Real-Time Video Services and Distributed Architectures

Abstract

As customer expectations for entertainment content evolve, service providers must transition from the role of basic access providers to full-fledged "experience providers." That means being able to deliver the full range of video, voice, and data services to any device, both inside and outside the home, whenever customers choose. To deliver an outstanding video experience, service providers need tools to effectively "define" the experience at the video headend, a next-generation, application-aware video network to "preserve" the experience as it is transported to subscribers, and superior equipment in the customer home to "realize" the experience exactly as intended.

The need to define a superior video experience for subscribers is driving a new generation of video technologies that offer consumers more personalized and interactive video content than ever before. In fact, as more and more service providers enter the TV market, the ability to offer personalized, on-demand entertainment has become a primary competitive differentiator. Today, service providers that offer on-demand content are seeing growth in average revenue per user (ARPU), increased customer loyalty, and reduced churn. In the future, as consumers gain the ability to access "anything-on-demand" content across the PC, the television, the networked home, and mobile devices, these capabilities will become a core requirement for service providers of all types. However, while new on-demand video technologies hold great potential for service providers, they also present new technological challenges.

Conventional solutions for real-time video services have focused on centralized architectures for both content storage and streaming. Whether the video servers are placed in a national, regional, or local headend location, the architectures are centralized, with the assumption that the application supports a non-real-time, hierarchical distribution of entire titles.

This paper outlines a more effective content management strategy, in which content is distributed in real time and on demand in smaller than whole-title units of granularity. This alternative architecture benefits from a more intelligent caching approach that allows service providers to benefit from both centralized content storage (reducing content management capital and operational expenses) and distributed streaming (which allows for much more efficient and cost-effective bandwidth management). At the same time, this distributed architecture preserves both the end-to-end short latency and unlimited content diversity that today's emerging generation of subscribers demands. In addition, the architecture allows service providers to implement a variety of advanced services, such as remixing of "long-tail" content (niche or user-generated content), which requires short-latency access to subsegments of content in deep content stores. Finally, this paper presents the Cisco® Content Delivery System, which embraces this distributed video architecture to provide service providers with a more efficient, scalable, and cost-effective solution for defining and delivering an extraordinary video experience.

Evolution of Real-Time Video Services

Today's service providers are undergoing a rapid transformation in the services they offer to subscribers, the networks that support those services, and the devices subscribers use to access them. As subscriber expectations for accessing and interacting with content evolve, service

providers are transitioning from the role of basic access providers to full-fledged "experience providers." That means delivering to subscribers a true "connected life," in which the full range of video, voice, and data services can be accessed on any device, both inside and outside the home, whenever customers choose.

Delivering video effectively is central to a service provider's ability to succeed in this evolving paradigm. However, to deliver a compelling video offering, carriers must manage video bandwidth consumption effectively, provide the broad array of content that subscribers demand, and deliver consistently high quality across all regions and devices.

For many years, video service has been defined as a service that distributes broadcast content. Within this paradigm, all elements of the service are centrally defined and controlled. A content owner creates the service – that is, a channel with a serialized delivery of "programs," with interspersed "commercials," with delivery timing that correlates to the time zone of the anticipated recipient of the content (for example, Pacific, Mountain, Central, and Eastern for North America). Times of the day are classified as "prime time" and "non-prime time," and content and commercials are customized accordingly. Content owners use services such as Nielsen Ratings to measure and classify channel popularity. Advertisers generate revenue based upon "popularity." More recently, technology innovations (such as the digital insertion standards SCTE35, SCTE30, and DVS 629) have allowed service providers to deliver more personalized advertising content at the neighborhood or regional level.

Although this paradigm has dominated video services for decades, a new type of video service called video on demand (VoD) has emerged in recent years. VoD is a personal service; each subscriber controls both the type and timing of content viewing. The amount of on-demand content provided by operators has rapidly increased from just a few thousand hours (both free and pay) to tens of thousands of hours today. Considering the amount of potential on-demand video content that has been created since the beginning of movies and TV, service providers should expect this content store to expand exponentially.

Most operators have traditionally thought of these two types of video services broadcast and VoD as separate, disconnected services. They do share a common edge distribution infrastructure (coaxial cable, air, digital subscriber line, fiber, and so on) and display device (the TV only, until recently). But otherwise, operators have created parallel infrastructures over the years to implement these two services.

