TANEJA GROUP
TECHNOLOGY ANALYSTS

T E C H N O L O G Y B R I E F

The Business Case for Centralized Storage Caching

December 2006

THE NET NET

Virtualization is rapidly reshaping data center infrastructure. By allowing physical server and storage devices to be deployed as resource pools, management and scalability practices change dramatically.

Server virtualization allows operators to create compute capacity pools and to rapidly provision and de-provision applications and their operating environments with shared CPU resources. Storage virtualization allows operators to pool storage capacity and drive higher utilization rates.

However, virtualization and consolidation of resources brings new challenges that must be addressed. Foremost among these challenges is scaling storage performance. The consolidation of application workloads and the drive to stack more and more virtual machine instances on a single server taxes the underlying storage systems and feeds I/O contention. Furthermore, the old approaches to provision and scale storage performance have significant drawbacks and are not cost-effective.

What is needed is a new approach to the storage performance problem. We believe that the disagregation of the storage cache from the physical storage device and the consolidation of that cache into a pooled network resource represent the next wave in this virtualization revolution sweeping the data center. The pooling and scaling of a separate shared caching resource will have far reaching impact on improving the overall performance and quality of service of a wide class of enterprise applications, like databases, email, and other line of business applications. In the following paper, we debunk the conventional wisdom regarding storage performance, introduce a new technology called centralized storage caching that addresses the storage performance bottleneck, detail the demonstrable ROI drivers fueling adoption of this technology, and finally spotlight Gear6, a promising vendor in the nascent field of centralized storage caching.

Storage Performance Conventional Wisdom

Understanding and characterizing performance for a given application or workload can be complex and highly specialized to the particular application and its I/O pattern. However, too often the industry dialog around storage performance lacks depth and insight. To that end, we examine the conventional wisdom regarding storage performance, storage costs, and storage performance scaling in an effort to reframe the discussion.

Copyright© The TANEJA Group, Inc. 2006. All Rights Reserved
87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com

Myth #1: IOPS Equals Performance

The most common published gauge of performance for a storage system is I/Os Per Second (IOPS). IOPS measures how much data can be streamed out from a single storage device or system at peak operating load. Typically, the number represents the sustained optimal high point of that device.

However, IOPS is only one dimension of performance. We believe that end users should think of storage performance in two dimensions – how much throughput (e.g. IOPS) can be sustained for a given workload and how long does it takes to respond to an I/O request (e.g. access time). For example, access time is a very critical factor in the overall response and performance of transactional applications, such as databases and email servers, and specific workloads like video streaming. On the other hand, IOPS can be a key determinant for applications like seismic processing and biotechnology research.

The key concern for the storage administrator is how he or she optimizes system wide performance for a given class of applications. In many cases, scaling IOPS is not the answer. End users and vendors alike must reframe their thinking on performance to not just focus on IOPS, but also factor in reducing access times in order to improve overall application performance.

Myth #2: Assess Storage Costs Solely on Capacity Metrics

End users and vendors alike tend to measure and compare the cost of storage purely based on a cost per capacity metric ($ per TBs). Although this is clearly an important dimension of storage costs, it does not take into consideration the cost to achieve a certain level of response time or throughput.

Taneja Group believes that end users would be well served to assess their storage investments not only on cost per capacity metrics, but also in terms of price performance metrics. Users must refine their view of storage costs and inject performance (particularly access times) into the storage cost calculus. From our conversations with forward thinking end users, they measure their systems on price-performance metrics, such as $/IOPS and $/access time, in addition to cost per capacity.

Myth #3: More IOPS Requires Provisioning More Capacity

Another significant conundrum facing end users is how to provision storage performance. Disk drives represent the only mechanical device in the data path. Due to the physical limitations of how fast a drive can spin and seek data, each disk can only deliver a finite amount of I/Os per second. However, disk I/O rates are extremely slow compared to the rate of modern day CPUs to process data and the network rates to transmit it. As a result, the disk becomes the key bottleneck to system and application performance.

End users have resorted to various workarounds for this fundamental data access issue. For example, end users typically create storage volumes that span multiple drives or spindles and stripe the data across those spindles in order to increase IOPS system wide.

Copyright© The TANEJA Group, Inc. 2006. All Rights Reserved

However, this common approach to provisioning performance has two negative repercussions. First, to increase IOPS performance for a particular volume, an end user must increase capacity. Capacity and performance cannot be scaled independently of each other. Thus, scaling performance leads to wasted space and increased administrator overhead. Second, adding spindles to a system does nothing to improve the response time of the system. Therefore, a storage administrator has no technique to address the second dimension of performance – access times.

