Sun Ray™ Multimedia Enhancements

Technical Brief

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.
650-960-1300

Version 2
November 2008

Contents

Introduction

Users have come to expect powerful video performance on all kinds of devices, including thin clients. However, video decoding is processor-intensive and transmitting video streams is bandwidth-intensive, so the fact that each server is expected to host many users can cause service levels to become untenable, especially when multiple users contend for resources as they view videos while performing other tasks. To deliver improved performance on thin client devices, developers need to resolve resource consumption issues and optimize multimedia delivery without compromising the cost and design advantages of the thin client computing model.[1]

The traditional method for handling multimedia streams has been to decode them on a server, rendering video into a virtual frame buffer and sending raw pixels and audio to the client asynchronously. Although this use of server-based decoding was a breakthrough when first introduced, it is now often characterized by sub-optimal frame rates, lack of audio/video synchronization, and high demand on shared processing, memory, and network resources.

Some vendors try to avoid server-based decoding by means of extra-cost options, such as operating systems, applications, and hardware add-ons, to enable video decoding on the client. But once a local operating system and applications are installed, the thin client model is compromised by the need for local updates and administration. Such client devices can no longer really be considered thin but fat or, at best, mixed. (It is axiomatic that true thin clients are stateless—they have no local operating system, applications, or storage.) On the other hand, the hardware-based decoding capabilities of a thin client device should be utilized when possible.

The challenge, then, is how to a provide a multimedia experience sufficiently rich to satisfy corporate and user needs on a mix of thin client devices, across operating system platforms, while retaining all the advantages of true thin client computing. This paper describes the Sun Ray solution.

1. Benefits of thin clients as opposed to fat or mixed clients include freedom from desktop hardware and software replacement cycles, elimination of the need for physical maintenance of desktop devices, and reduced vulnerability to theft of data.

Multimedia Redirection and Optimizations for Windows Users

For those who use the Sun Ray Connector for Windows OS to provide a Windows desktop environment, Sun Ray 2 series desktop units (*DTU*) remove the burden of decoding the video stream from the server. Video streams created in *H.264* or *VC-1* formats are streamed from the *Windows Terminal Server* to the Sun Ray Connector for Windows OS, using a separate channel, independent of the main *RDP* connection between the Windows Terminal Server and the Sun Ray Connector for Windows OS.

Piping video streams directly to the DTU gives a ten-fold decrease in CPU and bandwidth usage compared to the traditional *bit scraping* method. Decoding on the DTU takes full advantage of the H.264 or VC-1 compression and the on-board processing power of Sun Ray 2 class devices. This solution provides the maximum overall performance increase in terms of both server CPU usage and network bandwidth utilization and provides a high-quality viewing experience without requiring local hardware add-ons, operating systems, applications, or storage.

For formats other than H.264 and VC-1, sessions hosted via the Sun Ray Connector for Windows OS fall back to *YUV* when possible, providing optimization similar to that available for Solaris- or Linux-based Sun Ray sessions.

DirectShow Plug-in for Windows

Microsoft's standard RDP protocol, which connects Windows desktop sessions to the Sun Ray Connector for Windows OS, has no particular optimizations for media streaming; it must be extended to allow for a high-quality user experience. Sun Ray Server Software 4.1 and the Sun Ray Connector for Windows OS 2.1 use a DirectShow filter and new Windows components to implement the required extensions. The DirectShow filter intercepts audio and video data, which it sends to the Sun Ray Connector for Windows OS. From there, the audio is sent to the audio device, and the video is sent to the X server's *XVideo* extension, which then sends YUV directly to the DTU. H264 and VC-1 are sent to the DTU through the X server's *XVEnc* extension.

Fallback for Non-optimized Streams and Devices

When a stream is encoded with a *codec* that is not optimized by Sun Ray Server Software, or when non-Sun Ray 2 client devices, such as Sun Ray 1 DTUs, are used, the DirectShow filter falls back to YUV decoding and sends a YUV data stream to the client. When the YUV fallback is not available, for instance for WMV playback on Sun Ray 1 DTUs, the standard RDP path is used.