However, one new, increasingly popular feature in video services is changing everything: pause. When subscribers press the pause button, they transition from a broadcast service to an on-demand service. After pressing pause again, subscribers might now be consuming the content a few seconds or a half hour behind the live program. Although this pause capability has become immensely popular with subscribers, it also presents compelling opportunities for service providers. In the case of advertising, for example, if the play point is just a few seconds later (within the time frame of the live ad), then the system can play the same zoned ad. However, if the ad is no longer live, and if the subscriber has not been presented with the original zoned ad, then an opportunity exists to substitute a more personalized ad. In any case, the two previously disconnected services broadcast and VoD are now a single combined service. Events have been set in motion for a converged architecture that delivers all services.

Pause directly identifies the rapid mass-market adoption of time-shift TV, commonly called digital video recorder (DVR). In the past several years, the general population has become increasingly more familiar with the DVR phenomenon. The primary attraction of DVR is that it allows

subscribers to control when they can view content. This capability benefits both subscribers and content owners because, with the multitasking demands of today's society, subscribers often are unable to view content according to broadcast timelines. Thus, content owners get a larger audience, and subscribers get to view (and discover) more of the content they enjoy. Evidence is beginning to mount that time-shift TV advertising also gets watched more that subscribers are actually much more attentive while fast forwarding or rewinding and often stop and replay an ad on a topic of personal interest.

Time-shifting provides the capability to manipulate linear content streams, providing capabilities to pause or re-start linear program streams in real-time. Some service providers currently offer time-shifting services for select content, for example, allowing a user who tunes in late to a program to begin watching from the beginning of the program. Another advantage of the time-shifting phenomenon is the potential to complement traditional DVR services by delivering recorded content to PCs, handhelds, and other devices beyond the set-top box. Because there are some "fair use" legal questions surrounding time-shifting, service providers may need to negotiate content usage rights with content owners.

The recent announcements of content owners (for example, Disney and ABC) making content available within 24 hours of broadcast (or immediately, as is the case with Time-Warner Cable's "Start-Over" service) is a leading indicator that many content owners have begun to accept this new reality: customers are demanding and paying for content to be viewed when they want it, on the device of their choice. And, the time-shift experience suggests that subscribers want pause, rewind, fast forward (and even more subscriber-centric services such as previous chapter, next chapter, previous episode, and next episode) available for all content, both live and recorded. The days of pushing out 250 to 500 broadcast channels on a fixed schedule to the subscriber are gone. One might think about the emerging video paradigm more accurately as the era of billions of personalized channels.

Finally, service providers are seeing the emergence of Video 2.0 a new paradigm in which video no is longer downloaded from a central distributor using one-way connections to isolated customers, but instead becomes interactive and communal. Web 2.0 transformed the Internet from a repository of static published content to a dynamic, two-way exchange, in which consumers of content became creators of content, initiating new types of content, communities, and interactions. The rapid rise of YouTube, which went from zero to 120 million video streams downloaded daily in just a few months, highlights the new ways in which consumers are thinking about video content and the huge potential opportunity for the companies that can capitalize on this change in perspective.

What do these variants of real-time video service mean for content storage? Simple: storage requirements will grow exponentially, based on subscriber demand. And, from the perspective of the subscriber, when content is discovered, it is requested "now," with expectations for delay to picture of just fractions of a second.

As the technological and cost barriers for production of content continue to lower for both professionals and amateurs, service providers can expect rapid, continuous growth in the base of content producers. Text (blogging), audio (podcast), and photographs (FlickR) already have greatly increased in popularity. User-generated video is just beginning to become popular, but indicators of the rapid increase of video content already are present, as evidenced by the recent purchase of YouTube by Google. In the future, this rapid increase of mass-produced Video 2.0 content and

services likely will increase operators' storage requirements by several orders of magnitude further than the content growth already discussed.

Growing Storage Requirements

Two major technology innovations, time-shifting and VoD, are driving storage expansion. Either feature, but most probably both, will be present in any given operator/subscriber scenario. The two factors are not completely independent. Analysts expect that some broadcast content will be acquired into an operator's long-term VoD library, while some broadcast content will be discarded, since it already will be present in the VoD library.

