

Internet Engineering Task Force
Internet-Draft
CDNI Working Group
Intended status: Informational
Expires: Aug 05, 2016

Eui-Nam Huh
Kyung Hee University
Ga-Won Lee
Kyung Hee University
Yunkon Kim
Kyung Hee University
Jintaek Kim
Consortium of Cloud Computing Research
Feb 04, 2016

Software-Defined Storage Definition and Overview

Abstract

In accordance with rapid increase of data related to IoT and big data, techniques to control high capacity data is currently active and vibrant research field. Enterprises are trying to manage data on the cloud because of flexibility and capability. However, there are some limits to handle data intelligently in cloud. The SDS is considered as a good technique regarding this manner. SDS improves efficiency, scalability and flexibility in scale-out architecture, as well as, provides a cost effective solution using the existing storage resources efficiently.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>. Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress." This Internet-Draft will expire on August 05, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Related Work	2
2.1 Survey of existing technologies for SDS	2
2.2 Standardization activities to SDS	4
3. SDS	6
3.1 Definition of SDS	6
3.2 Overview of SDS	8
4. References	10

1. Introduction

IoT, Big Data and mobile paradigms are leading us to data-centric computing society. One of the objectives of data-centric computing paradigm is to realize ‘Data-as-a-Service’ in cloud storage perspective. A storage federation is good candidate technology to realize these paradigms. Software Defined Storage (SDS) is one of storage and data federation technology.

The SDS is a software-based technology that detaches the storage management from the physical storage and transforms it into a service. The SDS technology decreases storage management complexities by providing an automated and centralized management service to administer. Furthermore, the SDS improves efficiency, scalability and flexibility in scale-out architecture, as well as, provides a cost effective solution using the existing storage resources efficiently. Therefore, potential of storage and data federation on SDS is compelling.

The purpose of this contribution is to launch a new work item of storage/data federation by describing the overview, requirements and capabilities for SDS. This contribution will support to form a logical storage pool, to manage it by software, and finally to federate storage/data logically.

Through this contribution, various derivation, extension and combination of services can be created. Moreover, system interoperability improvement can give us new demand creation between different domain industries. Therefore, this contribution can contribute to invigoration of big data industry and expected to lead cloud and big data markets.

2. Related Work

2.1 Survey of existing technologies for SDS

2.1.1 HP StoreVirtual VSA

Transform your server’s internal or direct-attached storage into a fully featured shared storage array—without the cost and complexity associated with dedicated storage.

2.1.2 ViPR

EMC provides intelligent SDS solutions that help organizations drastically reduce management overhead through automation across traditional storage silos and pave the way for rapid deployment of fully integrated next generation scale-out storage architectures.

2.1.3 IBM Spectrum Scale (GPFS)

IBM Spectrum Scale is a proven, scalable, high-performance data and file management solution (based upon IBM General Parallel File System or GPFS technology, also formerly known as code name Elastic Storage) that's being used extensively across multiple industries worldwide.

2.1.4 GlusterFS

GlusterFS is a scalable network file system suitable for data-intensive tasks such as cloud storage and media streaming. GlusterFS is free and open source software and can utilize common off-the-shelf hardware.

2.1.5 Swiftstack

SwiftStack provides an enterprise-grade object storage system and an innovative storage controller that makes it simple for you to deploy, integrate and manage object storage clusters in your data centers.

2.1.6 NexentaStor

NexentaStor is our flagship Open Source-driven SDS (OpenSDS) platform, allowing thousands of customers all around the world to evolve their storage infrastructure, increase flexibility and agility, simplify management and dramatically reduce costs without compromising on availability, reliability or functionality.

2.1.7 SUSE Enterprise Storage

SUSE® Enterprise Storage is a highly scalable and resilient software-based storage solution, powered by Ceph. It enables organizations to build cost-efficient and highly scalable storage using commodity, off-the-shelf servers and disk drives.

2.1.8 yStor

yStor is a software-defined solution that allows you to build an enterprise-ready highly-scalable and distributed storage platform. yStor provides elastic provisioning and unmatched flexibility without the need for additional licenses, SAN hardware, or expensive infrastructure components

2.1.9 SANsymphony-V

SANsymphony-V10 software is a comprehensive and scalable storage services platform designed to maximize the performance, availability and utilization of your IT assets, no matter how diverse they may be, or what topology you've chosen.

2.1.10 VmWare (Virtual SAN)

VMware Virtual SAN is a radically simple, enterprise-class shared storage solution for hyper-converged infrastructure optimized for vSphere virtual machines.

