



DOES CLOUD STORAGE PORTEND “CREATION BY DESTRUCTION” OF TRADITIONAL STORAGE?



Mikhail Gloukhovtsev
Sr. Enterprise Architect
Comcast Cable

Table of Contents

| | |
|--|----|
| Introduction | 4 |
| Storage Solutions for Cloud-Based Information Services: Is SAN Dead for Cloud? | 6 |
| What Is Cloud Computing?..... | 6 |
| Cloud Computing from the Perspective of Storage Services..... | 6 |
| Cloud Friendly and Unfriendly Applications | 7 |
| Cloud Data Management Interface | 8 |
| Impact of Cloud Computing Requirements on Development of Storage Technologies | 9 |
| DAS Solutions for Cloud | 10 |
| Just DAS..... | 10 |
| High-Availability DAS | 10 |
| Distributed Host-Cache | 11 |
| SAN Solutions for Cloud | 12 |
| Transformation of SAN Solutions Driven by Cloud Requirements | 12 |
| “Cloudization” of Existing Storage Solutions..... | 13 |
| VCE: Integrated Infrastructure-as-a-Service Solution | 13 |
| Scale-up vs. Scale-out Solutions | 15 |
| Object-Based Storage | 16 |
| Atmos | 18 |
| Open Source Solutions and Cloud Computing | 21 |
| Storage Solutions in OpenStack..... | 21 |
| Comparison of Object-Based Storage with Features Required for Enterprise Class Storage Service | 24 |
| Software-Defined Storage | 26 |
| Symbiosis of HA DAS and SDS..... | 29 |
| NAS and Cloud | 29 |
| Data Mobility in Clouds | 29 |
| Cloud Storage Gateways and EMC Cloud Tiering Appliance..... | 30 |
| Big Data Meets Cloud Data..... | 31 |
| Isilon Storage for Big Data..... | 35 |
| EMC Greenplum Distributed Computing Appliance (DCA)..... | 35 |
| Object-Based Storage for Big Data..... | 36 |
| Fabric Storage for Big Data: SAN Functionality at DAS Pricing | 36 |

Conclusions37
References38

Disclaimer: The views, processes, or methodologies published in this article are those of the author. They do not necessarily reflect EMC Corporation's views, processes, or methodologies.

Introduction

How can a large company transform the storage services it currently uses to make them function as storage services for a private and/or hybrid cloud as well? If the ability to provide storage on demand is considered the key feature of cloud storage, what changes are required in the traditional storage infrastructure based on enterprise (i.e. EMC VMAX[®]) and midrange (i.e. VNX[®]) storage platforms to deliver cloud-based storage services? Companies have heavily invested in conventional storage systems and the question of how these systems can be re-used for cloud solutions becomes critical for development of the company IT infrastructure road map.

Should traditional storage systems oriented to management of structured data be kept only for transactionally-consistent applications? Or can they be upgraded to become storage components of a new cloud infrastructure? What features make storage “cloud storage”? What are the evolving cloud storage standards (Cloud Data Management Interface [CDMI], etc.)? This Knowledge Sharing article will try to answer these questions although the author understands that this search for answers may lead to more questions.

Recent years have witnessed the advent of two new phenomena in IT: cloud computing and Big Data. These two tectonic changes have been driven by business needs. They require new architectures and new service solutions and have affected all major components of IT infrastructure: server → server virtualization, network → network virtualization and software-defined network, storage → storage virtualization and software-defined storage, and security → security for cloud computing. In this article, I discuss cloud computing and Big Data only to the extent warranted by the discussion of the changes in storage services and technologies caused by these new phenomena. The reader can find many detailed reviews of cloud computing and Big Data storage in various publications, including books. I have included some of these publications in the References section.¹⁻⁴

While we all agree that cloud-based storage and Big Data require new storage architectures that may be radically different from traditional storage infrastructure, there are arguments about the scope of these changes. Does it mean a “rip & replace” of the existing storage ecosystems? A Forrester publications posed a question in its title: “Do You Really Need a SAN Anymore?”⁴ Has Direct-Attached Storage (DAS) come back in new incarnations such as HA-DAS, “application-centric storage”, “software-defined storage”, etc.?

In my opinion, all of these storage technologies and solutions based on them are complementary to each other as there is a very broad spectrum of applications and services requirements which existing and emerging storage solutions have to meet. From a philosophical perspective, the evolution of storage solutions is an example of a dialectic process driven by internal contradictions and negation (thesis, antithesis) that lead to further development by synthesis. Conflict between the thesis and antithesis is solved by the progress from quantity to quality, the negation of the initial development of the status quo, the negation of that negation, and then by reconciling the common features at a higher level.⁵ In terms of economics, it can be seen as an example of "creative destruction," the term used in economics to describe the way in which capitalist economic development arises out of the destruction of some prior economic order.⁶

The goal of this article is to provide a brief overview of the holistic guidelines that can help readers select storage technologies for implementation of cloud-based information services and Big Data solutions.

Storage Solutions for Cloud-Based Information Services: Is SAN Dead for Cloud?

What Is Cloud Computing?

We have to start with the definition of cloud computing. The definition offered by NIST is the most widely accepted: “Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model is composed of five essential characteristics, three service models, and four deployment models.”⁷

The essential characteristics are:

- On-demand self-service
- Broad network access
- Resource pooling
- Rapid elasticity
- Measured service

Note that the cloud-computing definition does not refer to any specific technologies. Rather, it is a model for the use of “a shared pool of configurable computing resources.”⁷ What is important for us when we consider the storage component of cloud services is rapid elasticity, resource pooling, measured service, and on-demand self-service.

Cloud Computing from the Perspective of Storage Services

The storage solution we choose for our cloud implementation should have the above characteristics. There are four deployment models—private cloud, hybrid cloud, public cloud, and community cloud⁷—and the relative importance of the cloud characteristics vary depending on the deployment model. We will focus our attention on the first two.

Rapid elasticity for storage resources means the ability to provide storage on demand and be able to scale it back as needed. This feature is important for public cloud service providers that face fluctuating workloads generated by various applications in a multi-tenant storage environment and provide services to many customers requesting quick storage allocations and releases/reallocations of the storage resources as needed. For private cloud implementations in which the number of customers is limited by business units and the I/O workload patterns for the

business applications used by the company are much better known, achieving rapid elasticity may not be as difficult as in public clouds but it is still a challenge.

Resource pooling is implemented as storage pools that can be created within a storage array, as an inter-array pool, and as a geographically stretched pool aggregating storage resources from several data centers.

On-demand self-service is the ability to automatically provide storage resources required for a given service requested by a customer. SNIA defines cloud storage as “the delivery of virtualized storage on demand.”⁸

Measured service provides control and optimization of storage resource use by implementing a storage service catalog including charge-back or show-back (internal) billing models. A storage service catalog offers multiple service classes with various standard service level options including operational support and charge-back. For example, storage service level can specify availability requirements, I/O performance metrics (latency, IOPS, throughput), and data protection requirements (DR RTO and RPO, backup RTO and RPO). The service can include other resources (network, security controls, etc.) and can be automatically deployed via the actionable service menu and managed.

Implementation of a private cloud can reduce CapEx growth by sharing a pooled commoditized resource that eliminates overprovisioning and underutilization of storage resources. Utilization is improved by provisioning storage for normal rather than peak loads, with greater agility. Using external cloud service providers such as in a hybrid cloud-based storage infrastructure enables transforming CapEx into OpEx by purchasing SLAs rather than storage hardware. OpEx will be reduced in private clouds by simplified management and automation that accelerates storage provisioning and service modifications. Data replication within a private cloud should simplify DR solutions.

