The Headend Revisited: A Multi-Service Video Data Center for the Modern MSO

S.V. Vasudevan, R. Wayne Ogozaly
Cisco

Abstract

Cable MSOs are under competitive pressure to deliver an increasing number of new services to effectively compete with rival service providers, as well as emerging over-the-top players. This paper proposes a strategy which evolves the traditional cable headend into a multi-service Video Data Center which is well positioned for the challenges of the modern era.

The Video Data Center combines the best of the traditional video headend with proven data center techniques. This paper discusses the intersection of these digital delivery systems and describes a next-generation data center model. Internet data center architectures are examined and compared with the data delivery requirements of a modern digital cable system. Specific insertion strategies will describe how cable delivery networks can evolve to a multi-service Video Data Center, spanning broadcast, on-demand, and a next generation IPTV services, while integrating the best of today's regional headend with advances in data center technologies.

INTRODUCTION

Cable MSOs face considerable operational challenges as they seek to smoothly integrate innovative new video services into their existing delivery platforms. As we witness the deployment of more interactive video services, IPTV overlays, internet streaming to set top receivers, and other advanced technologies, of equal importance is the smooth integration of these new functions with existing video, voice, and data services across headend, transport, and last mile networks. These new services must coexist with and even complement existing services such as linear broadcast, switched digital video, and on-demand video services. Unlike the passive linear broadcast delivery model, these advanced services implement much more dynamic traffic flows between subscriber and headend, resulting in increasing operational considerations.

But video delivery systems are not the only area that has recently experienced great infrastructure and traffic growth. Consumption of internet web content, information which is compositionally populated with a mixture of database content and increasingly rich media components, has fueled the evolution of data centers, super-evolved versions of the "computer room" of the 1970s. Traditional internet content and service providers have deployed data center computing systems to support the 24x7 operation of web, application, and database servers, supporting the efficient delivery of interactive web content to large user populations. In order to meet the increasing performance and scale demands of a global web audience, enterprise computing, network and storage architectures have evolved into organized and tiered architectures to support high-transaction information throughput.

Most modern cable headends and hubs already use IP as a common transport for data, voice, and video traffic. Accordingly, video delivery to consumers is rising in scale, growing in interactivity, and being distributed to an increasingly diverse set of video-capable devices. The increasing adoption of interactive services such as VOD, SDV and IPTV will present a similar increasing transactional traffic load on service delivery platforms. As these changes occur and content networks expand their reach to a national and global scale, and application server compute requirements grow to respond to the increasing transactional load, the headend of old begins to resemble a data center that is serving a variety of data types, led by video, to a large subscriber population.

The service provider industry has already benefitted from the leverage of technology that was originally developed for orthogonal markets. The very use of IP switching and routing for video delivery serves as an excellent example. The switches and routers that now carry triple-play traffic were originally developed for the enterprise networking market. All of the staff-years and capital investment in ASICs, software and hardware development in the first 10 years of switch and router development went towards products that were intended for the wiring closet, and essentially funded by the global enterprise IT market. By adapting these technologies to the performance, scale, and reliability requirements of the service provider market, and through clever media processing techniques that enable the safe carriage of media streams over IP networks, IP switching and routing of entertainment grade video is now seen as the norm. The service provider industry enjoyed a drastic reduction in per-bit transport costs as a benefit of this technology transfer.

The remainder of this paper examines the evolution of the enterprise/internet data center, and provides an overview of networking and computing requirements that have driven the evolution of data centers in their efforts to scale and server their requirements. We compare and contrast between internet data center information processing requirements and cable headend information processing requirements. Finally, we propose some appropriate data center techniques and best practices that can help cable headends more effectively scale the throughput of their delivery systems.

1.0 THE DATA CENTER EVOLUTION

Introduction

Much like the cable video headend, data centers in the enterprise are under increasing performance pressures. The enterprise data center is being influenced by shifting business pressures and increasing operational limitations. The new demands of the enterprise require enhanced video services, greater collaboration, near-instantaneous access to applications and information, and compliance with ever-tightening regulatory compliance. These strains have manifested as operational issues relating to power and cooling, efficient asset utilization, element and system management and monitoring, escalating security and provisioning needs, and increasing operational expense.

A fundamental metric of data center efficiency relates to cost, power, space, computational power, and ingress and egress throughput. Any data center that is designed, trades off between these metrics, depending on the underlying requirements of the offered service. Beyond this, there is increasing attention paid to the “-ilities”: [1]

- Scalability
- Reliability
- Availability
- Serviceability
- Manageability
- Security

The data center architecture has evolved to accommodate these demands in an efficient way. With each architectural shift, the network has become an increasingly important ingredient to enable the latest round of transitions. The transformation to the newest evolution of Data Center 3.0 technologies focuses on a service-oriented design, consolidation of server farms, virtualization across disparate networked elements, and enhanced automation to manage the heavy load of requests for content, applications, and services.