Server-based Decoding for Non-Windows Users

For non-Windows implementations, limited CPU and bandwidth resources call for a more efficient method of server-based video decoding. Sun's approach for Solaris and Linux installations that do not connect to Windows sessions relies on optimized server-based decoding and compression to improve multimedia performance.

XVideo and YUV

Computer output has historically occupied the *RGB* color space, but modern video formats such as H.264 and *VC-1*, and image storage formats such as *JPEG*, use *YUV* representation. YUV reduces the amount of color information in a signal that cannot be perceived by the human eye, allowing video streams to be compressed without a perceptible loss in image quality.

Before decompressed video can be rendered on a screen, it needs to be converted from YUV pixels to RGB. Sun Ray Server Software 4.1's *XVideo* extension, based on the *Xorg* X server, allows video in most formats to be decoded into YUV on the server and then sent directly to the Sun Ray DTU in this format.

Offloading two key parts of the video display process, color space conversion and upscaling, from the server to the DTU shifts the processing load downstream. Since videos are typically displayed at their encoded size or larger, transmitting the frames at their decoded resolution and then scaling them on the DTU relieves the server of the additional resizing overhead. It also reduces the network bandwidth load. YUV image representations are typically half the size of the equivalent RGB images, and they correspondingly use about half the bandwidth. Network traffic is not affected by users changing the size of the video output on the screen.

Using YUV and XVideo also increases frame rates, for example, 30 frames per second compared to 15 frames per second for RealPlayer video without using XVideo. This solution provides an improved video experience on both Sun Ray 1 and Sun Ray 2 series devices and accelerates video content encoded in all Solaris or Linux RealPlayer-supported formats.

Fallback for Other Formats

Videos encoded in formats other than H.264/AVC and Windows Media Video 9 fall back to YUV decoding when that option is available to the application on the server. If YUV decoding is not available to the application, the screen is bit scraped: Sun Ray Server Software finds the changed bits on the screen, encodes them, and then sends them to the DTU. This is the least optimal scenario.

Summary

Compressing and redirecting video streams based on file format, device type, and operating system platform helps to optimize Sun Ray multimedia performance. Sun Ray multimedia enhancements use an accelerated YUV path for Sun Ray 1 series DTUs and an accelerated H.264/VC-1 path for the Sun Ray 2 series, where H.264 and VC-1 codecs are supported in the DTU hardware. Other media formats use a standard RDP path. The Sun Ray multimedia redirection component supports Windows Media Player 10 and 11.