An interesting analysis[1] of DVD and book sales (which led to the now widely used concept of "long-tail" content) may be relevant for identifying a major growth factor for video content. The analysis studies the sales trends of "hit"-based inventory in "brick-and-mortar" stores versus "expanded" inventory in "low-overhead" Internet stores. The study considered Rhapsody, Amazon, and Netflix. The detailed analysis will not be repeated here, but in "low-overhead" inventory scenarios, the study suggested that 20 to 30 percent of total sales are from expanded inventory not available in offline "brick-and-mortar" competitors. This finding suggests that, when planning a time-shifting or VoD service, the availability of long-tail inventory holds the potential to generate substantial sales.

Time-shifting alone generates significant content. For a service provider offering 500 channels, content can be acquired at the rate of 12,000 hours per day (24 hours per day x 500 channels). If one assumes that 50 percent of this content is deleted because it is duplicate content, or because rights do not exist for long-term storage, the rate at which content is acquired it still significant. At 6000 hours per day over a two-week period, 84,000 hours of content would be acquired. If parallel versions of the content are generated upstream (for example, different feeds with different inserted zoned ads), then the storage requirement will multiply linearly. (As a baseline requirement for conserving storage space, it is assumed that a single copy of each broadcast channel will be retained in storage, not multiple zone-specific copies.[2])

While storage requirements are likely to grow substantially regardless of the underlying technology of the video infrastructure, new technological innovations are allowing carriers to manage content growth and delivery much more effectively. The rise of technologies such as switched digital video for cable operators and IPTV for wireline carriers is allowing service providers to offer far more content than was possible in the past. For example, many carriers now are able to offer ethnic tier programming and have found that such programming generates substantial revenues and is one of the most successful strategies for up-selling video subscribers. While offering all existing (and future) ethnic-tier content on demand also would drive additional storage requirements, it would likely be extremely attractive to the base of subscribers who use and are willing to pay for these services.

Once again, however, these capacity drivers do not include the coming Video 2.0 paradigm, in which subscriber-generated video will become just as common as conventional content (if not more so). This paradigm will likely escalate storage requirements even further beyond those drivers already discussed.

---

[1] Chris Anderson, http://www.thelongtail.com/the_long_tail/2005/08/index.html.
[2] The "n zone replicated channel" scenario is present in some early deployments because of the inability of existing streaming servers to dynamically splice content.

Conventional Content Distribution Architectures

A video infrastructure generally consists of a content storage subsystem and a streaming subsystem. The streaming subsystem streams the content in whatever format is compatible with the subscriber's device. Many variants of streaming systems exist. In some cases, the systems have local storage and require storage of whole titles. In other cases, the systems are arrays of servers which contain combinations of disk and random access memory (RAM) storage to cache content in small-segment granularities.

The content storage subsystem contains both live ingested (time-shifted) broadcast content (growing at the rate of number of content segments x resolution x formats per hour) and previously acquired content. This content can include movies, made-for-TV productions, and independent films, which also involve multiple resolutions and formats.

In conventional architectures, many operators manage storage in the streaming element by prepositioning the storage of whole titles, using the administratively determined popularity of a piece of content as the major criterion for storing that content. The majority of such solutions have a maximum upper storage limit from hundreds to thousands of hours (for standard-definition content). As long as the system does not exceed the storage limit, it can make content available within the streaming element. However, when the system exceeds its finite whole-title storage capacity, the subscriber experiences a "denial of service" situation (usually a "service not available" error message). Actually, the operator may experience a "denial of revenue" situation that is, the loss of the 20 to 30 percent additional revenue that can be realized by offering fee-based long-tail content.

Dynamic Content Distribution Architectures

Today, a more dynamic form of content distribution is becoming increasingly feasible. This approach allows streaming elements to dynamically to replace less-popular content based upon subscriber demand and to obtain newly requested content for short-latency streaming.

Figure 1 shows a conceptual diagram of this architecture. In this model, an array of content storage servers is connected to an array of streaming servers. The diagram assumes that the methods of interconnect are Gigabit Ethernet and 10 Gigabit Ethernet. This model specifies no specific content distribution protocol between the content and streaming subsystem elements.

