NEXT GENERATION VIDEO INFRASTRUCTURE: MEDIA DATA CENTER ARCHITECTURE

Gene Cannella, Cisco Systems
R. Wayne Ogozaly, Cisco Systems

Abstract

Service providers seek to deploy next-generation interactive, immersive, and profitable video services that reach any device with any content, anytime, anywhere. This requires a managed evolution to a much more powerful, scalable, end-to-end video infrastructure. Here we present a Media Data Center architecture for a scalable, video-aware delivery model that offers virtualization of services and subscriber applications, employs content delivery networks and adaptive bit rate methods, unifies support for MPEG and Web protocols, employs a cloud-based network design, and leverages video intelligence throughout the network, while interworking with legacy MPEG TV systems.

THE NEED FOR A MEDIA DATA CENTER

Video service providers are seeing competitive pressure to deliver the next wave of video services to more devices, in more formats, across a national footprint, and to millions of subscribers. These next-gen video services are more interactive and involve a wide range of subscriber applications, widgets, and streaming techniques that are more closely aligned with Internet protocols than they are with traditional MPEG TV systems. In many ways these new applications are very disruptive, because they impose a new set of operational requirements and a unified architecture that can handle both MPEG and Web-based protocols. Traditional TV and Internet services are converging and the video infrastructure must also evolve to support a new set of "TV Everywhere" services and cloud-based applications.

"TV Everywhere" services deliver premium entertainment to PCs, tablets, and mobile devices. Adaptive Bit Rate (ABR) streaming technologies such as Apple HTTP Live Streaming and Microsoft Silverlight Smooth Streaming are typically used to reach these additional Internet devices. This technology disrupts existing service provider infrastructure by creating dramatically increased compute and storage demand, multiplying file counts and file management requirements with fragmented file formats, and by increasing streaming load with an inherently unicast delivery model. The process to convert live content or on-demand titles into various ABR formats is very compute and storage intensive. ABR workflows typically include a number of processing stages, including acquiring content, transrating content into multiple streaming profiles at different bit rates, encapsulating each profile into the appropriate ABR formats, applying Digital Rights Management (DRM), and publishing the formatted content to origin servers. Storage and processing requirements can multiply as service providers begin to support several ABR formats, numerous streaming profiles, and diverse DRM requirements

In many cases, service providers are required to deploy more infrastructure to handle each of these new video service requirements. With each new service, operations and systems become more complex, channel lineups and VoD loads increase, and client types proliferate. Typically, infrastructure to support new services has been overlaid on legacy facilities and managed as individual service-delivery islands. Each video service requires new dedicated resources, network capacity, and drives increased operational costs. Replacing this siloed infrastructure approach with a more unified approach is advocated to evolve the traditional headend into a more powerful delivery system, namely the "Media Data Center."

FOUNDATIONAL TENETS OF THE NEXT-GENERATION MEDIA DATA CENTER

Support for a Multiservice Environment

The Media Data Center design should support continued scaling of traditional MPEG video services (linear broadcast, switched video, VoD, IPTV, etc.), but must also now support a broader set of Internet-based streaming and user applications. The design should support delivery of these services over any access network (cable, DSL, FTTx, mobile) to multiple end device types. For even if a service provider operates only one type of access network, the provider may offer its subscribers the ability to access their content as they move into or through other providers' access networks.

IP Early Content Acquisition

The Media Data Center design should provide a common content-acquisition process that moves video streams into an IP encapsulation early in the acquisition process. These encapsulated streams can then be shared across all service modes (MPEG, ABR, progressive download, etc.) with fewer ingest points and improved resiliency.

Service-Independent Scaling

Each video service, whether it be a legacy service such as linear broadcast, switched digital broadcast, and video on demand, or, one of the emerging next generation services such as IPTV, multi-resolution streaming, or cloud enabled subscriber applications, should independently accommodate scaling of infrastructure capacity. Existing services should not be disrupted when capacity is added, new services are deployed, or upgrade and maintenance operations are performed.

Cloud-based Network Infrastructure

The Media Data Center should support network assisted video delivery, proximity, load balancing, and cloud-based service orchestration. The design should optimize resiliency and content delivery over a distributed cloud enabled network. Video resources can then be managed, scaled, and secured independently using proven data center cloud innovations.

Virtualization of Content, Services, and Resources

Data center efficiencies are derived from the virtualization of servers, storage, and unified network elements. Content is abstracted, transcoded into different formats, and stored in multiple locations, for use across many devices.

Fault Containment and Resiliency

Video control plane and data plane faults should be contained to the Media Data Center. Fault identification and repair mechanisms are inherent to the design. Redundant data center locations should be used to provide high availability across a national and regional footprint.

Consolidation of Operations and Management

The consolidation of operations and management is achieved through the use of a unified computing platform and converged network technologies. Stream visibility and MPEG quality monitoring is provided across all services at each resource.

