around a content set, which requires a user to purchase access with an authorization/authentication system. The *Wall Street Journal*, for example, has a public facing site, freely available to anyone. It also has a “premium content” site for which individuals pay a monthly or annual fee. Almost every type of content has a walled garden option.

The value of walled gardens is intricately tied to longevity. Once inside, a permitted user can download, copy, and distribute whatever content he or she would like. The value for the content owner must be based on this assumption. The content gains value not because it can't be stolen, but because the owner refreshes the content regularly and rapidly. The *Wall Street Journal* has new content every morning. CNN's video content appears as stories develop. Financial and religious walled garden sites constantly add new audio streams, video clips, or images to entice return viewers and continued subscribers.

The “Cheap and Easy Access” Security Strategy

Apple Computer's phenomenal success with iTunes Music Service proves there is another option for content security: To prevent something from being stolen, make it available at a price that is less than the cost to steal it. The majority of consumers in the world are honest people. Give them what they want, at a price they can afford, and in an easily consumable manner, and they will purchase it rather than steal it.

This simple system is not effective, though, when consumers have already been conditioned in one direction; for example, free music. It is also not effective for content whose value is derived from scarcity, which is why Disney only re-releases its animated classics every seven years: A new generation can enjoy it, and in the intervening years, an artificial scarcity drives up the value of the content. At some point, the recording industry will have to write off the value of all music that has already been lost to file sharing and piracy. That music is out on the Internet, and regardless of how many individuals or file-sharing services are closed down, there will always be alternatives sprouting up. Consumers are also conditioned to get music for free. Until such time as it is either cheap and easy to purchase any song or album at any time the consumer wishes, or content security is affordable and secure enough to prevent new music from being copied and spread via the Internet, music distribution over the Internet will continue unabated.

PC Client Content/Download Managers

A relatively new approach to content security is to require consumers or users to install a small content or download management client on their PC. This application is connected to a dedicated service provided by the content owner or a content aggregator, who is the wholesale intermediary between a set of content owners and end users. Download managers can facilitate content distribution by using peer-to-peer or throttling strategies to move the content to the consumer's machine. More important, they can act as the DRM enforcement engine and audience monitoring and management application as well as other functions. The client can then enforce content security, whether or not the PC is connected to the Internet.

The disadvantage of this approach is that, like any piece of software, download managers can be hacked. With enough diligence or a smart enough cyber thief, even the most secure download manager can be broken.

Case Study

palmOne

palmOne, Inc. makes products used by customers around the world. To further strengthen relationships with its customers while improving its bottom line, palmOne turned to the Kontiki Delivery Management System (DMS). Kontiki (www.kontiki.com) gives palmOne better customer support at lower cost by automating the delivery of video-based set-up guides for its products. These video guides give customers easy-to-follow demonstrations covering common support inquiries. They are available around the clock, around the world, which represents a step up in service and a step down in cost.

For makers of handheld computing devices, the most necessary feature is ease of use. palmOne broke through in this marketplace in the 1990s with a combined hardware-software solution with an intuitive interface that revolutionized the market for personal digital assistants (PDAs). Competition and technology improvements have driven the price of the PDAs down while adding more features. This attracted countless new customers who were new to the world of PDAs and often technophobic. In turn, this drove up call volumes to palmOne's support center. Calls were often longer and hence more expensive. As competitive pressure and technological breakthroughs drove down the price of palmOne products, these calls represented an ever-increasing drain on the bottom line. palmOne considered and rejected remedies such as phone trees because of low customer satisfaction with such methods.

CONTINUED ►

CONTINUED ►

For palmOne, many customer calls focused on how to set up and perform a HotSync operation that transfers information between the PDA and a desktop computer. palmOne had previously implemented set-up and support guides on its website to help users with the set-up process, but the company felt they were still getting too many support calls on this and other basic questions. palmOne needed to add a new dimension to its customer support so that buyers with questions could really see and understand basic processes without needing to talk directly with a customer support representative. palmOne wanted to be able to deliver this information without requiring significant upgrades to its website and networking infrastructure.

palmOne has always emphasized value and simplicity in its products, and knows its customers expect the same in their communications with the company. For this reason, they embraced digital video as a customer communications medium, and chose the Kontiki DMS as the solution for managing and delivering its support videos to customers. The system provided palmOne with easy-to-use application tools for the publishing and management of support videos and also helped palmOne seamlessly integrate the videos into its customer support website. The DMS also made it easy for palmOne to reach large customer audiences by tapping computers throughout the palmOne Network to deliver video content faster without expensive network upgrades.

palmOne simply added the video delivery features to its existing set-up and support guides, which offered customers an additional easy-to-follow way to receive answers to their important support questions. The Kontiki delivery grid then delivered these digital videos to palmOne's customers. Customers received these full-motion videos by simply installing the Kontiki delivery manager—easy-to-use, lightweight Windows software, which enabled efficient, secure delivery of digital video to users' desktops. The delivery manager lets users view all types of digital media and documents. This content is easily integrated into existing websites, and users just click on links to have content automatically downloaded to their machines.

This digital media system also provides palmOne with custom response bars, which appear below videos playing on the user's desktop. Using the response bars and simple HTML programming, palmOne can lead users toward other materials that further reduce the need for telephone support. The response bars also introduce customers to new palmOne products and services, thus driving revenue. The success of the initial customer support videos has led palmOne to expand its approach to include videos on how to configure and use such advanced features as WiFi, Bluetooth, and more.