Cloud Friendly and Unfriendly Applications

As our goal is to create application-centric storage solutions, before discussing storage solutions for applications deployed in the cloud, we need to classify our applications to understand which of them can be moved to the cloud and what cloud deployment model should be used for that, which applications can have only some cloud-deployable components, and which cannot be moved to the cloud at all. Since this topic has been discussed in detail in many publications,¹⁻² I will briefly mention only key points related to storage solutions. Creating an

application taxonomy (“application consideration checklist for migration to cloud”) will help to identify the best candidates; typically, standard business applications such as e-mail and office productivity suites. Legacy applications that may be difficult to virtualize as well as highly customized applications are good candidates only for private clouds. Applications should be cloud-optimized, have loose interdependence, and be component-failure tolerant. There may be a need to change the application design leveraging the Normal Accident Theory to make the application design loosely coupled and as linear as possible.² Depending on the chosen cloud solution, the failover design can be transformed from enterprise-centric (centralized software on highly reliable hardware) to commodity-centric (distributed software on less reliable hardware). The latter is evolving into what is called a software-defined data center⁹ and software-defined storage,¹⁰ in particular.

When selecting a storage solution, it is important to understand whether a given application belongs to ACID or BASE type. ACID-type applications are characterized by Atomicity (all or no part of transactions are committed), Consistency (transaction creates valid states or returns to previous state), Isolation (non-committed transaction remains isolated from other transactions), and Durability (committed data is saved to ensure consistent state).¹¹ Conversely, the BASE model (Basically Available Soft-state, Eventual consistency) allows for variations in database consistency. As a result, the BASE model makes it possible to achieve levels of scalability that cannot be provided by the ACID model.¹¹

In mapping these requirements to storage solutions, it is likely we will end up having a set of solutions running the gamut from “traditional SAN”, to object-based storage, to “distributed storage”.

Cloud Data Management Interface

Cloud data should be managed and the management of very diverse types of data requires standards. The Cloud Data Management Interface (CDMI, Version 1.0.2 is the latest) defines the functional interface that applications will use to create, retrieve, update, and delete data elements from the cloud.⁸ CDMI allows clients to discover the capabilities of cloud storage offerings. It is also used to manage accounts, security access, and monitoring/billing information, even for storage that is accessible by other protocols.

More than two hundred vendor, end-user, and academic members belong to the Cloud Technical Working Group, many of whom have made significant contributions to the CDMI standard. CDMI addresses the lack of standardization in the three key problem areas related to cloud storage:

1. Client-controlled client-to-cloud data transfer
2. Client-controlled cloud data management
3. Client-controlled cloud-to-cloud data transfer

Impact of Cloud Computing Requirements on Development of Storage Technologies

Developing IT infrastructure to meet cloud computing requirements means changing from project-based infrastructure to service-oriented infrastructure. In traditional storage environments, the growth in storage capacity and performance is planned and storage assets are procured in advance according to the plan included in the budget. Taking into consideration the business dynamics leading to changes in the storage service requirements, such storage growth planning cannot be accurate even for traditional environments. Long term procurement planning for storage growth for cloud computing and Big Data is hardly possible. It means that rapid elasticity of resources is a requirement resulting in a need for adding new functionality to existing, mainly scale-up technologies and developing new cloud-oriented or Big Data-optimized technologies.

The relative importance of existing technologies is shifting. For example, Direct Attached Storage (DAS) transformed into HA-DAS is moving forward from being on the back burner. The purpose of storage virtualization is not consolidation—as it was in the recent past—but is storage agility as provisioning on demand and where storage capacity is needed at a given moment. As the application landscape is moving from “many applications on one server” to “one application on many servers” deployment, the role of scale-out storage technologies is gaining momentum. An important feature of this transformation of how storage is provisioned and consumed is that new cloud-oriented storage technologies should be seen not as a replacement of the existing technologies and processes but rather as a complementary approach.

DAS Solutions for Cloud

Just DAS

Direct Attached Storage (DAS) means that a server attaches directly to storage system ports without a switch. Internal drives in the server enclosure fall into this category. DAS can also mean external *shared* direct accessible storage when a few servers are connected to the same storage system using SCSI, FC, SAS, or iSCSI protocols. As DAS actually uses a point-to-point connection, it provides high bandwidth between the server and storage system.

While DAS solutions are assumed to be less expensive than SAN solutions, SAN storage with virtual provisioned pools can be more cost-effective.¹² This becomes clearer when all DAS cost factors are considered (including DAS overhead costs such as CPU workload that could be moved to a specialized processor or controller, backup servers, disaster protection infrastructure, many extra spindles for performance, growth, utilization, copies, etc.).

Despite well-known limitations (low scalability, low utilization levels, no replication or tiering solutions, no centralized management), DAS has survived the era of networked storage. While some reasons for this survival may be just inertia and conservatism, others are low cost and acceptable performance. These advantages of DAS led to its renaissance in a new form enriched with new functionality such as high-availability DAS.

As an example, DAS-based storage solutions are offered by Rackspace¹³ and other hosting providers. A DAS-managed storage solution is recommended for applications requiring a lower-cost, entry-level cluster to maintain availability.¹³

High-Availability DAS

Other examples of the DAS comeback are the storage solutions implemented by Facebook, Google, and others. These systems consolidate server and storage in a single node connected to other nodes with storage accessible from every node. Writing data locally and later copy/sync them to other nodes is acceptable for BASE-type applications. Such a shared high-availability DAS (HA-DAS) model provides the simplicity and cost-effectiveness of local storage combined with some beneficial features of SAN solutions as shown below. Use of inexpensive commodity storage hardware along with storage management implemented at higher software levels of the storage technology stack creates a new paradigm that is called software-defined storage which will be discussed in detail later.

An example of HA-DAS is storage clustering, in which controller-to-controller connectivity is provided through a redundant SAS connect. Servers share storage and can balance application loads and provide failover capability in case of server failure. HA-DAS that is easy to deploy offers the features of traditional high-availability (HA) with the benefits of Direct Attached Storage and meets the cost and availability needs of data centers and small/medium businesses.¹⁴

For example, the Nutanix Complete Cluster consolidates DAS with compute resources in four-node Intel-based appliances called “Compute + Storage Together”.¹⁵ The internal storage—a combination of PCIe-SSD (Fusion-io) and SATA hard disks from all nodes—is virtualized into a unified pool by Nutanix Scale-out Converged Storage and can be dynamically allocated to any virtual machine or guest operating system. A Nutanix Controller Virtual Machine (VM) on each host manages storage for virtual machines on the host. Controller VMs work together to manage storage across the cluster as a pool using the Nutanix Distributed File System (NDFS).

Distributed Host-Cache

Using DAS as an extension to SAN, storage systems can intelligently pre-stage the most active data within the PCIe flash cards. As a comeback of DAS, EMC has recently introduced a server-based caching solution: VFCache product is a PCIe generation two-flash card installed in a server and supporting software in the form of an installable device driver (EMC Project Lighting). VFCache accelerates reads and protects data by using write-through cache to the networked-storage. VFCache can also be used as a DAS device storing the application data (“the split-card” feature). VFCache offers the following advantages over a traditional DAS device:

- It can provide performance combined with data protection as data can be stored on an array in the backend, in contrast to a traditional DAS server-based device.
- While DAS solutions are limited by the size of the installed flash capacity, VFCache integration with the backend storage systems allows it to adapt to different workloads and data sets by dynamically moving data to/from the array.
- VFCache data can be centrally managed.

