Powering Microsoft Smooth Streaming with the Cisco Media Processor

March 2009 was the first time anyone saw live streaming content at high-definition (HD) resolutions. The technology that enabled such video delivery came from Microsoft's next-generation media delivery, Smooth Streaming. Cisco has been there since the beginning. In fact, the Cisco® Media Processor live streaming appliance was the power behind the first live Smooth Streaming events, and this role continues today.

Prior to the innovation that Smooth Streaming enabled, video on the web had always been plagued by two major problems: the quality of the video and the reliability of the playback. Viewers are disappointed if the video quality is poor (either too small or too blurry) or frustrated if the playback is unreliable (constantly rebuffering or stuttering). Both of these problems persisted with digital video for many years. Smooth Streaming solves both of these problems by making the experience of playing video on the web just like watching TV: the video plays instantly and at the highest possible quality that will reliably play back.

Figure 1. Smooth Streaming with Cisco Media Processor enabling live Smooth Streaming events

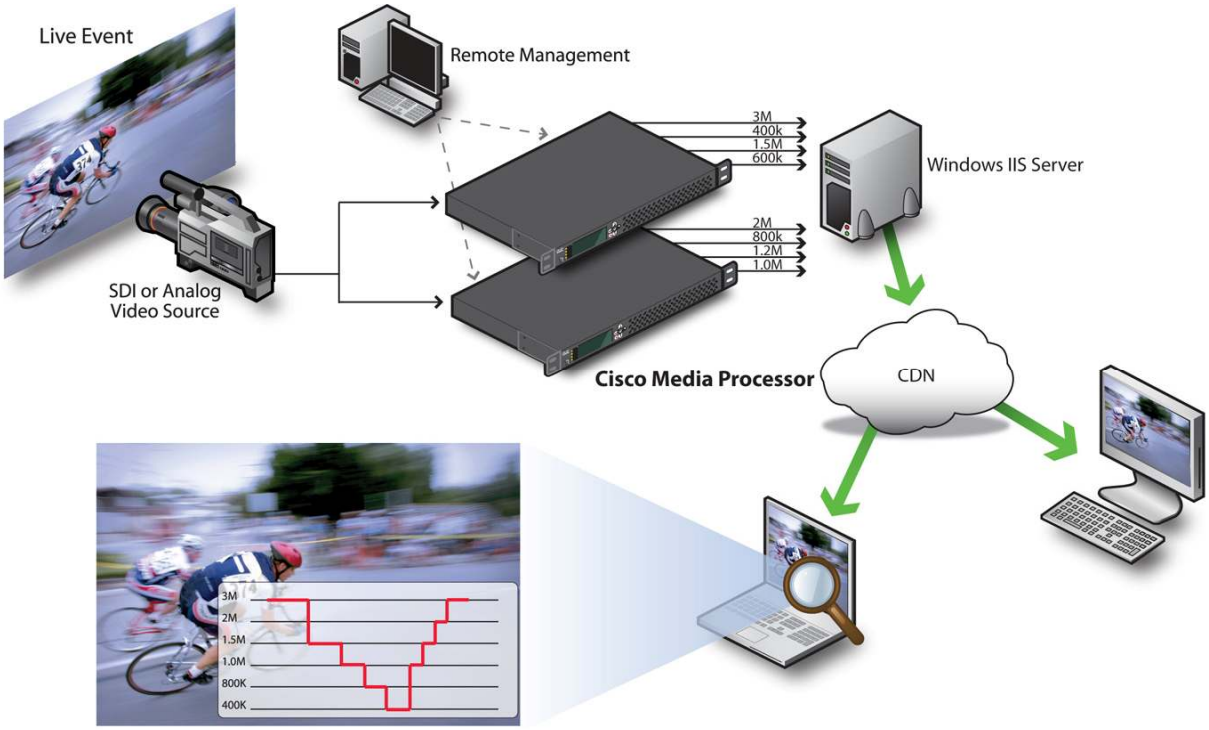

How It Works

The Smooth Streaming platform has built-in technology that dynamically detects local bandwidth and CPU conditions and transparently switches—in near real time—the video quality of a media file that a player receives. Thus, consumers with high-bandwidth connections can experience HD quality streaming while those with lower bandwidth receive the appropriate stream resolution for their connectivity. The intelligence of the stream switching is inherent in the "conversation" between the server and the player, and requires no intervention, activity, or decision from the viewer. Smooth Streaming allows consumers across the board to enjoy a compelling, uninterrupted video experience, without guessing at their bandwidth availability or locking in to a single high or low resolution. It also alleviates the need for media companies to cater to the lowest common denominator of quality level within their audience base.

Content owners can boost brand awareness and advertising revenue by extending average viewing times through higher-quality, true HD experiences (resolution of 720p or greater). They can also benefit from unprecedented network scalability using distributed HTTP-based web servers and offer better quality to more customers.

Cisco products support Smooth Streaming, providing content creators with Smooth Streaming options for both live and video-on-demand (VoD). With the Cisco Media Processor and the Cisco Media Processor Management Console, you can set up and schedule live Smooth Streaming events. With Cisco Transcode Manager, you can generate the necessary assets to support on-demand options.