FIGURE 1 Overview of Redirection of Windows Media Input by Device Type

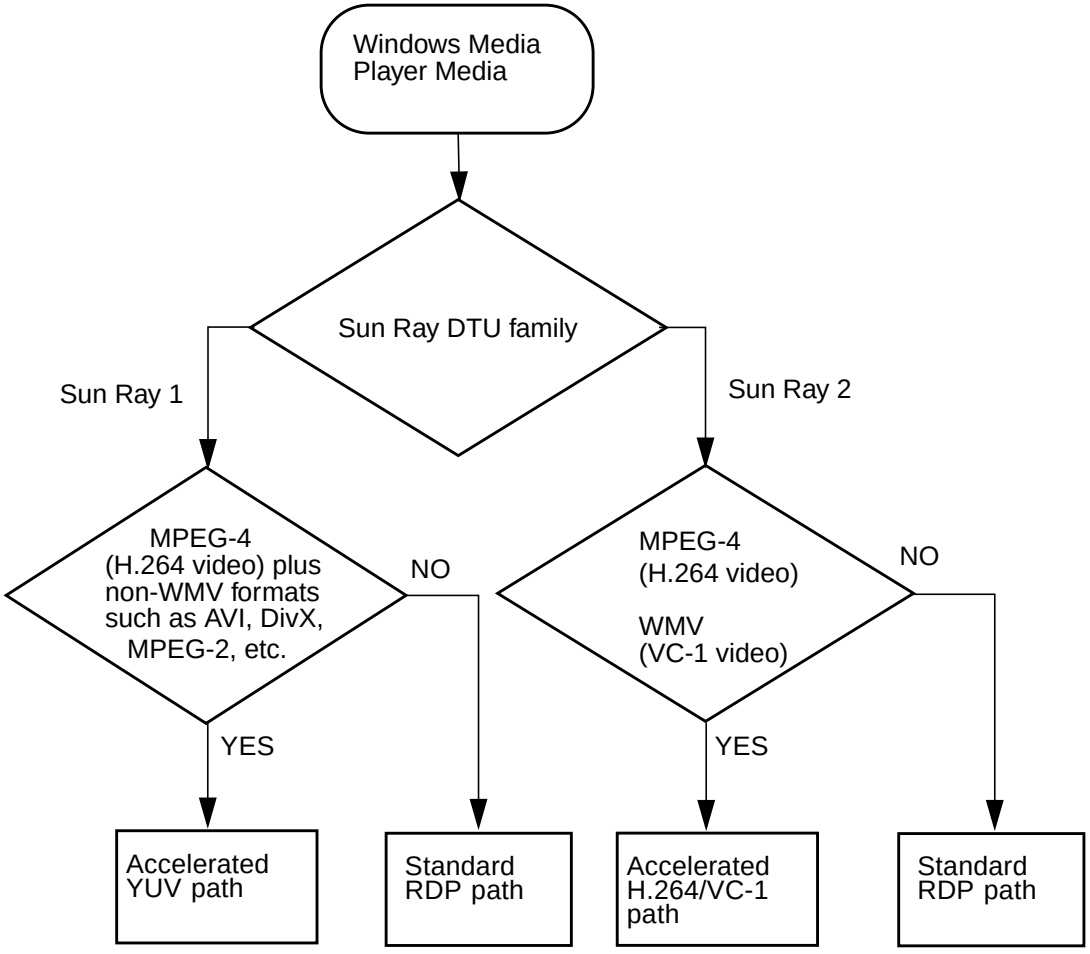

YUV Video

An accelerated path for YUV video delivery enables improved playback of video formats such as MPEG-1 and MPEG-2 by reducing the bandwidth required to deliver decoded video to the Sun Ray DTU. The accelerated YUV path is used automatically so long as the correct software decoders for the particular video format are available and the software is configured to make use of the XVideo extension. The following YUV formats are supported:

- Planar: YV12, I420
- Packed: UYVY, YUY2

H.264 or VC-1 video playback on a Sun Ray 1 DTU, which does not have a hardware decoding capability, uses software decoding and the accelerated YUV path.

VC-1 Video

Sun Ray 2 family DTUs support all Simple and Main VC-1 profiles, up to the following levels:

TABLE 1 Supported VC-1 Video Levels

Profile	Video Level
Simple Profile, Low Level =	176x144 pixels (QCIF) at 15 frames per second
Simple Profile, Main Level =	352x288 pixels (CIF) at 15 frames per second
	320x240 pixels (QVGA) at 24 frames per second
Main Profile, Low Level =	320x240 pixels (QVGA) at 24 frames per second
	352x288 pixels (CIF) at 30 frames per second

VC-1 Simple/Main profiles are compatible with the Windows Media Video 9 (WMV9) format and also use the hardware decoding in Sun Ray 2 series DTUs. At this time, `.wmv` files that are not VC-1 encoded cannot take advantage of accelerated playback.

H.264 Video

Baseline, up to level 2.0, is the best profile supported by the current Sun Ray hardware. Certain videos encoded in the Main profile may play; however, the Sun Ray decoder does not support *CABAC* encoding or data partitioning, and cannot decode high-profile streams. Unsupported streams result in a black window or an error reported to the player. For best results, video files should be encoded in Baseline profile at up to 352x288 pixels (CIF) and 15 frames per second (fps).

Conclusion

Sun Ray multimedia enhancements use accelerated video streams to provide fast, reliable, high-quality multimedia performance without compromising the cost, maintenance, and security benefits of the Sun Ray true thin client computing model. In cases using older model DTUs which lack video decoding capability, the standard display mechanism offers a lower-quality, but still acceptable, user experience.

