

SNIA.
CSI | CLOUD
STORAGE

Hyperscaled Enterprise Storage

SNIA.
DPCO | DATA PROTECTION &
CAPACITY OPTIMIZATION

September 2016

Abstract: *Hyperscaler storage customers have always built their own storage systems from commodity components. Now enterprises are following suit, disrupting the traditional enterprise storage market. This paper explores that trend and highlights real-world methods enterprises are using to build their own hyperscale storage environments to ensure scalability, improve agility, avoid vendor lock-in and lower costs.*

Foreword

This paper is the work of a joint effort by the SNIA Cloud Storage Initiative (CSI) and the Data Protection and Capacity Optimization (DPCO) committee. It includes results from a discussion with the head of a major global bank that is building its own datacenters and filling them with commodity components built into custom racks.

Executive Summary

Enterprises have traditionally purchased highly available storage systems with built-in redundancy and dedicated storage controllers with proprietary firmware. However growth of these types of systems has slowed and revenues are declining.

*Total worldwide enterprise storage systems factory revenue declined 2.2% year over year to \$10.4 billion during the fourth quarter of 2015 (4Q15), according to the International Data Corporation (IDC) **Worldwide Quarterly Enterprise Storage Systems Tracker**. Total capacity shipments were up 10.7% (year over year) to 35.5 exabytes during the quarter. Revenue growth declined within the group of original design manufacturers (ODMs) that sell directly to hyperscale datacenters. This portion of the market was down 14.4% year over year to \$1.2 billion. Sales of server-based storage were up 6.1% during the quarter and accounted for \$2.2 billion in revenue. External storage systems remained the largest market segment, but the \$7.0 billion in sales represented a 2.3% year-over-year decline. IDC March 11, 2016*

Global spending on public cloud services is expected to grow 16% to \$204.2 billion this year, compared with 13.8% growth in 2015 and a 17.7% rise in 2014, according to data from researcher [Gartner](#)

*Global IT spending overall is expected to be \$3.49 trillion in 2016, a 0.5% decline from last year. Many companies embrace a **“hybrid cloud” model** that integrates cloud services and private data centers. CIOs often say that not all applications belong in the cloud, and that they are looking for ways to bridge their multiple IT environments. [Wall Street Journal](#)*

Some of the data is obviously moving to the cloud and that seems to be growing significantly. But few (Netflix mainly) have made the transition to total cloud based IT. They principally cannot get all the services they need from today's public cloud.

Our research has indicated that many of these enterprises, in addition to utilizing cloud storage are also building their own storage systems using software defined storage (SDS) and best in class commodity components, assembled in racks.



Introduction

This paper describes the overall trend of enterprises adopting methods from hyperscalers such as Amazon, Facebook, Google and Microsoft Azure to build their own storage systems while achieving the required service levels for internal projects. These private cloud storage systems are used to host the data for business critical applications and comply with regulations from multiple government jurisdictions.

We profile an actual use case from a large (in the top 30 of the Fortune 500) bank that allowed us access to the details of their own efforts in this space. Other enterprises are following suit and impacting the revenue from traditional enterprise storage systems.

1 Hyperscale Storage Overview

This section briefly describes key storage technologies with the intent of setting the stage for the enterprise use of these methods.

There is a new category of storage vendors called original design manufacturer (ODM) direct who package up best in class commodity storage devices into racks according to the customer specifications and who operate at much lower margins. They may leverage hardware/software designs from the Open Compute Project (OCP) or a similar effort in China called Scorpio, now under an organization called the Open Data Center Committee (ODCC), as well from as other available hardware/software designs.

Hyperscalers built huge datacenters with a minimum of 5,000 servers and a footprint of at least 10,000 square feet, but generally much larger. They need to scale their storage at the speed of their business growth. They are very sensitive to the cost of acquisition, thus the emphasis on best in class commodity components. They typically use their own or open source software (software defined storage) to create manageable storage systems in these datacenters. The designs are tailor-made for the application/service requirements.

Large enterprises still building their own datacenters include large banking institutions such as Goldman Sachs, Capitol One, Fidelity and Bank of America, as well as telcos such as AT&T. Even though they may be increasing their use of the public cloud, they also have requirements which are difficult to meet with the services currently available. Banking institutions in particular are subject to many audits from multiple jurisdictions of regulators. They essentially create their own private storage cloud to ensure regulatory compliance with the many and varied requirements each government mandates.

Hyperscalers have never been fond of enterprise storage systems with built-in redundancy and high margins. Enterprises are just now mimicking this trend and the results are illustrated in the market numbers.