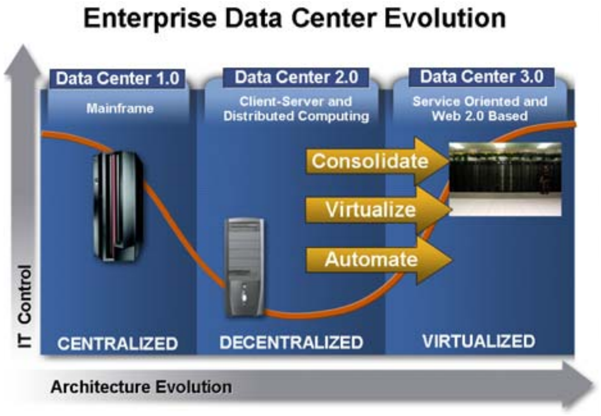

Figure 1 - The Enterprise Data Center is undergoing an extreme makeover, which is focused on server consolidation, service virtualization, and controlled automation.

The net result of this transformation is a massively scalable architecture; some data centers are built to contain upwards of 100,000 servers, stitched together as a fabric of virtualized elements.

Many design elements of the enterprise data center are applicable to the next generation cable Video Data Center. To better appreciate the similarities and differences, it is useful to present a brief overview of enterprise data center architecture.

Attributes of the Modern Data Center

Most modern data center architectures are based on a proven layered approach, which has been deployed in many of the largest data centers in the world. This layered approach provides the foundation of the contemporary data center and includes 3 major components: a ***core*** layer, ***aggregation*** layer, and ***access*** layer [2].

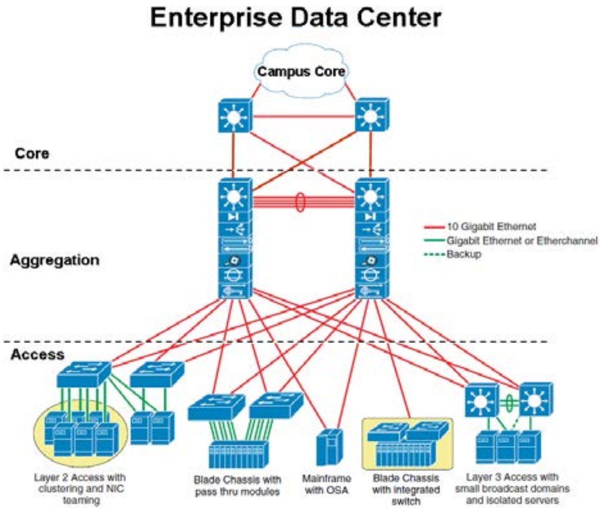

Figure 2 – The Data Center architecture implements a layered approach designed to diffuse massive loading across thousands of server elements [2].

It should be first noted that in enterprise data center terminology, the use of the descriptors *core*, *aggregation*, and *access* are used in an almost opposite manner to the way these terms would be used an interpreted by a cable headend engineer.

The core layer provides the high-speed packet switching backplane for all flows entering and leaving the data center. The core layer provides connectivity to multiple aggregation modules and provides a resilient Layer 3 fabric for this redundant system. This layer runs an interior routing protocol, like OSPF or EIGRP, and load balances traffic between the core network and the aggregation layer.

The aggregation layer includes multi-service switches which manage the interaction between servers through a Layer 2 domain. Server-to-server traffic flows through the aggregation layer and use features such as firewall and server load balancing, to optimize and secure applications.

The access layer is where servers physically attach to the network. Unlike the cable world, “access” in the data center refers to core of server design, where the network provides access to thousands of servers. Server components consist of 1RU servers, blade servers with integral switches, blade servers with pass-through cabling, clustered servers, and mainframes. The access layer network typically consists of modular switches and integral blade server switches. These switches manage Layer 2 and Layer 3 topologies which distribute flows across various server pools in different broadcast domains.

Distributing Requests Across Server Pools

Many applications run inside an enterprise data center. Each application may have one or more IP addresses associated with it, providing an identity to which Internet users send their requests. As Internet requests are received through the core and aggregation layers, specialized load balancers distribute these requests across a pool of targeted servers, using a combination Layer 2 VLANs and internal virtual IP addresses. A series of complex algorithms sift through requests, associate a virtual IP address and VLAN domain, and distribute the requests to the private addresses of physical servers [3]. These complex load balancers may confine requests for specific applications to a pre-allocated pool of servers, which are segregated into manageable clusters. In that way, server performance, redundancy, and security can be managed and scaled in a predictable way.

The multi-tier model in today's data center is dominated by HTTP-based applications. Scalable web-based services are built with multi-tier processing layers which include web, application, and database tiers of servers. Web servers field http:// requests and pass transactional information. Application servers fulfill these transactional requests by consulting one or more databases or sub-services. This information is presented back to the web server, which composes the properly formatted HTML page containing all requested information. This model enables the independent scaling of presentation, application or database resources as the service requirements dictate.

The multi-tier model may use software that runs as separate processes on the same machine using inter-process communication (IPC), or on different machines with communications over a switched network.

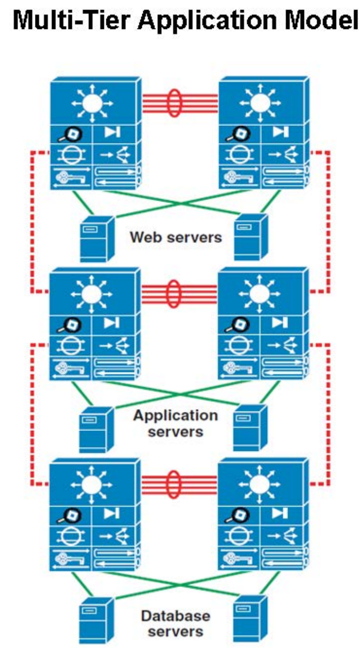

Figure 3 – For most HTTP-based applications, Web, Application, and Database servers work across clearly defined server pools using a tiered model, to process HTTP-based service requests [2].

