# BUILDING A WEB SERVICES-BASED CONTROL PLANE FOR NEXT-GENERATION VIDEO EXPERIENCES

Yoav Schreiber, Sunil Mudholkar, Nadav Neufeld
Cisco

*Abstract*

*Next-generation video services require solutions that are aware of real-time data and granular business rules encompassing identity, location, policy, etc., and can make decisions based on that data for a multitude of applications. In conventional video systems, however, this collection of data and business rules resides on disparate elements linked by closed, proprietary connections. This reliance on closed, "siloed" video systems impedes an operator's ability to develop new services and features, or to effectively scale cloud-based video services.*

*This paper presents a scalable, open-standards approach to orchestrating video services to allow for video control plane extensibility in multi-vendor ecosystems. It describes the architectural foundation for a loosely-coupled, modular video control plane with service provider-grade high availability and scalability. This approach enables video end-points to discover cloud functionality and external systems to expose services and communicate with video endpoints. Drawing on Internet communication methods, the proposed architecture enables a more flexible and scalable video services platform.*

## INTRODUCTION

An array of market forces is driving demand for new kinds of video services and, ultimately, profound changes to the service provider video systems delivering them. The Cisco Visual Networking Index (VNI) projects that more than 90 percent of consumer IP traffic and two thirds of the world's mobile traffic will be video by 2015.[1] The same study projects 10 billion mobile Internet-connected devices connected by the following year. And, the Cisco Global Cloud Index projects cloud IP traffic to reach 133 exabytes per month by 2015.[2]

These trends point toward a massive shift in consumer viewing habits from traditional, closed TV video systems to an open, cloud-based model. Consumers want the ability to access multiple types of content on multiple types of devices, regardless of the users' location or of the network over which they are connecting (e.g., the managed service provider footprint, an unmanaged Wi-Fi network, a cellular network, etc.). Consumers also seek new kinds of video experiences that integrate conventional video content with cloud services and interactive applications, and extend intuitively across multiple screens.

Service providers are well aware of these industry changes, and many are already moving to expand their video services to new screens and devices beyond the traditional set-top box (STB). However, the traditional service provider video architecture – designed to deliver legacy broadcast and on-demand video content, over a closed network, to a managed STB endpoint – is simply not equipped to support cloud and multi-screen delivery. These next-generation video services demand a level of service orchestration that conventional video platforms do not address. Consider: to deliver a personalized next-generation video experience to a subscriber, the video system must account for:

- Identity
- Device
- Content entitlement

- Location
- Bandwidth availability
- Past user activity
- Social network connectivity
- And much more…

Yet in today's video architectures, this information resides on several disparate, closed systems, including OSS/BSS, content management systems, session resource managers, client software, applications, middleware, etc. In addition, the isolated service "silos" on which conventional service provider control planes rely (i.e., treating managed and unmanaged clients, wired and wireless networks, etc., as entirely separate environments) further impede the service orchestration necessary to deliver a seamless, personalized multi-screen video experience. Conventional video system architectures are also ill-equipped to address the complexity inherent in serving multiple and changing consumer devices connecting to the service, or in optimizing the quality of experience (QoE) based on changing conditions.

Meeting the requirements of modern, cloud-based video delivery will require a new architectural model: a control plane for loosely coupled video systems that is designed specifically to meet the technical requirements of next-generation video services. This paper presents such architecture.

The proposed architecture is a scalable, standards-based approach to orchestrating video services to allow for video control plane extensibility in multi-vendor ecosystems. Drawing on proven Internet communication approaches used by some of the largest web companies in the world, it provides an architectural foundation for a loosely coupled, modular video system with service provider-grade high availability and scalability. Chiefly, this architectural approach:

- Enables video endpoints to discover cloud functionality in a loosely coupled system, and allows external systems to access exposed web services for communication with video endpoints
- Provides a platform to dynamically manage sessions, resources, and workflow in a loosely-coupled system incorporating both managed and unmanaged devices

## ARCHITECTURAL ELEMENTS AND KEY CAPABILITIES

The proposed video control plane employs an Internet-based architecture, and as such, represents a significant departure from conventional video systems. Effectively, this approach applies the proven architectural model and design principles used by major web companies like Google and Facebook to contend with massive amounts of data and users, and applies them to video services. This gives service providers a more scalable video services platform, and affords them the same degree of flexibility and speed as web companies when designing, testing, and rolling out new features and applications.