Adoption of SSDs using the PCI Express bus for enterprise computing applications has accelerated after the release of NVM Express, NVMe, or Non-Volatile Memory Host Controller Interface Specification for accessing SSDs on a PCI Express bus. While previous solutions required proprietary software drivers that complicate OEM qualification, standardizing the interface of the SSDs means that operating systems only need one driver to work with all SSDs adhering to the specification. It also means that SSD manufacturers can reduce the production costs because they do not have to design specific interface drivers.¹⁶

SAN Solutions for Cloud

Transformation of SAN Solutions Driven by Cloud Requirements

Since 2000, SAN has quickly become the mainstream storage solution because of its advantages over DAS solutions: higher storage utilization, scalability, centralized management, high-availability, performance, data mobility (SAN and storage array-based migrations), and storage system-based replication. Now we will review how SAN solutions evolve to be able to offer the features required for cloud computing.

Resource Pooling and Rapid Elasticity. Implementation of storage pools and ability to increase or shrink the pool size transparently to the users makes storage resource pooling possible. Storage pools can be created not only within a storage array but also between arrays. For example, using EMC Federated Tiered Storage, storage pools can span across EMC VMAX frames and other EMC arrays including third-party storage systems. Storage resource pools can be expanded beyond a single site with VPLEX[®] Metro and VPLEX Geo. Implementing storage auto-tiering along with virtual provisioning such as EMC FAST[®] VP allows for delivering high performance and, at the same time, cost-effective storage solutions.

On-demand self-service can be provided using an actionable storage service catalog.

Measured service can be implemented with charge-back or show-back options. Storage utilization can be reported back to the users using SRM tools such as EMC ControlCenter[®] Storage Scope[®] or ProSphere[®].

“Cloudization” of Existing Storage Solutions

Many enterprises plan to implement private cloud solutions first, then create a hybrid cloud as the next step.^{1,17} Use of existing SAN solutions can continue in private cloud environments but they should be augmented with cloud-aware management tool sets to provide automation and orchestration for storage service delivery. Virtual storage requires implementation of dynamic pools of compute and storage; EMC storage solutions such as VMAX and VNX lead the market (Fig. 1).¹⁷




Figure 1: Market Shares of Storage Vendors for Virtualized Server Environments.¹⁷

VCE: Integrated Infrastructure-as-a-Service Solution

One of the challenges of private cloud implementations is a seamless integration of storage, compute, and network resources in a scalable solution. In 2009, Cisco and EMC, with investments from VMware and Intel, formed VCE (Virtual Computing Environment) coalition with the goal to develop products enabling businesses to transform traditional data centers into virtual data centers, following up with transition to cloud computing.

Offered by VCE as an integrated and tested unit of virtualized infrastructure¹⁸, Vblock[®] systems provide the integration of server virtualization, network, compute, storage, security, and management products along with predictable performance and operational characteristics. Unified support and end-to-end vendor accountability are Vblock features that make this solution very attractive to customers. There are a few base Vblock configurations:

- Vblock System 300 is a mid-sized configuration with the use cases including e-mail, file and print services, and virtual desktops. It includes high density, compact fabric switches (Cisco Nexus 5548 Series IP switches, Cisco MDS 9148 Series storage switches (Fibre Channel)), integrated fabric-based blade servers (Cisco UCS 5108 blade server chassis and Cisco UCS B-Series blades), and unified storage (EMC VNX Series unified storage with Unisphere®).
- Vblock System 700 is an enterprise-level infrastructure package and its use cases include business-critical ERP and CRM applications. System 700 consists of EMC Symmetrix® VMAX 10K or 20K storage component with optional EMC VNX VG2 Gateway or EMC VNX VG8 Gateway, Cisco Nexus 1000V VSM and VEM virtual switch, Cisco MDS 9148 Multilayer fabric switch, and Cisco UCS 5108 blade server chassis with Cisco UCS B-Series blades.

The use of Vblock solutions for SAP leads to significant CapEx and OpeX savings as a result of dynamic provisioning of compute resources, simple resizing and reconfiguration, integrated consistent backups, reduction of manual labor in daily IT operations, storage-assisted provisioning of multi-tiered/system SAP application landscapes within minutes, and without any post-provisioning configuration.¹⁸

| Areas of TCO Savings | Savings Resulted from Using Vblock Platform |
|---------------------------------|---|
| Hardware/Software Investment | 20% |
| Implementation | 30% |
| Hardware/Software Ongoing Cost | 2%-5% |
| Operations | 5% |
| Continuous Improvement Projects | 20% |
| Upgrade Projects | 30% |

Table 1: Savings in SAP implementations Using Vblock¹⁷

EMC Unified Infrastructure Manager simplifies management of the Vblock Infrastructure by providing¹⁹

- Provisioning
- Configuration Management
- Compliance Management

Unified Infrastructure Manager




Figure 2: Unified Infrastructure Manager for VCE¹⁹

The cloud solutions offered by CSC—CloudCompute and BizCloud—are examples of using Vblock in implementing cloud-based services.²⁰ CSC operates 41 data centers and provides both off-premises and on-premises solutions to companies worldwide, offering CloudCompute as Infrastructure-as-a-Service (IaaS) in all CSC global cloud data centers. CloudCompute serves as the IaaS layer of CSC BizCloud, an out-of-the-box private cloud, billed as a service and designed for rapid deployment on a customer's premises.

Scale-up vs. Scale-out Solutions

While SAN-based solutions can be transformed for providing cloud storage services—as we have discussed above—the role of SAN-based storage gradually decreases. A recent EMC review shows a decline in the importance of SAN technologies from 81% to 71% in 12 months (2011-2012).²¹

While deployment of SAN-based solutions has seen a 4% reduction, it is accompanied by 6% growth in deployment of storage technologies for cloud environments. SAN-based technologies are mainly scale-up solutions providing high IOPS per GB and low response time, whereas the scalability requirements for cloud services can be met in a more cost-effective manner by scale-

out storage solutions by adding nodes as the data amount grows (Fig. 3).²² As capacity is added to a scale-up storage system, the increase in the compute power and memory is limited by the configuration maximums. Therefore, the scale-up storage system reaches a peak level of performance at the maxed-out compute and memory configuration and then the performance starts to degrade as more capacity is added.




Figure 3: Preference Domains for Scale-up and Scale-out Storage Architectures²²

Scale-out solutions use object-based storage and NAS scale-out architectures which will be discussed below.

Object-Based Storage

While block-based storage stores data in groups of blocks with a minimal amount of metadata storage with the content, object-based storage stores data as an object with a unique global identifier (128-bit Universally Unique ID (UUID)) that is used for data access or retrieval. The Object-based Storage Device (OSD) is a new disk interface technology being standardized by ANSI T10 technical committee (Fig. 4). Metadata that includes everything needed to manage content is attached to the primary data and is stored contiguously with the object. The object can be any unstructured data, file, or group of files; for example, audio, document, email, images, and video files. By combining metadata with content, objects are never locked to a location, enabling automation and massive scalability required for cloud and big data solutions.

Incorporation of metadata into objects simplifies the use of data management (preservation, retention, and deletion) policies and, therefore, reduces the management overhead. With flat-address-space design, there is no need to use file systems (file systems have an average overhead of 25%) or to manage LUNs and RAID groups. Access to object-based storage is provided using web protocols such as REST and SOAP.