Emergence of Centralized Storage Caching Architectures

Based on our conversations with end users and several promising startups, Taneja Group is tracking a new trend in storage infrastructure architectures – centralized storage caching. From our vantage point, cache that has traditionally been included in storage systems is being disaggregated and consolidated into a network resource. This complements and augments the cache in existing storage devices. Effectively, storage caching is about creating a huge pool of memory cache that stores not only the most recently accessed blocks across the storage network, but also has a policy engine that determines the optimal distribution of data across the caching system.

End users can now tackle the storage performance challenge of scaling both IOPS and reducing latency without increasing backend capacity. Ultimately, storage caching represents a new tool in the storage administrator's toolkit to deliver optimal performance across the storage infrastructure.

Centralized storage caching offers three compelling advantages to complement traditional storage systems:

1. **Order of Magnitude Better Access Times Than Disk.** Fundamentally, the cache resource accelerates performance because retrieving a block of data from RAM is 10,000 faster than retrieving a block from a disk. In fact, caching allows access times to be compressed from milliseconds to microseconds.

2. **Non-disruptive**. A network cache is completely transparent to the applications and underlying storage infrastructure. Network cache devices are in-band network resident devices. No host agents are required and backend storage systems require no additional configuration or tuning to take advantage of a network cache.

3. **Scale Cache Capacity Independently of Storage Capacity.** By separating the cache from the storage system, the centralized storage cache allows the cache size to grow in proportion to the needs of the application and its I/O profile. Administrators can simply add more cache without provisioning storage capacity to increase IOPS or reduce access times.

At present, the vendor community championing centralized storage caching is focused on optimizing NAS response times

Copyright© The TANEJA Group, Inc. 2006. All Rights Reserved

and overall IOPS performance, but we have every reason to believe that this same technology will ultimately be applied to Fibre Channel and iSCSI block access.

The Implications of Centralized Storage Caching for Tiered Storage

We believe that the emergence of a scalable caching resource has far reaching implications for how end users conceive of their tiered storage environments and ILM strategies. Traditionally, end users have defined storage tiers based on the performance and SLAs required for their application data. To that end, the characteristics of the drives, the speed of the drives (RPMs), drive types (SAS, FC, SATA), RAID levels (RAID 0, RAID 5, RAID 6, etc), and advanced data protection techniques (frequency of snapshots, replication, etc) taken together form a given class of service or tier. Typically, mission-critical transactional application data, like databases and email are assigned tier 1 storage, while long term archival, backup and restore, and regulatory compliance data is assigned tier 2 storage.

We believe that memory based caching architecture adds a new dimension or variable for users to consider when they define storage tiers. We believe that end user must reframe their thinking about application performance to include the concept of access time and the notion of quality of service as they plan storage tiers. In fact, with caching, users will be able to provision storage to deliver a given response time or ensure that performance characteristics (IOPS and access times) can scale to meet usage spikes. Caching allows users to marry the concept of delivering better response times and quality service with tiered storage classes. As a result, network-based caching has a chance to change the current calculus of storage tiering.

Ultimately, we believe that memory based caching systems will meld with existing storage infrastructure to complement the current storage tiering architecture. Users will use caching and its ability to independently scale IOPS and decrease access time as a mechanism for augmenting existing storage tiers and further tailoring them to the I/O profiles of applications in their environments. We foresee a broad usage for centralized storage caching technology whereby applications ranging from transactional applications, like databases and email, to High Performance Computing applications leverage the caching technology to boost performance. Going forward, network caching technology should have far reaching ramifications for tiered storage deployments.

The Business Case for Centralized Storage Caching

In addition to turbo charging the raw performance of traditional storage infrastructure, a shared network cache delivers a demonstrable Return On Investment (ROI) and a quick investment payback. Centralized storage caching reduces costs and improves efficiency in five key dimensions:

1. **Compressed Time to Revenue.** The earliest adopters of centralized storage

Copyright© The TANEJA Group, Inc. 2006. All Rights Reserved

caching technology have been companies that need to speed up business critical processes. These firms typically have revenue generating applications, such as video rendering, drug discovery, and oil and gas exploration that are bottlenecked on I/O access. By speeding access to the data, they can shave hours and days off these business critical processes and complete their revenue generating projects quicker. Ultimately, these firms can gain competitive advantage and compress the time to revenue. The ROI and payback of centralized storage caching technology is immediate and demonstrable for this type of application.

