

Cisco Virtual Video Infrastructure

Managing Complexity and Scale in a Next-Generation Video Network

Executive Summary

Consumers have never had higher expectations for TV and video services. They want more video, in more formats, on more screens, and with more personalization and control than ever before. Service providers are aggressively striving to meet these expectations, and to deliver more video content and personalization than their competitors. However, while operators have made major investments in adding real-time, time-shifted, and on-demand video capabilities to their service offerings, the video distribution model on which they have relied can become a significant barrier to success.

As operators look to add tens or hundreds of thousands of hours of video to content libraries over the next several years, they need new tools to manage growing complexity, deliver new services, and efficiently scale video offerings to more customers. The Cisco® Virtual Video Infrastructure (VVI), built using proven Cisco Content Delivery System (CDS) technology, provides a highly scalable foundation for storing, managing, and distributing personalized TV and video services across national and transnational video networks.

The Cisco VVI delivers many of the essential capabilities service providers need to evolve their networks to become a medianet—an intelligent network optimized for rich media. Part of the broadest, deepest video portfolio in the industry, the Cisco VVI integrates media intelligence throughout the network to deliver a more visual, social, and personal experience for customers, while helping service providers accelerate the rollout of new services and scale to millions of subscribers.

Introduction

As more service providers enter the TV and video entertainment market, competition has never been fiercer. With more video content available than ever before and an ever-widening array of personalized options for viewing that content, consumer expectations are growing. To succeed in this environment, service providers must do the following:

- **Offer more content:** The amount of high-definition content, local content, foreign language content, and on-demand titles a service provider can offer is a key competitive differentiator. Increasingly, differentiating a service also means incorporating online, user-generated, and other nontraditional video content into the service offering.
- **Deliver a personalized user experience:** The days of building a service offering on basic broadcast TV are over. Today's consumers want the specific content they want whenever they want it, on whichever device they choose. For service providers, that means dramatically expanding on-demand video libraries and time-shifting capabilities, and putting capabilities in place to deliver video to a variety of consumer devices.

- **Reduce operational costs:** At the same time that consumers are demanding more content and capabilities, they are also demanding lower prices. As service providers add next-generation content and complexity, they will need next-generation efficiencies from their distribution systems.
- **Improve scalability and service velocity:** Operators need tools to not only create new types of service offerings, but to rapidly roll those offerings out to national and transnational audiences. They need a scalable, manageable video network that allows them to roll out multiple services simultaneously in months and reach millions of subscribers immediately.
- **Develop new revenue streams:** Facing significant pricing pressure for consumer subscriptions, service providers need other options for growing revenues. For example, they need new advertising models that offer more diverse and targeted methods for reaching customers.

These requirements are driving the need for more flexible and scalable distribution systems for TV and video entertainment. The current video distribution and video-on-demand (VoD) approaches employed by most service providers, however, have major limitations. Chiefly, today's model relies on a market-based approach which, while often effective at the local level, does not allow for easy coordination across the entire subscriber base and does not capitalize on the economies of scale that typically are available for national or transnational networks.

For most service providers, the TV/video distribution network is divided into separate geographical areas, each with its own video delivery infrastructure. Video content offered in each market is typically acquired, stored, and managed locally. This model gives operators a great deal of control in customizing offerings for individual markets, but it is extremely inefficient. A typical Tier-1 operator in a large country like the United States may manage more than 100 separate markets, and service providers in other parts of the world may serve a similar number of markets across multiple countries. With the need to ingest and store a single title hundreds of times in hundreds of separate local infrastructures, the cost of expanding on-demand libraries and services in each local market quickly becomes exorbitant.

More than simple inefficiency, however, most local architectures also lack important capabilities. For example, there is usually very little linkage between local TV/video systems and the online video offerings in which many operators are currently investing as they build out national and transnational backbone networks. As operators move to expand their content libraries exponentially over the next several years and roll out a variety of new video services, the inherent inefficiencies of a localized distribution model will become increasingly problematic.