ENTERPRISE DATA CENTER DESIGN AS A MODEL

Data center computing, storage, and networking technologies are undergoing rapid change. Virtualization practices are creating a more scalable and secure data center. Massive remote storage, migration to 10-Gigabit Ethernet for edge connectivity, and higher-performance servers are changing design standards. These data center industry innovations can now be applied foundationally to create an advanced video infrastructure.

Enterprise data center networks are built hierarchically to provide high availability and continuing scalability. Hierarchical data center design provides the operator with the flexibility to create logical topologies that use traditional Layer 2 and Layer 3 network configurations, service module insertion, and network device virtualization to create architectures that scale orthogonally at each layer within dynamic application requirements.

As is shown in Figure 1, the enterprise data center network is typically modeled as Core, Aggregation, and Access network layers, with a Services layer for both security and network services.

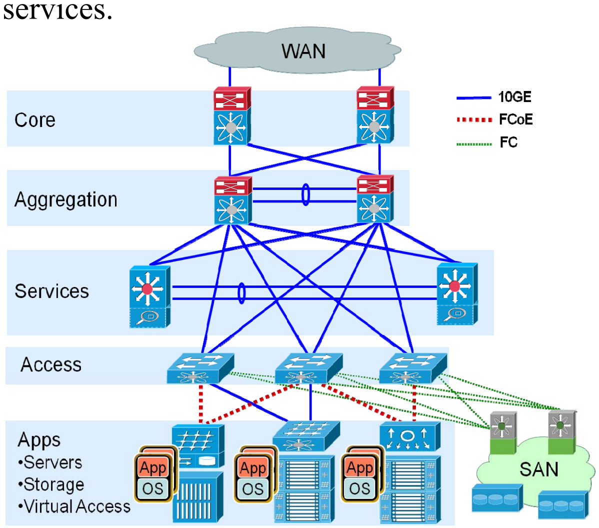

Figure 1. Enterprise Data Center Layered Network Design

This hierarchical, layered data center model employs redundant switches at each layer of the network topology for device-level failover, creating a highly available transport between end nodes over the network. The Services layer appears in this model because data center networks ordinarily require additional services beyond basic packet forwarding. These include, for example, server load balancing, firewalls, and intrusion prevention. These services can be introduced as standalone appliances or as modules that populate a slot of one of the switching nodes in the network. For each of these services an independent decision may be made whether to deploy redundant hardware based on the availability requirements associated with that service.

Experience with enterprise data center operations has shown the "Core, Aggregation, Services, and Access" layered model to be an enabling paradigm for scalability, performance, flexibility, resiliency, and maintenance. Here we will adapt and extend this model to next generation service provider video infrastructure, focusing on the Media Data Center while recognizing many of the architectural principles are equally applicable to other service provider facilities, such as next generation Hubs or Pod deployments.

ADAPTING ENTERPRISE DATA CENTER INNOVATIONS TO THE NEXT GENERATION SERVICE PROVIDER MEDIA DATA CENTER

In Figure 2, video resources are folded into the hierarchical data center model, providing functionality for acquisition and other traditional MPEG workflows as well as next generation IP-based subscriber apps and content distribution. Subsequently, we will develop the ideas of virtualization of infrastructure at each layer and the unification of compute and storage networking into a logical representation of this physical topology model.