At a high level, the proposed architecture encompasses the following building blocks (Figure 1):

- **Base platform:** The foundation of the video architecture is an open-source operating system, on top of which resides a distributed cache that acts as a shared data store accessible to all loosely coupled elements and workflows in the system.
- **Common messaging infrastructure:** A standards-based, highly-scalable, real-time messaging bus provides the communication framework over which distributed endpoints communicate.

- **Service infrastructure:** The architecture employs a standards-based, hardware-agnostic service infrastructure and workflow engine that enables complex service orchestration with the necessary performance for video services.
- **Session and resource management:** The architecture provides real-time session and resource management capabilities across multiple networks and devices, and is designed to support flexible policy enforcement and dynamic business rules.
- **Applications:** All video applications (i.e., service assurance, device authentication, emergency alert services, etc.), are built upon this platform.
- **Application programming interfaces (APIs) and web services:** the video control plane brokers communication among all elements of the system, including both legacy and cloud-based video applications, via APIs and web services.

Together, these building blocks create a next-generation video control plane that represents a fundamental shift from traditional video systems. Unlike legacy video systems, which rely on tightly coupled client/server communications, the proposed architecture employs a distributed model, similar to that used in web applications. The Internet has solved many of the problems of software resiliency, performance and scale and taking advantage of those is key to building this distributed control plane. In this distributed control plane, functionality and intelligence is allocated to various loosely coupled endpoints (i.e., clients, virtual machines, network elements, etc.), which then advertise, discover, and consume functionality from a shared cache of data.

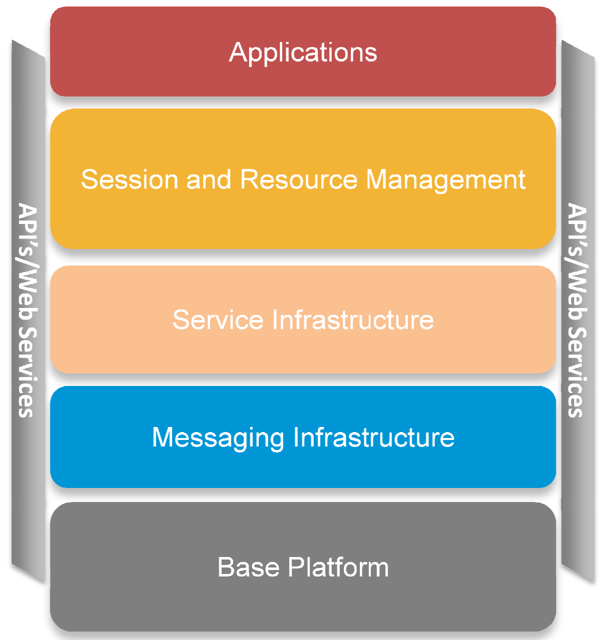


*Figure 1. High-level architecture for next-generation video control plane*

This data includes all of the essential information that applications need to deliver a video service via the cloud, including presence, state, entitlement, resource availability in the network, etc., all of it updated in real time. The data are stored in a high-speed shared cache, from which they are accessible to distributed applications in real time.

Clients, network elements, and applications can access data stored in the shared cache via a standards-based messaging bus. This common messaging infrastructure connects all elements to the cache via an encrypted, authenticated connection, and facilitates messaging back and forth among the various applications. Cisco has designed the architecture using a widely adopted communication protocol known for its scalability and performance in social presence and instant messaging applications.

Once this core framework is in place – open-source operating system, high-performance data store, and real-time messaging infrastructure – operators can build applications on top of it. These can include everything from straightforward core service functions like device authentication and service assurance, to advanced cloud service offerings such as cloud- or network-based DVR, social TV experiences, and synchronized companion device experiences.

Effectively, the intelligence in the proposed architecture is decentralized – residing in the applications. The proposed video control plane merely acts as a broker for these applications, providing all of the essential information they need to make real-time decisions and deliver cloud-based video services. Together, the shared data cache and common messaging bus functions almost like a web-based news feed: various elements throughout the system publish events or information, and every other element in the system can subscribe to any relevant information. This web services-based communication framework provides inherently more flexibility and scalability than a client/server model, and represents a significant departure from traditional closed video systems, and even some contemporary IP video systems. It should also be noted that the operator need not store every piece of data in the system in a single, centralized real-time cache. Less frequently used data can easily be stored elsewhere in the infrastructure, and remain accessible via the same messaging bus.