What if there was a way to ingest video content once for all markets, store that content in a centralized content library, and then efficiently distribute it to viewers across a large national or transnational network as it is requested, in whichever format and for whichever device customers choose? This is not merely an idealistic vision of the video system of the future. Service providers can take advantage of these capabilities right now with the Cisco Virtual Video Infrastructure.

The Cisco Virtual Video Infrastructure

The Cisco VVI provides an intelligent platform for delivering any type of real-time, time-shifted, or on-demand video content to any screen, across a national or transnational infrastructure. Based on a highly scalable network design and the unique dynamic caching model currently employed in Cisco CDS on-demand video systems worldwide, the Cisco VVI allows operators to streamline

operations, cost-effectively scale services across vast geographies, and deliver a compelling and differentiated user experience.

With the Cisco VVI, operators can:

- **Exponentially expand content:** The Cisco VVI's innovative architecture, proven Cisco CDS content distribution technology, and industry-leading low latency mean that service providers can ingest video content once and distribute it everywhere, in real time, as it is requested. This allows operators to exponentially expand content libraries, making 100,000 titles or more available to any subscriber, anywhere, while reducing local ingest and storage requirements.
- **Dramatically cut costs:** With the ability to serve an entire national or transnational customer-base with a single content library, service providers can collapse headend and content storage infrastructures across the network and consolidate operations centers, reducing both capital and operational expenses.
- **Deliver new and differentiated services:** The Cisco VVI allows service providers to support a broad range of new services. For example, with the ability to distribute content from anywhere to anywhere, operators can provide user-generated and online video just as easily as any other on-demand title. The ability to deliver content with sub-second latency also lets service providers dramatically expand the video library that can be made immediately accessible to customers, allowing them to access content that resides in a different state or country virtually instantly.

Operators can also support popular real-time and time-shifted services, such as letting viewers tuning into a program in progress and restart it from the beginning, or providing network-based personal video recorder (nPVR) functions such as the ability to pause, fast forward, and rewind live TV. The Cisco VVI's centralized storage/localized streaming architecture also distributes screen-formatting processes to the network edge.

- **Minimize video network bandwidth:** The Cisco VVI's innovative content caching and propagation model dynamically distributes content based on real-time popularity and ensures that only a fraction of content traverses a segment of the network. As a result, the Cisco VVI can reduce an operator's national or transnational backbone bandwidth consumption by as much as 90 percent.
- **Optimize existing infrastructure investments:** Service providers are already investing in national or transnational IP backbone networks to deliver voice and Internet services. Now, operators can use that existing backbone infrastructure to support national/transnational distribution of real-time, time-shifted, and on-demand video. The Cisco VVI can also interface with third-party VoD systems, allowing operators to expand and centralize their VoD service offering while still benefiting from investments they made previously in local video streaming and storage infrastructures.
- **Rapidly scale new and existing services:** The Cisco VVI lets operators easily distribute content to a large and continually growing customer base. It facilitates the rapid rollout of current and next-generation video services in multiple markets simultaneously or on a national/transnational basis.
- **Uncover new revenue streams:** The Cisco VVI creates a national presence for advanced advertising services, offering advertisers access to national audiences, local audiences, targeted personal audiences, and anything in between. The ability to deliver real-time and

on-demand video, user-generated and online content, and content tailored for a variety of screens and devices also offers many new advertising and revenue opportunities.

All of these capabilities depend on the ability to ingest real-time or on-demand content once, to dynamically propagate and cache that content across the network, and to stream it to the customer as it is requested without any noticeable delay. Only Cisco can provide the highly specialized video caching and distribution techniques necessary to achieve sub-second latency times, meet aggressive service-level agreements (SLAs), and deliver a superior customer experience.

