

White Paper

Akamai HD Network

Encoding Best Practices for Akamai HD for iPhone and iPad



Table of Contents

EXECUTIVE SUMMARY	1
Step 1: Planning a Mobile Content Strategy	2
Step 2: Deciding on a Business Model	2
Step 3: Deciding on a Delivery Model	2
VIDEO STREAMING FOR THE IPHONE	3
Video Support	3
Audio Support	4
Resolution	4
Delivery	4
User Experience	4
IPHONE & IPAD STREAMING CAPABILITIES	4
Introduction	4
Architecture Overview	4
BEST PRACTICES	6
Video Encoding and Bit-Rates	6
CONTENT CREATION	7
Encoders	7
Stream Segmenter	7
Content Origin	8
Ingest to NetStorage	8
Storage Considerations	8
Client Playback	8
HTML 5 Video Support	8

Executive Summary

The consumption of mobile content around the globe is growing rapidly, thanks to the introduction of more open, Internet based smartphones like the iPhone. These devices expand the opportunity for content publishers to leverage the public internet (rather than closed-walled garden approach of carriers) to extend audience reach to new platforms.

Video capabilities on mobile handsets have increased significantly. According to a recent study, 75% of worldwide data-capable handsets are able to download video clips, and almost 50% support streaming video. Apple's iPhone represents a large and growing share of this market when measured by number of handhelds. But iPhone users are much more active on the network than others, generating a disproportionate amount of mobile web and application traffic. With the introduction of the iPad, it's anticipated that this trend will continue as consumers leverage this new device to consume a variety of rich media content such as video, audio, electronic books etc.

In addition to screen size and ease-of-use, Apple's integral support of WiFi appears to be a big contributor to increased mobile internet usage. WiFi handheld traffic is growing rapidly, from 3% in late Q3 2008 to 8% in mid Q4 2009. Apple products represent almost 80% of worldwide WiFi requests by handhelds, 50% from iPhones and another 30% from the iPod Touch.

The opportunity for mobile digital media and applications is growing as manufacturers add WiFi functionality to this leading edge of smartphones. With iPhones out in front, now is the time for your business to explore and plan a mobile content strategy.

Since the iPhone/iPad offers the optimal user video experience at this time, this guide will focus primarily on delivering content to iPhone/iPad users. However we will discuss other mobile devices in passing.

This paper is intended as a best practices guide to optimizing video streaming content to the iPhone & iPad 3.0+ OS, and leveraging Akamai HD for iPhone streaming solution.

Step 1: Planning a Mobile Content Strategy

The iPhone has demonstrated that consumers are willing and able to download and view a wide variety of content and applications on their handsets, a trend that will be reinforced as the iPad penetration grows within the consumer market. Windows of opportunity range from two to three minute breaks while waiting for something or someone, to an hour or more while commuting or other travel. Prior to the iPhone & iPad access to mobile content was primarily available through 'on-deck' services offered by the carriers. Carriers would determine which brands and content services are available to consumers from their mobile internet home pages. However, there is now a growing shift to 'off-deck' or 'direct-to-consumer' content through brand owned sites and applications. This has enabled new business models, revenue streams and delivery options for both 'on-deck' and 'off-deck' content.

Step 2: Deciding on a Business Model

There are three primary business models in today's mobile content market, those being;

1. Free content: if the content provides direct support for a product or service, then you might distribute it for free.
2. Paid for content: another model is to charge a fee directly to the customer for the content or the application. Variations on this include subscriptions, rental charges for viewing for a limited time, or an outright purchase.
3. Advertising supported content: supporting content distribution with advertising is a frequently considered business model. Advertising options include banner ads, 'pre-roll' of ads before the main video or more complex integration such as graphics overlays on the video content.

Choosing the appropriate business model for your content is the first step.

Step 3: Deciding on a Delivery Model

After choosing your business model, the next step is content delivery. 'On-deck' delivery requires partnership with the carrier, but can provide more details on user location, demographics, subscription, and billing options. 'Off-deck' delivers in two ways: using standard web protocols and browser; or with a custom application that gives you greater control over quality and user interaction.