Figure 1. Tiered Content Storage and Streamer Arrays
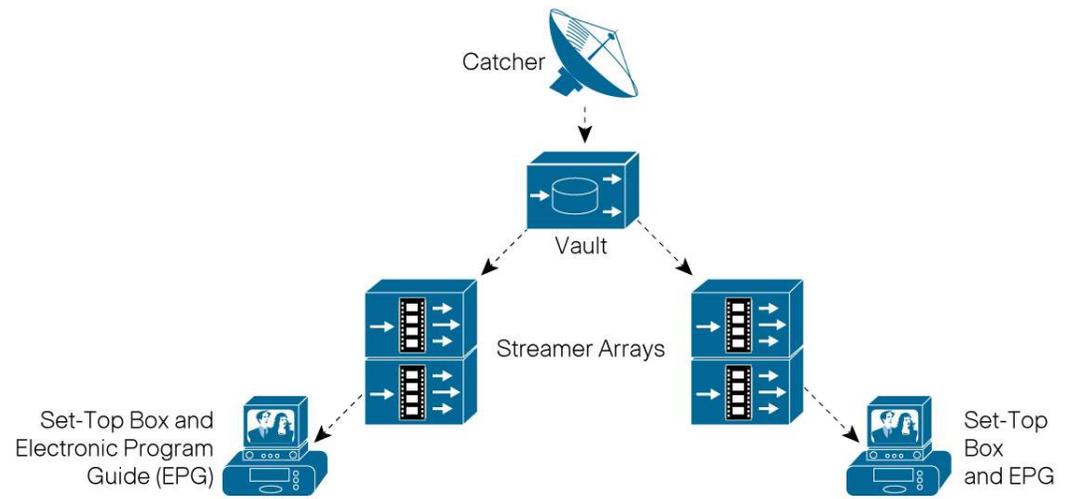
This dynamic content distribution model can employ one of three different methods of fetching and streaming requested content, depending on the capabilities of the architecture and the streamers. However, the first two methods accelerated whole-title fetch and progressive fetch both have limitations that prevent them from providing an ideal approach to real-time video delivery. The third method, segmented cache-fill, can avoid many of the latency and other issues presented by the other approaches and provides a much more efficient, scalable strategy.

Method 1: Accelerated Whole-Title Fetch

For video systems that are unable to stream during ingest, a tradeoff must be made between ingest latency and content bandwidth. If, for example, the worst-case requirement is to be able to play a piece of content no later than two minutes from time of request, then a one-hour movie must be retrieved at 30 times the stream rate, that is, 120 Mbps for standard-definition MPEG-2-encoded content or 60 Mbps for standard definition MPEG-4.

In systems where the whole content title must be ingested prior to streaming, ingestion tends to rapidly overrun the link. For example, in such a system a maximum of eight one-hour movies can be retrieved within two minutes over a Gigabit Ethernet link to a storage system. In a nationwide, regional, or even local network, the probability of eight or more concurrent requests for long-tail content is very high, as is the probability of concurrent requests for 80 movies over a 10 Gigabit Ethernet link in a larger network. Streaming of high-definition movies is not even feasible under this model.

Method 2: Progressive Fetch

For those systems that are able to begin steaming concurrently while ingesting, content can be provided with shorter latency and requiring less bandwidth. Assuming that content is obtained at stream rate, streaming can begin after the streaming system's minimum ingest latency has elapsed. However, whenever the streamer begins streaming, it is dependent on the ability of the content storage server array to continue providing the content at stream rate, without jitter. If the content server array does introduce jitter, a gap also will be introduced in the streaming output. Typically, the streaming system will be required to introduce additional delay in order to handle any potential jitter. Therefore, while a streaming system might be able to ingest content with a 10-second latency, the system might end up introducing even more delay. And, since pausing the display when the buffer underruns would be considered unacceptable in a real-time video service, the delay might have to be padded further.

In this scenario, the system could partially ingest the title with longer delay (perhaps one to two minutes) and then begin playing. Typically, today's PC-based streaming implementations employ this approach, in which the delay occurs first and then streaming starts. However, with this method, there is still a possibility that the network will be unable to maintain stream rate, and thus the display will freeze. (In PC-based streaming applications, a "refreshing the buffer" icon typically appears when this occurs.) In addition to this problem, operators also will be challenged to scale such solutions over burst-mode networks, input/output (I/O) systems, and disk subsystems.