Block-based vs. Object-based Storage Models




Figure 4: Block-Based vs. Object-Based Storage Models (from Ref 11)

Object-based storage was brought to the market first as content-addressed storage (CAS) systems such as EMC Centera[®]. The main goal has been to provide regulatory compliance for archived data. Cloud-oriented object-based systems appeared in 2009 and have become the next generation of object-based storage. These cloud-oriented systems have to support data transfer across wide-area geographies, such as global content distribution of primary storage, and to function as low-cost storage for backup and archiving in the cloud. The new cloud- and Big Data-oriented object-based storage solutions represent the synthesis of new features of object-based storage for cloud through negation of the traditional CAS systems (even their names point to their heritage – for example, CAS^tor offered by Caringo).

The advantages of using object-based storage in cloud infrastructures are exemplified in solutions offered by the biggest cloud service providers such as Amazon and Google since object-based storage can simplify many cloud operations.

The software role becomes critical, as it should automatically recognize and configure new nodes into a wider system to meet data protection and performance requirements, typically set as policies based on SLOs. The software manages resource allocations to meet varying service levels and manage hardware failures.

Atmos

Atmos[®]—EMC's software solution for object-based storage—uses the Infiniflex HW platform to deliver cloud storage services. Atmos is the first multi-petabyte information management solution designed to help automatically manage and optimize the delivery of rich, unstructured information across large-scale, global cloud storage environments. New applications can be introduced to the Atmos cloud without having to specifically tie them to storage systems or locations.

While Atmos is designed to be scalable and globally distributed, Atmos management is easy since it is operated as a single entity. In Atmos, metadata is used to trigger policy that automatically distributes the information across the nodes. As all Atmos internal services run on virtually every node, there is no single point of failure.

Everything in the Atmos system is treated as an object, including tenants, policy, metadata, files, etc. Objects stored in Atmos are divided into two parts. Metadata is stored on disks managed by Meta Data Service (MDS). The metadata is further divided into system metadata including filename, file size, modification date, timestamps, access-control lists, a layout storage object (LSO), and user metadata that include arbitrary key/value pairs. The second part is user data, which is traditional application data such as Word files, text files, movies, and MP3 files. User data is managed by storage service. There are several advantages to storing metadata and user data in different places, as they are managed by different services (MDS and storage service) that can be optimized separately.

The Atmos cabinet contains a combination of front-end nodes to service I/O requests, which communicate with storage nodes where data is maintained. Atmos front-end nodes run the Common Application Program or "CAP", which is an EMC Linux distribution based on a Red Hat kernel. The Atmos software is layered on top of CAP and is considered a closed appliance

model. While primarily sold as an appliance configuration, this same Atmos software can also be deployed within a VMware environment leveraging VMware-supported servers and storage as the underlying hardware resources.¹¹

Data protection is accomplished with data replication on multiple servers. Atmos uses two types of replicas for objects—synchronous and asynchronous—that are defined as part of policies. A given object may have both types of replicas. In general, synchronous replicas are bit-for-bit copies of each other that are identical at any point in time. A successful write acknowledgement may not be returned to the client until writing all synchronous replicas is complete. It results in performance impact directly related to the number of synchronous replicas, their distance from each other, and the type of connections that exist between those locations. Asynchronous replication done at a best-effort level offers some level of disaster recovery while not affecting performance as much as synchronous replication. The system tries to keep asynchronous replicas up to date, but while such updates will be made eventually, there is no guarantee of when they will occur. This may be acceptable for BASE-type applications.

Atmos provides multi-tenancy support, so administrators can more easily provision and meter capacity, bandwidth, and usage across tenants from a single system. Tenants can be provided with self-service access and management of their storage.




Figure 5: Atmos Deployment Options

The physical Atmos platform is a purpose-built appliance consisting of Intel processors and high performance disk drives. Atmos may be deployed in three initial configurations: 120, 240, and 360 spindles (Fig. 5) and can be upgraded incrementally. Each configuration supports 1TB, 2TB, and 3TB drives, so the total capacity of the WS2-360 fully populated with 3TB drives, is 1,080 TB in a single cabinet. There is no limit to how many cabinets may be combined to form a single Atmos storage space. The system does not require or support RAID for customer data. Servers contain a LSI 3801 simple SAS/SATA HBA for customer data. Since there is no caching, there is no potential for data loss.

Access to Atmos includes Web Services APIs, file services interfaces, and packaged offerings as well as traditional file-sharing protocols such as NFS and CIFS. The Atmos Cloud Delivery Platform (ACDP) is optional software that layers orchestration functions on top of Atmos to deliver turnkey Storage-as-a-Service (StaaS). ACDP provides public cloud-like functions through private cloud centralized controls for both enterprises and service providers.

In December 2012 EMC announced new products and improvements around the entire EMC Atmos portfolio, significantly extending Atmos value for enterprises as they transform to a cloud. The company introduced a new Atmos G3-Dense 480, 10 GbE network connectivity and options for 60 disks or 480 disks with 3 TB drives. Atmos is now available in complete, light, and virtual

editions. Virtual Edition (Atmos VE) provides the full functionality of Atmos Complete Edition deployed on any VMware-certified third-party storage, including EMC Celerra®, CLARiiON®, Symmetrix, VNX, VNXe, and Isilon®.

Other vendors of object-based storage products include Amazon S3, NetApp (StorageGRID), Dell (DX Object Storage Platform), Caringo (CAStor), Cleversafe (Dispersed Storage), DataDirect Networks (Web Object Scaler), NEC (Hydrastor), and Amplidata (AmpliStor).

Open Source Solutions and Cloud Computing

There is growing interest toward open-source cloud solutions in the storage industry. The appeal of such solutions is the ability to use open standards to provide interoperability and avoid vendor lock-in. Open source cloud platforms such as OpenStack, CloudStack, and Eucalyptus are gaining momentum with companies that are building private cloud environments and have decided to avoid using proprietary cloud management platforms such as VMware vCloud. Usually these companies have strong development teams in-house or engage consultants for implementing open source cloud solutions.

Storage Solutions in OpenStack

In this review, we consider OpenStack that is supported by large IT vendors (HP, Cisco, IBM, Intel, Dell, and VMware—more than 150 companies in total —Microsoft is actively involved in OpenStack supporting Hyper-V integration) as well as by service providers (Rackspace, which uses it for its Rackspace Cloud offering). OpenStack seems to be the most popular among open source cloud platforms. EMC joined the OpenStack Foundation in December 2012. NASA and Rackspace were the first primary OpenStack supporters.

OpenStack is positioned to become an open standard for cloud implementations with an aim to achieve interoperability and avoid or at least reduce proprietary lock-in. Freely available under the Apache 2.0 license and often referred to in the media as "the Linux of the Cloud", OpenStack officially became an independent non-profit organization in September 2012. It should be noted that interest in other cloud management platforms such as Eucalyptis, Nimbula, and Apache CloudStack is also growing fast.

OpenStack is defined by OpenStack.org as:

“OpenStack is a cloud operating system that controls large pools of compute, storage, and networking resources throughout a data center, all managed through a dashboard that gives administrators control while empowering their users to provision resources through a web interface.”