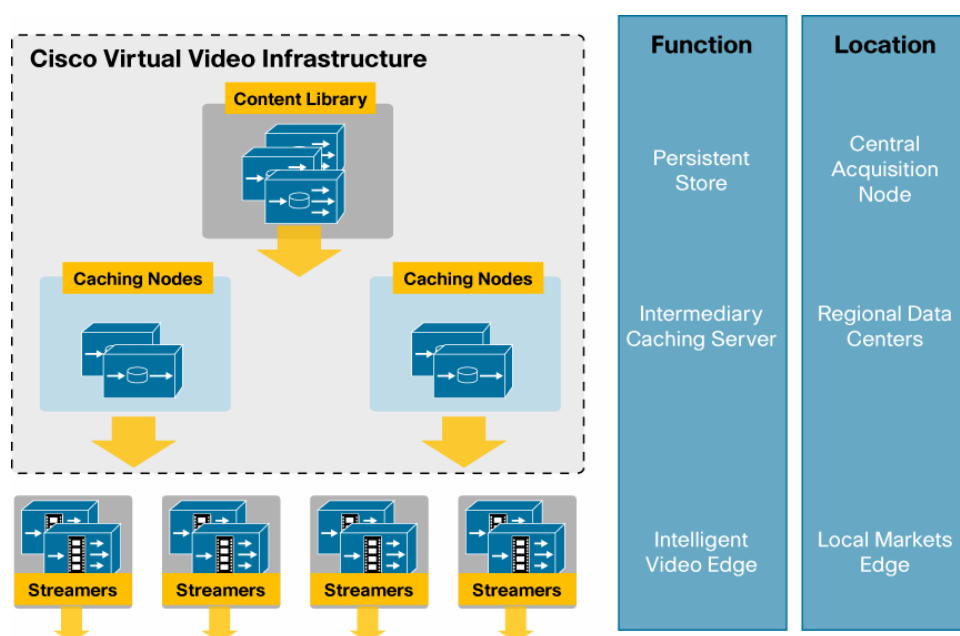
Building Blocks of the Cisco Virtual Video Infrastructure

The unmatched capabilities and efficiencies of the Cisco VVI are based on its highly scalable architecture and its dynamic caching and distribution techniques. This architecture separates content ingest/persistent storage, temporary storage (caches), and streaming functions to provide unparalleled efficiency, resiliency, and scalability.

Dynamic Video Caching and Distribution

In a traditional localized video distribution model, operators pre-position content manually in each market. This requires them to predict which content will be popular and where, position content in local streaming and storage infrastructures, and set content to expire at a given time. Customer preferences, however, are constantly changing and notoriously difficult to predict. If operators guess wrong, the content customers want is not available when they want it. The only alternative is to waste significant resources by storing huge amounts of content locally, much of which may rarely be viewed. This model also makes it economically impractical to offer niche “long-tail” content that appeals to smaller audiences—even though there is ample demand among customers for such content.

The Cisco VVI offers an alternative model based on a hierarchical architecture consisting of a centralized content library that serves all markets, a layer of regional caching nodes that each serve a subset of those markets, and localized video caching/streaming platforms in each market (Figure 1). Designed specifically for the next generation of IP-based video networks, this model applies the same scalability and caching efficiencies used to deliver web pages over the Internet to the distribution of video across national or transnational networks.

Figure 1. Cisco Virtual Video Infrastructure High-Level Architecture

In the Cisco VVI, content is ingested once for persistent storage in the content library, from which it is made instantly available to any streamer in any market. (Note that this national or transnational content library is a logical entity; in practice, the library might actually be distributed across multiple sites for higher resiliency and availability.) Video content is not transmitted across the network, however, until it is actually requested by a subscriber. The first time a title is requested, it is cached both in the regional caching node and in the local streamer serving the requesting customer. The title then resides both in the local cache serving all customers in the requesting customer's market and in the regional caching node serving all of the markets in that region. In this way, service providers can maintain a single content ingest and storage point (dramatically reducing local infrastructure costs), while fulfilling most requests for video content from a local or regional caching node (reducing backbone bandwidth consumption). In addition, the Cisco VVI propagates video content in 64-kilobit segments. As a result, the system distributes and caches only those portions of a title that are actually viewed (for example, omitting closing movie credits or any portions of the video skipped by the user that go unwatched), further optimizing bandwidth across the network.