2.0 THE VIDEO DATA CENTER

Comparing and Contrasting Enterprise and Video Data Centers

An important distinction between enterprise and video data centers can be understood by analyzing the traffic pattern of each service.

For a 3-tired web site, the amount of information offered is relatively small (an http GET method), the computational demands are moderate, and the amount of information returned is also relatively small (the average web page size is about 300kB). For a search engine web site, while the information offered is also small, the computational demands are very large (a search request may be dispatched to dozens of index servers working in concert) [4], and the information returned is also very small. For a video service provider with a service such as VOD, the information offered is relatively small, the computational demands are moderate, but the information returned is very large (a 2 hr movie in HD moves 14GB of data). Thus we can see that there is an immediate distinction in both computational as well as egress throughput requirements between the internet data canter and the video data center. Accordingly, the architecture can be optimized for the specific service requirements. But it cannot be mistaken that both types of data center have shared goals in the areas of scalability, reliability, availability, serviceability, and manageability, and security.

In this spirit, a video headend architecture implementing selected data center best practices is proposed.

Combining the Best of Data Center and Video Headend Technologies

As MSOs continue to expand the number and scale of video services, the resulting permutation of content, formats, interactive features, and end-devices has placed a burden on the traditional cable headend. This burden may be felt within the headend in a number of ways. First, services such as VOD, SDV, and IPTV require a resilient 2-way communications network, including outside plant and home connections. Upstream traffic will organically grow with the number of subscribers and applications. Correspondingly, the performance of the application servers that process these upstream requests will need to improve. In many cases the end result of the subscriber request will be the switching or streaming of video content. This information must be delivered via a resilient IP network with strict packet loss requirements. In the process of being delivered, the stream may be spliced, encrypted, or otherwise transformed by another cascading device. Additionally, sophisticated processes and systems are needed to manage the increasing rotation of licensed on-demand content, as well as ad content.

This evolving mix of new services, expansive capacity, and more interactive services, has triggered a evolution of the video headend. To accommodate this new service mix, the *Video Data Center* aims to combine the best of data center and traditional headend models into a more scalable and manageable system.

Video Data Center Design Goals

A number of key components have emerged to address the needs of the Video Data Center in the modern era:

- ***Multi-Service Support:***

  The next generation Cable Video Data Center not only needs to support core cable video services including broadcast, Switched Digital Video, and VOD, but it must also modularly accommodate new services such as IPTV, internet streaming to the TV, and video streaming to the PC or handheld device.

- ***Service Independent Scaling:***

  Each video service (broadcast, SDV, VOD, IPTV, etc) should be able to independently grow in capacity, without disrupting existing services. The IP network and multicast control plane will play a key role to both identify and independently manage each video service.

- ***Resource Modularity and Demarcation:***

  Each video resource within the Video Data Center (acquisition, grooming, encryption, etc.) is organized as a modular component with clear demarcation points. A pair of multicast addresses are used to identify the input and output points of a video stream as it traverses the delivery network. The IP network and multicast control plane are used to identify, grow, and independently manage each video resource. This methodology will prepare the way for service virtualization and improved scaling. It also enables the development of IP appliances to provide future stream processing functionality, or super-appliances that consolidate more than one stream processing function.

- ***Service Virtualization:***

  Service virtualization has many meanings, both in enterprise and service provider information processing. From the service provider perspective, it describes the implementation of a service that is logically viewed as a single client-server entity, but may physically be realized by more than one application instance. In content networking, virtualization refers to the abstraction of a video object from its physical attributes or location. For example, a movie may be stored in multiple locations, and be transcoded into multiple formats to suit multiple receiving devices, but in a virtualized implementation it need only be known as a single entity.

- ***Fault Containment and Resiliency:***

  The combination of embedded quality monitoring and enhanced redundancy techniques allow video faults and outages to be rapidly identified, circumvented, and contained to the Video Data Center. The Video Data Center provides a fully redundant system in which most failures are contained and not propagated throughout the video network.

- ***IP Early Acquisition:***

  The next generation Video Data Center will acquire large volumes of content, both real-time and asset-based, sourced from diverse locations in a wide range of formats. Much of this content will be shared across multiple video services. The "IP Early Acquisition" module ensures that this content is available to all services in an IP format, at the earliest point in the content acquisition process.

- ***Improved Operations and Management**:*

  Stream visibility and quality monitoring across all services at strategic measurement points is always important. Since video streaming/processing devices and networking devices comprise the video delivery chain, it is important to have monitoring tools that can present a synthesized and unified view of the delivery system. The Video Data Center integrates video service monitoring across the delivery network.