2. **CAPEX Savings on Underutilized Capacity**. As discussed above, conventional provisioning of storage performance leads administrators to buy and provision more capacity than they need. This approach leads to underutilized storage capacity. Centralized storage caching devices fix this counter-intuitive provisioning model because they allow the independent scaling of cache size without increasing capacity. Administrators no longer need to deploy more spindles to grow IOPS. In fact, they can deploy and grow a network cache to improve both IOPS and access time. Therefore, end users can dramatically decrease the amount of storage capacity that they need to purchase for high performance applications.

3. **Power, Cooling, and Space Savings**. A spillover effect of reducing capacity is that centralized storage caching also reduces the amount of power, cooling, and floor space needed in the data center. In short, each spinning disk that is eliminated results in less power consumption to spin it, less power needed to cool it, and less rack space needed to house it. Given that power, cooling, and space represent one of the largest line item expenses in data center operations, reducing these costs can be a critical ROI driver for centralized storage caching technology.

4. **OPEX Savings from Less Storage.** A second spillover effect of reducing deployed capacity is the reduction of administrative cost. There is a direct relationship between the number of TBs deployed and the number Full Time Equivalents (FTEs) needed to manage that capacity. Most IT departments staff their storage teams based on a fixed staff ratio (e.g. Managed Capacity in TBs to FTEs). By reducing the need to grow capacity for performance, IT no longer needs to increase headcount aggressively. Moreover, critical personnel are now available for more strategic IT projects and tasks. As a result, centralized storage caching technology allows companies to contain their OPEX spending and achieve better productivity from their current staffing levels.

5. **Improved Quality of Service.** Centralized storage caching offers the ability to improve response times for key transactional applications such as databases, where there is a high ratio of peak to normal usage. Traditionally, in these workloads, an administrator must

Copyright© The TANEJA Group, Inc. 2006. All Rights Reserved

provision the storage environment to support peak loads, even if the peak loads are only hit once a month or once a quarter. As a result, a storage administrator may need to purchase faster storage, more storage controllers, and even more spindles to ensure adequate response time. Centralized storage caching provides a demonstrable ROI because it can boost responsiveness of critical applications without undue purchases in networking, storage controllers, faster storage, and additional spindles. Furthermore, delivering rapid service response times will become increasingly important as more and more business processes become automated and are online 24x7.

Spotlight on Gear6

Headquartered in Menlo Park, California, Gear6 is a pioneer in the burgeoning field of scalable centralized storage caching. Gear6's flagship product, CACHEfx, is an in-band, network resident caching device that can attach to any Ethernet network. The product supports data access via NFS today, but we expect the company to expand its offerings to iSCSI, Fibre Channel, and CIFS protocols as it matures its offering.

CACHEfx can accelerate client response time and IOPS by 10-50X depending on the workload. Moreover, CACHEfx can scale out to support millions of IOPS. The company's product is particularly well suited for applications with large data sets, random data access, high number of concurrent clients and transactions, significant peak loads, and strict service level requirements.

CACHEfx utilizes an innovative clustered architecture that presents the cache capacity as a single shared pool of cache. Each CACHEfx module within the appliance contributes a fixed amount of RAM to the overall distributed cache. As a result of this clustered architecture, both capacity and IOPS of the CACHEfx appliance can be scaled out by adding additional modules on an "as needed" basis. To this end, users can start with a relatively modest network cache and grow it over time to meet increasingly storage performance requirements.

Taneja Group Opinion

We believe that centralized storage caching offers a killer value proposition – an order of magnitude better performance (for both IOPS and access times) coupled with a non-disruptive approach that complements existing storage infrastructure. In short, we believe that centralized storage caching has the potential to change how end users think about scaling performance and deploying tiered storage environments.

Centralized storage caching has wide applicability across a wide range of enterprise workloads and applications. We foresee centralized storage caching being adopted in the core of the data center in support of mission-critical transactional applications, such databases, email, and file serving. As server virtualization continues to feed I/O contention, we believe that centralized storage caching could be a crucial element in alleviating that inevitable bottleneck. Furthermore, we expect that centralized storage caching technology will

Copyright© The TANEJA Group, Inc. 2006. All Rights Reserved

be applicable to particularly latency sensitive applications, like video streaming and other I/O intensive workloads outside the core data center.

The implications and ROI of centralized storage caching technology are compelling and demonstrable. To end users grappling with performance bottlenecks, we recommend that you consider a new approach to the current storage performance provisioning model. To that end, we believe that Gear6 is an emerging leader in developing memory-based cache architectures. We recommend that you give them a look. We don't think that you'll be disappointed.

NOTICE: The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document.

Copyright© The TANEJA Group, Inc. 2006. All Rights Reserved
87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com