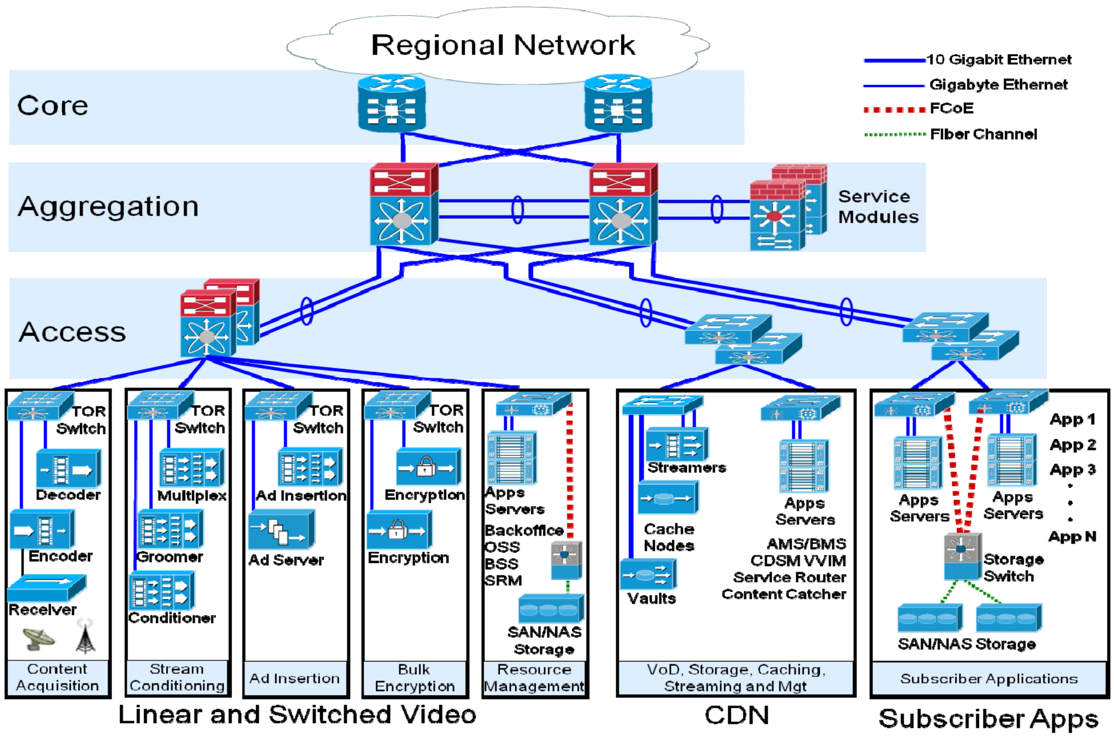

Figure 2. Media Data Center Components

Hierarchical Networking

As in the enterprise data center, the Media Data Center is implemented via hierarchical network layers. The design provides distinct layers for Core, Aggregation, Network & Security Services, and Access. With this layered approach, the Media Data Center eases the operational challenges of managing resource access, security, and scaling across a diverse set of video elements.

The Core network layer provides a Layer 3 connection to the national or regional transport network.

The Aggregation network layer performs load balancing and aggregation of video flows and management data across various service contexts.

The Network and Security Services Layer implements firewalls, server load balancers, intrusion prevention systems, application-based firewalls, and network analysis modules.

The Access Layer (a switching layer in the Media Data Center not to be confused here with the service provider access network or outside plant) provides efficient access to a range of modular video resources.

Compute, Network, and Storage Unification and Virtualization

Emerging best practices in the data center will unite compute infrastructure, storage access, virtualization, and the networking of both compute elements and storage in cohesive systems that simplify operations and reduce total cost of ownership.

Flexible and advanced x86 architecture server technology integrated in blade form factor enables the running of compute infrastructure as virtual machines. Server virtualization increases the total utilization factor for the compute infrastructure, enables application mobility, and increases service availability, all while generating power savings and total cost reduction. Virtualization becomes the new data center best practice with the emergence of hypervisor applications such as VMware's ESX/ESXi Server, Microsoft's Hyper-V, and the open source hypervisor Xen, leveraging the availability of powerful, multi-core x86 CPUs and virtualization technology from both AMD (AMD Virtualization) and Intel (Intel Virtualization Technology).

Stateless computing is used to enable seamless server mobility. Service profiles are used to abstract server attributes and decouple them from physical hardware attributes. Applications will run on virtual machines that are instantiated in response to dynamic conditions. Applications and processes can rapidly and easily migrate between hardware instantiations. This will allow server capacity planning that leverages the statistical characteristics of both application and service load and improves overall utilization accordingly.

As mentioned briefly above, virtualization of servers is achieved using virtualization applications known as hypervisors. These are virtual machine monitors that enable multiple operating systems to run concurrently on a host computer. Native hypervisors are software systems that run directly on the host's hardware as a guest operating system monitor, and provide the ability to scale virtual machines to the degree required in the Media Data Center. Multiple virtual machines can be created to run in isolation, side by side on the same physical server. As shown in Figure 3, each virtual machine has its own set of virtual hardware (RAM, CPU, NIC) upon which an operating system and applications are loaded. The operating system when loaded sees a consistent and normalized set of hardware without regard to the actual physical hardware components upon which it is instantiated. This ability to run multiple virtual machines abstracted on a single set of hardware sees it greatest benefit in applications with high peak to average CPU load, or simply low average load. Examples of these in the Media Data Center might include management or control applications, or video on demand applications, but likely extend to a surprising number of other applications.

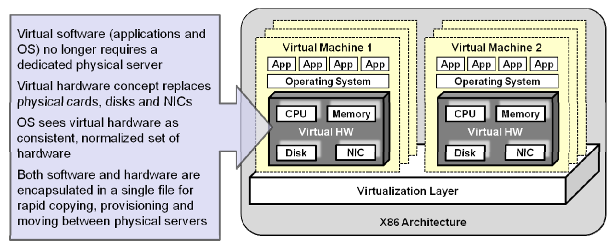

Figure 3. Virtual Machine Concept

Yet another feature of the virtual machine approach is the ability to encapsulate virtual machines into files for rapid saving, copying, and provisioning. Full systems, including operating systems, BIOS, virtual hardware, and fully configured applications can be moved within seconds from one physical server to another in response to changing service loads or for zero down time maintenance. This "*service orchestration*" feature gives the server infrastructure the ability to dynamically apply additional resources to new service loads, yielding operational and utilization benefits across a wide range of use cases, but offering benefits of particular interest in applications with high average CPU load but variable service load. In the Media Data Center examples of such applications might include video processing operations for ingest, live or on-the-fly transcoding, time-shifted streaming, or many others. Together, the two capabilities of 1) running multiple virtual machines on the same physical hardware, and 2) dynamically moving file-encapsulated virtual systems onto new hardware, give the virtualized server infrastructure the simultaneous ability to dynamically scale in *two* dimensions in response to both dynamic CPU loads and dynamic service loads. Improvements in utilization factor can be dramatic. In a use case deploying a home security monitoring application, server hardware requirements were reduced by a factor of 4:1. In another use case, a hardware reduction factor of 8:1 was seen for a session & resource management application such as is commonly deployed by service providers to manage QAM sessions for VoD, switched broadcast, or data services.