3.0 BROADCAST SERVICES USING VIDEO DATA CENTER TECHNOLOGIES

Implementing Switched Digital Video within the Video Data Center

While there are some fundamental differences between the volume and type of information exchanged from a headend versus a database-driven website, there are also many similarities between each data center's design goals. Many data center architecture and networking practices can be applied in the cable headend. As an illustrative example, we

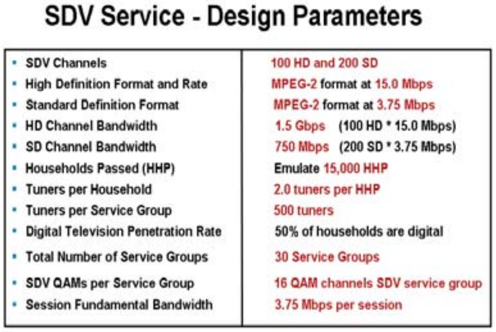

SDV Service - Design Parameters

SDV Channels	100 HD and 200 SD
High Definition Format and Rate	MPEG-2 format at 15.0 Mbps
Standard Definition Format	MPEG-2 format at 3.75 Mbps
HD Channel Bandwidth	1.5 Gbps (100 HD * 15.0 Mbps)
SD Channel Bandwidth	750 Mbps (200 SD * 3.75 Mbps)
Households Passed (HHP)	Emulate 15,000 HHP
Tuners per Household	2.0 tuners per HHP
Tuners per Service Group	500 tuners
Digital Television Penetration Rate	50% of households are digital
Total Number of Service Groups	30 Service Groups
SDV QAMs per Service Group	16 QAM channels SDV service group
Session Fundamental Bandwidth	3.75 Mbps per session

Figure 4 – The Switched Digital Video prototype service is implemented using Video Data Center techniques and includes 100 HD and 200 SD channels, mapped to 30 Service Groups, in a 15,000 household hub site.

present a hypothetical system design for a Switched Digital Video (SDV) service implementation. The SDV service described in this paper was designed and tested using the design parameters presented in Figure 4.

The advent of Switched Digital Video (SDV) technology provides a fundamental change in the way the industry delivers digital video entertainment. With SDV, service providers have the ability to offer a wider variety of programming while managing HFC network bandwidth in a sustainable way. The SDV architecture switches only selected content onto the HFC based on channel change requests from users within a service group. Thus, content that is not requested by any user in a particular service group does not occupy HFC bandwidth.

MSOs considering the deployment of SDV technology face three operational challenges. The first includes the integration of a large number of operational components required by the SDV service. Unlike the linear broadcast model, SDV implements a much more dynamic service in which SDV Servers dynamically map requested content to various QAM channels across hundreds of service groups. This increases operational complexity.

The second challenge for Cable MSOs considering SDV is how to smoothly integrate an SDV service with existing video services across the headend. The SDV service must co-exist and even complement existing services such as linear broadcast, VOD, IPTV, and other streaming techniques. The SDV system must also be managed and scaled concurrently with these adjacent video services.

A final challenge is in capacity planning for future growth. Since the SDV system has the ability to admit a great number of linear programs with only a modest increase in required HFC stream resources, operators need to rely on tools to provide visibility of system resource utilization. When the opportunity is presented to augment the amount of switched programming, the Video Data Center should be able to accommodate capacity expansions without any major disruption to the system in place.

Layered Network Approach

The Video Data Center implements a layered network approach to manage resource access, security, and scaling across a diverse set of video elements. Figure 5 highlights the Core, Aggregation, and Video Resource Access layers within the Video Data Center. Similar to the Enterprise Data Center, these layers provide a Layer 3 core connection to the Regional Network, aggregation and load balancing of video flows across various cable services, and access to a range of modular video resources. Individual video resources are linked together using the multicast control plane to create a SDV cable service.

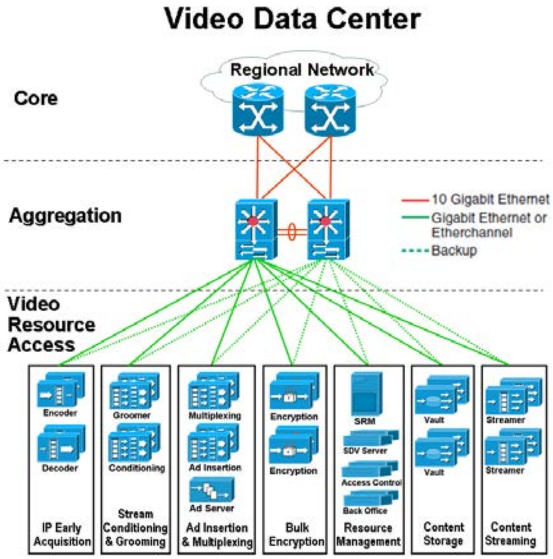

Figure 5 – Similar to the Enterprise Data Center, the Video Data Center implements a layered network approach, including Core, Aggregation and Access layers.

The Aggregation Layer includes redundant multi-service switches which provide GE access to individual resource elements. Since broadcast services within the Video Data Center operate on a much smaller scale than the traditional data center, a fewer number of multi-service switches are required. In this SDV example, two high density switches provide dual-homed GE inter-connects to all video resource elements.

