



Evolving Requirements for On Demand Networks

Motorola's On Demand Video Solutions



Introduction

Video on demand (VOD) was originally introduced by the cable industry in the late 1990s to differentiate its video product offering from that of the direct-to-home (DTH) satellite competitors. In today's competitive environment for video services, the role of the on demand platform has only grown in importance. Not only has VOD become a "must have" service for cable, IPTV, and even DTH providers, but the on demand infrastructure has moved beyond VOD to serve as the launching pad for a range of critical advanced services. These include Time-Shifted TV, Network-based Digital Video Recording (nDVR) and its close cousin, Remote Storage DVR (RS-DVR), and On Demand Ad Insertion.

Weekly VOD users are twice as likely to be "extremely satisfied" with their digital cable service.

Parks Associates, 2008

The growth of VOD and the advent of new services have dramatically changed the requirements for an on demand solution. Today, a flexible, scalable, and reliable on demand network is required for:

- Maintaining high levels of customer satisfaction with existing VOD services
- Reducing churn
- Generating revenue from new services and premium content

At the same time, in order to maximize the profitability of these services, operators must optimize the resources that make up their on demand network, most notably streaming, storage and network resources.

Achieving these goals is no simple task however, as the success of VOD and related services has increased the demands on the platform and added complexity to the solution. While there are many factors, there are three primary trends that have changed the requirements for an on demand network.

Major Trends in On Demand

Streaming & Ingest Growth

In the past several years, on demand usage has grown considerably. In the United States the number of on demand views grew 50% between 2006 and 2007¹. More importantly, as VOD usage grows, so do concurrency rates, or the percentage of homes using VOD at any one time, increases as well. This is the key driver in determining the amount of streaming resources required to service the subscriber base.

Realtime services such as Time-Shifted TV further increase concurrency rates by converting traditional linear TV viewing into on demand streams. Adding to the challenge is the rapid increase in high definition (HD) content, which in most of today's cable networks represents the equivalent of approximately four standard definition (SD) streams.

HD and realtime content also increase the need for resources to ingest content into the on demand platform. In the case of a Time-Shifted TV deployment, potentially hundreds of SD channels and dozens of HD channels may need to be ingested into the on demand network and turned around in seconds so they can be offered to subscribers as on demand programs.

Library Growth

In parallel with the growth in streaming requirements, the amount of content offered is expanding dramatically. On demand libraries have more than tripled over the past several years, and are projected to grow by more than nine times by 2011².

The growth in library sizes has been driven by:

- Greater willingness of content owners to distribute content in the on demand tier
- Shorter theater to VOD windows, including some "day and date" efforts where movies premier on VOD simultaneously with, or even before, the theatrical premier
- Growth in real-time content from services such as Time-Shifted TV
- HD content, which today requires approximately 4 times the storage of SD content
- An increase in local on demand content

On Demand share of total viewing will grow from 9% to 38% by 2012.

Pike & Fisher

¹ Rentrak, 2008

² Motorola data and estimates.

Flattening Concurrency Curves

Concurrency can also refer to the number of viewers simultaneously viewing a particular piece of content. As the concurrency curve flattens, the equivalent number of subscribers accesses a larger number of unique titles in the library.

For example, in the early days of VOD, the “80/20 rule” held that, due to the popularity of certain content, 80 percent of on demand streams could be served with the top 20 percent of the content.

Today, based on detailed analysis of more than 160 deployments worldwide, Motorola estimates that in a 24-hour period the same 80 percent of streams actually access 50 percent or more of the content viewed in that period.

As we will see, this has dramatic implications for the design and optimization of an on demand network.

The Evolution of the On Demand Network

Phase 1: Distributed, Autonomous Networks

Early deployments of VOD services provided subscribers with access to a small content library with relatively low concurrency rates.

The infrastructure used to support these early deployments included proprietary back office software systems for asset management, resource management and billing, as well as proprietary interfaces to disk-based computer hardware systems used for content ingest, storage and streaming functions.

At the time, many cable operators enabled early VOD services by utilizing a highly distributed architecture with VOD servers located at the majority of network hubs. Due to their small size, the entire library was stored at each VOD server. An example of this architecture is displayed in Figure 2.