In the next generation Media Data Center design, the formerly separate LAN, SAN, and high performance computing networks are consolidated onto a single network and unified fabric which uses 10G Ethernet as the fundamental link interconnect. Legacy data center storage network designs typically use Fibre Channel interfaces for storage and associated switching. These Storage Area Networks (SANs) are built and operated independently of the data center Ethernet switching networks, creating a duplication of cabling, switching, and other resources. The Fibre Channel over Ethernet encapsulation, as shown in Figure 4, below, is a standards based approach that converges these two networks onto a single more cost effective Ethernet infrastructure.

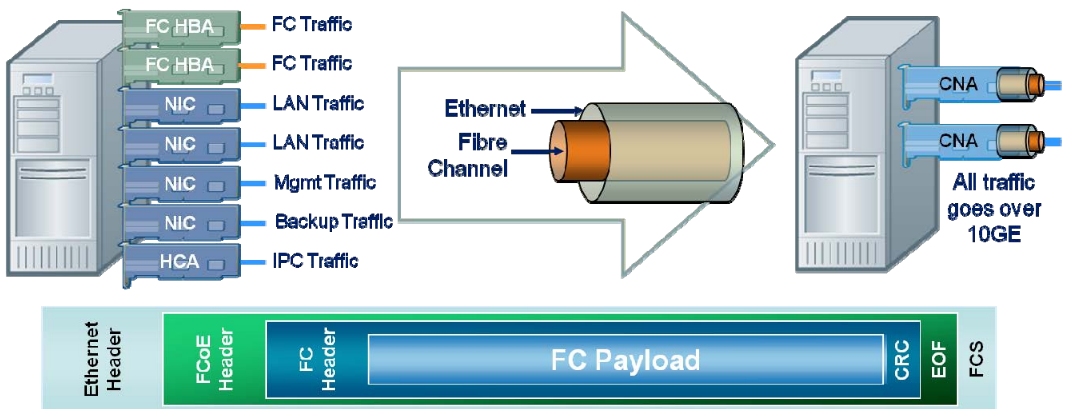

Figure 4. Consolidating LAN/SAN Networks using Fibre Channel over Ethernet

Devices that formerly attached to Fibre Channel networks using Host Bus Adapters (HBAs) will now attach to Ethernet networks via Converged Network Adapters, which integrate the functionality of both the HBA and an Ethernet Network Interface Card (NIC). These Converged Network Adapters, or CNAs, perform the Fibre Channel over Ethernet encapsulation and allow storage systems to use native storage protocols over the unified fabric at lower cost than traditional SAN switching. Figure 4 depicts a single pair of 10 GigE CNA ports replacing the multiple adapter types and multiple links of the legacy configuration. This scheme provides a fully redundant transport that supports both Fibre Channel SAN and Ethernet LAN topologies over a single physical link,

The unified fabric approach eases requirements for cabling, switching, power, cooling, and network adapters, with particular savings derived as a result of the migration from 1 Gbps Ethernet and Fibre Channel cabling to 10 Gbps Ethernet connections. In addition, the unified fabric provides a simplification of the switching network architecture. Compute servers will use a FCoE connection to the access layer that supports both LAN and SAN interfaces. This allows a single pair of switches that support FCoE with 10G Ethernet interfaces to replace four switches (two LAN and two SAN) that were present in the legacy topology. As a result, a complete layer of Fibre Channel switches and cabling can be removed from the topology as is illustrated in Figure 5.

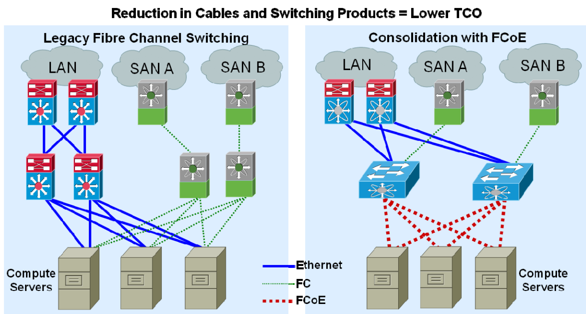

Figure 5. Unified Fabric Switching Topology

Another dimension of virtualization available in the next generation Media Data Center is virtual access switching. In a dynamic environment, individual and total application requirements for compute and storage will scale up and down based on subscriber behavior. Network access for the responsible virtual machines will need to scale correspondingly. Providing a virtual access switching layer is a powerful method to virtualize the interfaces for virtual machines and unified storage networks.