2.1.11 Maxta (MxSP)

The Maxta Storage Platform (MxSP) provides organizations the choice to hyper-converge on any x86 server, the ability to run on any compute abstraction layer, and the flexibility to support any combination of storage devices eliminating the need for complex and expensive NAS and SAN devices.

2.1.12 Scality RING

The Scality RING is a proven software storage solution that enables customers to build petabyte scale storage infrastructures leveraging industry-standard servers.

2.1.13 Solution Comparison

Table 1 SDS Solution Features Comparison

Solutions	Automated Policy-based	Centralized	Heterogeneous hardware	Storage Type	Scale-in/out	Elastic /Resilience	Self-service	Tiered
HP StoreVirtual VSA		✓	✓	F, O	✓	✓		
ViPR	✓	✓		O, B	✓	✓	✓	
IBM Spectrum Scale (GPFS)	✓	✓	✓	F, O, B	✓	✓		✓
GlusterFS		✓	✓	B	✓	✓		
Swiftstack		✓	✓	F, O	✓	✓		
NexentaStor		✓	✓	F, B	✓	✓		
SUSE Enterprise Storage	✓		✓	O, B	✓	✓	✓	
yStor		✓		B	✓	✓		✓
SANsymphony-V		✓	✓	B	✓	✓		✓
VmWare (Virtual SAN)	✓	✓	✓	B	✓	✓	✓	
Maxta (MxSP)	✓	✓	✓	F, B	✓			
Scality RING		✓	✓	F, O	✓	✓		

2.2 Standardization activities to SDS

The standardization activities in various organizations are at an early stage. They are establishing SDS related working groups and define SDS technology. In particular,

open-source based storage technologies, integrated management tool, interoperability standardization discussion is active in many organizations.

This section provides various standardization activities related to SDS, which is brief summary of each organization's description.

2.2.1 DMTF (Distributed Management Task Force, Inc.)

Cloud Management Initiative (CLOUD) is working to address management interoperability for cloud systems. The following DMTF working groups produce the SDS related standards and technologies promoted by the Cloud Management Initiative:

- Cloud Management Working Group (CMWG) has developed DMTF specification entitled "Cloud Infrastructure Management Interface (CIMI)." The CIMI specification describes the model and protocol for management interactions between a Cloud Infrastructure as a Service (IaaS) provider and the consumers of an IaaS service. The basic resources of IaaS (machines, storage, and networks) are modeled to provide consumer management access to an implementation of IaaS and facilitate portability between cloud implementations that support the specification.
- Cloud Auditing Data Federation Working Group (CADF) defines the CADF standard, a full event model anyone can use to fill in the essential data needed to certify, self-manage and self-audit application security in cloud environments.
- Open Virtualization Working Group (OVF) produces the Open Virtualization Format (OVF) standard, which provides the industry with a standard packaging format for software solutions based on virtual systems.

Open Software Defined Data Center Incubator (OSDDC) is aim to develop standard architectures and definitions to describe the Software Defined Data Center (SDDC). The incubator is developing SDDC usecases, reference architectures and requirements for industry standardization.

In addition, various related activities are ongoing such as Common Information Model (CIM), Configuration Management Database Federation (CMDBf), Systems Management Architecture for Server Hardware (SMASH), etc.

2.2.2 SNIA (Advanced Storage and Information Technology)

The standardization activities in SNIA are mainly performed in the Cloud Storage Technical Work Group. The Cloud Storage TWG acts as the primary technical entity for the SNIA to identify, develop, and coordinate systems standards for Cloud Storage. This group aims to produce a comprehensive set of specifications and drives consistency of interface standards and messages across the various Cloud Storage related efforts. Representatively, Cloud Storage TWG promotes cloud storage adoption with open standards such as "Cloud Data Management Interface (CDMI)". CDMI is an ISO/IEC standard that enables cloud solution vendors to meet the growing need of interoperability for data stored in the cloud. There are currently more than 20 products that meet the CDMI specification. SDS requires

a standardized storage management interface, such as “Storage Management Initiative Specification (SMI-S)” developed by Storage Management Initiative Specification (SMI-S) Core TWG, in order to automate the management of the storage resources and discovery of their capabilities for use in various pools. Besides, Object Drive TWG, Disk Resource Management TWG, etc. are collaborating for Software Defined Data Center realization.

2.2.3 OASIS (Organization for the Advancement of Structured Information Standards)

OASIS Topology and Orchestration Specification for Cloud Applications (TOSCA) Technical Committee is enhancing the portability and management of cloud applications and services across their lifecycle. TOSCA standards aim to enable Software Defined Environments (SDEs) by optimizing the underlying cloud infrastructure.