Traditional Content-Delivery Platforms

Prior to Smooth Streaming, media delivery on the web used two delivery methods: progressive download and traditional streaming. Popular video-sharing websites such as YouTube, Vimeo, and even Facebook use progressive download almost exclusively. Progressive download is nothing more than a simple file download from an HTTP web server. The term "progressive" stems from the fact that most player clients allow the media file to be played back while the download is still in progress—before the entire file has been fully written to disk (typically to the web browser cache). If you pause a progressively downloaded video at the beginning of playback and then wait, the entire video will eventually download to your browser cache, allowing you to smoothly play the entire video without any problems.

HTTP is known as a stateless protocol, meaning if an HTTP client requests some data, the server responds by sending the data, but does not remember the client or its state. Each HTTP request is handled as a completely standalone one-time session.

Real-Time Streaming Protocol (RTSP), on the other hand, is an example of a traditional streaming protocol, though only one of many versions of streaming protocols for the web. RTSP is defined as a stateful protocol, meaning that from the first time a client connects until the time the client disconnects, the streaming server keeps track of the client's actions. The client communicates its actions, or states, to the server by issuing commands such as PLAY or PAUSE. After a session between the client and the server is established, the server begins sending the media as a steady stream of small information packets (the format of these packets is known as Real-Time Transport Protocol [RTP]).

A New Content-Delivery Option

Combining the best of both technologies, Smooth Streaming is a hybrid delivery method that acts like traditional RTSP-style streaming but is based on HTTP progressive download. Although it is a more evolved, advanced technology, it requires new protocols. It relies on HTTP as the transport tool and performs the media download as a long series of very small progressive downloads rather than one big progressive download. It is one type of what we call adaptive bitrate (ABR) delivery, a new and innovative way to deliver digital media and solve the problems of reliable playback and quality.

In a typical adaptive streaming implementation, the video or audio source is cut into many short segments ("chunks") and encoded to the desired delivery format. Chunks are typically 2- to 4-seconds long. At the video codec level, this length typically means that each chunk is cut along video group-of-pictures (GoP) boundaries (each chunk starts with a key frame) and has no dependencies on past or future chunks or GoPs. This scenario allows each chunk to later be decoded independently from the other chunks, but when collected and played back by the end user it is viewed as an uninterrupted video experience.

The encoded chunks are hosted on an HTTP web server. A client requests the chunks from the web server in a linear fashion and downloads them using plain HTTP progressive download. As the chunks are downloaded to the client, the client plays back the sequence of chunks in linear order. Because the chunks are carefully encoded without any gaps or overlaps between them, the chunks play back as a smooth video.

The "adaptive" part of the solution comes into play when the video or audio source is encoded at multiple bitrates for each 2- to 4-second chunk. The client can now choose the chunks that suit its needs best—the best bitrates and resolutions that it can handle at that moment. Web servers usually deliver data as fast as network bandwidth allows. The client can easily estimate user bandwidth and decide to download larger or smaller chunks ahead of time. Additionally, the size of the playback and download buffer is fully customizable.

Summary

With Smooth Streaming, the experience of watching video online now has the quality and reliability of the experience traditionally associated only with a TV environment. Smooth playback starts nearly instantly, never pauses or has to buffer, and offers the highest possible quality video without the users having to guess their connection speed.

Whether you need to deliver live or on-demand content, Cisco advanced encoding solutions produce exceptional quality video quickly and efficiently

Give your viewers the quality, scale, and interactivity they deserve. For more information, visit Cisco at Cisco Media Processor Family page or contact your local Cisco account representative.

Selecting the Right Product for Your Smooth Streaming Needs

For Video-on-Demand (VoD)

If you have larger volumes of content, such as extensive video or movie archives, or those who wish to leverage the additional benefits of Inlet's Pre-Encoding Analysis and Post Encode features to help sort, manage, automate and quality control (QC) your content, the Cisco Transcode Manager portfolio is the ideal solution.

For Live

Cisco has a wide variety of solutions for delivering live content. The Cisco AS5100 Series Media Processor supports standard-definition (SD) Smooth Streaming. The Cisco AS7100 or AS8100 Series Media Processors are required for high-definition (HD) Smooth Streaming. If you have an IP Multicast video infrastructure, you should select the Cisco AS6000 Series Media Processor, which also supports high-definition (HD). The add-on Cisco Media Processor Management Console also allows you to manage and synchronize multiple media processors with one easy-to-use interface.

Advantages of Adaptive Bitrate (ABR) Delivery

For the Content Owner

- Lower costs to deploy because ABR can use generic HTTP caches or proxies and does not require specialized servers at each node
- Better scalability and reach, reducing “last-mile” concerns because it can dynamically adapt to inferior network conditions as it gets closer to the user’s home
- Adaptability, because it adapts content to the user, rather than requiring content providers to guess which bitrates are most likely to be accessible to their audience

For the End User

- Fast startup and seek times because startup and seeking can be initiated on the lowest bitrate before moving to a higher bitrate
- No buffering, no disconnects, and no playback stutter (as long as the user meets the minimum bitrate requirement)
- Transparent bitrate switching based on network conditions and CPU capabilities
- Consistent and smooth playback experience

Encoding Smooth Streaming Assets with Cisco

Live Events and 24-Hour Feeds

The Cisco Media Processor was the first in the industry to support Smooth Streaming, and it continues to power most live events that use the Smooth Streaming platform. The Cisco Media Processor and Smooth Streaming delivered the Vancouver 2010 Winter Olympics for NBC and CBC, as well as the weekly Sunday Night Football Extra product for NBC Sports.