The Cisco VVI continually tracks content popularity among actual users to make intelligent decisions regarding when to release particular content from cache. By caching and distributing video dynamically based on real-time viewer behavior, the Cisco VVI eliminates the need to try to predict the popularity of video content. In fact, it is because of this unique caching intelligence, as well as the Cisco VVI's industry-leading low latency, that Cisco alone can deliver a true national or transnational on-demand video network.

- Innovative dynamic caching:** Without the ability to dynamically propagate content out toward the network edge for streaming, all of the benefits of maintaining a single, centralized content library would be impossible – the bandwidth requirements of serving a national or transnational audience would quickly overwhelm the core network backbone. The Cisco VVI's unique caching and propagation intelligence, however, continually propagates content based on real-time popularity and ensures that the vast majority of viewers can be served from a local or regional cache. The result is the ability to centralize

permanent content storage, while actually reducing backbone bandwidth consumption by as much as 90 percent.

- **Industry-leading low latency:** Whether content is drawn from the centralized library or from a local or regional cache, the Cisco VVI delivers it to any subscriber, anywhere, in less than a second, allowing for virtually instantaneous playout of even the less popular long-tail content. This unparalleled low latency is also what positions the Cisco VVI as a foundation for a new generation of video services, since it allows service providers to transparently integrate real-time, on-demand, online, and user-generated content, as well as a variety of broad-based and targeted advertising, into a single video experience.

Based on widely deployed Cisco CDS technology, the primary elements of the Cisco VVI (content library, caching nodes, and video streamers) employ the same content distribution protocols that are used in existing Cisco CDS deployments. The Cisco VVI can deliver the most capabilities and cost savings when deployed with Cisco CDS Streamers at the local level. However, the Cisco VVI solution can also work with third-party streaming equipment. As a result, it also allows operators to continue benefiting from the video storage and streaming infrastructure investment they have already made in local switching offices, if they choose.

Nonstop Service Availability

In addition to providing a more efficient video distribution system, the Cisco VVI's IP intelligence and auto-failover capabilities also makes it inherently more resilient than conventional video systems. The architecture employs intelligent resource pooling and load leveling across the network, allowing multiple distributed caches and streamers to work together as a single logical resource pool, which can be automatically redistributed across the network in response to changing conditions. This strategy not only allows for unprecedented capacity, it creates a "self-healing" video network.

In a traditional video system, if a video server fails, all customers who were watching a program streamed from that server see a service interruption. In the Cisco VVI, if a video streamer fails (or is taken offline intentionally for a scheduled maintenance), the functions it was performing are immediately and transparently taken over by other resources in the network – without any impact to subscribers and without the need for idle standby units that other solutions require. These self-healing capabilities apply not only to real-time video and VoD streams, but to every other service running on the system, including time-shifted TV and ad insertion. The Cisco VVI also allows operators to distribute content ingest and storage for national or transnational content libraries across multiple physical locations to ensure availability and resiliency.

Content Library Virtualization

Service providers may be able to realize dramatic savings with a centralized content library, but they still need the ability to maintain separate content offerings tailored to each market. The Cisco VVI provides library virtualization tools that allow operators to maintain separate policies for each market and treat each market as a distinct entity, even as they are benefiting from the economies of scale afforded by a centralized national or transnational content library.

With the Cisco VVI's virtualization tools, each market can continue to operate as if it had its own copy of the assets on offer to that market. Policies governing the management, addition, and cancellation of content can still be handled on a local market-by-market basis, even though all actual content is maintained in a single physical library. The Cisco VVI also includes asset-distribution tools that allow operators to manage permissions, mirroring policies, and availability of

titles in the content library. For example, a service provider that is backing up titles to a secondary content library for redundancy can specify the times of day in which copying occurs to avoid overtaxing network links when they are being heavily used by customers.