The OpenStack platform consists of several service components (projects). The most important are:

- Compute services (Nova)
- Object-based storage service (Swift)
- Block-based storage service (Cinder)
- Image service (Glance)
- Network service (Quantum)
- Identity service (Keystone)
- Dashboard self-service portal (Horizon)

Swift (object-based storage service) which is used by cloud providers such as Rackspace (that provided the code), is considered to be the most stable and broadly adopted component of OpenStack. In Swift, data is accessed through an HTTP interface, typically with a REST API. All client data access is done at the user level: the operating system is unaware of the presence of the remote storage system. The OpenStack storage service is implemented using clusters of standardized servers and, therefore, is capable of storing petabytes of data.

The OpenStack object storage architecture includes storage nodes that run Account, Container, and Object services, Proxy node running Proxy services, and Auth node which is an optionally separate node running the Auth service separately from the Proxy services. A zone is a group of nodes isolated as much as possible from other nodes (separate servers, network, power, even geography). Each storage node is configured as a separate zone in the ring that is a set of mappings of OpenStack object storage data to physical devices. A minimum of five zones is recommended. The ring goal is to ensure that every replica is stored in a separate zone. Figure 5 shows a possible OpenStack storage configuration for a minimal installation.²³




Figure 5: OpenStack Object-Based Storage (from Ref. 23)

For OpenStack Object Storage, an external network should connect the outside world to the proxy servers. The storage network (at least 1 Gbps is recommended) is intended to be isolated on a private network or multiple private networks. OpenStack Object Storage currently runs on Ubuntu, RHEL, CentOS, or Fedora; the large scale deployment at Rackspace runs on Ubuntu 10.04 LTS.

As the data access interface is at a low level of abstraction, it is possible to build file-based applications on top of object storage to provide a higher level of abstraction. For example, the OpenStack Image service can be configured to use the Object Storage service as a backend.

OpenStack also provides persistent block-level storage devices for use with OpenStack compute instances. With block-based storage, files are exposed through a low-level computer bus interface, such as SCSI or ATA, accessible over the network. OpenStack block-based storage (Cinder) is synonymous with SAN. In OpenStack, the Nova-volume service that forms part of the Compute service uses iSCSI to expose remote data as a SCSI disk that is attached to the network. As the data is exposed as a physical device, the end user is responsible for creating partitions and formatting the exposed disk device. Block storage volumes are fully integrated into OpenStack Compute and the Dashboard allowing for cloud users to manage their own storage needs. In OpenStack Compute, a device can only be attached to one server at a time so block storage cannot be used to share data across virtual machine instances

concurrently. EMC is working on developing interfaces that are relevant to EMC's storage business (e.g. for OpenStack block storage [Cinder]).

OpenStack has various backup solutions. For example, Rackspace Private Cloud Software v2.0 supports the Folsom release of OpenStack and Amanda Enterprise (commercial edition of Amanda) has been integrated with HP Cloud Object Storage using the open REST-based API of the OpenStack Swift storage service.

OpenStack API modules are compatible with other APIs such as those for Amazon EC2, Xen, and KVM. Tools exist to move virtual servers from OpenStack to Amazon EC2.

Comparison of Object-Based Storage with Features Required for Enterprise Class Storage Service

The hype around OpenStack may make one believe that as OpenStack is open source, it is an open and widely-adopted standard, with broad interoperability and free from commercial interests. However, as concluded by Lydia Leong (Gartner):

“In reality, OpenStack is dominated by vendor interests, where they want customers to adopt their own offerings, potentially to include proprietary lock-in. Some of the participants, notably Rackspace and other service providers are afraid of the growing dominance of AWS in the cloud IaaS market and do not believe that they have the ability to muster, on their own, the engineering resources necessary to successfully compete with AWS at scale, nor do they want to pay an ongoing license fee for a commercial CMP like VMware's vCloud stack. Do not plan the future of your data center with the assumption that OpenStack will be at its core. While it is a promising project, that is no guarantee of future success.”²⁴

Does OpenStack have features required for enterprise-class cloud platform? Is the OpenStack Object Storage service (Swift) an ideal solution for cost effective, scale-out storage? Vendors involved in OpenStack typically anticipate that it will take until at least late 2013 before it reaches a level of stability and maintainability comparable to typical commercial software sold to the enterprise.²⁴

OpenStack provides a basic feature set and is not likely to replace full-featured commercial solutions. The missing key features are monitoring, high availability (HA), business continuity (BC), and integration with existing user management systems (AD and LDAP). Users are encouraged to integrate existing open source tools, such as the HA Pacemaker utility or DRBD (Distributed Replicated Block Device), an HA clustered storage file system but some users may be not comfortable with these tools. Storage management and reporting capabilities also have some gaps. For example, to get the average age or size of objects in a Swift storage container, users have to run individual queries.

While the set of cloud platform features offered by OpenStack will grow over time, it is unlikely, according to Gartner,²⁴ that OpenStack will displace VMware's vCloud suite or Microsoft's System Center, or competing products such as BMC Cloud Lifecycle Management, over the next five years. OpenStack is supposed to be hypervisor-neutral, but in reality it supports KVM or Xen and its support for VMware's vSphere and Microsoft's Hyper-V is inadequate for enterprise deployments. It is also unlikely that OpenStack will lead to the displacement of vSphere or Hyper-V, in favor of KVM or Xen, in the enterprise.²⁴

As discussed above, the selection of a cloud solution depends on the business goals—if implementation of a private cloud is a planned step in transforming a traditional data center to a virtual data center and, eventually, to a cloud computing environment, proprietary solutions such as VMware vCloud are likely to be chosen for mission critical applications. In my opinion, we will see development of interoperability of OpenStack modules and distributions (including commercially supported) with proprietary cloud solutions and as solution choice is driven by business requirements, there will be a broad spectrum of cloud implementations integrating open source and commercial components and users will benefit from more options to choose from. This trend is exemplified by the recent addition of native support for the Amazon S3 API to Atmos 2.1. It enables customers to easily migrate their S3 applications to any of the more than 40 Atmos-powered public clouds worldwide or any internal private cloud they are deploying.

VMware vCloud suite competes with the OpenStack platform. It is positioned by VMware as the foundation for building software-defined data centers having software-defined storage as the key component. Let us briefly review this new concept of software-defined storage.

Software-Defined Storage

I am sure readers will agree with Christos Karamanolis (VMware, Office of the CTO) that 2012 has been the year of the “software-defined data center”.¹⁰ Indeed, this buzzword could be found in many online and print trade magazines and heard at almost every technical briefing on the next generation data center. Christos predicts 2013 to be the year of “software-defined storage” (SDS).

SDS can follow the development and acceptance of “software-defined networking” (SDN) that gained popularity as a component of the software-defined data center. A number of emerging storage vendors (Nexenta, Nutanix, ScaleIO, Zadara, CloudByte, GridStore, Virsto, and others) are bringing SDS products to the market. While these products are based on various technologies, what puts them in the same category is decoupling of storage service features (provisioning, dynamic tiering, cloning, snapshots, replication, deduplication, etc.) from the hardware layer and moving them up to software-based management. This abstraction allows for using commodity hardware in scale-out architecture. The OpenStack Storage Object service (Swift) and Block-Based storage service (Cinder) discussed above are other examples of software-defined storage.

Can legacy storage be converted into software-defined storage? The traditional storage frame has a lot of hardware functionality built in and transformation into an SDS system would require a complete re-design. But what is the goal of this conversion? It hardly makes sense since traditional storage and SDS are like two different animals, each with its own ecosystem as seen in Table 2.