A Virtual Ethernet Module is installed into each server hypervisor kernel, while a distributed virtual switch is created by installing a Virtual (Switch) Supervisor Module on an additional virtual machine. Network LAN policies are implemented at this layer in port profiles, much as server profiles were created for the stateless computing mode of operation of the virtual machines. A typical port profile identifies, for example, MAC, World Wide Name (the Fibre Channel unique element ID), Boot Order, and Firmware, as well as network (LAN) and storage (SAN) policies. In this way operators may dynamically and easily apply policy-based configuration and operation of network services for each virtual machine; they can also easily manage virtual machine connectivity across physical servers to balance server workloads for optimized application performance or to provide improved availability during routine hardware maintenance.

Each virtual machine connects to its own Virtual Ethernet port on the access switch. This creates a key capability; the network administrator is given traffic visibility and policy control on a *per virtual machine* basis. Virtual machines can be managed as though they were physical servers from the perspective of network connectivity. If applications need to be moved or to grow, port profiles will follow the application to the new resources. Operators will simply modify service profiles to add or subtract capacity for video applications, eliminating multiple complex configuration operations and the need to separately touch the configurations of servers, adapters and interfaces, LAN/SAN switches, and storage devices.

The Media Data Center design should also include embedded management systems that enable management of the unified compute infrastructure (servers, storage, network, virtual machines) as a single entity. To fully leverage the virtualized and unified environment of the new Media Data Center, management GUIs, CLI interfaces, and robust programmatic interfaces will allow the definition of service profiles that logically encapsulate desired physical configurations and provisioning of unified resources.

<u>Device Virtualization and Virtual Device Contexts</u>

Increasingly, routers are offering advanced features allowing the virtualization of network device features. The possible dimensions of device virtualization include: forwarding plane, control plane, management plane, and the partitioning of software and hardware components into those planes. The Media Data Center architecture here described leverages each of these dimensions of virtualization in the aggregation switch layer, creating partitions that we will subsequently call “Virtual Device Contexts.” In this way, aggregation switches become virtual devices that present as independent and unique logical entities to users connected within each Virtual Device Context. Each Virtual Device Context will maintain its own set of software processes, have its own configuration, and be administered via its own management context.

The Media Data Center uses Virtual Device Contexts to partition families of services and applications *according to the operational practices of the service provider.* An example partitioning might delineate the resources associated with legacy linear, switched digital broadcast, and IPTV in one Virtual Device Context from the resources associated with CDN-based services such as VoD and Internet streaming in a second context, while creating yet a third context for management applications and advanced subscriber interactive applications.

USE CASE: INSERTING THE “TV EVERYWHERE” MEDIA POD

This section describes an insertion strategy for a "TV Everywhere" Media Pod, or design module. This business use case describes how the Media Data Center architecture can accommodate the insertion of new features such as adaptive bit rate video streaming. The following description provides a vision of how new services can be integrated smoothly into the overall architecture.

Common Cloud Infrastructure

Service providers have embraced "TV Everywhere" as a way to bring TV-quality video programming to devices beyond the home television. The "TV Everywhere" initiative develops the infrastructure and applications to deliver premium live and on-demand content to many IP devices, including the PC, tablet, and mobile devices. Service providers are in various stages of system deployment and system trials.

With the Media Data Center and the Media Pod modular design, service providers can implement a distributed cloud based infrastructure that supports “TV Everywhere” services. As shown in Figure 6, this infrastructure includes a combination of National Media Data Centers, Regional Media Pods, and a distributed Content Delivery Network to support streaming video services. The “common-cloud” infrastructure utilizes data center infrastructure that has proven performance and known scaling for the relevant video workflows. These workflows can be replicated across regional Media Pods and National Media Data Centers. Given this approach, service provider operations teams can develop a level of certainty in their infrastructures, increase service velocity, and focus on the deployment of video applications and subscriber services.

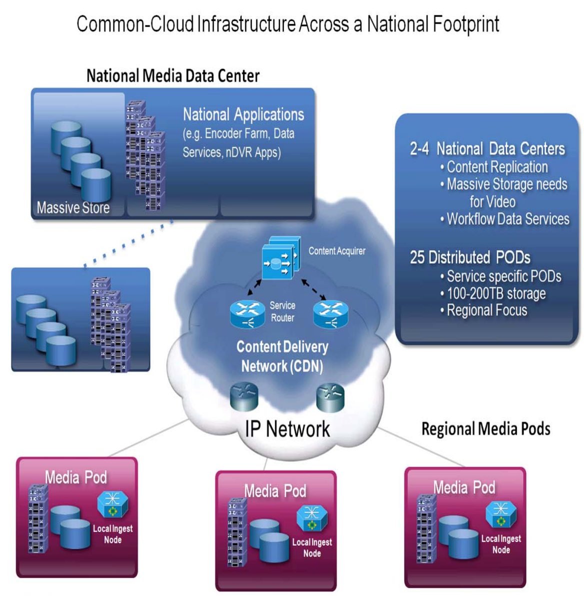

Figure 6. Common Cloud Infrastructure

Adaptive Bit Rate Workflows

Adaptive Bit Rate video delivery plays a critical role, because video content is likely to traverse an unmanaged or congested network to reach devices in the home or on the road.