The Cisco Media Processor allows content owners to control all the settings necessary to encode and publish smooth streams to a Microsoft Internet Information Server (IIS) or content delivery network (CDN). Additionally, the Cisco Media Processor Management Console, the Cisco dashboard for managing multiple encoders, allows customers with large-scale broadcasting needs to control multiple encoders simultaneously. This function enhances the Smooth Streaming experience by allowing you to use multiple processors for one event to provide a wider range of resolution options, guaranteeing your viewers the best possible consistent playback and highest quality. The following sections outline the core components. Similar to the Cisco Transcode Manager VoD products, the Cisco Media Processor creates manifest files and sends them to the IIS server along with the video streams, letting the server and eventually the client player (for example, Silverlight) know which resolutions are available for playback.

Video-on-Demand

Cisco Transcode Manager, the Cisco Video-on-Demand (VoD) solution, supports Smooth Streaming asset creation and allows you to create jobs and templates that identify all the necessary bitrates and resolutions desired, to quickly generate Smooth Streaming VoD assets. The Cisco Transcode Manager also allows you to manage Smooth Streaming encodes across a farm of dedicated encode servers, resulting in a significant increase in productivity and faster time to market.

These VoD solutions convert media assets—both audio and video—into all the resolutions and formats required by your specific application. In addition, the Cisco Transcode Manager automatically create manifest files, which are also required to support Smooth Streaming. The Cisco Transcode Manager even allows you to configure customized manifest files, enabling content creators to make unique user experiences without reencoding the collection of assets.

A typical Smooth Streaming media asset, therefore, consists of the following files:

- MP4 files containing video and audio.
  - *.ismv: Contains video and audio, or only video.
    - One .ismv file per encoded video bitrate.
  - *.isma: Contains only audio.

  Note: In videos with audio, the audio track can be muxed into an .ismv file instead of a separate .isma file.
- Server manifest file.
  - *.ism.
    - Describes the relationships between media tracks, bitrates, and files on disk.
    - Used only by the IIS Smooth Streaming server—not by the client.
    - Client manifest file.
  - *.ismc.
    - Describes to the client the available streams, codecs used, bitrates encoded, video resolutions, markers, captions, etc.

Smooth Streaming Manifest Files

As stated earlier, the Cisco Media Processor creates manifest files and sends them to the IIS server along with the video streams, letting the server know which resolutions are available for playback.

The Smooth Streaming Wire/File Format specification defines the manifest XML language as well as the MP4 box structure. Because the manifests are based on XML, they are highly extensible. Among the features already included in the current Smooth Streaming format specification is support for:

- VC-1, WMA, H.264, and AAC codecs.
- Text streams.
- Multilanguage audio tracks.
- Alternate video and audio tracks (that is, multiple camera angles, director’s commentary, etc.).
- Multiple hardware profiles (that is, same bitrates targeted at different playback devices).
- Script commands, markers or chapters, and captions.
- Client manifest Gzip compression.

- URL masking.
- Live encoding and streaming.

  Note: Both manifest file formats are based on XML. The server manifest file format is based specifically on the Synchronized Multimedia Integration Language (SMIL) 2.0 XML format specification.

A folder containing a single Smooth Streaming media asset might look something like the picture shown in Figure 1.

Figure 2. Folder with a Single Smooth Streaming Media Asset

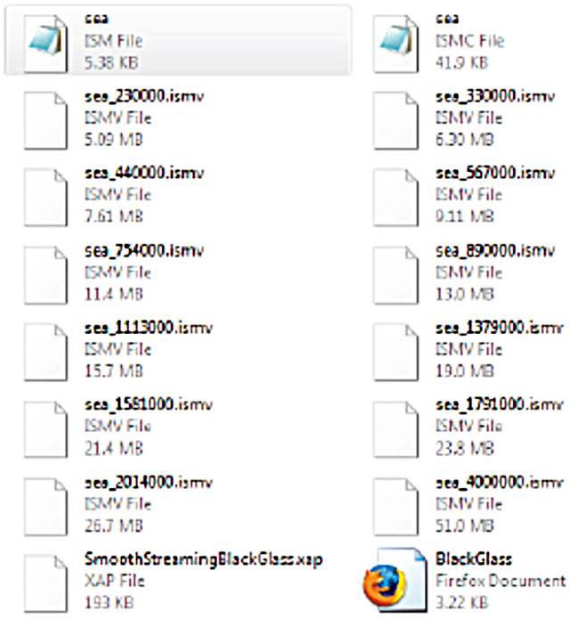

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C22-675929-00 06/11