Enabling Service Virtualization

Service virtualization is a critical principle upon which many of today's advanced data centers are built. The goal of service virtualization is to provide a more efficient delivery platform which is better aligned with the massive scale and complexity of today's networked environment. Although service virtualization in the data center continues to evolve, different components are directly applicable to the Video Data Center:

- ***Network Enabled Modularity***

  Resources are defined as modular components with clear network demarcation points, providing a more scalable and flexible design. Capacity can typically be added with minimal impact to existing services. (Example: SAN-based network storage which is independently managed and scaled from the server farms)

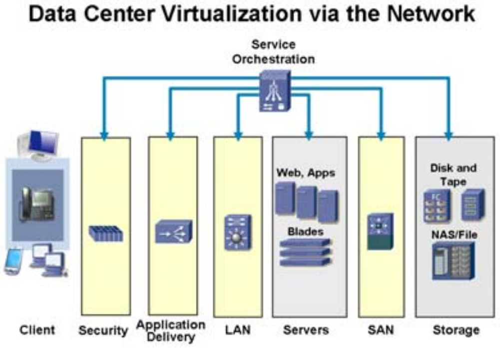

Figure 6 – Data Center storage is a modular resource, which is networked via the SAN, and can be independently managed and scaled [2].

- ***Distributed Service Models***

  Providers have the flexibility to distribute various processing stages across multiple servers, regional networks, and national backbones. Resource managers that are present for interactive services will play a principal role in the Video Data Center. These managers will optimize the allocation of resources to satisfy a number of parameters, including cost, latency, even potential revenue opportunities.

- ***Masking Complexity within Multi-Service Environments:***

  Improved resource management systems, advanced load balancing, and multi-tiered access networks hide the underlying technology from the end user. Content is available in multiple formats to a myriad of consumer devices. Exactly how this content is acquired, formatted, and delivered is hidden from the end user.

Similar techniques are implemented within the Video Data Center. As described in Figure 7, video resources operate within a modular architecture, defined by clear network demarcation points.

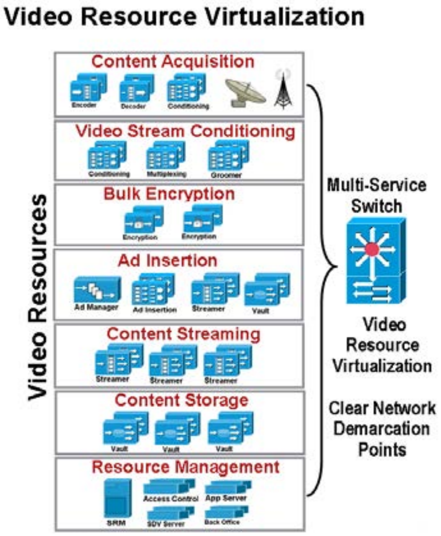

Figure 7 – Similar to the Enterprise Data Center, video resources are split into modular components and available as distributed network elements

These video resources are shared by multiple video services including linear broadcast, SDV, time-shift TV, PCTV/IPTV and other streaming services. This combination of resource modularity and network demarcation enables specific functions like Ad Insertion or VOD Streaming to be contained within a centralized Headend or distributed throughout a regional network.

SDV Service Creation

The multicast control plane plays a vital role in the Video Data Center design. Figure 8 describes how the multicast control plane strings together modular resources to create video content for the SDV service. In our example, 300 SDV SPTS streams are mapped to unique IP multicast group addresses. Devices along each processing stage issue IGMPv3 JOIN messages to draw specific multicast streams to each device over the layer 3 GE network. Redundant multi-service switches provide the routed control plane to manage the handoff of these streams between each processing stage.

Each of the major processing stages utilizes a unique multicast source address to identify primary and backup video sources, and obtain the correct content. Source-Specific Multicast (SSM), a feature of IGMPv3, also allows each stream to maintain a single multicast group address (but varying source addresses), spanning all of the video processing stages.

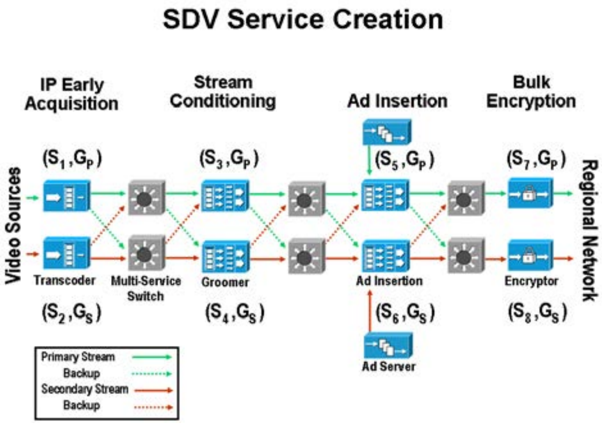

Figure 8 – The 300 channel SDV service is created by stringing together processing elements, using the SSM multicast control plane.

In this example, the same multicast group address is used for each SPTS stream across the four processing stages. There is one group address "**G_P**" mapped to the 300 primary video streams in the SDV tier, and a second group address "**G_S**" mapped to the 300 secondary or backup streams. This SSM feature greatly reduces the number of multicast groups used throughout the Video Data Center since a single group address (with varying source addresses along each stage of the delivery chain) can follow a stream throughout the Video Data Center. This technique also provides a clear network demarcation point between each stage and offers MPEG monitoring tools complete visibility into all SDV multicast streams.