Glossary

AAC Advanced Audio Coding, a "lossy" compression format capable of delivering relatively high quality at relatively low bit rates.

ALP The Sun Appliance Link Protocol, a suite of network protocols that enable communication between Sun Ray servers and DTUs.

CABAC Context-adaptive binary arithmetic coding, a "lossless" entropy coding technique used in H.264/MPEG-4 AVC video encoding.

bit scraping The traditional server-based method for handling multimedia streams, in which the server software finds changed bits on the screen, then encodes them and sends them to the thin client device.

client Normally, this term refers both to the physical hardware, such as a Sun Ray thin client desktop unit, and the process that accesses resources such as compute power, memory, and applications from a server. The server may be located remotely or locally. In the present context, the Sun Ray DTU is a client of the Sun Ray server; the Sun Ray Connector for Windows OS is a Windows Terminal Server client.

client-server A common way to describe network services and the user processes of those services. Although this term can apply to a wide range of interactions between desktops and larger computing facilities, the thin client model suggests that all, or nearly all, computing be performed on the server.

codec A device or program capable of encoding and/or decoding a digital data stream or signal.

DTU See *Sun Ray DTU*

H.264 A standard for video compression developed by MPEG and VCEG for a wide range of bit rates and resolutions. Also known as MPEG-4 AVC (Advanced Video Coding) and MPEG-4 Part 10.

JPEG A common, “lossy” method for compression of photographic images. The acronym stands for the Joint Photographic Experts Group that originated the standard.

MPEG A series of standards for compression of video images. The acronym stands for the Moving Picture Experts Group of ISO/IEC.

RDP Microsoft Remote Desktop Protocol.

RGB A device-independent color space (red, green, blue), used for representing images.

server Generically defined as a network device that manages resources and supplies services to a client. This paper refers in particular to the Sun Ray server(s), which host Sun Ray sessions as well as DTUs, and to Windows Terminal Servers, which act as hosts for Windows applications that can be reached by RDP clients, of which the Sun Ray Connector for Windows OS is an example. The Sun Ray DTU is a client of the Sun Ray server; the Sun Ray Connector for Windows OS is a Windows Terminal Server client.

server farm A cluster of servers linked with load balancing software.

Sun Ray DTU The desktop unit, originally known as the desktop terminal unit, is the physical appliance used to transmit keystrokes and mouse events to and receive display information from the Sun Ray server. The Sun Ray DTU hardware has a built-in smart card reader, and some models also contain a flat-panel display. Sun Ray DTUs are also known as Sun Ray ultra-thin clients and Sun Ray virtual display clients.

Terminal Server client The client software used to access remote sessions hosted on a Windows Terminal Server, in this case, the Sun Ray Connector for Windows OS.

thin client Thin clients remotely access some resources of a computer server, such as compute power and large memory capacity. Sun Ray appliances rely on the server for all computing power and storage. Within the client-server computing model, thin clients are distinguished from fat clients by the absence of local operating systems, applications, disc drives, fans, or other devices that fat clients need in order to operate.

VC-1 A video codec specification developed by Microsoft and others and implemented as WMV9 (Microsoft Windows Media Video 9) to deliver higher performance than earlier standards such as MPEG-2.

Windows terminal Any device used to access Windows applications residing on a Windows Terminal Server.

Windows Terminal Server A server that hosts Microsoft applications for remote terminals or clients.

WMA Windows Media Audio data compression file format and codec developed by Microsoft.

WMV9 See *VC-1*.

Windows Media Player A digital media player and media library application developed by Microsoft. Windows Media Player runs on Mac OS, Mac OSX, and Solaris operating systems as well as on Microsoft Windows.

XVEnc An ActiveX control used to compress audio files.

XVideo An extension to the X server used to rescale video playback. XVideo helps to reduce server-side processing and network load by enabling users to resize video on Sun Ray 2 DTUs.

Xorg The X server in the X.Org Foundation's open source public implementation of the X Window System (X11). The *XVideo* extension in Sun Ray Server Software 4.1 is based on Xorg.

YUV The color encoding system used for analog television.