A few examples highlight these choices:

Long-playing videos, such as movies or full TV episodes, have typically been downloaded to the handset and stored until the user is ready to play them. This involved a two-step process that goes first to a desktop and then to the handset when docked (for example, via iTunes). With the availability of wifi enabled smartphones and faster wireless 3G data speeds, direct to handset streaming delivery over standard web protocols is becoming increasingly popular for full TV episodes. In this model, the video begins to play as soon as enough content has been received.

Short videos and time-sensitive video clips, such as sports or news highlights, are often delivered by streaming, which means they begin to play as soon as enough content has been received. These typically have to fit into brief windows of opportunity, with selection, download, ad display (if any), and main video completed in less than two or three minutes.

Simulcast of live sports, news and events to mobile handsets requires special considerations around scalability and performance of the delivery mechanism. Prior to the release of the 3.0 OS, live streaming required a custom proprietary solution. iPhone/iPad 3.0+ OS now has built in support for live video streaming that dynamically adapt to available bandwidth, optimize full screen video playback and increase battery life.

Video Streaming for the iPhone

Video and audio capabilities of the iPhone is geared towards creating and providing the best media experience on a mobile device. iPhone 3.0+ supports several video formats for progressive download playback, including H.264 and MPEG-4, in .m4v, .mp4, and .mov file formats.

Video Support

iPhone

Full-screen video playback of movie files with the .mov, .mp4, .m4v, and .3gp filename extensions, using the following compression standards:

- H.264 video, up to 1.5 Mbps, 640 by 480 pixels, 30 frames per second, Low-Complexity version of the H.264 Baseline Profile with AAC-LC audio up to 160 Kbps, 48kHz, stereo audio in .m4v, .mp4, and .mov file formats
- H.264 video, up to 2.5 Mbps, 640 by 480 pixels, 30 frames per second, Baseline Profile up to Level 3.0 with AAC-LC audio up to 160 Kbps, 48kHz, stereo audio in .m4v, .mp4, and .mov file formats
- MPEG-4 video, up to 2.5 Mbps, 640 by 480 pixels, 30 frames per second, Simple Profile with AAC-LC audio up to 160 Kbps, 48kHz, stereo audio in .m4v, .mp4, and .mov file formats

Akamai recommends the use of H.264 encoding in an MP4 container file for iPhone video content. This combination has excellent compatibility and delivers higher quality at lower bitrates than MPEG-4.

iPad

- Support for 1024 by 768 pixels with Dock Connector to VGA Adapter; 576p and 480p with Apple Component AV Cable; 576i and 480i with Apple Composite AV Cable
- H.264 video up to 720p, 30 frames per second, Main Profile level 3.1 with AAC-LC audio up to 160 Kbps, 48kHz, stereo audio in .m4v, .mp4, and .mov file formats; MPEG-4 video, up to 2.5 Mbps, 640 by 480 pixels, 30 frames per second, Simple Profile with AAC-LC audio up to 160 Kbps, 48kHz, stereo audio in .m4v, .mp4, and .mov file formats; Motion JPEG (M-JPEG) up to 35 Mbps, 1280 by 720 pixels, 30 frames per second, audio in ulaw, PCM stereo audio in .avi file format

Audio Support

iPhone

- AAC, Protected AAC, MP3, MP3 VBR, Audible (formats 2, 3, and 4), Apple Lossless, AIFF, and WAV


iPad

- Frequency response: 20Hz to 20,000Hz
- Audio formats supported: HE-AAC (V1), AAC (16 to 320 Kbps), Protected AAC (from iTunes Store), MP3 (16 to 320 Kbps), MP3 VBR, Audible (formats 2, 3, and 4), Apple Lossless, AIFF, and WAV
- User-configurable maximum volume limit


Resolution

The native video display format of the iPhone is 480 x 320 pixels, or about half the resolution of standard TV. When viewed sideways the aspect ratio of the screen is 3:2, which is slightly wider than standard TV, and narrower than widescreen/HDTV. The native video display format of the iPad is 1024x768.