Within this Video Data Center architecture, the SDV video path contains four processing stages. The IP Early Acquisition stage provides redundant access to all SDV video content. During this stage, SDV video content is acquired from redundant sources including satellite, off-air, and terrestrial links. SDV content is transcoded to an IP MPEG format at this early stage. Where necessary, video streams are converted from multi-program transport streams (MPTS) to single program transport streams (SPTS), as used by SDV.

This second stage provides Stream Conditioning of all SDV channels by redundant groomers. Since redundant copies of the SDV video streams are available, this stage performs a stream selection process in which each groomer independently selects the best available copy based on ETR-290 MPEG quality measurements. After stream selection, groomers rate cap the resultant SPTS streams from Variable Bit Rate (VBR) to Constant Bit Rate (CBR), as used by SDV.

The third stage of processing provides Ad Insertion, delivered by ad servers and standards-based SCTE 30/35/130 MPEG insertion techniques.

In the final stage of SDV processing, two independent encryption devices bulk-encrypt all SDV SPTS channels. The redundant encryption devices independently create two instantiations of the SDV channel lineup. The complete SDV channel lineup is now available for use by remote hub sites. In this fully redundant system, 300 primary streams are identified by multicast group address “$\mathbf{G_P}$“, in addition to 300 secondary video streams branded by group address “$\mathbf{G_S}$“.

Throughout this data center design, processing elements compare stream quality between the primary and backup SDV copies. In many cases, redundancy mechanisms identify MPEG level faults, and perform a cutover to the backup stream to contain the fault to the data center. These techniques prevent the propagation of video faults throughout the video network, pre-empting QAM-level cutovers to backup streams in many cases.

Source Specific Multicast (SSM) Load Balancing

A key element of the enterprise data center is the ability to load balance millions of flows across hundreds and potentially thousands of servers. These complex load balancing techniques typically use a creative combination of L2 VLANs, virtual address schemes, and server pools to evenly distribute requests across thousands of servers. This design is driven in part by the commoditization of server hardware in which enterprise-class servers are being replaced with thousands of low-cost blade servers. With the assistance of distributed computing and distributed systems management, load balancing across thousands of server elements is both efficient and reliable in the traditional data center design.

By contrast, most elements within the video headend remain specialized and far from commoditization. These specialized processors require their own sophisticated session and resource management systems. These management systems string together a number of video resources to generate each properly encoded, conditioned, ad spliced, and encrypted video stream. Fortunately, for broadcast video services, the scale of today’s Video Data Center is typically much smaller than its enterprise counterpart with respect to number of elements. Given these differences, Video Data Center designs continue to employ resource managers and the multicast control plane to provide flow control and load balancing for most broadcast services.

For the SDV design, the SSM multicast address scheme provides a stable and predictable load balancing technique, which is suitable to the size and scale of the SDV service. Specific SDV streams are mapped to individual GE ports via IGMPv3 JOINs. Figure 9 illustrates how a subset of multicast group addresses are mapped to specific GE ports within the Stream Conditioning stage of the Video Data Center. In this example, 300 SDV SPTS streams are evenly distributed across 3 GE ports. The streams are divided into 3 groups of 100 streams (e.g. $\mathbf{G_{1-100}}$, $\mathbf{G_{101-200}}$, $\mathbf{G_{201-300}}$). SSM allows the same group address to be reused across different processing stages, simplifying stream management. This processing stage employs different multicast source addresses to acquire content (from **S3)** and handoff the groomed streams (from **S5**) to the next stage.

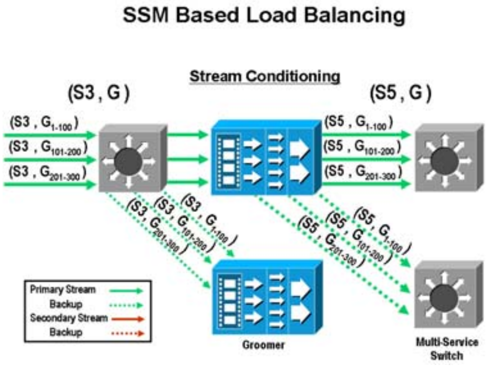

Figure 9 – SSM multicast provides a well-defined load balancing technique to distribute 300 SDV video streams across 3 GE ports.

This technique achieves a well-defined load balancing of SDV video flows. At each processing stage, multicast groups associated with video programs are joined to predefined GE ports. Each processing element is assigned a work load of multicast groups by the resource management system.

Gaining Control over Massive Scaling

The SDV system should be able to scale to a larger number of channels without impact to the current SDV deployment. Figure 9 provides the estimated bandwidth if we doubled the number of SDV channels. Based on these calculations, the total bandwidth required to deliver 600 channels of SD and HD content in the SDV tier is 4.5 Gbps.