Employing this web-based infrastructure for brokering information between applications, the proposed architecture can:

- Orchestrate cloud-based services across multiple devices in real time
- Perform end-to-end session and client/device management for both managed STBs and cloud-connected endpoints
- Provide an interface between multivendor systems and applications
- Achieve service provider-grade scale and availability
- Allow for fully customizable user experience and applications

### An Open System, Incorporating Standards and Web Design Approaches

An essential characteristic of the proposed next-generation video control plane is that it is based on open standards to allow for maximum flexibility. As a result, it integrates with legacy systems and with any standards-based third-party technology or application. It is also designed to allow operators to change technology vendors, equipment, systems, etc., as they choose. By avoiding proprietary standards that can age quickly, it provides a more future-ready architecture.

Along the same lines, the proposed architecture is designed to increase service velocity by incorporating web services and design approaches. It is modular and loosely-coupled, allowing for phased introduction of services and technologies. As described later in this paper, the architecture also facilitates service velocity and flexibility through its ability to dynamically manage workflows. The following sections describe this architecture in greater detail.

## OPEN, STANDARDS-BASED SOFTWARE PLATFORM

In a legacy QAM-modulated video system, applications and middleware reside on STBs, and video content acquisition and provisioning systems populate back-offices. In this legacy environment, deploying and maintaining software, especially when proprietary, requires a substantial investment of time, skill, and financial resources. As operators transition to cloud distribution based on IP, however, they can take advantage of existing Internet-based standards to achieve greater scalability, and afford greater flexibility and service velocity.

The core technology envisioned in the proposed distributed video control plane has several facets. It is based on an open-source programming language such as Java, it uses web services-based SOA design principles, and it functions as a workflow engine.

Open-Source Technologies

As stated, the proposed architecture is fundamentally an Internet-based approach to video service delivery. As such, it should be based on an open-source or standards-based language such as Java. By using open standards, an operator has its choice of additional open-source libraries and frameworks that simplify the process of building, testing, piloting, and enhancing new services and features.

Java in particular is an excellent fit for a distributed video system. Known for its portability and openness, Java is an apt programming language for an extensible multi-screen video platform with automated workflow, session, content, and other video control plane functionalities. Because Java language code can be represented in the intermediate form Java bytecode, Java can run on a multitude of operating systems that otherwise would require platform-specific machine code. Other reasons for Java's popularity include its efficient memory management and relatively simple object model.

Service-Oriented Architecture

The proposed architecture is designed to operate on a service-oriented architecture (SOA) platform, and should be operable on any standards-based SOA platform. SOA principles allow for loose couplings between clients and servers, and facilitate the development of services independent of the client or underlying platform. As such, SOA design principles help facilitate the systems' ability to share information and functions among multiple distributed applications and system elements in a widespread and flexible manner.

Workflow Capabilities

At the top of the SOA stack is the workflow engine. Users can create workflows by using the workflow engine's graphical editor or by editing XML files directly. Workflows are not hard coded, are not compiled, and do not require any kind of system downtime for modification.

The system supports the real-time execution of multiple workflows, which can be easily extended to add new features rapidly in a controlled manner. These workflows can be characterized as:

- **Atomic:** Once a request begins executing a workflow, it will continue executing that workflow, without interruption.
- **Extensible:** Workflows are defined according to standards and can be modified using standard tools.
- **Flexible:** Workflows are built with a series of nodes that support the basic concepts of sequential operations (IF statement, multi-threading, etc.)

As with the rest of the proposed architecture, the control plane is designed to use a standards-based workflow engine. However, the workflow engine must be fine-tuned for speed to function effectively in a video system. After all, while a lag of a few seconds may be acceptable when initiating a video-on-demand (VOD) session, such a delay in more advanced real-time cloud applications (i.e., pausing or rewinding content in a cloud DVR service) would render the service unusable. Cisco worked to optimize open-source workflow engines to meet these demands, and contributed those gains back to the open-source community.

The communication with external or third-party application procedures also requires a new kind of workflow invocation service. This allows operators to deploy and test features to various overlapping subsets of subscribers, based on criteria such as set-top media access control (MAC) address, account number, service group, or requester IP address range. In the proposed control plane, options for selecting the workflow to execute also include service endpoint and designated market area (DMA) code.