Issues with the progressive fetch approach include:

- Difficulty of scaling the content distribution subsystem to match increasing numbers of parallel concurrent fetch operations
- Aggregate performance of the ingest function with the streaming element is limited
- Trick modes (fast forward, rewind, and so on) typically generated in real time, causing scaling limitations

Operators can employ the progressive fetch method in scenarios in which the frequency of externally pulled content is extremely low (likely much less than 10 percent, or perhaps 1 percent), as long as the content storage subsystem can maintain stream-rate delivery of such content (avoiding excessive contention for a shared spindle and efficiently implementing protocols), and as long as the delay variations are accepted by subscribers.
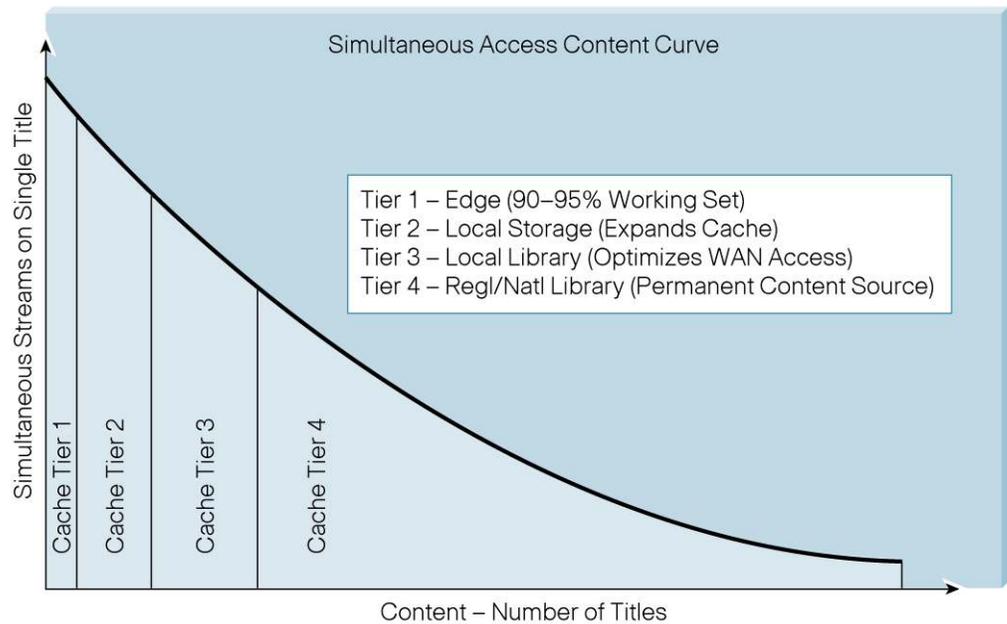
Method 3: Segmented Cache-Fill

An alternative content delivery method that provides both low latency (less than 300 ms) and minimum bandwidth utilization is segmented cache-fill. In this model, segments are defined as small portions of content, with a representative size per segment of 64 Kilobytes. This model requests and transmits segments in all portions of the system, including storage, distribution, and streaming.

In the segmented cache-fill approach, caching provides a complementary mechanism in support of segments, since the model requires permanent storage of content only within the storage array itself. Cache is defined as transient storage, which benefits the system by conserving all types of resources (disk, bus, memory) in the system. Caching is a well-known and widely deployed technique for services such as serving Web pages. If a segment of content is in a local cache resource, then it does not need to be requested using a more expensive resource such as network I/O. Cache management techniques significantly reduce resource consumption, especially network I/O. In a large regional or nationwide storage configuration, where it is likely that two or more people will request access to a piece of long-tail content concurrently, long-tail content distribution also can benefit from caching.

Figure 2 shows a multitier caching hierarchy. Tier 1 is the storage closest to the network edge, from which content is played out in real time. This storage accounts for a small percentage of the hardware (RAM) because of the caching efficiencies delivered by the dynamic content popularity algorithms that the system employs. Tier 2 is local storage that may be located within a single system or in a peer system connected by high-speed local bandwidth. Tier 3 is a local content library that allows popular content to be held locally and accessed using inexpensive networking resources. Tier 4 is the permanent source of content and can be located anywhere, even in a national or super headend. With tiered caching, operators can build local, regional, and even national video networks with as few as a single copy of all available content, both short-tail and long-tail.

Figure 2.   Caching Hierarchy

Simultaneous Access Content Curve

Tier 1 – Edge (90–95% Working Set)
Tier 2 – Local Storage (Expands Cache)
Tier 3 – Local Library (Optimizes WAN Access)
Tier 4 – Regl/Natl Library (Permanent Content Source)

Caching is a popular method for optimizing the mix of high-cost and low-cost components in systems. The most expensive but highest-performance storage is RAM, while the least expensive storage is disk drives. RAM storage is ten times more expensive than Small Computer System Interface (SCSI) storage, and 100 times more expensive than Serial Advanced Technology Attachment (SATA) storage.