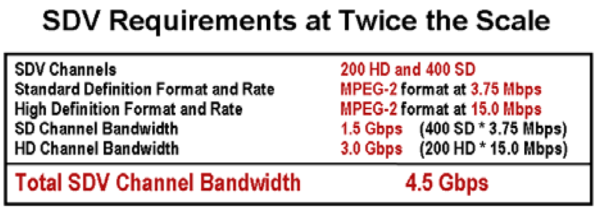

SDV Requirements at Twice the Scale

SDV Channels	200 HD and 400 SD
Standard Definition Format and Rate	MPEG-2 format at 3.75 Mbps
High Definition Format and Rate	MPEG-2 format at 15.0 Mbps
SD Channel Bandwidth	1.5 Gbps (400 SD * 3.75 Mbps)
HD Channel Bandwidth	3.0 Gbps (200 HD * 15.0 Mbps)
Total SDV Channel Bandwidth	4.5 Gbps

Figure 10 - Bandwidth required if the SDV Service was doubled in capacity to include 200 HD and 400 SD programs.

To accommodate this additional capacity, extra video resources are required by the Video Data Center. Figure 10 highlights this growth in capacity, without interfering with the currently deployed SDV service. At each stage, additional multicast group addresses are assigned to the extra SDV channels. These new SDV multicast groups are drawn through the different switch ports, groomers, ad insertion, and encryption systems, ultimately providing an efficient and scalable technique to double the SDV service. Since each video processing stage is modular and utilizes clear network demarcation points, video resources within each stage can scale independently. The multicast control plane and resource management system simply directs new SDV streams to the new processing elements.

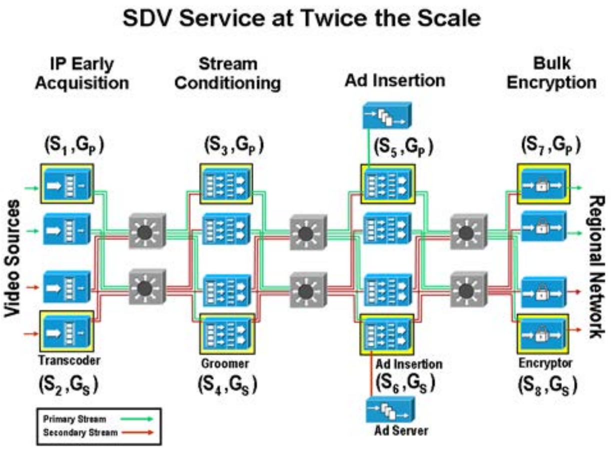

Figure 11 - Modular video resources and extra multicast groups are simply added to each processing stage to accommodate 600 SDV video channels.

4.0 IPTV INSERTION INTO THE VIDEO DATA CENTER

IPTV Insertion Strategy

The power and modularity of the Video Data Center are highlighted with the insertion of a future IPTV service. As described in Figure 12, IPTV video streams may require unique stream conditioning, ad insertion, and bulk encryption. To accomplish this unique processing, a separate set of multicast group addresses (**G_{IPTV}**) are used to draw the new IPTV streams through additional video resource elements. Similar to other cable services, a single group address for the IPTV streams will follow those streams throughout the video network. The combination of resource modularity, network based demarcations, and SSM based flow control allow this insertion technique with minimal disruption to existing video services.

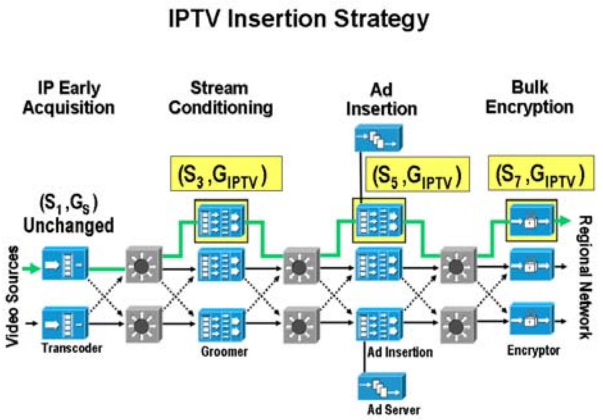

Figure 12 – The Video Data Center provides an IPTV insertion strategy with minimal disruption to already deployed services.

The IP early acquisition stage provides a common source of content for all cable services, including IPTV. As improvements to resource and element managers continue, next generation session-resource management systems can support both traditional cable services and newer IPTV systems, enabling further consolidation. This IPTV insertion strategy delivers a manageable insertion process while minimizing the impact to existing cable video services.

IP Early Acquisition

The next generation Video Data Center will acquire large amounts of video content, sourced from diverse locations in a wide range of formats. Much of this content will be shared across multiple video services. The "IP Early Acquisition" model insures that this diverse content is made available to all services in an IP encapsulated format, at the earliest point in the acquisition process.

The goal of the IP Early Acquisition stage is to convert all streams acquired from diverse sources to an IP MPEG SPTS format. Figure 13 provides a typical Headend video acquisition process in which content is acquired from satellite, off-air, and terrestrial video sources. RF video content from satellite and off-air sources is processed by a series of encoders, and groomers to convert RF, SDI, and ASI video into IP MPEG SPTS streams. Similarly, terrestrial video sources from remotes sources will be converted to an IP MPEG SPTS format.

In this example, common MPEG video streams, acquired from diverse sources are available for linear broadcast, SDV, VOD, and IPTV services.