Effectively, this workflow engine provides the tools to create rules, and supports tremendous customization. It also allows for greater flexibility, resiliency, and velocity when rolling out new functions or applications. Traditional systems require operators to shut down the system and come back up when implementing changes, and are unable to execute multiple workflows in parallel. The proposed control plane architecture supports multiple workflows, allowing operators to modify a workflow dynamically, without affecting connections already in use. This enables greater innovation and service velocity by giving operators web-like workflow capabilities, such as A/B testing, where a service provider can use multiple workflows that differ from one another in order to target a specific set of customers, endpoints, or video assets. For example: "Apply workflow (WF) 1 if customer lives in Massachusetts. Apply WF2 if customer is also a high-speed data customer." The engine can even target individual endpoints, which is useful in beta-testing to familiar customers. Operators can test multiple similar versions of a function or application, and expand or roll them back with relative ease – providing a major boost in their ability to innovate, in less time, at a lower cost. This is typically not possible with a traditional video system.

In the same way, because the video control plane operates via virtualized software instances in a data center rather than tightly coupled hardware systems, the proposed architecture also gives operators greater ability to scale services dynamically. For example, the datacenter can dynamically spin up resources on the East Coast as prime time approaches, and shift those resources to the West Coast as the evening progresses.

## VIDEO CONTROL PLANE

The next-generation video control plane is built upon the workflow system described above. It performs three key functions:

- Real-time session management
- Resource management
- Business policy management

### Session Management

A next-generation video service must provide a framework for session management across multiple screens, in a variety of video applications. This can include "session-shifting" across devices and networks (i.e., beginning playback on the a TV via the STB, pausing, and then resuming playback later from a smartphone connecting over a cellular network), as well as more advanced applications. One example is a "companion screen" experience that integrates both a managed service (e.g., video content delivered to a managed STB over the service provider's managed video network) and unmanaged services (e.g., IP data services that complement the STB video service and may be synchronized with it, but are delivered to an unmanaged device such as a tablet or smartphone).

The session management function of the proposed control plane architecture performs the majority of the core functions of a traditional video system back office, but includes support for both QAM and IP environments, and operates according to open, web-services practices.

The traditional approach to video services has been to perform all session processing within a proprietary environment, based on a client/server model with the STB endpoint tightly coupled with back-end servers controlling session and state. In other words, all of the logical software components (resource management, business policy, billing, entitlement, etc.) communicate with each other using closed protocols. The video control plane envisioned here functions differently: The client drives session and state, dictating the format and streaming bit rate required for a specific viewer using a specific device. As discussed, this model is based on the way software works on the web, where diverse applications and devices from multiple vendors share a common delivery language and services framework. By applying this model of session management to video services, operators can gain more flexibility and control.

Self-contained or siloed legacy video back office systems have also had trouble scaling to handle growing volumes of traffic. For a large service provider running a popular VOD service with millions of concurrent users may require dozens of separate session management systems just to handle the load. Additionally, if one part of the system was resource-constrained, the system needed to be expanded as a whole, rather than simply adding an additional node or virtual machine to support that function. Since the distributed control plane envisioned here uses SOA design principles, a shared data cache to store real-time state information, and a messaging infrastructure designed for Internet scale, a single session management system can theoretically scale to serve unlimited clients.

Resource Management

The objective of resource management is to manage video objects in the video distribution system and balance the load among video servers/streamers and networks. To achieve this, the proposed architecture employs resource management tools that intelligently and automatically consider such factors as allocation, video server selection, replication, and cache management to help ensure optimal load balancing among video servers, and to minimize the delay for video requests to be served.

In addition, the resource management function of the proposed architecture relies on the same shared cache as the session management function, and affords the same degree of scalability. The distributed control plane also uses a common resource manager for both legacy and cloud services.

The session and resource management functions of the proposed next-generation video control plane are implemented as logical individual components. Multiple component instances may be deployed throughout the operating environment as virtualized applications in varying degrees.

Flexible Policy Management and Dynamic Business Rules

A traditional TV video distribution system applies a variety of business rules to the delivery of content and services to consumer endpoints. These rules can encompass specific times content can be distributed, specific markets barred from receiving content (for example, blackout rules governing some sports broadcasts), rules barring a device from receiving content without the right content security, etc. For a cloud-based, multi-screen video service, creating explicit rules governing how content can be delivered to every possible IP endpoint in every possible location is simply not practical. The proposed distributed video control plane therefore includes a more flexible business rules engine capable of dictating rules for delivering content through the cloud (for example, allowing streaming of entitled content to any authenticated IP device that supports a particular digital rights management [DRM] system).

The business policy management function of the proposed architecture is powered by a next-generation business rules engine that brings deeper sophistication and intelligence to the process of delivering video services. The business policy management function gives operators a greater level of detail about how the content is going to be consumed. Effectively, it allows operators to store all business rules and policy logic in a centralized script or table, where it can be accessed by other elements in the system. This globally accessible data repository provides the system with vital information on sessions, devices, business rules, etc., and facilitates automated decision-making by the network in applying policy.