The segmented cache-fill function is a much lower-overhead function than the ingest function in the progressive and whole-title accelerated fetch methods. In a cache-fill architecture, the ingest function is performed by the content storage element itself. The cost of ingest (metadata, trick files, resiliency) is amortized over all future plays of content both popular and long-tail.

With a segmented cache-fill content distribution model, operators also can meet much lower latency requirements. For example, operators can implement a system in which the streaming element is aware of the bandwidth needs of all streams, and can aggregate its segment I/O with the content element to help ensure that service is maintained for all streams. Because of the benefits of caching for more popular content (cached popular content does not generate network I/O), this model reserves sufficient network bandwidth to dynamically retrieve long-tail content.

The segmented cache-fill content distribution model also allows for quick switching among different content titles. For example, in a targeted advertising scenario (in a switched broadcast or time-shifted environment), an embedded ad insertion, which requires seven seconds to react, can be handled with a sub-300 ms reaction time. Likewise, some of the newer services such as remixing of segment clusters from multiple content titles (for example, offering "The Best Scenes from 'Da Vinci's Inquest'") also can be handled without managing whole-title units of granularity, greatly increasing storage and network bandwidth efficiencies. And, in the future, as subscribers control the remixing of long-tail content (for example, skipping intros or trailers), the congestion of all resources, including network I/O, also will be reduced.

A More Intelligent Video Architecture: The Cisco Content Delivery System

More and more operators are now reaching the conclusion that a collection of scalable content storage arrays is an integral part of an end-to-end, real-time video services system. With the

emerging requirements for long-tail content and advanced Video 2.0 personalization services, service providers need a method of content access that can support both short-latency startup of streams and new, advanced services such as latency-intolerant remixing. Service providers also are increasingly recognizing that a distributed architecture, in which streamers with transient cache storage are placed nearer to the edge of the network, can offer substantial scalability and bandwidth efficiency advantages over conventional, centralized video systems.

The ideal video system should employ the cache-fill segment distribution method, which distributes content on demand in less than whole-title units of granularity and can effectively scale real-time video services to regional and national configurations. By employing segment granularity caching, such a video system can reduce the bandwidth consumed for video services transport by several orders of magnitude. And, with centralized storage arrays, it can amortize the growth of content from n x 1000 hours to n x 100,000 hours (to n x 1,000,000 hours and beyond) across entire regional and national footprints.

The Cisco Content Delivery System (CDS) incorporates all of these strategies into a single platform that allows operators to provide subscribers with all of the content they demand in broadcast, time-shifted, and VoD mode whenever and however they request it. The Cisco CDS transcends conventional VoD systems by providing operators with an intelligent, network-based platform for delivering the next generation of entertainment, interactive media, and personalized advertising services to their subscribers. This platform the latest addition to the Cisco IP Next-Generation Network (IP NGN) Service Exchange Framework combines video ingest, storage, distribution, personalization, and streaming capabilities into a solution that operators can use to deliver localized, interactive, and personalized content across a growing portfolio of heterogonous devices.

Unlike early VoD solutions, which functioned as large, centralized video servers and were extremely difficult to scale as the subscriber base and content libraries grew, the Cisco CDS functions as a true video network not just a video server. By fully embracing the capabilities of IP networking and the content delivery strategies outlined in this paper, the Cisco CDS embraces an entirely new paradigm for the delivery of subscriber video services, representing a much more cost-effective, flexible, and future-ready video solution.

As service providers plot out their strategies to transform into experience providers, to offer customers a true "connected life" in which content transcends the access device, and to deliver on the promise of Video 2.0 to democratize video content production and distribution, they need a video solution that was designed to support all of these future requirements. They also need a solution that provides the full range of tools to deliver the interactive, personalized, and high-quality video experience that customers demand right now, and that is part of a larger IP video framework that integrates all aspects of defining, preserving, and realizing a superior customer experience. The Cisco CDS provides this platform. As a core component of the Cisco IPTV solution, the Cisco CDS builds on the IP and VoD expertise of Cisco and the video leadership of Scientific Atlanta in the headend and in the customer home to provide a truly comprehensive video solution.

For more information about the Cisco CDS, visit www.cisco.com/en/US/products/ps7191/Products_Sub_Category_Home.html.

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

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

C11-395592-01   06/07