



The New Streaming Media Economy

Why Performance Matters

by
Curtis Franklin, Jr.

Streaming media is growing in importance as a business tool for both internal and customer-facing purposes. New media and entertainment applications have been joined by education, corporate and government information, business training and travel replacement in the list of factors driving streaming media's growth.

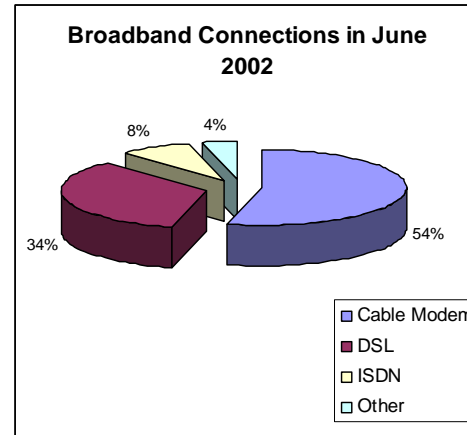
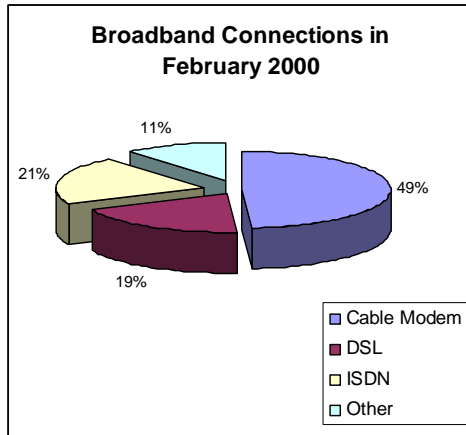
Along with the growth, however, have come significant challenges as streaming media points out to corporations and individuals the limits of their information infrastructure. These limits crash against the users' demands for smooth video, clear audio, and performance levels specified and guaranteed by contract. Media providers seeking to provide consistent quality of service face a daunting array of performance issues, caused by lack of last-mile broadband build-out to service interruptions and carrier quality of service problems.

One key to dealing with quality of service issues is accurately measuring streaming media performance from the customer perspective. With the lack of industry standardization, customer confidence rests on providing information from a trusted source. Keynote Systems provides such performance management tools, measuring performance from the crucial customer perspective in order to give both providers and customers crucial feedback on performance levels and their causes. This paper will examine some of the reasons businesses are approaching streaming media with renewed interest, the business and technical difficulties of economically delivering content via streaming media, and the tools available for measuring streaming performance from the most important perspective—the perspective of the customer.

Driving the New Streaming Economy

The streaming media industry is undergoing dramatic market growth today. After a period of exuberant expectations in the late 1990s, markets have settled into a pattern of high growth rates fueled by a variety of applications and business factors. Ananthakrishnan, Case, Cowgill and Menon in their paper "Using Streaming Media for Corporate Training" found that forty percent of enterprises report that they have instituted or are planning to institute on-line streaming media-based training programs. A growing number of people are looking to streaming media on news-oriented web sites for their news during the business day as they keep up with world events from their workstation, rather than their television set.

Based on this increasing consumer acceptance of streaming services, news and sports organizations have turned to streaming media as a "premium" service in the continuing quest for content that consumers will pay for on a regular basis. Major league baseball has scheduled the majority of its games in the 2003 season to be available via streaming media. Concerns about the cost and complication of business travel has lifted the visibility of streaming media as an alternative to in-person business meetings and briefings.



Growth in DSL and cable modem installations have helped increase the market for streaming content

Government institutions and agencies have also begun looking to various streaming media to present briefings, to train individuals in a broad range of subjects and to allow members of different government agencies to share information without requiring travel. Information on preparing for emergencies, for example, is streamed from federal sources to regional and local emergency services agencies on a regular basis. As governments at varying levels (federal, state and local) cope with the effects of reduced revenues due to recession, alternatives to travel have been sought, with streaming media through point-to-point, multi-cast and access grid topologies.

Business Challenges in Streaming Media

The streaming media market was supposed to be different. In the late 1990s, projections were put forward predicting that most news, much entertainment, the convergence of advertising and order placing, and a revolution in personal communications would have taken place by now, all built on the foundation of near-universal streaming media access. What happened?

The biggest thing that happened was disappointed customers. Individuals and companies which had been led to believe that TV-quality video would quickly come to their web browser found that broadband access appeared much more slowly than had been advertised, and that dial-up access could simply not support the promised streaming media programming. The broadband buildout was, like so many other plans, a victim of the declining economy of the new millennium. Even when buildout began to catch up with demand, companies had begun looking for real, measurable benefits to broadband services—benefits that were hard to measure without industry-standard tools.