The response from customers has been increased satisfaction and fewer "live" calls to support staff. In a survey of more than 900 palmOne users, two-thirds

CONTINUED ►

CONTINUED ►

of the respondents cited two features as important to them in receiving video: a high-quality, full-screen picture and the ability to watch offline and stop or pause the video. Kontiki provides palmOne with both these features, which many other video delivery systems don't offer. More than one-third of survey respondents said the video answered their questions, thus avoiding the need for personal help by customer support staff. The benefit of this system alone is saving palmOne an estimated \$370,000 per year in customer service costs. User satisfaction is also very high for delivered video support. The majority of users prefer video demonstrations over phone, e-mail, or an online knowledge-based search.

This success led palmOne to expand its use of digital media to another key part of its business—delivering new software programs and updates to the desktops of palmOne owners for installation into their products. More than 600,000 such downloads, frequently topping 20 MB each, are delivered every year. Kontiki lowers palmOne's costs for these downloads by making the most efficient use of existing bandwidth, which allows palmOne to cancel an expensive contract for additional peak-usage bandwidth. Adding it up, palmOne saves nearly two times its investment in Kontiki every year. That's why palmOne is planning additional uses for employing digital video at different points in the product life cycle for demonstrations, reseller training, and media briefings.

Case Study

Ernst & Young

Ernst & Young, one of the world's largest professional services firms, is committed to restoring the public's trust in professional services firms and in the quality of financial reporting. In more than 140 countries its employees face challenges internally as well as externally as they work with colleagues to understand the implications of recent legislation. They also face challenges as they work with clients to ensure that corporate governance and adherence to strict ethical standards form the cornerstone of their businesses.

Today, a business video-on-demand solution is allowing Ernst & Young to cost-effectively deliver high-quality video communications to all employees. Examples of these types of communications include:

- Effectively communicating critical directives in corporate governance, regulations, and strategic practices from the very top of the organization to every employee.

CONTINUED ►

CONTINUED ►

- Ensuring that members of globally distributed practice groups can stay on the same page and more effectively complete their jobs with targeted communications.
- Providing effective education of new hires, including recent college graduates, so that they meet the high standards of the firm.
- Fostering organizational knowledge and adherence to evolving best practices as determined by leaders within each business unit or industry group.
- Deploying effective “high-touch” communications and training to all employees without requiring expensive upgrades to the existing network infrastructure.

The Sarbanes-Oxley Act (SOA), passed in mid-2002, was meant to address the crisis in corporate confidence. SEC regulations are now beginning to spell out more precisely what SOA will mean to companies. But one thing is clear: Businesses must be more transparent, more accountable, and pay more attention to what is happening internally if trust is to be restored. This includes managing intangibles such as the corporate culture and organizational ethics.

Ernst & Young believes a “principles-based approach” to corporate governance requires that key changes in policies and procedures be thoroughly defined by those at the top and communicated quickly and effectively to the ranks. How should they communicate effectively with impact and power? To John McCreadie, chief information officer (CIO) of Ernst & Young, the answer was obvious—business video. As John McCreadie commented: “We believe that by using business video-on-demand for corporate communications and training, we can maximize the return on our greatest asset, our employees.” He asked the Ernst & Young IT department to investigate solutions that would allow him to utilize business video for globally distributed communications and training.

Searching for a solution

Ernst & Young had a number of technical “must haves” when it began searching for a solution. They felt the solution needed to:

- Provide sufficient functionality and flexibility to universally reach all employees in a timely manner through on-demand and “push” video broadcasting.
- Enable “narrowcast” communications that are targeted and secure so that geographically dispersed practice groups could efficiently communicate.
- Enable flexible video delivery options, including push (automatically sending video to selected recipients), pull (fulfilling user on-demand requests for video), subscription (periodic video delivery), and e-mail.
- Deliver video globally without requiring expensive network upgrades or impacting country network capacities.

CONTINUED ►

CONTINUED ►

- Support a highly mobile workforce where employees often move between offices or need to access information from locations outside of the corporate network.
- Provide the capability to protect sensitive corporate information by leveraging DRM technology and existing directory services.

The solution

Kontiki is providing Ernst & Young with a complete end-to-end solution for business video-on-demand with network-friendly delivery for all employees worldwide. The solution includes:

1. **Kontiki Delivery Management System 3.0:** The DMS 3.0 is a software suite that allows Ernst & Young to publish, protect, deliver, and measure video and other rich media. This enables video communications and training for all Ernst & Young employees worldwide. It also includes grid delivery technology, which enables secure and centrally managed delivery of video while greatly minimizing the amount of bandwidth and hardware required for full-screen, high-quality video communications.
2. **System integration:** Using XML Web Services, Kontiki is seamlessly integrated with Ernst & Young's enterprise-wide list management technology—the Global List Management System (GLMS)—to allow flexible narrowcast delivery of video to targeted groups as well as broadcast communications to the entire workforce.
3. **Business media services:** Kontiki worked together with Ernst & Young to identify opportunities that easily and cost-effectively created video in-house using standard software tools and inexpensive digital video cameras. For example, Kontiki helped Ernst & Young set up several “self-service” video production studios for less than \$10,000 each.

How has the system worked so far? The deployment has been completed to over 106,000 desktops/laptops worldwide with high-quality video communications and training successfully delivered to achieve the stated business goals.