ABR delivery models can adapt the streaming rate in response to real-time network conditions, in order to maintain the video experience. The data center infrastructure needed to acquire, create, manage, and deliver ABR video streams includes a mix of traditional video components along with advanced data center products. ABR workflows contain multiple processing stages in which video is acquired, transrated, encapsulated, packaged in a DRM, and stored in an origin server. These stages must support linear and on-demand titles, as well as various ABR formats and bitrate profiles. Sample workflows for an adaptive bit rate Linear, On-Demand, and Time-shift services are provided in Figure 7 below. As video content passes through each stage of the workflow, a combination of compute, network, and storage resources will be consumed. Data center infrastructure elements must support these dynamic processing requirements and scale to vast content sources and wide ranging delivery formats.

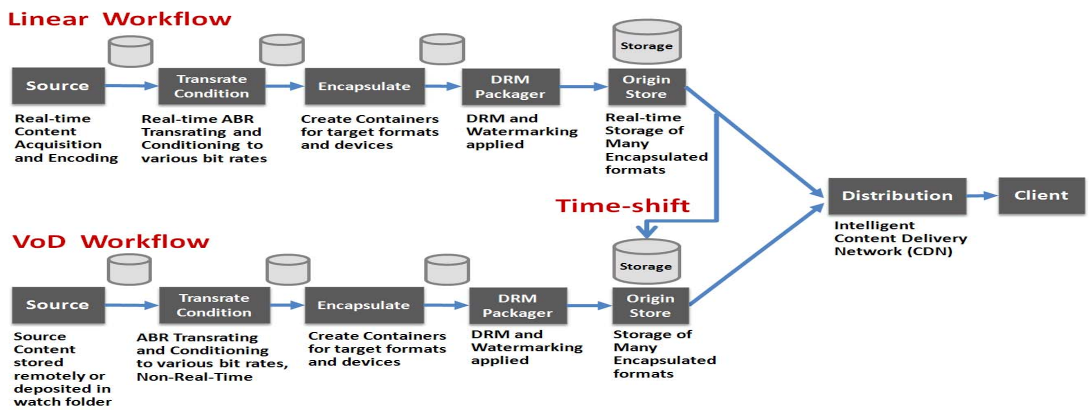

Figure 7. Workflows for Adaptive Bit Rate Services

Media Pod Components and Applications

The Media Pod includes the processing elements and infrastructure required to manage and produce ABR content. A number of Media Pod components and applications required to support a "TV Everywhere" service are shown in Figure 8. These ABR applications include content transcoders, packagers, DRM and watermarking applications, as well as a wide range of content management and publishing tools.

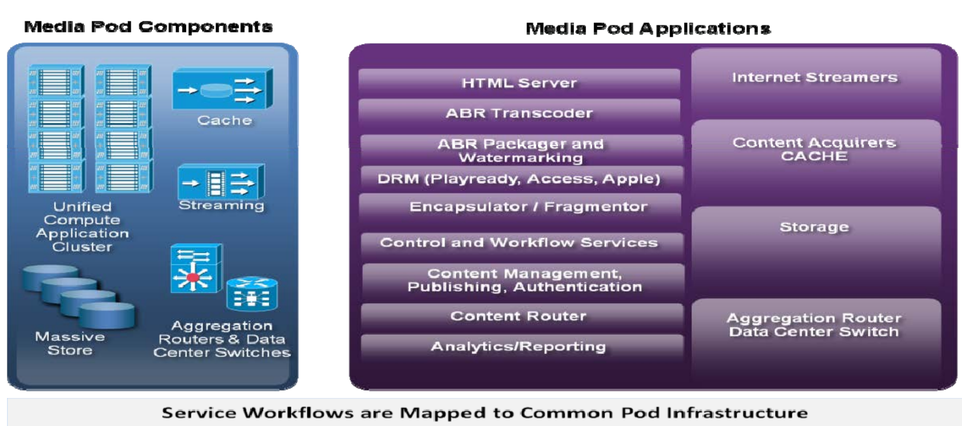

Figure 8. Media Pod Applications and Components

Service providers can map video data plane workflows and content management functions to the Media Pod infrastructure. The Media Pod is designed to provide a scalable compute, network, and storage architecture required by next-gen video services and cloud-based applications. The Media Pod is depicted in Figure 9 below and utilizes the Media Data Center architectural principles and innovations to significantly reduce operational complexity and reduce total cost of ownership.

Combined SAN / NAS storage devices supporting virtualized storage with block and NFS storage technologies are used in the Media Pod. Unified storage scales easily to support massive centralized storage or smaller scale regional storage.

The Media Pod incorporates unified storage switches delivering 10GE based FCoE technology. These switches enable 10GE FCoE network connections, saving significant cabling costs and achieve a consolidation of switching devices.