Streaming Media – Gaining Momentum

In 2002 and 2003, broadband access is finally being installed at rate that is bringing optimism back to streaming media proponents. In November of 2002, Gartner Dataquest reported that 28% of U.S. households connected to the Internet through some form of broadband connection. Cable modem technology continues to be the dominant access medium, with 54% of the market. After its rocky start, DSL is growing rapidly,

increasing its share from 19% of the market in February 2000 to 34% in June 2002. Even with its growth, however, Gartner Dataquest says that lack of access to broadband technology remains one of the leading inhibitors to faster growth. An additional inhibitor is the perceived lack of compelling content requiring broadband access. Businesses are forming with the understanding that the market is sufficiently immature to allow total market growth to provide individual company sales growth. There is also the clear understanding that a compelling case must be made for any new media venture after the excesses of the 1990s.

Measuring Return on Investment

Part of the new discipline of Internet media companies is carefully measuring the effectiveness of any technology investment. In order to be useful, measurements must provide information on two broad topics. One is the experience of the end-user. How much delay is a user experiencing between requesting a stream and seeing it begin? How smooth is the playback once it begins? The second broad topic is the effect on the provider network. How much bandwidth is being delivered? Do any delays experienced by users represent faults within the systems controlled by the vendor, or are there issues with the Internet at large?

Another issue in analyzing the results of measurements is that there have been no uniform, accepted standards for various aspects of streaming media. In basic web pages, the “8-second rule” –end-users will wait 8 seconds for a page to load before clicking to another option – has been accepted by many as a reasonable rule of thumb for acceptable performance. No similar *de facto* standards exist within streaming media.

Streaming Media Speedbumps—Technical Challenges

Delivering a high-quality user experience is often a significant enough obstacle to prevent organizations from moving forward with a streaming media strategy. Most of the challenges revolve around acceptable performance, either achieving it, or understanding the impact on the performance of the enterprise network as a whole of delivering adequate performance.

First, IT Managers must understand the technical distinctions between streaming media and other technologies implemented through their networks. Second, there are solutions that exist today to help address these distinctions such as multicast and Content Distribution Networks (CDNs). Finally no technology implementation is complete if you can't measure what you manage. Given the additional challenges and complexities associated with streaming, an intuitive streaming management tool measured from the user's perspective is needed to validate performance.

Streaming Media —Not Your Typical Web Bottleneck

The problems begin with the fact that streaming media is fundamentally different than most files delivered across a network. Minor delays in receiving packets won't be noticed, while delays in a video stream most certainly will be. Most streaming media providers try to work around the problem of delays by using a transport protocol that allows smooth streaming performance at a cost of leaving out all the protocol pieces which work to ensure that a connection and complete transmission actually occur.

Using this protocol, called the User Datagram Protocol (UDP), removes a number of technology impediments to smooth media streaming, though it is “working without a net” for the stream itself. That risky nature can lead to issues when it comes to quality. Maintaining the Quality of Service (QoS) necessary for acceptable video has proven a special challenge as each client has requested a separate stream, each taking up a large portion of the overall bandwidth available within an enterprise network.

CDNs – a streaming solution or a performance challenge?

One of the methods used to improve streaming media performance and reduce the impact of streaming media delivery on enterprise networks is the development of Content Delivery Networks (CDNs), which distribute caching content servers between the central streaming media servers and receiving clients, to reduce total hop count and lower the stress on networks between originating server and client. Some large networks place thousands of servers worldwide in order to move content closer to consumers. Akamai, for example, placed 12,000 servers in 54 countries in building a world-wide CDN. Large infrastructure buildout did not, however, lead to agreements in which streaming content delivery providers committed to specific quality of service levels for their customers.

Quality levels are important to CDN customers because the same technologies that move content closer to the viewing workstation may act to insulate the streaming content provider from their customers. Questions of customer experience and comparisons between CDNs can be crucial to content providers working to justify expenditures, but there are few tools that allow insight into these issues. Tools that do exist, such as Keynote Streaming Perspective, are used both by streaming media providers—to determine how well CDNs move their programming to customers—and by the CDNs themselves, to support delivered performance claims.

Managing Customer Expectations – You Can’t Manage what you Don’t Measure
Service-level Agreements (SLAs) became a standard way of doing business through the turn of the century. Customers wrote agreements in which providers were contractually obligated to provide specific levels of network performance or face financial penalties. Network providers of various sorts were driven to SLAs by customers, but streaming media providers have resisted, basing their resistance on the difficulty of providing consistent, measurable performance from the server to the client.