Exploring Cisco Virtual Video Infrastructure Advantages

The Cisco VVI provides all of the service capabilities, architectural capabilities, and operational capabilities that service providers need to streamline video operations and deliver an exceptional customer experience.

Key Service Capabilities

- Supports multiple content formats (high-definition and standard-definition content, multiple video codec formats, multiple media file types, etc.)
- Supports ingest and streaming of real-time video services, VoD services, and Internet video
- Supports advertising content distribution and streaming
- Supports nPVR capabilities to provide a digital video recorder (DVR)-like experience with the network
- Provides a single content delivery network for serving set-top boxes (STBs), PCs, and mobile devices
- Supports content security and encryption
- Supports narrowcast service such as VoD, time-shifted TV, and switched digital video (SDV) sharing the same infrastructure
- Supports both traditional and next-generation STBs and headends
- Delivers scalability, reliability, and high performance
- Provides data warehouse reporting, logging, and network management
- Provides a standards-based solution that can support multi-vendor environments

Key Architectural Capabilities

- Supports sub-500 millisecond latency from ingest to the network edge for both standard and interactive content
- Supports millions of titles and concurrent VoD sessions
- Supports content storage and forwarding on a local, regional, and national/transnational scale
- Places no constraints on the number of ingest points that can be supported
- Supports guaranteed SLAs across WANs for real-time services with industry-leading ingest-to-play latency
- Optimizes bandwidth utilization with intelligent asset propagation to the network edge
- Distributes content library and asset-location functions
- Allows for independent scaling of storage and streaming resources
- Helps guarantee high availability of service across a WAN

Key Operational Capabilities

- Allows for individual control of content libraries and billing management systems (BMS) in distinct markets, even with a common storage architecture
- Supports market-specific control methods and content location/resiliency rules

- Delivers the improved economics of centralized content ingest, while maintaining the flexibility to support ingest in local markets
- Provides hierarchical management capabilities at national/transnational, regional, and market levels

Conclusion

As service providers strive to deliver a more compelling and personalized video experience than their competitors, they need a video distribution system capable of managing the extraordinary complexity and scale of the next generation of video services. They need the capacity to exponentially expand content libraries, the flexibility to integrate new types of services and advertising models, and the efficiency to deliver any type of content anywhere, under aggressive SLAs.

As the worldwide leader in networking, Cisco has unparalleled experience employing networks and IP intelligence to deliver superior scalability and efficiency. Cisco has also invested heavily in dynamic caching intelligence and is a full generation ahead of its competitors in optimizing the delivery of video over IP networks.

The Cisco VVI is not built on brand new, unproven video distribution techniques. It is an evolution of one of the most successful video network architectures in history, the Cisco CDS, which is currently used in many of the largest video infrastructures in the world. Through the ongoing development and support of real-world video networks, Cisco has gained vital experience dynamically caching and distributing video content at the local and regional level. More than any other technology provider, Cisco understands the requirements of tomorrow's large-scale video architectures and can offer the ideal foundation for a scalable national or transnational video system.

In addition to innovative technology, Cisco can also offer extensive implementation and operations expertise, including in-depth system integration services to help service providers deploy the Cisco VVI quickly and successfully.

Cisco's revolutionary video technologies extend beyond the Cisco VVI's video delivery capabilities. The Cisco VVI is just one major building block in Cisco's overarching strategy to support an intelligent end-to-end medianet that delivers a new generation of media services with unprecedented scalability and efficiency. With the Cisco VVI and other medianet-optimized Cisco technologies, service providers can deliver a true Connected Life to their customers—linking them to the human network and delivering experiences that are more visual, more social, and more personal.

For more information, visit <http://www.cisco.com/go/sp/medianet>.



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.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)