This repository is flexible, allowing the operator to define business rules within an XML-based searchable workflow. The workflow can incorporate not just native services associated with the control plane, but allow operators to make “off-board” calls to existing or third-party services. For example, a mobile operator could configure the workflow to make calls to an existing location service, instead of having to recreate that service for the next-generation video control plane. Furthermore, the operator can call entire off-board workflows (not just services), and take advantage of existing third-party business logic instead of having to recreate it. This is yet another example of the value of using an open workflow definition language.

Contrast that with current video systems, which give operators only the most rudimentary knowledge of where the content is being delivered. Essentially, operators know only whether content will be delivered on-net or off-net. With the business policy management tools in the proposed next-generation video control plane, operators can see beyond that, to know exactly what type of device the content is going to and the subscriber consuming it. This allows them to define more finely grained, sophisticated rules for delivering content, giving them an opportunity to generate additional outlet revenue and reduce capital expenditures.

## APPLICATIONS AND USE CASES

Once the next-generation video control plane architecture is in place, operators can deploy all applications involved in the video service on top of this framework. This includes essential applications such as device authentication, service assurance, emergency alerts across multiple devices, etc. The proposed architecture is also well-suited to enabling the unique capabilities essential to delivering a cloud-based, multi-screen video service, including the ability to authenticate users and devices among multiple back-end systems (both legacy and IP), and the ability to manage and assure services across multiple devices and networks. The proposed video control plane architecture can also support more advanced applications that take full advantage of cloud capabilities, such as a synchronized companion device experience and a cloud DVR service.

### Authentication Among Multiple Back-End Systems

A next-generation video system must communicate with varying types of back-end billing systems, from mainframe to web-based. For cable operators especially, authentication systems are based on a tightly coupled, usually proprietary authentication process between the STB and the back-end billing system, where the STB boots and queries the system for each subscriber’s/STB’s entitlements. These legacy authentication systems will likely remain in place for the foreseeable future, but they present a significant barrier for newer IP media delivery systems and endpoints, which authenticate in very different ways.

Some IP video delivery systems in use today have attempted to bridge this divide. Typically, however, this entails invasive changes to the core of the IP application to support communication with legacy authentication systems. This is to be expected: closed video systems communicate via closed, proprietary mechanisms. Exposing core software elements for authentication (or billing, or other services that require communication with a conventional video back-end system) is typically a significant development project, undertaken at a significant cost. In addition to the potentially onerous costs of this custom integration, this process also impedes an operator's ability to quickly design and deploy new service offerings that interconnect with legacy billing and authentication systems.

As discussed, the proposed distributed video control plane architecture is an open system. It provides a common infrastructure that can unify legacy authentication and billing systems with newer IP distribution services. In the proposed architecture, this is accomplished via protocol conversion mechanisms that broker this communication and provide the interface to various back-end systems. Rather than incorporating communication with legacy systems into the core of the video control plane, protocol conversion mechanisms deployed at the "edges" of the system communicate with legacy systems, while the core of the distributed services platform remains purely IP-based, and highly scalable. Effectively, this model preserves a kind of stateless, web-aware application core, even as the software communicates with older billing and authentication systems, and frees the video control plane from having to adapt to legacy systems. And, since these protocol conversion mechanisms are open and standards-based, operators need no proprietary intelligence to develop applications to communicate with the IP system.

<u>Service Assurance and Management Across Multiple Device Types</u>

An adaptive bit rate (ABR) streaming capability – the ability to optimize video streams for the specific connecting endpoint based on real-time network conditions – is an essential requirement of a next-generation video system. However, operators cannot rely on a standard, generic ABR functionality. More sophisticated ABR management is required for next-generation services, augmenting basic information about client and network with an additional layer of business rules and priorities, based on a more sophisticated awareness of the network and the subscriber.

Consider a premium cable customer streaming a high-definition program to a TV in her living room via a "smart" TV. Upstairs, her children begin streaming a movie on an iPad. The cable operator would not want an automated ABR function to downgrade the living room TV stream to standard-definition video halfway through the program. The video control plane envisioned here draws on network intelligence to inform ABR decisions, and uses the software control plane to effect quality-of-service (QoS) prioritization and bandwidth reservation in the network. In the scenario described above, the system can draw on operator-defined rules, as well as customizable rules defined by the subscriber, to assure that the large-screen TV in the living room takes priority over a mobile device, and that the primary subscriber takes priority over secondary users.