Both CDNs and their customers have vested interests in seeing SLAs proliferate. From the customer side, SLAs can help ensure that they are receiving solid benefit from the investment they make in CDN service. In the event of performance problems, SLAs can be a valuable tool in forcing CDNs to deliver anticipated levels of streaming delivery. The CDNs themselves also benefit from agreements that allow them to offer guarantees to customers and differentiate themselves from competitors. Both sides of the equation benefit from third-party confirmation that the performance terms of an SLA are being met on a consistent basis.

Establishing a Performance Standard

One major part of the performance management problem with Web technologies such as streaming media is the nature of the Internet itself—it is impossible to say from one moment to the next precisely which links will be traversed in sending data from one computer to another via the Internet. Another significant part of the problem is in defining precisely what “acceptable performance” means from the end-user perspective. Performance management tools and an accurate measurement methodology exist today that can help overcome these obstacles.

There is significant difficulty, in quantifying both the expectation of the media consumer and the points of degradation that cause the media to fall below the expectations. The questions for those attempting to set agreements based on the user experience include precisely what to measure, where to obtain the measurements, and what limits to place on the measurements to indicate non-compliance.

There are many criteria that may be used, but initial time to startup, dropped packets, rebuffering counts (the number of times the played stream stops so that buffering can occur), and total packet counts are among some of the most important metrics used to determine streaming performance. Just as it’s important to measure a number of criteria, it’s important to put the interpretation of those measurements in terms that can be understood by executive management. Keynote Stream Quality Score and the Maximum Stream Quality score are two such interpretive tools. In the Stream Quality score, every streaming clip is measured on a scale of 1 – 10, with 10 representing DVD quality. The score gives managers the received score and the theoretical maximum, known as the Maximum Stream Quality score, given the clip’s encoding and available bandwidth. A “3.2 out of a possible 4.6” has meaning that is communicated quickly and can be compared against results at other times, or from other networks.

Measuring the End-User Experience

The “User Experience” became the watchword of Internet business in the first rush of Internet success. Because Internet media must compete with other media types, user experience is still the most useful consideration when trying to understand whether a streaming media project will succeed. Delays in stream startup, gaps in playback and errors in reproduction are all factors that can lead to a disappointing user experience. It’s not enough, though, to merely test for the existence of these factors—corporate customers need quantifiable, verifiable measurements for technical staff to act upon, or to be carried to vendors and partners for problem remediation.

Minimum Meaningful Measurement

For measurements to be useful, they must possess four minimum qualities. They must reflect, as accurately as possible, the condition being analyzed. Measurements must be reliable. Measured results must be comparable to one another. Finally, the measurements must not present information that isn’t useful to the analysis being performed. It’s quite easy for network managers to be overloaded with information

In a field without recognized, industry-based standards, the credibility of the testing methodology must be respected, especially if the results will be used to drive vendors or

partners to take action on performance problems. In many cases, a customer or service provider will be able to lean on the reputation of the impartial third party that is seen by all parties as an effective tool for resolving questions about true performance levels. Companies including as Keynote Systems, and others have released services designed to measure the performance of various aspects of streaming media delivery. The remainder of this paper will deal with one of these options, Keynote Systems' Streaming Perspective.

Keynote Systems' Streaming Perspective

Streaming Perspective is designed to give companies that provide streaming media content an understanding of the performance of their service from the user perspective. This is accomplished by measuring media clips via *agents*, measurement computers distributed around the world. The agents request media clips on a predetermined schedule and measure the most relevant parameters of the stream from a customer's perspective as it is served in response. The data generated from these measurements are then used to calculate four major and two minor target scores which then provide an overall picture of user experience quality. The hierarchy of scores, from low-level scores such as packets, bandwidth, frames, image resolution and codec, through intermediate scores such as relative packets, relative bandwidth, relative frames, video and audio rendering, and video and audio encoding, to the highest levels of Startup, audio and video, and finally the Stream Quality Score, provides customers with a range of measurements that can be used by different internal consumers for the data. The highest scores are useful for "big picture" understandings of the customer experience while the lower-level, internal scores provide technicians with useful analysis tools for problem resolution. The major target scores are:

Stream Quality – Represents the overall quality of the performance in a single number. Based on measurements at the media client program, this number is built from Video Score, Audio Score and Startup Score. Every stream receives a Stream Quality score based on a scoring scale of 0-10. A score of 10 is equivalent to full-screen, full-audio DVD quality.

Maximum Stream Quality Score—Represents the quality of the overall stream experience compared to the maximum possible score based on clip quality and encoding, and bandwidth available. This score is a useful tool for getting an immediate understanding of the user experience.