Table 2: Comparison of Traditional Storage and Software Defined Storage

| Features | Traditional Storage Frame | Software-Defined Storage (SDS) |
|------------------------|---|---|
| Reliability | Mature technologies. Stable software. Reliability is built into the storage hardware, software protects against rare cases that are not handled by hardware protection features. | SDS is developed under an assumption that hardware is not necessarily reliable, and the software is responsible to continue providing storage services in case of hardware failures. |
| Performance/ QoS | While implementation of storage pools shared by apps with very diverse I/O profiles makes delivery of predictable performance a challenge, a large installation base and many years of use in various industry verticals have resulted in development of excellent best practice guides and performance tuning processes. Resource partitioning (cache, ports, drives) is used for providing QoS. | Linear scalable performance. By definition, an SDS platform has to incorporate a variety of storage devices of different generations with widely varying performance and reliability characteristics, as they are found in a typical data center. The value of the SDS is in delivering a predictable quality of service for the workloads running on heterogeneous hardware. |
| Scalability | Scale-up architecture. The growth is constrained within the frame. | Scaled-out architecture. Hundreds and thousands of nodes. HA and data redundancy are managed by software. |
| Management /Automation | Proprietary CLI. | HTTP/REST API. Automation for all apps across all types of storage. |
| Interoperability | Limited. Third-party storage support is implemented through storage virtualization (array controller, SAN, appliance-based). | SDS systems use standard x86 hardware and standard OSs; interoperability of those has already been tested by their vendors. |
| Agility | Limited by high acquisition CapEx and procurement lead time. | Elasticity: add, move, remove nodes “on the fly.” |
| TCO | High because of scale-up design and expensive development of specialized HW and SW and the interoperability testing. | Low. Investment protection; can leverage the existing server infrastructure. |

Time and again, SDS can be seen as the evolution of legacy storage virtualization through creative destruction. Traditional block-based storage virtualization solutions are based on a virtualizer running on an array-controller, appliance, or SAN-switch blade. In the SDS realm, a storage virtualizer evolves into a storage hypervisor such as the VM-centric storage hypervisor developed by Virsto and SANsymphony-V from DataCore. The storage hypervisor provides a higher level of software intelligence capable of delivering storage services on commodity hardware without relying on ASIC-built storage functionality.

VMware's vision for SDS was presented at VMworld in 2012 as a series of technology previews. Requirements for SDS include the following capabilities:²⁵

- Can converge with the compute platform
- Can be managed as resource along with CPU and memory
- Can scale on demand in lock-step with application needs
- Can provide per-VM automated SLA management

The previewed technologies include an ability to pool local disks on servers for creating shared storage resources for VMs (Distributed Storage, vSAN), a new method for storage arrays to communicate with virtual or cloud environments (virtual volumes, vVol), and better ways to leverage Flash-based storage. VMware sees SDS delivery implemented through

- Storage Policy-Based Management for automation that spans across all types of storage
- Enabling traditional SAN/NAS systems using Virtual Volume and Virtual Flash frameworks to provide integration with existing tiers of storage as well as to introduce new tiers of storage such as local DAS or Flash (SSD or PCIe card)
- Enabling DAS (Server Disks) using VMware's Distributed Storage technology

According to VMware,²⁵ the use of storage policy-based management, with per-VM settings for capacity, availability, and IOPS will eventually shift the role of the key storage management unit from LUN to VM.

Symbiosis of HA DAS and SDS

ScaleIO ECS is an example of symbiotic HA DAS and SDS solution.²⁶ It is a software-only solution that uses application hosts' local disks to realize a virtual SAN that is comparable to or better than external SAN storage, but at a fraction of the cost and complexity.²⁶ The lightweight software components of ECS are installed on the application hosts alongside applications such as databases and hypervisors, as well as shared-everything applications. ECS natively supports all the leading Linux distributions and hypervisors. The product includes encryption at rest and quality of service (QoS) features.

NAS and Cloud

While cloud computing enters enterprise environments, traditional scale-up NAS solutions continue to be used but are augmented with scale-out NAS systems deployed in private and/or hybrid clouds.

In January 2013, EMC announced beta availability of EMC's Syncplicity cloud-based online file sharing service that offers customers an option to combine cloud-based service with the use of on-premises storage that can be either EMC's Isilon[®] scale-out NAS or Atmos[®] object-based storage.²⁷ This solution gives IT the ability to control where managed files should reside and which users can share them. Regulatory requirements can be met by managing data placement with security policies so that sensitive data always resides on the internal storage. At the same time, users can benefit with a secure solution for file sync and sharing with IT retaining complete control over data and storage resources. Syncplicity enables enforcing of data retention policies that can be automated and be able to delete shared corporate data from user devices if needed.

There are also various so-called Cloud NAS device offerings targeting small businesses and home users; for example, Red Hat Storage Server for Public Cloud, CloudStor from Buffalo Technology, and File Station from Synology. Red Hat Storage Server for Public Cloud is based on scale-out NAS and works with Amazon Web Services enabling deployment of storage in the Amazon Cloud.

Data Mobility in Clouds

Implementing an agile storage infrastructure leveraging clouds require ways to move data between cloud environments. Such data mobility solutions rely on interoperability between traditional storage system and cloud-based storage. Possible solutions include cloud storage gateways that we will consider now.

Cloud Storage Gateways and EMC Cloud Tiering Appliance

Cloud storage gateways link the past and future of storage by providing integration between traditional storage and emerging cloud-oriented storage systems. The main thread of this article is that the interactions between developments of traditional technologies and new ones represent creation of a new “synthesized” ecosystem evolving in a spiral-form progress through contradiction and negation of preceding technologies. Instead of replacing traditional storage with new technologies, these two technology “generations” will co-exist and complement each other by allowing for data placement according to the data value and I/O profile.

To move data between these technology domains requires cloud storage gateways. As standardized cloud storage APIs are not yet available, cloud storage gateways are deployed to function as protocol translators to avoid rewriting applications to support Web service interfaces. Since traditional NAS systems use NFS and CIFS protocols and cloud-based object storage services use web-service APIs, storage gateways play the role of translator.

Vendors such as EMC, Riverbed, CTERA, Nasuni, Panzura, StorSimple, TwinStrata, and Zetta have developed cloud gateways either as hardware appliances or software that customers can install on their on-premises devices. Recently, Amazon announced its software-based AWS Storage Gateway but it is applicable only to Amazon services. Nirvanix and Rackspace also offer file-based gateways. Many of these gateways offer both block-access (iSCSI) and file-access (SMB/SMB2/CIFS) interfaces to cloud storage, intelligent caching architecture that delivers local performance, data compression/deduplication, at-rest/in-flight data encryption, and bandwidth optimization and scheduling. They can be deployed as a virtual storage appliance (VSA) or as physical hardware appliance.

Use cases for cloud storage gateways include extending Information Lifecycle Management (ILM) processes onto the cloud environment by implementing tiering and archiving, remote office backups, and balancing I/O workload for unstructured data (file services) in hybrid clouds by moving data to/from public cloud.

EMC Cloud Tiering Appliance (CTA), announced at EMC World 2012, is an example of a storage gateway. It can move files from NetApp arrays and Celerra® arrays to VNX® and Isilon, and from VNX to Atmos® and Data Domain®. The CTA can also send files from VNX to Data Domain or to Atmos storage.