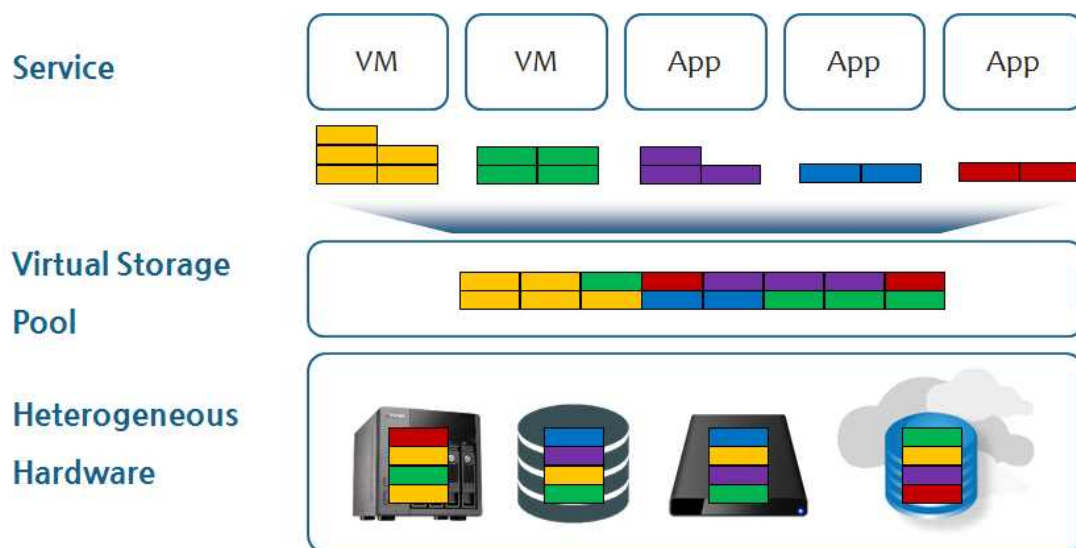
Cloud Application Management for Platforms (CAMP) Technical Committee defines interfaces for self-service provisioning, monitoring, and control. Based on REST, CAMP is expected to foster an ecosystem of common tools, plugins, libraries and frameworks, which will allow vendors to offer greater value-add.

3. SDS

3.1 Definition of SDS

3.1.1 Definition

The SDS can be a software-based model that detaches the storage management from the physical storage hardware and transforms it into a service. The SDS as a logical storage pool can provide an automated and centralized management service to administrator. The SDS might also decrease storage management complexities. Moreover, the SDS is able to improve efficiency, scalability and flexibility using commodity hardware as well as to provide a cost effective technology using existing storage resource. In this SDS environment, users can request to allocate storage to applications by their requirements.



3.1.2 Objectives

The objectives of the SDS are to improve efficiency and to reduce wasting cost of storage. A facility cost was a problem of traditional storage environment. Service providers had to construct big storage service to provide service because they couldn't expect how much storage space was exactly needed. Additionally, the traditional storage had some problem such as vendor lock-in, difficulty of scale out storage, and so on. On the other hand, the SDS can provide storage to application unit and be managed by single control point. It can support block, file, and object type storage using existing hardware infrastructure. Moreover, it is possible to scale out storage nodes with system disruption.

Table 2 Traditional Storage and SDS Comparison

Features	Traditional Storage	Software Defined Storage
Scale up	Deploy entire new shelves of disks	Add new shelves of disks
Lock-in	Use specialized hardware	Use heterogeneous hardware
Service unit	System component unit (Shared Container)	Application component unit (Individual Container)
Flexibility	Static	Dynamic
Management	Storage administrators' intervention	Single control point

3.1.3 Benefits

- Simplicity
 - Automated policy-based: To store data in the right place, at the right time, with the right performance, and at the right cost based on defined policies.
 - Centralized management: To convert the whole storage hardware in a seamless storage pool and offer a single control point to manage all the resources
- Flexibility
 - Heterogeneous hardware: To allow the use of commodity hardware and the implementation over an existing infrastructure
 - Supporting block, file and object storage: Integration of the three basic storage types
- Scalability
 - Scale-out architecture: To incorporate more storage nodes to increase the capacity and improve the performance
 - Elastic: Enable increase capacity as needed without disruption the availability or performance
- Efficiency
 - Component unit: To make traditional, big storage pool into application unit storage to reduce wasting expense

3.2 Overview of SDS

3.2.1 Key Characteristics

3.2.1.1 Automated Policy-Based Management

- Data Provisioning: Provide users with access to data and resources.