Startup Score—Represents the total initial delay experienced by a streaming media client between the request for a stream and the beginning of play. This score is based on parameters which include initial connection time, redirect count and time, and initial buffering time.

Audio Score—Represents the audio portion of the streaming media clip. This score includes parameters such as audio encoding, delivery and rendering.

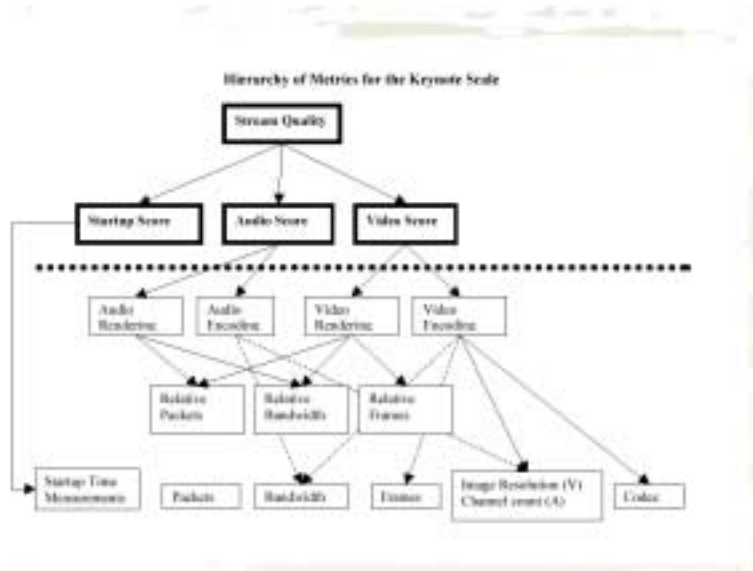
Video Score—Represents the video portion of the streaming media clip. This score includes parameters such as video encoding, delivery and rendering.

In addition, two scores are used as components of other scores. These are reported as components of Stream Quality

Rendering Score—Represents the performance of the end-to-end infrastructure in reproducing the clip at the client system. This score is designed to rate what is received, relative to what is sent. Individual score are calculated for audio and video.

Encoding Score—Represents the quality of the media clip streamed. Calculated separately for audio and video, these scores are designed to rate what the media server attempts to send.

Keynote makes all these scores, and others, available on the MyKeynote web site for each customer. A variety of graphs and charts can be customized to focus on the parameters that are of primary importance to the individual organization.



Individual scores are reported, then brought together in higher-level primary scores within Streaming Perspective.

Modeling User Experience

Streaming Perspective has been designed to model end-user experience in several key ways. First, it is based on results gathered from actual streaming media measurement agents in operation at locations around the world. This approach means that Streaming Perspective does not try to extrapolate from laboratory conditions, it measures stream quality at the place, and in the manner, that it would be viewed by a user.

By using technology for streaming that is based on the technology used in Windows Media and Real player software, as part of the measurement process, Keynote enables Streaming Perspective to avoid low-level technological obsolescence that would be an issue if Keynote had directly programmed measurement software to interact with a client protocol stack. If, for example, a streaming technology is introduced making use of

TCP rather than UDP, IPv6 rather than IPv4, or an alternative to any of the other protocols involved in streaming media, Streaming Perspective should be able to adapt as rapidly as the client software itself, and maintain the level of accuracy and repeatability present in the original service.

Multiple parameters also contribute to Streaming Perspective's accurate reflection of user experience and insulation from changing technology. If, for example, there is a particular streaming application which doesn't rely on audio—an Internet silent film festival, perhaps—the parameters can be tuned to reflect the user's indifference to that portion of the total stream. In the same way, a change in underlying protocol that makes an aspect of the stream's performance less germane to correctable concerns, the parameter weighting can be adjusted to take the fact into account.

Improvement Tool

Streaming Perspective requests a clip every hour for standard streaming files (though customers may request measurements as often as every 15 minutes), every second for live events (data points are rolled into one minute and then five minute intervals), and on individually-tunable intervals for sites with varying needs. The regular, consistent monitoring of streaming media performance allows companies to profile the performance of servers, networks and other components across time-of-day and day-of-week levels. This sort of consistent, regular monitoring is valuable when discussing problem resolution with suppliers and partners, or trying to pin down the source of performance issues within the organization.

Streaming Perspective provides a repeatable, consistent, respectable tool for measuring the performance of streaming media as it reaches the client. The combination of qualities means that the results will carry weight with third-parties that might be involved in performance issues, and that the key determinant of top-line revenue—user experience which leads to repeat customer interaction—is the key factor in streaming content performance analysis.

For more information about Keynote Streaming Perspective, please go to <http://www.keynote.com> or call 1-800-KEYNOTE.