Cloud storage can be used for data archiving and CTA uses automated file tiering to transparently move infrequently accessed files to Atmos or service provider clouds and thereby reclaim primary storage space. CTA delivered as an appliance or software installable as a virtual machine is deployed non-disruptively and the archiving process operates out-of-band. When files are moved to Atmos or Data Domain, stubs are left behind to be used for recall. As there is no persistent metadata that controls the location of files, no database must be maintained.

Some gateways use WAN-optimization techniques to make the data transfer between the data center and the cloud more efficient. EMC and Riverbed offer a joint solution consisting of Atmos as a cloud-optimized storage platform and the Whitewater cloud storage gateway. Whitewater gateways reduce data sets by 10 to 30 times on average, substantially reducing cloud storage costs while accelerating data transfers by using deduplication, compression, and optimization technologies.

In most cases today, network connections to cloud storage cannot provide low latency and high-bandwidth access. Consequently, cloud-based storage cannot compete against enterprise class storage and storage gateways should not be used for mission-critical or very high-transaction-rate applications. However, storage gateways can be suitable for latency-tolerant low IOPS applications. Challenges for integrated traditional/cloud storage environments include regulatory compliance and storage resource management that should be able to support such hybrid systems comprised of two very different components.

Big Data Meets Cloud Data

There are various definitions for “Big Data”. Sometimes Big Data is defined as the data that cannot be managed using existing technologies: applications and infrastructure. Therefore, new applications and new infrastructure for those new applications as well as new processes and procedures are required to use Big Data. The storage infrastructure for Big Data applications should be capable of managing large data sets and providing required performance. Development of new storage solutions should address the following characteristics of Big Data.²⁸

- Huge volume of data (for instance, billions of rows and millions of columns)
- Speed or velocity of new data creation
- Variety of data created (structured, semi-structured, unstructured)

- Complexity of data types and structures, with an increasing volume of unstructured data (80%-90% of the data in existence is unstructured)

These Big Data characteristics determine the selection of storage solutions to manage them. The massive volume of data are in various formats with a wide range of data generation and processing velocities and access latencies. As discussed above, large data volume and the variety of data types and complexity are common features that we also see in cloud storage. Therefore, storage architectures are the place where Big Data meets Cloud Data and the storage solutions for the cloud described above can apply to Big Data storage:

- Object-based storage
- DAS and HA-DAS
- Scale-out NAS with single Namespace
- NAS gateway to SAN and Scale-out SAN

The Apache Hadoop platform,²⁹ an open-source software framework supporting data-intensive distributed applications, has two core components: the Hadoop Distributed File System (HDFS) which manages massive unstructured data storage on commodity hardware and MapReduce which provides various functions to access the data on HDFS. The storage function consists of HDFS that provides a redundant, distributed file system optimized for large files. MapReduce consists of a Java API as well as software to implement the services that Hadoop needs to function. Hadoop integrates the storage and analytics in a framework that provides reliability, scalability, and management of the data.

Hadoop supports four different node types. The NameNode and the DataNode are part of the HDFS implementation. Apache Hadoop has one NameNode and multiple DataNodes. The NameNode manages that name space by determining which DataNode contains the data requested by the client and redirects the client to that particular DataNode.


Hadoop is “rack aware”—that is, the NameNode utilizes a data structure that determines what DataNode is preferred based on the “network distance” between them. Nodes that are “closer” are preferred (same rack, different rack, same data center). HDFS uses this when replicating data, to try to keep different copies of the data on different racks. The goal is to reduce the impact of a rack power outage or switch failure so that even if these events occur, the data may still be readable.

HDFS uses “shared nothing” architecture for primary storage—all the nodes have direct attached SAS or SATA disks (DAS). No storage is shared as the disks are locally attached and there are no disks attached to two or more nodes. The default way to store data for Hadoop is HDFS on local direct attached disks. However, it can be seen as HDFS on HA-DAS because the data is replicated across nodes for HA purposes. Compute nodes are distributed FS clients if scale-out NAS servers are used. Pros and cons for various storage options^{30,31} for HDFS are presented in Table 3.

| | Pros | Cons |
|--|--|---|
| DAS | Writes are highly parallel and tuned for Hadoop jobs; Job Tracker tries to make local reads. | High replication cost compared with shared storage. NodeName keeping track of data location is still a SPOF (can be addressed in dispersed storage solutions [Cleversafe]). Scalability bottleneck, as “everything has to be in memory” |
| SAN | Array capabilities (redundancy, replication, dynamic tiering, virtual provisioning) can be leveraged. As the storage is shared, a new node can be easily assigned to a failed-node data. Centralized management. Using shared storage eliminates or reduces the need for three-way data replication between datanodes. | Cost, limited scalability of scale-up storage arrays. |
| Distributed File System /Scale-out NAS | Shared data access, POSIX-compatible, and works for non-Hadoop apps just as a local file system, centralized management and administration. | While HDFS is highly optimized for Hadoop, it is not likely to get the same level of optimization for a general Distributed File System (DFS). Strict POSIX compliance leads to unnecessary serialization. Scaling limitations, as some DFSs are not designed for thousands of nodes. |

Table 3: HDFS Storage Options

A tightly coupled Distributed File System (DFS) for Hadoop is a general purpose shared file system implemented in the kernel with a single name space.³¹ Local awareness is part of the DFS—no need for NameNode. Compute nodes may or may not have local storage. Remote storage is accessed using a file system-specific internode protocol. If DFS uses local disks, compute nodes are part of DFS with data spread across nodes.



Source: Gartner (February 2012)

Figure 5: Technologies Evaluated or Being Deployed to Meet Big Data Requirements.³²

Figure 5 presents technologies in the priority order in which companies are evaluating them or have deployed them to meet Big Data requirements.³² EMC Isilon is an example of shared storage as primary storage for Big Data Analytics whereas EMC Data Domain/VMAX and Greenplum[®] exemplify shared storage as secondary storage for Big Data Analytics. A Big Data “stack” like the EMC Big Data Stack presented in Table 4 needs to be able to operate at multi-petabyte scale, handle structured and unstructured data, and be increasingly collaborative across the enterprise.

| Technology Layer | EMC Product |
|------------------------------|-------------------------------------|
| Collaborative - Act | Documentum xCP, Greenplum Chorus |
| Real Time - Analyze | Greenplum + Hadoop. |
| Structured/unstructured data | |
| Storage, Petabyte Scale | VMAX, Isilon, Atmos |

Table 4: EMC Big Data Stack

Isilon Storage for Big Data

Isilon is an enterprise storage system that can natively integrate with the HDFS.³³ EMC Isilon storage uses intelligent software to scale data across vast quantities of commodity hardware, enabling explosive growth in performance and capacity. Isilon OneFS[®] uses the InfiniBand back-end network to allocate and stripe data across all nodes in the cluster automatically. OneFS eliminates single point of failure (see Table 5) by distributing the Hadoop NameNode to provide high-availability and load balancing. Isilon NAS storage provides a single file system/single volume scalable up to 15 PB.³³ Data can also be staged from other protocols to HDFS by using OneFS as a staging gateway.