The Media Pod employs a powerful x86 based compute environment that can implement a "bare metal," or a virtualized compute environment. Virtual compute servers provide more efficient use of compute resources, and deliver superior operational management. Video workflow applications, content management, and subscriber applications reside on this compute environment.

Access switching in the Media Pod uses virtual access-switch technology that supports cloud based operations, application movement across different compute resources, and profile based configuration of network and storage requirements.

Media Pod supports Next-Gen Video Services and Cloud-Based Applications

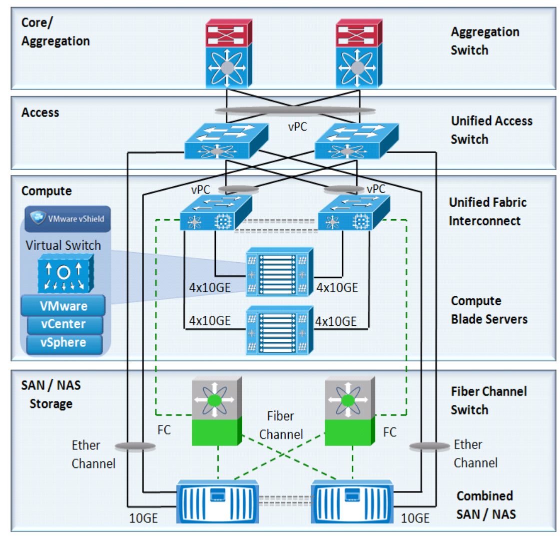

Figure 9. Media Pod Architecture

Video applications are mapped to virtual machines and unified storage elements in the Media Pod. Different ABR workflow stages (transrate, fragment, DRM, watermark) are overlaid onto blade servers and storage elements within the Pod architecture. Application profiles that define the operating system, compute requirements, network interfacing and QoS policies, and storage requirements are created for each workflow application. These video application profiles are overlaid onto Media Pod resources as described in Figure 10 below.

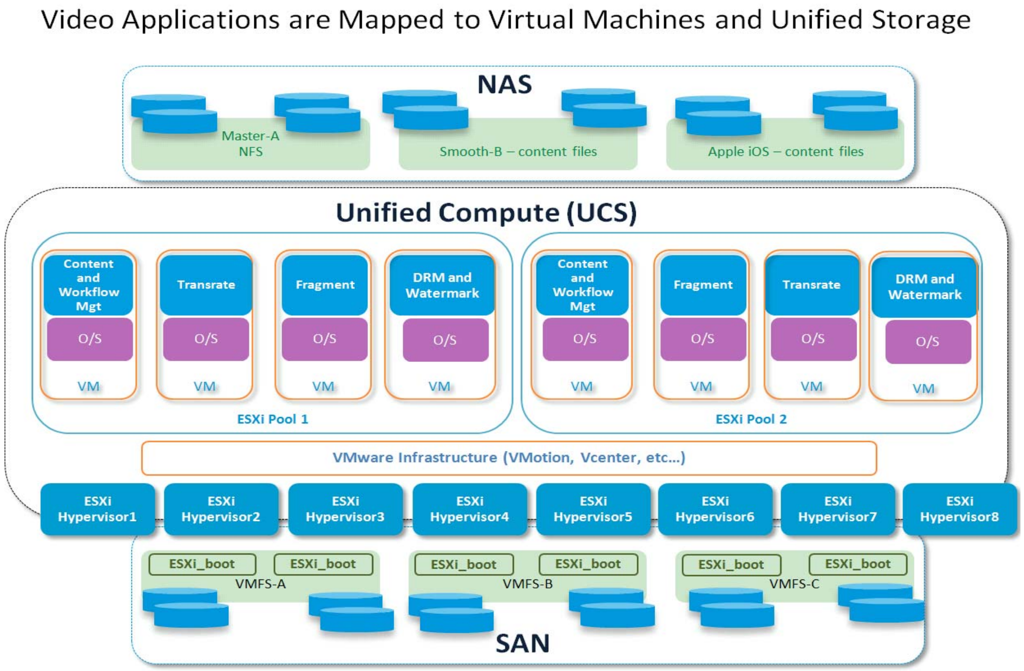

Figure 10. Video Applications are mapped to Virtual Machines and Unified Storage

<u>Media Pod Insertion</u>

The Media Data Center has been architected to allow the smooth insertion of design modules such as the Media Pod. Figure 11 depicts how the Media Pod can be inserted into the architecture to manage and deliver ABR content for "TV Everywhere" applications. It is likely that the Media Pod and other data center design modules will share some certain processing workflows such as the live content acquisition workflow. The content acquisition workflow is designed to support this shared service environment. In the example below, both traditional TV and the ABR Media Pod share the content acquisition workflow. In addition, the CDN system will be updated to include ABR Internet streaming components. Because the Media Data Center has been designed to leverage enterprise data center principles and innovations in accordance with the foundational tenets discussed earlier, the Media Pod and Internet CDN components insert easily into the architecture.