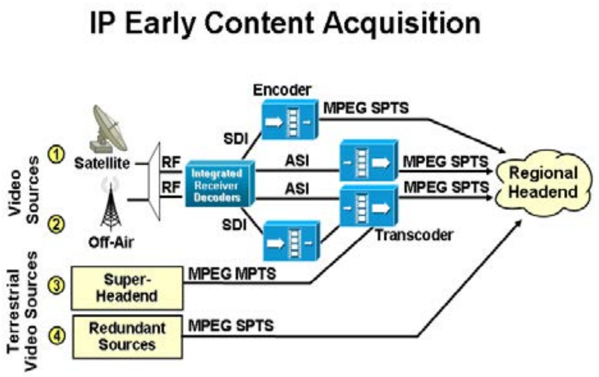

Figure 13 – IP early content acquisition collects and transcodes video content to an IP MPEG format for use by Broadcast, SDV, VOD, and IPTV services.

5.0 ON-DEMAND SERVICES USING VIDEO DATA CENTER TECHNOLOGIES

Second generation video-on-demand systems take advantage of Data Center content caching techniques and distributed video streaming to more efficiently absorb the massive scale and unpredictable loading of on-demand services. Figure 14 provides an example of an advanced on-demand network.

The intelligent IP infrastructure allows VOD content to be stored in large vault arrays at strategic locations across the network. On-demand vaults can provide a centralized and consolidated repository for large quantities of content, which are acquired from a wide range of sources to support both live and on-demand applications [5]. The ability to grow content storage independently from content streaming is a critical enhancement, and was inspired from the Web caching architectures.

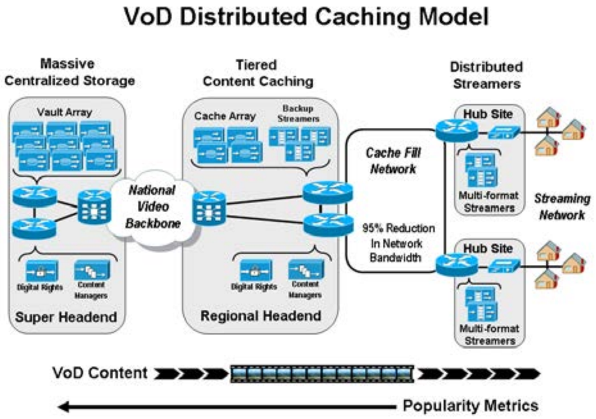

Figure 14 – The VOD content caching model is similar to Web-based cached systems in which the most popular video content is efficiently cached near the network edge, and delivered to the user via multi-format streamers.

In this distributed caching model, on-demand titles are distributed upon request using a tiered caching scheme. Cached video makes its way through the intelligent transport to video streamers located near the edge of the access network. As content streamers are moved closer to the edge, providers benefit from improved scaling, video quality, and reduced transport costs. This distributed caching model has been shown to reduce transport network bandwidth by as much as 95% since the most popular content is delivered once, cached, and reused across many requests for the same title.

Multi-format video streamers in the hub sites handle traditional MPEG video formats, but also support internet streaming formats such as Flash or Windows Media. These multi-format streaming improvements coupled with in-home gateway devices can enable a single video infrastructure to now support multiple streaming services to various in-home devices. Backup video streamers located within the regional headend can be positioned to handle unexpected peak periods, when streamers in a specific hub site are saturated with requests.

When combined with a national network, an on-demand caching scheme allows for a significant consolidation of content ingest points. Advanced caching protocols enable real-time ingest and the delivery of content throughout a national footprint within a 250 ms period. HD bit rates are also supported at a massive scale. Based on these data center design improvements, terrestrial content distribution will likely supplant the traditional on-demand "pitcher/catcher" distribution techniques of today. This networked infrastructure also provides a flexible platform to implement robust resiliency schemes, where multiple copies of content are stored and accessed from distributed locations. All of these elements work together through an intelligent network to more efficiently deliver the next massive wave of on-demand content.

Figure 15 – Advanced content caching protocols enable real-time ingest and real-time distribution of video content across a national footprint.

SUMMARY

Enterprise computing is the midst of another evolutionary transformation, as data center architectures give way to "cloud computing" systems such as Amazon Dynamo [6]. Advances in this area are driving investments in advanced and highly-scalable networking and distributed computing architectures. The expected increase in interactive video traffic requirements is placing higher requirements on cable headends to deploy similarly scalable video delivery systems that are easily adaptable to growth and performance. From this standpoint, the evaluation and application of data center technologies and practices can greatly benefit service providers' efforts to scale and grow their video delivery systems.

REFERENCES

[1] ***"Datacenter Reference Guide",*** Sun Microsystems, July 2007.

[2] "***Cisco Data Center Infrastructure 2.5 Design Guide",*** Cisco Systems, 2007.

[3] ***"Towards a Next Generation Data Center Architecture: Scalability and Commoditization",*** A. Greenberg, P. Lahiri, D. Maltz, P. Patel, and S. Sengupta. Microsoft Research, 2008.

[4] ***"The Anatomy of a Large-Scale Hypertextual Search Engine",*** L. Page and S. Brin, 1998.

[5] ***"Climbing Mount Everest",*** S. Kotay (Comcast) and J. Pickens (Cisco). NCTA Technical Papers, 2008.

[6] ***"Dynamo – Amazon's Highly Available Key-Value Store",*** G. DeCandia et. al., ACM Symposium on Operating Systems Principles, 2007.