The benefits of using Isilon for Big Data are shown in Table 5.

| Hadoop/DAS Challenges | Hadoop with Isilon Solutions to Address Them |
|--|---|
| Dedicated storage infrastructure. One-off for Hadoop only. | Scale-out storage platform. Multiple applications and workflows. |
| Single-point of failure | No single point of failure – Distributed namespace |
| Lack of enterprise class data protection – no snapshots, backup, replication | End-to-end data protection – SnapshotIQ, SyncIQ, NDMP backup |
| Poor storage efficiency – three-way mirroring | Storage efficiency > 80% storage utilization |
| Manual import/export | Multi-protocol support. Industry standard protocols: NFS, CIFS, FTP, HTTP, HDFS |
| Fixed scalability – rigid compute to storage ratio | Independent scalability: decoupling compute and storage – add compute and storage independently |

Table 5: Benefits of Using Isilon for Hadoop

EMC Greenplum Distributed Computing Appliance (DCA)

Combining Isilon and Greenplum HD provides the best of both worlds for Big Data Analytics. Greenplum Database[™] (GPDB) with Hadoop delivers a solution for analyzing structured, semi-structured, and unstructured data.³⁴ The Greenplum DCA is a massively parallel processing (MPP) architecture and the GPDB is the industry’s most scalable analytic database. It features “shared nothing”, in contrast to Oracle and DB2. Operations are extremely simple—once data is loaded, Greenplum’s automated parallelization and tuning provide the rest: no partitioning is

required. To scale, simply add nodes (Greenplum DCA fully leverages the industry standard x86 platform); storage, performance, and load bandwidth are managed entirely in software.

Users can perform complex, high-speed, interactive analytics using GPDB, as well as streaming the data directly from Hadoop into GPDB to incorporate unstructured or semi-structured data in the above analyses within GPDB. Hadoop also can be used to transform unstructured and semi-structured data into a structured format that can then be fed into GPDB for high speed, interactive querying.³⁴

Object-Based Storage for Big Data

Cleversafe announced plans to build the Dispersed Compute Storage solution by combining the power of Hadoop MapReduce with Cleversafe's Dispersed Storage System.³⁵ The object-based Dispersed Storage systems will be able to capture data at 1 TB per second at Exabyte capacity. Combining MapReduce with the Dispersed Storage Network (dsNet) system on the same platform and replacing HDFS—which relies on three copies to protect data—will significantly improve reliability and allow analytics at a scale previously unattainable through traditional HDFS configurations.

Fabric Storage for Big Data: SAN Functionality at DAS Pricing

Scale-out Fabric storage offered by AMD SeaMicro as a Big Data storage solution provides massive scale-out capacity with commodity drives.³⁶ Decoupling from Compute and Network to grow storage independently enables moving from DAS with a rigid storage-to-compute ratio to flexible scale-out fabric storage up to 5 PB.

According to AMD SeaMicro,³⁶ the SM15000 Server Platform is optimized for Big Data and the cloud can reduce power dissipation by half and be able to supply SAN functionality at DAS pricing by coupling data storage through a "Freedom Fabric" switch that removes the constraints of traditional servers. Unlike the industry-standard model, where disk storage is located remotely from processing nodes, SeaMicro has worked out a networking switched fabric that connects servers to the "in rack" disk drives and is extensible beyond the SM15000 rack frame allowing construction of cumulatively very large systems.

Conclusions

Just as DAS survived networked storage, likewise SAN will survive and get new features and use cases in the era of cloud computing and Big Data—negation of negation leads to synthesis of a new phenomenon with a new level of enhanced functionality.

When we map cloud and Big Data requirements to storage solutions, it is likely we will end up having a set of solutions running the gamut from “traditional SAN” to “distributed storage” to cloud- or Big Data-optimized storage. Extremes like “No SAN” slogans should be avoided.

Cloud computing and Big Data will affect infrastructure design, storage service road maps, and budget priorities. We as users need to work with both established storage vendors and emerging vendors bringing innovative technology solutions to the market so that we can understand their cloud computing and Big Data solution road maps and develop our own storage service strategy to meet dynamic business requirements.

References

1. G. Schulz. Cloud and Virtual Data Storage Networking. CRC Press, Boca Raton. 2012.
2. IT-as-a-Service. Planning and Design. EMC Educational Course for Cloud Architect Experts. EMC, 2011.
3. P. Goyal. Big Data: The Big Challenge - A Paper For Every CIO. EMC Proven Professional Knowledge Sharing, 2012.
4. A. Reichman, S. Balaouras, C. Kane, B. Echols, L. E. Nelson, and A. Crumb. Do You Really Need A SAN Anymore? Forrester Research, February 2009.
5. B. Russell. A History of Western Philosophy. Simon & Schuster, N.Y., 1972.
6. E. Pol, P. Carroll. An Introduction to the Creative Economy (2nd edition). McGraw-Hill, Australia, 2007.
7. P. Mell, T. Grance. The NIST Definition of Cloud Computing (NIST Special Publication 800-145), September 2011.
8. <http://cdmi.sniacloud.com>
9. <http://www.vmware.com/solutions/datacenter/software-defined-datacenter/index.html>
10. C. Karamanolis. "2013 predictions: The year of software-defined storage?" VMware Chief Technology Office. <http://cto.vmware.com/2013-predictions-the-year-of-software-defined-storage>.
11. Virtualized Data Center and Cloud Infrastructure. EMC Educational Course for Cloud Architects, EMC 2011.
12. D. Merrill. Is DAS really cheaper than SAN? HDS, 2009.
13. http://www.rackspace.com/managed_hosting/services/storage/das/
14. T. Gibbons. High Availability-DAS Continuous Uptime Meets Volume Server Deployment. SNW, October 2012.
15. <http://www.nutanix.com/products.html>
16. <http://www.nvmexpress.org>
17. P. Gelsinger. The Perfect Storm That Makes IT Transformation Possible. EMC IT Leadership Council, 2011.

18. Vblock Infrastructure Platforms in the Real World. VCE. EMC World, 2012.
19. B. Klotzman, F.-S. Svingen, Introduction to EMC UIM Provisioning and EMC UIM Operations. EMC World, 2012.
20. Vblock Drives Enterprise Cloud Growth for CSC. CSC Case Study.
<http://www.vce.com/products/vblock/system-700>.
21. A. Shrivastava. Managing Information Storage – Trends, Challenges, and Options 2012-2013. EMC World, 2012.
22. K. Ghosh. Scale-out and Scale-up Storage Architectures. EMC World, 2012.
23. OpenStack Install and Deploy Manual. Red Hat Ubuntu. November 2012.
24. L. Leong. Don't Let OpenStack Hype Distort Your Selection of a Cloud Management Platform in 2012. Gartner. September 2012.
25. S. Bhat, A. Farronato. Software-Defined Storage - VMware's Storage Vision & Strategy. INF-STO2143, VMWorld 2012.
26. <http://www.scaleio.com/>
27. <http://www.syncplicity.com/company/press/syncplicity-on-premise-storage-isilon-atmos.aspx>
28. Data Science and Big Data Analytics. Student Guide. EMC, 2012.
29. <http://wiki.apache.org/hadoop/ProjectDescription>.
30. J. Webster. Storage for Hadoop: A Four-Stage Model. SNW, October 2012.
31. S. Fineberg. Big Data Storage Options for Hadoop. SNW, October 2012.
32. S. Childs, M. Adrian. Big Data Challenges for the IT Infrastructure Team. Gartner, 2012.
33. Hadoop on EMC Isilon Scale-out NAS. EMC White paper H10528. EMC, 2012.
34. EMC Big Data Storage and Analytics Solution. Solution Overview. EMC, 2012.
35. <http://www.cleversafe.com/overview/how-cleversafe-works>

36. S. Nanniyur Fabric Architecture: A Big Idea for the Big Data infrastructure. SNW, April 2012.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.