Comparing Video Resolution



Comparing Video Aspect Ratio



This means that images in unmodified standard-aspect videos will appear slightly squashed (shorter and wider). Widescreen videos will either display with black bars on the top and bottom, or the picture will have the left and right sides clipped

Video content in unmodified standard-aspect videos will appear slightly squashed (shorter and wider). Widescreen video will either display with black bars on the top and bottom, or the picture will have the left and right sides clipped.

Delivery

All media delivery to the iPhone is over the Akamai HD Network, a HTTP based network. Connectivity is available through EDGE, 3G, and Wi-Fi networks – optimal encoding takes available bandwidth into account, with 2 to 4 bit-rates that are optimized to different connections.

User Experience

The iPhone OS uses its built-in Media Player application to play media; even when video is embedded in a web-page. Video automatically expands to the size of the screen and rotates when the user changes orientation. The controls automatically hide when they are not in use and appear when the user taps the screen.

With the iPad, video is by default played within a window in the browser/application. The user can choose to playback in full screen mode or not, at their discretion.

iPhone & iPad Streaming Capabilities


Introduction

Today you can send streaming audio and video over HTTP from an ordinary web server for playback on iPhone, iPad, iPod Touch, or other devices, such as desktop computers, without the limitations of Progressive Downloads.

The streaming protocol made available in 2009 supports Multiple Bitrates and automatically switches to the optimal bit-rate based on network conditions for a smooth quality playback experience. The goal here is to reduce the amount of rebuffering events, and time spent rebuffering, that has traditionally plagued video playback experience particularly for Live streaming. This implementation also provides for media encryption and user authentication over HTTPS, allowing publishers to protect their work.

Architecture Overview

The encoder (or a separate segmenter process) will produce H.264/AAC content in a sequence of small content segments, in MPEG-2 TS format (.ts). There is also an M3U8 index file that references the segments; in the case of live content the M3U8 is continuously updated to reflect the latest content.



Best Practices

Video Encoding and Bit-Rates

User's accessing content on the iPhone, iPad or iPod Touch have widely varying download speeds depending on whether they accessing the content over Edge, 3G or WiFi. These settings are applicable to both live and VoD content.

Figure 1 below contains recommended encoding settings to use when creating HTTP Live Streaming media. These settings apply to both live and prerecorded (video on demand, or VOD) encoding.

The provided settings are grouped according to whether the content is intended to be streamed over the Cellular or Wi-Fi network, whether the content is for iPhone/iPod Touch or iPad, and whether the content is 4:3 or 16:9 aspect ratio.

The following audio and video formats are supported:

- Video: H.264 Baseline Profile Level 3.0
- Audio: HE-AAC or AAC-LC up to 48 kHz, stereo audio OR MP3 (MPEG-1 Audio Layer 3) 8 kHz to 48 kHz, stereo audio

Note that the bit rates are general suggestions, and should vary depending on target audience, content subject matter, and specific encoder behavior.

H.264 Encoding should be single-pass Baseline Profile, frame re-ordering disabled. Key frames are suggested every 5 seconds, ideally an even divisor of the chosen segment length.

iPhone/iPod Touch

	Dimensions	Frame Rate*	Total Bit Rate	Video Bit Rate	Audio Bit Rate	Audio Sample Rate	Keyframe	Restrict Profile to:
Aspect 16:9								
Cell	480x320	N/A	64	N/A	40	22.05	N/A	N/A
Cell	400x224	10	150	110	40	22.05	30	Baseline, 3.0
Cell	400x224	12-15	240	200	40	22.05	45	Baseline, 3.0
Wifi	400x224	29.97	440	400	40	22.05	90	Baseline, 3.0
Wifi	400x224	29.97	640	600	40	22.05	90	Baseline, 3.0
Aspect 4:3								
Cell	480x320	N/A	64	N/A	40	22.05	N/A	N/A
Cell	400x300	10	150	110	40	22.05	30	Baseline, 3.0
Cell	400x300	12-15	240	200	40	22.05	45	Baseline, 3.0
Wifi	400x300	29.97	440	400	40	22.05	90	Baseline, 3.0
Wifi	400x300	29.97	640	600	40	22.05	90	Baseline, 3.0