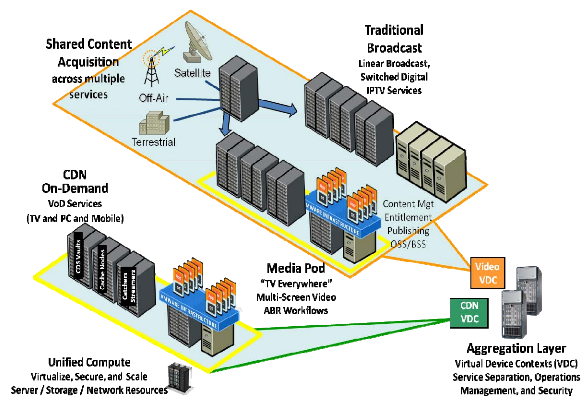

Figure 11. Media Pod Insertion

Figure 12 depicts the infrastructure for the new TV Everywhere ABR service overlaid on the Media Data Center logical topology. The Media Pod folds into the design as a secure partition, and can efficiently use new computing, storage, and network resources. Similarly, the CDN partition can accommodate new storage, cache, and streaming components as needed for the Internet streaming function.

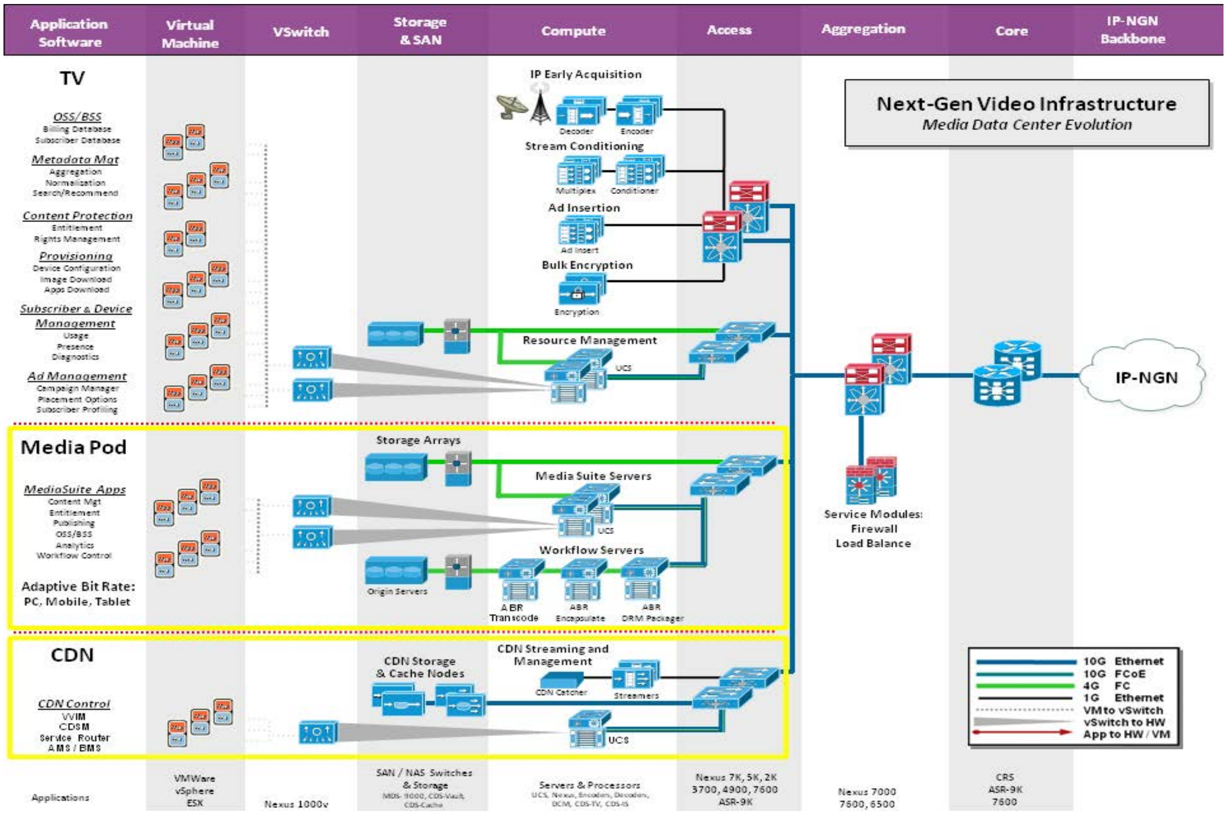

Figure 12. Media Pod Partition

CONCLUSION

An architecture for a Media Data Center is developed that leverages enterprise data center principles and innovations for a scalable video infrastructure. Virtualization is leveraged in the compute domain as well as in all layers of network connectivity and administration. Storage, compute and network resources are unified to achieve powerful operational benefits and cost reductions. Finally, a use case is demonstrated, with the implementation of a Media Pod design module that provides scalable infrastructure for adaptive bit rate streaming and the emerging “TV Everywhere” service.