While adequate for the requirements of the time, as on demand usage and libraries grew, these systems faced significant challenges:

- Operational complexity, including the management of several copies of large content libraries across distributed server complexes
- Reliability, including a large number of disk failures in multiple remote locations
- Scalability
- Inability to support real-time services

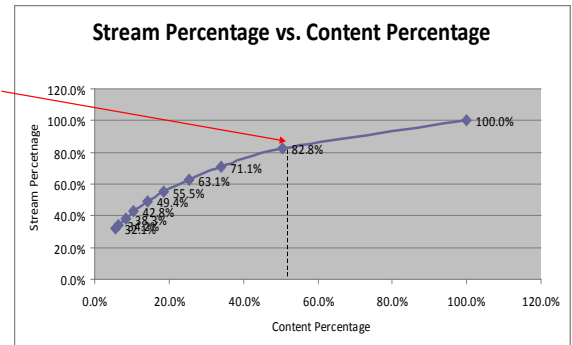


Figure 1. Today, the 80% of streams that access the most popular content reach more than 50% of the on demand content.




Figure 2. An early VOD deployment with highly distributed disk-based servers, each with the entire library.

Phase 2. Consolidated, Autonomous Networks

To better address these challenges, on demand deployments moved to a model in which streaming and storage libraries were aggregated in more central locations, as shown in Figure 3.

Two important developments were necessary for the shift to this model:

1. The large-scale buildout of high-bandwidth transport networks, which allowed for dynamic data capacity and the centralization of on demand streaming servers.
2. The emergence of high-capacity, extremely reliable solid-state video servers with decoupled streaming and storage resources.

The combination of these two factors led to an architecture in which:

- Libraries could be managed more easily with less duplication of content
- The on demand platform became much more reliable, scalable and cost-effective
- Storage and streaming resources could be scaled independently
- High-performance solid-state servers could enable real-time services




Figure 3. A Consolidated model for VOD, where streaming and content libraries are more centralized for better performance, reliability and scalability.

Phase 3. Hybrid Model

As successful as the model in Phase 2 has been, the challenges facing operators with regard to on demand networks is driving further evolution. In particular, recent massive growth in content libraries has made it financially untenable to continue to maintain the entire library at each streaming server, even in this more consolidated model.

Increasing network connectivity between streaming sites, systems and even regions makes it possible to further centralize content libraries while maintaining flexibility in the placement of streaming resources. Thus, while in Phase 2 streaming and storage were separated in terms of the hardware dedicated to each, today they can also be separated geographically. This entails moving the master content library to a central location that feeds multiple local streaming servers, and positioning a subset of that library as an edge cache in an attempt to handle the bulk of streaming requests locally. Streaming servers can either be centralized or distributed based on the constraints of the network, subscriber density and other factors.

In this model, the delivery platforms are separate from the systems that distribute the content, but are interrelated as can be seen in figure 4.

This model builds on the advantages of Phase 2 by:

- Reducing capital expenditures for storage by minimizing the amount of content duplicated at each streaming server
- Increasing the number of choices available to subscribers, including virtually unlimited HD, time-shifted television and real-time content
- Leveraging centralized solid-state servers for reliability, scalability and performance

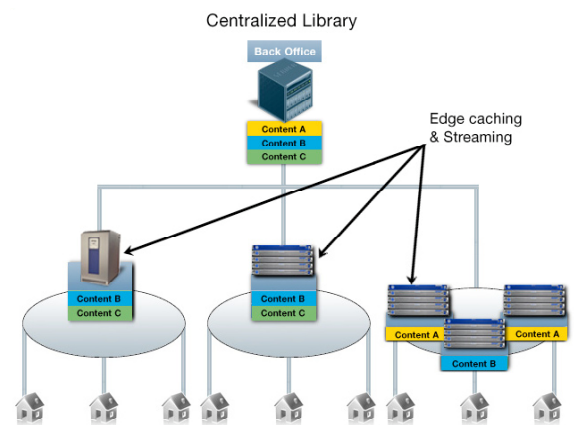


Figure 4. A hybrid model where storage is centralized in a regional library and streaming resources can distributed as required by network and subscriber constraints. A subset of the library is cached with each video server.

Challenges with the Hybrid model

As discussed earlier, increasing library size has led to a flattening of the concurrency curve, meaning subscribers are accessing a larger percentage of the library. Looking at any instant in time, the number of assets being viewed may be relatively small. This has enabled the successful use of solid-state servers that separate streaming and storage resources at the same location. But over the course of a 24 hour period or longer, the number of assets being requested grows dramatically. In the hybrid model, where the storage and streaming resources are separated by a network with finite bandwidth, this fact creates significant challenges.

Given that not all content is located at the edge cache in the *hybrid model*, the objective is to position content **at the edge**, and therefore serve as many requests as possible. Because of flattening concurrency curves, this is no simple task. Moreover, most solutions utilize simple algorithms for determining which content to send from the central library to the edge cache—and when to send it.

The example in Figure 5, which is based on Motorola analysis of actual field data, shows the impact on the network of moving to a hybrid model when using simple algorithms.

As concurrency curves flatten resource consumption shifts to the network. In this fairly typical scenario, 3.6TB of content would be sent to this server per day. In other words, the equivalent of the entire library is sent through the network more than twice every 24 hours.

The Motorola Solution

Based on the data gathered from these live systems around the world, the analysis suggests that the requirements for an effective on demand network include:

- Solid-state servers that provide the performance, scalability and reliability required for growing concurrency rates, HD content and real-time services.
- A broad range of platforms that can support multiple network architectures, scale from low stream counts to very high stream densities, and support very large regional or even national libraries
- Intelligent placement and distribution of content and streaming resources in order to optimize streaming, storage and network resources

Adaptive Media Management

One of the keys to the success of the solid-state server has been the ability to intelligently manage the movement of content between the storage library and the streaming server over a network connection. This architecture is an ideal proxy for analyzing the performance of servers in the hybrid model, which follow the same general configuration.

Over the last 5 years, Motorola has leveraged the experience from 160 deployments encompassing 900,000 streams worldwide to model this network topology. The company has conducted in-depth analysis of concurrency rates and content usage patterns and is able to simulate how they are impacted by such changes as increasing library sizes and the introduction of real-time services.

The result of the research is a highly sophisticated set of algorithms that can be shown to reduce the impact of content movement on the network by up to 85 percent.

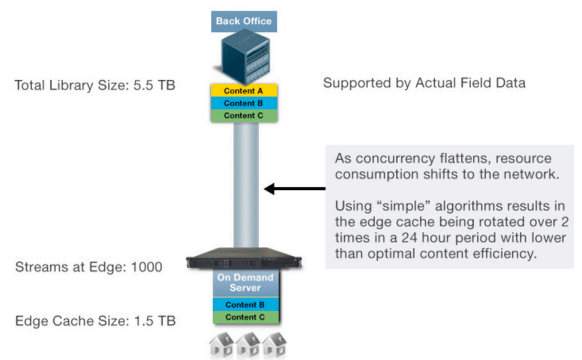


Figure 5. In this example, the use of simple content placement algorithms results in the equivalent of the entire library being sent over the network every 24 hours.

Based on the approaches used to move content to and from the solid-state cache in the B-1 Video Server, these algorithms are used to:

- Intelligently place content in a cluster of on demand servers
- Decide which servers in a cluster are optimal for responding to a stream request
- More efficiently move content to edge servers when content must be propagated

These technologies and methods are referred to as Adaptive Media Management, and are used throughout the Motorola on demand portfolio to optimize streaming, storage and network bandwidth.

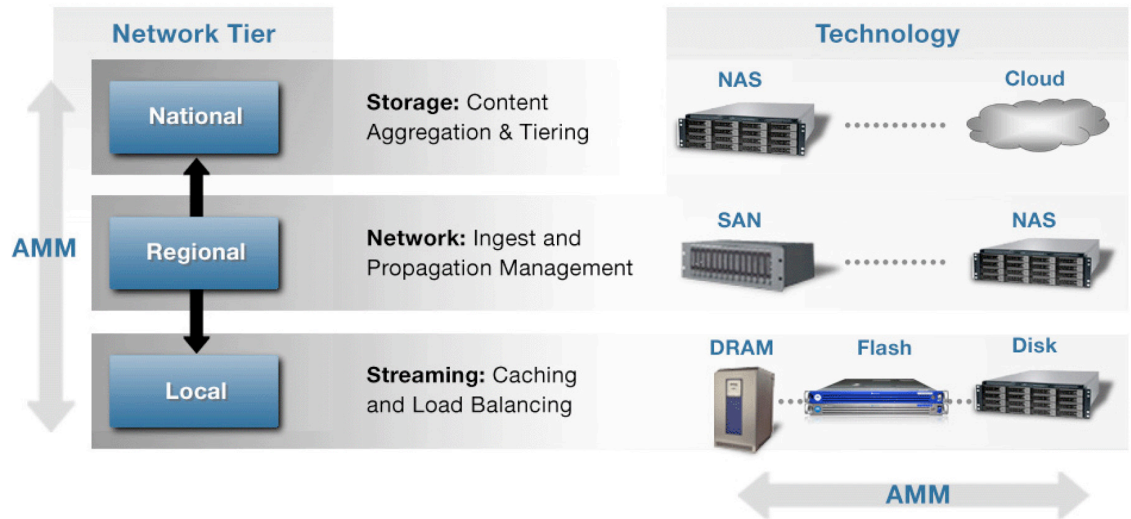


Figure 6. Motorola's Adaptive Media Management

Delivery Products

At the foundation of the on demand network are the servers that reliably ingest and deliver enormous amounts of high-quality video to subscribers. Motorola offers the broadest, most flexible lineup of solid-state servers in the industry, designed to scale from very low stream counts to high-density streaming and ingest applications with high performance, reliability and scalability.

The Motorola B-1™ Video Server

The world's most widely deployed solid-state on demand platform, the Motorola B-1 powers more than 900,000 streams in 160 deployments with leading operators.

Targeting medium to very large streaming sites, the Motorola B-1 Video Server is a highly scalable platform that provides high-performance streaming of MPEG-2 and MPEG-4 (H.264) content directly from Dynamic Random Access Memory (DRAM). This carrier-class, purpose-built hardware and software platform was designed to deliver advanced video services, including VOD, Time-Shifted Television, On-Demand Ad Insertion and nDVR.

The unique design of the B-1 offers video service providers a high-performance, reliable and scalable platform that also optimizes streaming resources and reduces total cost of ownership (TCO).




Figure 7. The solid state B-1 video server with decoupled streaming and storage.

Motorola B-3 Video Server

Designed for small- to medium-sized deployments, the B-3 Video Server melds open, standards-based software with industry-standard hardware to create a powerful on demand platform. The result is a high-performance, highly scalable, fault-tolerant server that delivers premium support for VOD, Time-Shifted TV, On Demand Ad Insertion, nDVR, and other advanced services.



Figure 8. The solid state B-3 video server targeting small- to medium-sized deployments.

The key to the B-3 Video Server is Motorola Video Grid Technology. Leveraging the grid computing paradigm, Video Grid Technology stripes video content across all blades in the server cluster, treating the cluster as one scalable, fault-tolerant server. The result is a scalable, high-performance, self-healing platform for video delivery, distribution and processing.

The B-3 offers flexible configuration options, and can be deployed as a solid-state edge cache served by a central content library, or as a standalone video server with integrated library.

Distribution

CPS1000 Cluster Manager

The CPS1000 Cluster Manager (CPS1000-CM) is the first component of Motorola's Content Propagation System. It enables a hybrid, on demand network architecture by supporting a central library and the intelligent, real-time distribution of content to edge servers. CPS1000-CM provides cluster management between the on demand servers and the operator's back office system, managing the distribution of content and stream allocation. The CPS1000-CM utilizes Adaptive Media Management and sophisticated content management algorithms to optimize storage, streaming and network resources and improve fault resiliency for a cluster of co-located or geographically-distributed cluster elements.

Integrated Solution

Together, these products provide a complete, fully integrated system for the creation of a flexible, reliable on demand network. Each component is designed to allow video service providers to adapt the architecture to their particular needs. Figure 9 shows how the components fit together to deliver a hybrid on demand architecture.

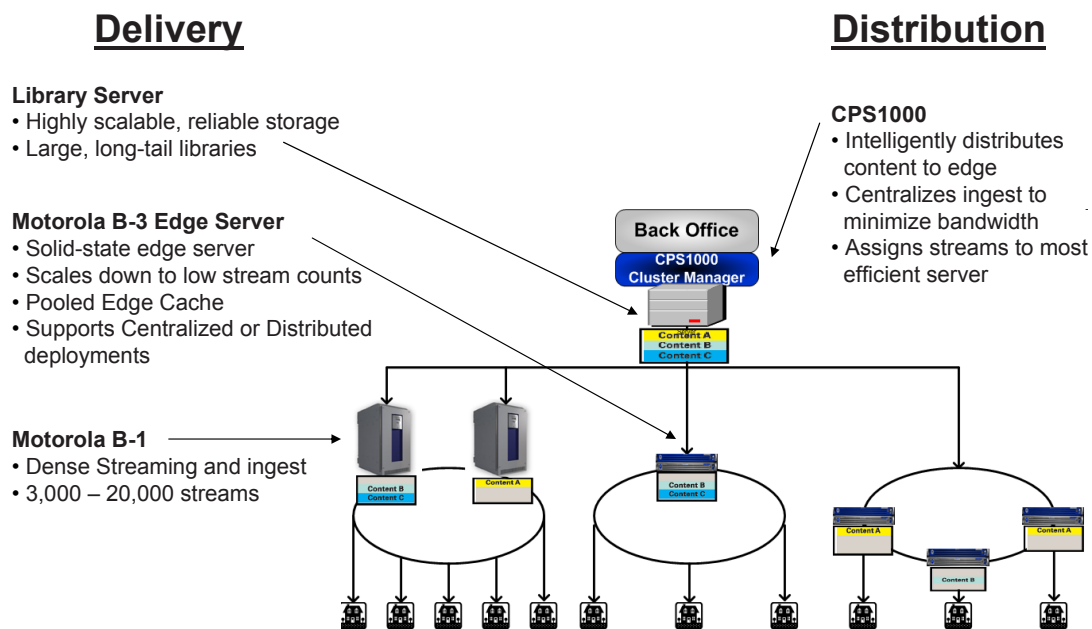


Figure 9. A hybrid on demand network as enabled by Motorola's delivery and distribution products.

Conclusion

The requirements of an on demand network have continued to evolve in response to the trends of streaming growth, increasing content libraries and flattening concurrency curves. These trends, in turn, have been driven by factors including the introduction of new services such as Time-Shifted TV, increased high-definition programming and changing viewing patterns.

For the past 5 years Motorola's on demand products have led the way not only by responding to these trends but enabling new models for how on demand networks are designed and built. As existing trends accelerate and new ones emerge, Motorola is leveraging its experience to enable the next shift in architecture and a new wave of revenue-generating services.

The Motorola on demand portfolio offers the following benefits to video service providers facing these evolving challenges:

Efficiency

- Optimized Storage, Streaming, and Network Resources
- Adaptive Media Management based on real-world data
- Solid-State servers that lower total cost of ownership

Flexibility

- Supports multiple on demand architectures
- Features the industry's broadest range of solid-state servers
- Provides low and high density streaming optimization
- Offers support for Cable, IPTV, and Internet video

Future-Proofing

- "Pay-as-you grow" scalability minimizes CAPEX
- Supports rapidly expanding content libraries
- Offers support for advanced services, codecs, and multiple platforms

Global Leadership

- Recognized as the world's most widely deployed solid-state server
- Permits live deployments of advanced services with leading operators
- Delivers proven integrations with third-party products
- Is an integral part of Motorola's world-class, end-to-end video portfolio

For more information, please visit www.motorola.com/ondemand.



MOTOROLA

Motorola, Inc. www.motorola.com

The information presented herein is to the best of our knowledge true and accurate. No warranty or guarantee expressed or implied is made regarding the capacity, performance or suitability of any product. MOTOROLA and the Stylized M Logo are registered in the U.S. Patent and Trademark Office. All other product or service names are the property of their respective owners. © Motorola, Inc. 2009