iPad

	Dimensions	Frame Rate*	Total Bit Rate	Video Bit Rate	Audio Bit Rate	Audio Sample Rate	Keyframe	Restrict Profile to:
Aspect 16:9								
Cell	480x320	N/A	64	N/A	40	22.05	N/A	N/A
Cell	400x224	10	150	110	40	22.05	30	Baseline, 3.0
Cell	400x224	12-15	240	200	40	22.05	45	Baseline, 3.0
Cell	400x224	29.97	440	400	40	22.05	95	Baseline, 3.0
Wifi	640x360	29.97	640	600	40	22.05	90	Baseline, 3.0
Wifi	640x360	29.97	840	800	40	22.05	90	Main, 3.1
Wifi	640x360	29.97	1240	1200	40	22.05	90	Main, 3.1
Aspect 4:3								
Cell	480x320	N/A	64	N/A	40	22.05	N/A	N/A
Cell	400x300	10	150	110	40	22.05	30	Baseline, 3.0
Cell	400x300	12-15	240	200	40	22.05	45	Baseline, 3.0
Cell	400x300	29.97	440	400	40	22.05	95	Baseline, 3.0
Wifi	640x480	29.97	640	600	40	22.05	90	Baseline, 3.0
Wifi	640x480	29.97	840	800	40	22.05	90	Main, 3.1
Wifi	640x480	29.97	1240	1200	40	22.05	90	Main, 3.1

**Frame Rate: Assumes current frame rate is 29.97. For other frame rates see below:*

For	For 150k Use	For 240k Use	For All Else Use
30	10	12 to 15	30
60	10	12 to 15	30
29.97	10	12 to 15	29.97
59.94	10	12 to 15	29.97
24	8	10 to 12	24
23.98	8	10 to 12	23.98

Content Creation

Encoders

Delivering live video to the most popular mobile device on the market doesn't have to be complicated. That's why Akamai works with best-of-breed partners like Inlet Technologies. Inlet Technologies has taken the difficulty, complexity and headaches out of the mix with their Spinnaker™ encoding product.

Apple's iPhone 3.0+ OS and iPad enables a new approach to delivering video to handsets. Both the Apple iPhone and iPad now supports next generation adaptive delivery for live video - with or without an application on the iPhone. Inlet has worked directly with Apple on this solution and Spinnaker is the de facto choice for broadcasters, media & entertainment enterprises and sports leagues or delivering content to the iPhone. Spinnaker supports adaptive bitrate (ABR) delivery technology, allowing viewers to move seamlessly from one resolution and data rate to another without buffering, pausing, or missing a moment of their video. This approach provides the best viewing experience possible, while offering exciting features we've come to expect in other devices, like the ability to pause/rewind/fast forward through content, as we do on our DVRs.

Stream Segmenter

Apple has created the Media Stream Segmenter, a command-line tool that segments Live MPEG-2 transport streams (containing H.264 video and AAC, HE-AAC or MP3 audio) for delivery, and also creates the associated M3U8 index. The Segmenter receives an MPEG-2 transport stream over UDP. The index file and media files can then be deployed using common web server infrastructure. The Media Stream Segmenter can also encrypt the segments using AES-128 encryption.

This approach should work with many encoders, as MPEG-2 transport stream support is wide-spread. Some encoder vendors are building this functionality into their product, so no separate segmentation process is needed.

In addition to the Media Stream Segmenter, Apple has also created the Media File Segmenter. The Media File Segmenter (mediatilesegmenter) divides a .mov, .mp4, .m4v, .m4a, or .mp3 file into small media segments and creates an index file. The index file and media segments can be deployed using almost any web server infrastructure for streaming to iPhone and Mac OS X 10.6 Snow Leopard. The Media File Segmenter only produces VOD streams.

Content Origin

As this content is delivered over HTTP, it can be served from nearly any origin web server, including Apache or IIS. The segmenter can be run locally, or configured to output content to a shared folder on the web server, or you can use a variety of methods suitable to your OS to move the content.

Akamai HD Network supports serving content from NetStorage, a scalable cost-effective storage solution. NetStorage is transparently replicated across the globe, making it the best option for managing and delivering high-demand content.

Ingest to NetStorage

Akamai NetStorage is an ideal place to ingest the content, making it available to the Akamai HD Network. HTTP is the preferred protocol for live ingest of segmented content, as it is a simple protocol with low overhead, wide support, and compatibility with most firewalls and network configurations. The content can be sent to NetStorage through an HTTP PUT or POST, and old segments can be deleted with HTTP DELETE. Several encoder vendors have added support for this mechanism and can directly publish encoded content to NetStorage in real-time. For On-demand content, customers typically choose FTP for transferring files to NetStorage.

Storage Considerations

Dealing with many small files can be a challenge for any storage system. Apple and Akamai both recommend 10 second segments as the best length for balancing performance and user experience. Shorter chunk sizes will incur additional transfer overhead and result in more file entries. It is also recommended that each content item is stored in its own sub-directory, making performance issues from excessive file entries less likely. For long-running events consider rolling to a new folder periodically.

Client Playback

The iPhone and iPad directly supports this segmented streaming approach and will launch when a user enters the M3U8 link into Safari. Publishers may also choose to write their own native iPhone and iPad apps that are able to seamlessly browse content, leveraging the Media Player for content rendering.

HTML 5 Video Support

HTML 5 is the latest revision of the HTML standard. Of relevance is the new video tag, which allows video content to be directly embedded into a page, without the need for proprietary runtimes or plug-ins. It is still in the last stages of the design process, however some browsers are beginning to add support for it, starting with Apple's Safari 4 browser, which includes Safari on the iPhone 3.0+ and iPad.

Like HTML object tags, the video tag has a 'fall-through' mechanism. If the video tag is not recognized, the browser will attempt to render the next child tag. This is extremely useful, allowing a runtime like Flash to handle playback on browsers that do not yet support HTML 5 video playback. The video tag can also fall-through across multiple sources, in the event that a browser does not support a particular file format or compression. The HTML 5 specification defines the syntax for the video tag, but it is up to the browser which media formats to support.

The Akamai Difference

Akamai® provides market-leading managed services for powering rich media, dynamic transactions, and enterprise applications online. Having pioneered the content delivery market one decade ago, Akamai's services have been adopted by the world's most recognized brands across diverse industries. The alternative to centralized Web infrastructure, Akamai's global network of tens of thousands of distributed servers provides the scale, reliability, insight and performance for businesses to succeed online. Akamai has transformed the Internet into a more viable place to inform, entertain, advertise, interact, and collaborate. To experience The Akamai Difference, visit www.akamai.com.

Akamai Technologies, Inc.

U.S. Headquarters

8 Cambridge Center
Cambridge, MA 02142
Tel 617.444.3000
Fax 617.444.3001
U.S. toll-free 877.4AKAMAI
(877.425.2624)
www.akamai.com

International Offices

Unterfoehring, Germany	Bangalore, India
Paris, France	Sydney, Australia
Milan, Italy	Beijing, China
London, England	Tokyo, Japan
Madrid, Spain	Seoul, Korea
Stockholm, Sweden	Singapore



©2010 Akamai Technologies, Inc. All Rights Reserved. Reproduction in whole or in part in any form or medium without express written permission is prohibited. Akamai and the Akamai wave logo are registered trademarks. Other trademarks contained herein are the property of their respective owners. Akamai believes that the information in this publication is accurate as of its publication date; such information is subject to change without notice.