<u>Synchronized Companion Device Experience</u>

A session-aware control plane for both IP and QAM traffic can be used to deliver complementary viewing services consumed on two devices at the same time, such as viewing a live TV broadcast while using a synchronized application on a tablet. Synchronized companion services can include push applications for polling, alternative content, instant replay and other time-sensitive experiences. A viewer watching "American Idol," for example, could receive bios of contestants pushed to the companion screen when contestants come on stage, and a real-time voting application to vote for the winner during the show.

The proposed video control plane architecture is an ideal platform for deploying these types of synchronized multi-screen applications. The real-time messaging infrastructure allows the operator to synchronize the web applications, companion screen, and live TV stream to enable these and other real-time interactive multi-screen applications. The same messaging infrastructure can also support social applications, such as the ability to allow a subscriber watching a TV show to identify and chat with friends who are watching the same show. These and other synchronized multi-screen applications can benefit service providers by differentiating their services, enhancing subscriber loyalty, and driving customers to higher subscription tiers – and all are enabled by the proposed video control plane architecture.

### Cloud DVR Service

Many operators are now seeking to move content storage back into the network, rather than on DVR appliances in the subscriber's home. These nDVR or cloud DVR services offer operational advantages over traditional DVR service offerings – most notably, the ability to reduce capital expenditures on costly DVR appliances. For a cloud DVR service to function effectively, however, operators need to manage those cloud-based resources and enforce business logic across unmanaged devices to deliver a seamless transition of experience for the subscriber. A cloud-based DVR service requires a system with extremely high performance that can scale sessions well beyond normal VOD utilization behavior, with hundreds of millions of assets.

The proposed video control plane architecture with its next-generation, high-performance workflow engine meets all of these requirements. In addition, because the architecture employs a web-services based control plane, it simplifies the process of extending the DVR service to other IP devices. These capabilities allow service providers to roll out differentiated, revenue-generating services like cloud DVR itself, but also allow for additional revenue-generating services. For example, operators can invoke a workflow that detects when a subscriber's DVR storage is nearly full, and pushes a message out to the user's companion device to ask if the subscriber would like to purchase additional cloud DVR storage space.

## CONCLUSION

For many service providers, the shift from traditional QAM-based video architectures to an open, cloud-based distribution model is a five- to 10-year transitional exercise. But the video services landscape is rapidly evolving, and this transition already is well underway. Video service providers are faced with the need to adopt cloud distribution capabilities even as they serve existing customers over legacy infrastructure.

A key to navigating this transition is handling video session, resource, policy, and other functions from a common software control plane. The proposed next-generation video control plane architecture can manage this transition, offering a range of overlapping and mutually reinforceable benefits. These include:

- **Service velocity:** Transitioning to a video control plane based on open programming languages and web services offers a competitive advantage. By providing web services management capabilities for video services, the platform helps enable rapid testing and deployment of new features and applications to a large number of end devices and targeted sets of customers, while allowing operators to continue taking advantage of legacy systems in a graceful manner.
- **Flexibility:** A core value of SOA-based systems, flexibility is exemplified in the proposed control plane architecture and its real-time execution of multiple workflows. Operators have the ability to develop and test all manner of features and applications, in both targeted and large-scale deployments. The ability to dynamically manage multiple workflows in parallel also allows operators to roll out changes and upgrades with no downtime.
- **Scalability:** While legacy systems could expand to handle growing volumes of traffic only with difficulty and cost, the proposed architecture scales simply by adding another node or virtual machine to support any given function. Based on Internet design and communication principles, it provides a platform for video services at massive scale, capable of supporting theoretically unlimited clients.

Ultimately, the proposed video control plane fills a critical gap in the ongoing evolution of service provider networks into a cloud-based delivery framework. Legacy transport technology has many years of life remaining. This platform allows providers to continue supporting those services, while adapting to emergent market realities and using the most efficient cloud-oriented software architecture, techniques, and methods to deliver compelling new subscriber experiences.

## REFERENCES

[1] Cisco. (2011). *Cisco Visual Networking Index: Forecast and Methodology, 2010-2015.* Cisco. Retrieved April 1, 2012, from Cisco.com: http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf

[2] Cisco. (2011). *Cisco Global Cloud Index: Forecast and Methodology, 2010-2015.* Cisco. Retrieved April 1, 2012, from Cisco.com: http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf