SeaChange®

SOLUTION BRIEF

SeaChange® MediaCluster® releases the power of flash memory to drastically reduce operator service costs.

A FAST Architecture Reduces Total Cost of Ownership for On Demand Content Delivery

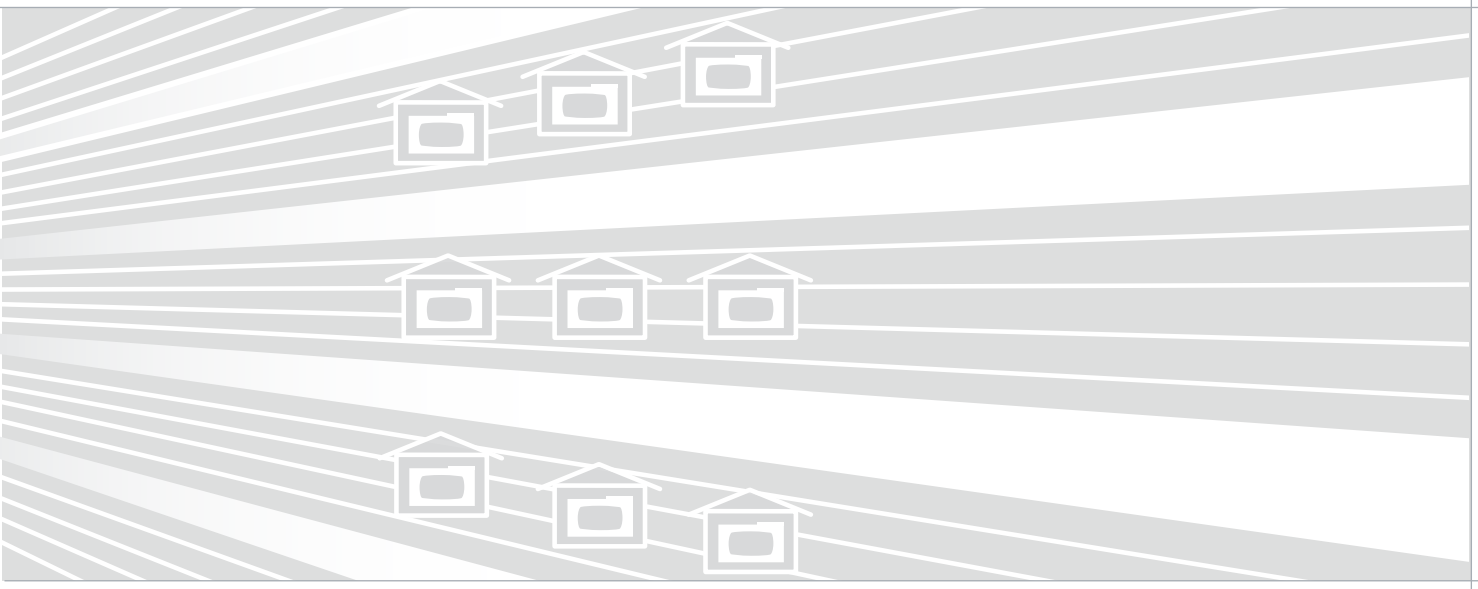

4/2007

In the 1990s, SeaChange® MediaCluster® technology revolutionized Video on Demand (VOD) by replacing large monolithic enterprise servers with clusters of small inexpensive disk servers. Now MediaCluster does it again, this time replacing expensive DRAM and spinning disks with flash memory.

Operators who deliver content on-demand over a network — whether cable, telecom, direct satellite, or Internet — have faced one inescapable barrier to growth. The more content they serve, the more expensive it gets to serve the content. What's wrong? The culprit is storage, and the need to deploy storage to the edge of the network so that subscribers get the quality of service (QoS) they expect.

As demand for particular content goes up, so does the number of copies that must be stored near subscribers at hundreds of network edge locations. And as the number of subscribers grows, so does the number of those edge locations. The result is compound growth in storage expense. Disks cost money to buy — but they cost more money to operate. The costs of operating disks include replacing failed disk drives, providing space, power and cooling, and spreading IT personnel across distributed edge locations.

An alternative approach is to replace some disks with DRAM (dynamic random access memory). But DRAM offers severely limited storage capacity and is volatile, i.e., it does not retain content if the server fails or is taken offline for maintenance. That's why every DRAM-based server comes with an abundance of locally attached disk drives.

The SeaChange solution is to deploy a relatively slow but also relatively large and inexpensive media — flash memory — in a way that enables multiple flash modules to outperform DRAM. It is an approach similar to what SeaChange did when it created MediaCluster. At that time, SeaChange replaced high-cost enterprise servers with clusters of low-cost media (inexpensive disks). Now it's taking that strategy one step further — replacing disks in MediaCluster with flash memory — an approach called FAST, or Flash Assisted Storage Technology.

Serving On Demand Applications

The performance test for any server architecture is how well it serves the usage profile of the applications served. Figure 1 shows a sample usage profile that is representative of most VOD deployments. The x-axis shows the assets in decreasing order of popularity and the y-axis contains the number of concurrent access of these assets during peak hours. As expected, a small number of popular assets contribute most of the work-load. Engineers often use an 80/20 rule — meaning 20% of the content generates 80% of the demand — when allocating server resources against content. The actual ratio, of course, will vary from site to site and will also vary from moment to moment in a given site.

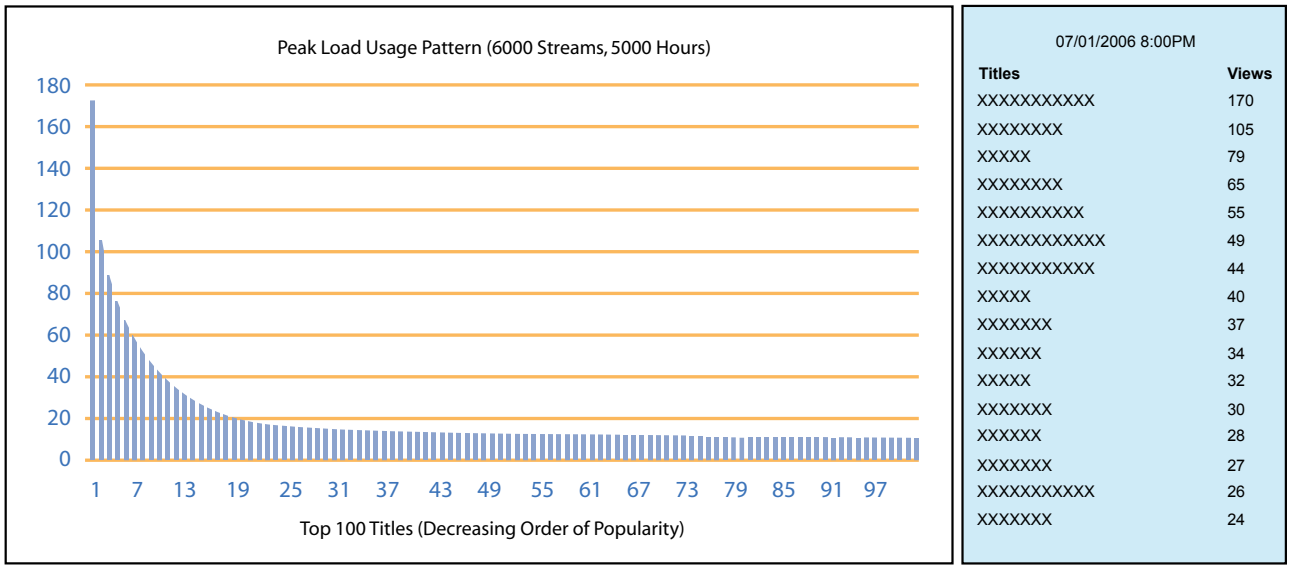

Figure 1: Small Number of Assets Contribute to a Majority of the Sessions

The 80/20 rule means that not all content needs to be replicated to the edge close to subscribers — only the 20% (more or less) hot content. The less popular 80% can be stored on, and served directly off, a disk-based system at the central location. QoS can be assured for both hot and cold content — assuming sufficient bandwidth exists between the centralized location and any edge location to support 20% of the workload.

So the question facing server vendors is this: If you only need 20% of content at the edge to serve 80% of streams, why are there so many disks at the edge — given disks' high ownership costs? Can't they be replaced with something less expensive?

Problems with DRAM

One approach that has gained some attention is to replace some of the disks in edge servers with DRAM (see Figure 2). However, that introduces three key issues. First, DRAM is severely limited in the amount of information it stores. A commodity server supports 16GB of DRAM, which stores eight hours — less than 0.1% of total content for most operators. Apart from its high price tag, a proprietary server can only support 128GB of DRAM or 64 hours of cached content, which is less than 1% of total content. That means you would need to purchase many servers in order to store 20% of content in DRAM.

A second issue is DRAM's volatility. Any server failure or maintenance downtime would cause the cached DRAM content to be lost. That would require a re-population of the cache — which in turn would overwhelm the network or the locally attached disk drives. Time-delayed applications would be especially vulnerable, where real-time ingested content would be lost forever.

High power consumption is DRAM's third issue. A 2GB DRAM module consumes as much power as a 500GB disk drive, further worsening the power and cooling problems at the distributed edge locations.

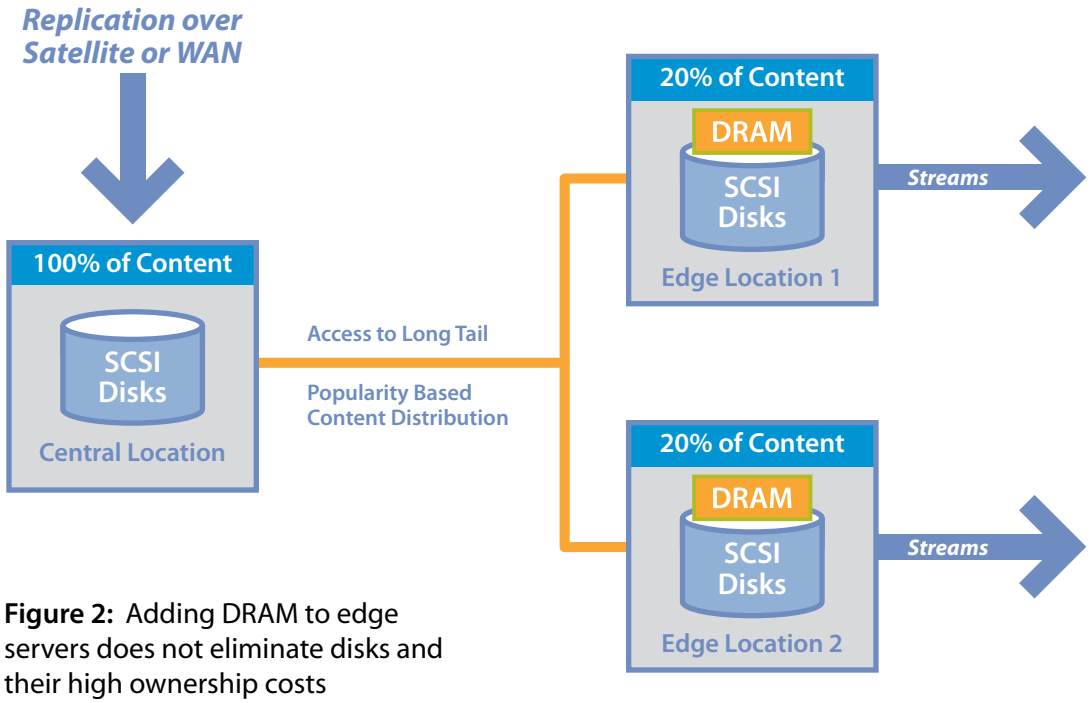

Figure 2: Adding DRAM to edge servers does not eliminate disks and their high ownership costs

What about Flash Memory?

Flash memory, which is used widely in consumer electronics, avoids DRAM's issues of storage capacity, volatility, and power consumption. A 2RU commodity server can support up to 760GB flash memory or 380 hours of content. That is 40 times more content than for DRAM in a comparable server. And since flash is non-volatile, the hot 20% content can be replicated to edge servers that have no locally attached disk drives.

Flash also consumes 5 to 10 times less power than disk, saving significant power and cooling costs. And, since it has no moving parts, flash memory is 100 times more reliable than disk.

The real test, however, is how flash performs against DRAM, which is inherently faster than both flash and disk. Eliminating disks at the edge — and their high ownership costs — requires DRAM-class performance. How can flash memory offer that?

The Answer: Flash Assisted Storage Technology

A similar question was once posed about another technology — inexpensive disk drives — that long ago became the building block of enterprise class storage. SeaChange MediaCluster adapted that technology to on-demand applications — striping data both across all the disks in a server and across all the servers (nodes) in a cluster. Because clusters output contiguous "chunks" of content in parallel across all the disks and all the nodes, effective bandwidth is magnified by a factor that scales with the number of disks and the number of nodes. The same effect is also achieved, however, when flash modules replace disks in a MediaCluster. Bandwidth of individual modules and individual nodes is combined — bringing DRAM performance to commodity flash.

By striping data across all the flash modules and all the servers, MediaCluster guarantees fault resilience in the case of flash module failure as well as in the case of complete server failure or service downtime. The method of RAID$^{2®}$ (RAID-squared) striping eliminates the expensive option of mirrored or backup servers. Missing data is recovered in real time without affecting streaming performance. Additional servers and flash memory can be added to an existing MediaCluster online without service disruption.

MediaCluster technology is also the key to extending flash memory's life expectancy. A flash module has a limited number of write cycles. Write hot spots — where writes are more intensive on one flash module than the others — would dramatically reduce flash memory's life expectancy. MediaCluster eliminates write hot spots by perfectly load balancing writes across all flash modules in a node and across all nodes in a cluster. By eliminating write hot

spots, MediaCluster extends flash memory's lifetime to more than 10 years versus five years for hard disk drives.

MediaCluster forms the key building block for SeaChange's FAST (Flash Assisted Storage Technology) architecture for on-demand content delivery. As shown in Figure 3 the FAST architecture consists of two tiers, a streaming tier that contains one or more flash-based MediaClusters and a storage tier that contains one or more disk-based MediaClusters. The streaming tier and the storage tier can be co-located or geographically distributed. In a co-located model, the streaming tier serves as an independently scaling streaming booster to the storage tier. In a distributed model, the streaming tier enables diskless operation at the edge locations while the storage tier enables centralized storage and storage management. At the heart of the FAST architecture is SeaChange® Axiom™ content distribution software, which automatically distributes high demand content to the streaming tier and keeps the less popular, or long tail, content at the storage tier.

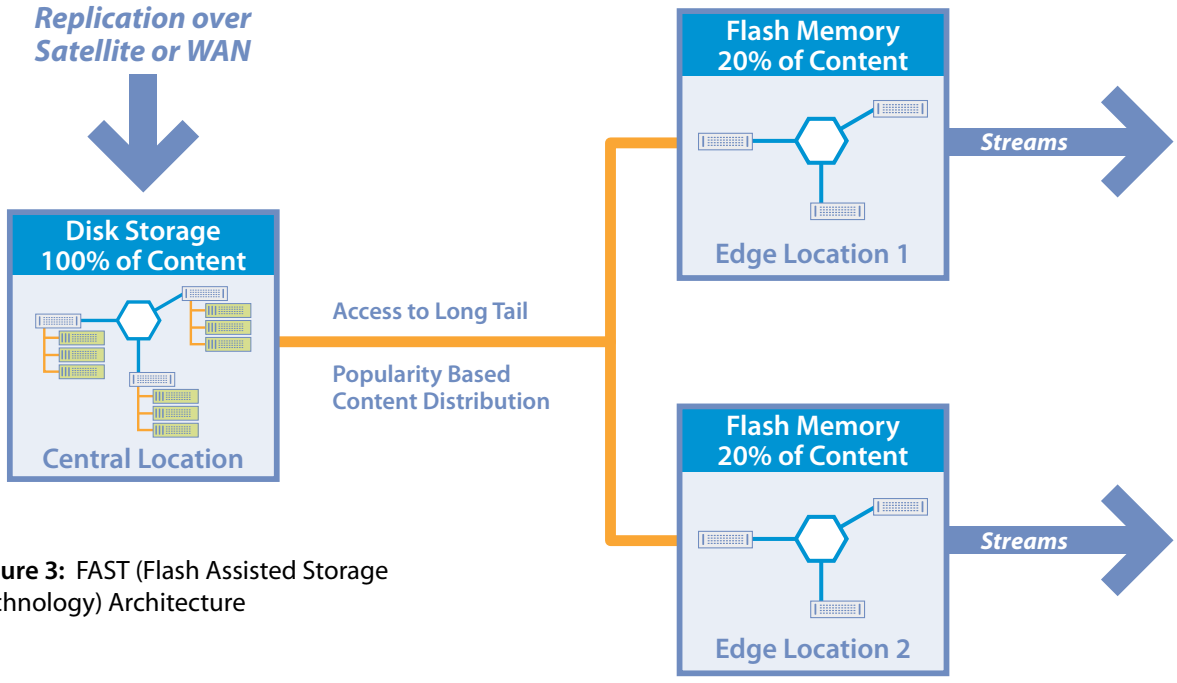

Figure 3: FAST (Flash Assisted Storage Technology) Architecture

Architectural benefits include:

- **Reduced Ownership Costs.** Superior reliability, longer life, lower power usage, reduced administrative overhead, and reduced space requirements all result in lower ownership costs than with DRAM-based approaches.

- **Increased Scalability.** By completely decoupling storage and streaming, FAST brings a new level of scalability and flexibility. An operator can scale the streaming tier for increased subscriber/edge locations and scale the storage tier for increased content.

- **Open Architecture.** Third party applications ingest, transcode, and preview content on the storage tier via open storage interfaces.

- **Investment Protection.** Operators can leverage SeaChange's widely deployed disk-based MediaClusters as the storage tier while boosting streaming capacity by adding the flash-based MediaClusters.

The FAST architecture radically moves the curve when it comes to providing cost-effective, scalable, and reliable storage for on-demand applications. Operators can now bring much more content to the network edge to more subscribers without having to bring disk storage with it. That means more content for consumers with higher quality of service for greater market differentiation and lower cost of technology ownership.

SeaChange®

For more information please contact your SeaChange Representative
or visit our website at www.schange.com

SeaChange International is the leading provider of end-to-end and best-of-breed solutions for the world's growing on-demand and IPTV industry. Its powerful, open video on demand and advertising software and scaleable hardware enable broadcast, cable and telecommunications operators to provide new on-demand services and to gain greater efficiencies in advertising and content delivery. With its Emmy® Award-winning and patented technology, thousands of SeaChange deployments are helping broadband, broadcast and satellite television companies to streamline operations, expand services and increase revenues. Headquartered in Acton, Massachusetts, SeaChange has product development, support and sales offices around the world.

SeaChange®

SIMPLY THE BEST.