Hyperscaled Enterprise Storage



1.1 Best in Class Commodity Components

Hyperscale storage is built from best in class commodity storage components such as SSDs and hard drives. These drives are typically delivered populated into racks and purchased as a “pod” with a dozen racks making up one storage “system”. Around half of the racks are populated with server “heads” to perform the software defined storage (SDS) in a scale-out manner. The disk trays are JBODs – just a bunch of disks, and the SSD trays are JBOFs – just a bunch of flash. Another technique is essentially densely populated compute nodes where all the storage is directly attached.

The SDS and other custom built software handle the resiliency by replicating or using erasure code across the datacenter and geographically between datacenters. Performance scales almost linearly at least within each pod. They also have a “fail in place” policy, letting failed components exist powered down in the datacenter until a certain percentage of trays or racks is unusable, and only then replacing these larger units.

1.2 Software Defined Storage

As mentioned above, software defined storage is used to provide the higher level resiliency needed when using these components. This cuts out the margin and markup typical in enterprise storage arrays. Yet they can still fulfill redundancy, availability and performance goals. SDS allows them to scale out incrementally with identical nodes. SDS also provides a platform for further automation and may be part of an internal or external storage cloud.

The operations personnel work closely with the SDS vendor, or may even maintain an open source SDS implementation. This may require development as well as operation teams called “DevOps” – a trending concept. The storage provided to internal customers might be tuned to specific application requirements.

The focus is on datacenter scale monitoring and management and over time new features may be required in the drives to accommodate some of the unique requirements.

2 Use Case: Large Global Bank Example

This section provides a description of the methods a large global bank uses to build its own storage systems and private storage cloud.

2.1 Scale of Deployment

After consulting with an IT vice president of a large multinational bank, we learned something of the scale and operation of their internal infrastructure. It matches that of the large Internet hyperscalers both in size and the trend toward using the same procurement and operation techniques.

Hyperscaled Enterprise Storage

The bank has over 20 datacenters around the world. They create an internal private cloud for the entire bank's IT project usage. They cannot use the public cloud for this data as they currently need to comply with over 200 different country government regulations. Their storage budget dwarfs the revenue of most medium sized storage vendors.

Enterprise companies that embrace a blended value model, offering software defined storage, long term retention, data life cycle management and traditional work horse storage are better suited to benefit from this shift in the industry.

Their deployed storage has 10s of thousands of nodes with around 200 petabytes of active data and ½ exabyte of inactive data. Their overall data footprint is growing at 45% annually. They process 10s of trillions of transactions daily, and downtime is very expensive.

They are big enough that vendors will custom build for them, but they also have a policy of no single source for any of their hardware. They buy storage in 6 PB pods that are ½ CPU and ½ storage drives, pre-assembled by the ODM direct vendor. They installed their first such pod in 2015. Their next pod will be all flash.

Most importantly their cost savings is projected to be 50% over traditional storage.

2.2 Software

As previously mentioned, they use SDS with the best in class commodity hardware to create a private cloud for internal IT projects and customer facing services. They are currently deploying about 11% of their storage as SDS today, but have a goal of 20% by the end of 2016. Their goal is to grow this to 50% by 2020.

They license their SDS from a major vendor (site license) and are training up their staff in the new approach. They virtualize the hardware to abstract away any differences between the multiple vendors. From the storage service they serve up an S3 compatible interface for new projects and mainly provide block services for existing applications.

They plan to look at Ceph for SDS and NVMe for SSD interfaces in the future.

Summary

The trend to build hyperscale storage infrastructures involves utilizing best in class commodity hardware, and providing all the differentiation in the software to manage all aspects of the storage. Thus, all the "intelligence" is in the software, such that the infrastructure behavior is programmable via that software. Look for platforms running intelligent software on general purpose hardware in the future. This is a trend for enterprises that bears watching.

Hyperscaled Enterprise Storage



For More Information

Additional information on SNIA Cloud Storage activities, including the Cloud Storage Technical Work Group, can be found at <http://www.snia.org/csi>. Additional information on SNIA Data Protection and Capacity Optimization can be found at <http://www.snia.org/forums/dpco>

Suggestion for revision should be directed to <http://www.snia.org/feedback/>.

About the SNIA

The Storage Networking Industry Association (SNIA) is a not-for-profit global organization, made up of member companies spanning the global storage market. SNIA's mission is to lead the storage industry worldwide in developing and promoting standards, technologies, and educational services to empower organizations in the management of information. To this end, the SNIA is uniquely committed to delivering standards, education, and services that will propel open storage networking solutions into the broader market. For more information, visit <http://www.snia.org>.



Storage Networking Industry Association

4360 Arrows West Drive • Colorado Springs, CO 80907 • Phone: 719-694-1380 • Fax: 719-694-1389 • www.snia.org

© August 2016 Storage Networking Industry Association. All rights reserved.