- Data Protection and Availability
 - Backup/Recovery: backup the data and restore it from backups.
 - Snapshots: copy of the state of a system or to a capability provided by certain systems at a particular point in time.
 - Replication: the same data is stored on multiple storage devices, to improve reliability, fault-tolerance, or accessibility.
 - Clustering: linking many computers together to act like a single computer and all de computers has access to all the data.
 - Mirroring: the act of copying data from one location to another in real time.
 - Self-healing: the ability to perceive that it is not operating correctly and make the necessary adjustments to restore itself to normal operation.
 - Data migration: Transfer data between storages.
 - Data Performance
 - Caching: a component that stores data so future requests for that data can be served faster.
 - Thin provisioning: a virtual provisioning mechanism that allows addressable storage capacity to be provisioned without consuming or reserving physical capacity.
 - Auto-tiering: put data in the appropriate class of storage based on how frequently the data is accessed in real-time.

- Event management and alerting

- Thin Provisioning
 - This virtual disk form is very similar to the traditional format with the exception that it does not pre-allocate the capacity of the virtual disk from the datastore when it is created. When storage capacity is required the virtual disk will allocate storage in chunks equal to the size of the file system block.

3.2.1.2 Centralized Management

- Pool Management: create, deliver and manage storage pools.
- New Resources: manage new resources.
- Policy Settings: allow administrators to set policy for automate managing the storage and data services.
- Service Levels: determine or set service level to the system resources.
- Monitoring:
 - Track capacity consumption.

- System Health.
- Monitor and report on performance trends from host to storage.
- View the physical resources and relationship dependencies.
- Event logs and system alerts management.
- Reports
- Troubleshoot: systematic search for the source of a problem so that it can be solved, and so the product or process can be made operational again.

3.2.1.3 Self-Service

- Users can easily subscribe to storage resources that meet their workload demands.
- Users are not required to know or care about the underlying hardware and software that is providing the storage to their application.
- Automatically provisions the right hardware and software to meet the users' needs based on policies pre-defined by the storage administrator.

3.2.1.4 Heterogeneous Hardware/Storage Type

- Support Heterogeneous Hardware as follows:
 - Direct-attached storage (DAS)
 - Network-attached storage (NAS)
 - Storage area networks (SAN)
 - Cloud Storage
- Support Block/File/Object Storage
 - Block-based storage stores data on a hard disk as a sequence of bits or bytes of a fixed size or length (a block).
 - File-based storage systems, such as NAS appliances, are often referred to as "filers" and store data on a hard disk as files in a directory structure
 - Object-based storage systems use containers to store data known as objects in a flat address space instead of the hierarchical, directory-based file systems that are common in block- and file-based storage systems.
- Support implementation in virtual, physical environments or a mix of them.

3.2.1.5 Scale-out Architecture

- Distributed computing
- Active-Active Architecture
- The data and the metadata are distributed across all nodes
- Allow Scale-Up

3.2.1.6 Scalability

- Scale while continuing to manage storage as a single enterprise-class storage system.

- Provides massive, virtually limitless scalability.
- Scale without disruption the availability or performance.

4. References

- [1] DMTF DSP-IS0501, Software Defined Data Center (SDDC) Definition, 2015
- [2] SNIA, Software Defined Storage (White Paper), 2015
- [3] VMware, The VMware Perspective on Software-Defined Storage (White Paper), 2014
- [4] EMC, Transform Your Storage For The Software Defined Data Center (White Paper), 2015

Appendix A. Acknowledgements

This draft is supported by Institute for Information & communications Technology Promotion(IITP) grant funded by the Korea government(MSIP)(R-20150223-000247, Cloud Storage Brokering Technology for Data-Centric Computing Standardization)

Authors' Addresses

Eui-Nam Huh

Computer Science and Engineering Department, Kyung Hee University
Yongin, South Korea

Phone: +82 (0)31 201 3778

Email: johnhuh@khu.ac.kr

Ga-Won Lee

Computer Science and Engineering Department, Kyung Hee University
Yongin, South Korea

Phone: +82 (0)31 201 2454

Email: gawon@khu.ac.kr

Yunkon Kim

Computer Science and Engineering Department, Kyung Hee University
Yongin, South Korea

Phone: +82 (0)31 201 2454

Email: ykkim@khu.ac.kr

Jintaek Kim

Consortium of Cloud Computing Research, Seoul, South Korea

Phone: +82 (0)2 2052 0156

Email: jtkim@cccr.ir.kr