



Object Storage: What, How and Why?

Live Webcast
February 19, 2020
10:00 am PT

Today's Presenters



Moderator:
John Kim
Mellanox
@Tier1Storage



Presenter:
Rick Vanover
Veeam
@RickVanover



Presenter:
Chris Evans
Architecting IT
@ChrisMEvans



Presenter:
Alex McDonald
NetApp
@AlexTangent

SNIA-at-a-Glance



185
industry leading
organizations



2,000
active contributing
members



50,000
IT end users & storage
pros worldwide

Learn more: snia.org/technical



Technologies We Cover

- ✓ Ethernet
- ✓ iSCSI
- ✓ NVMe-oF
- ✓ InfiniBand
- ✓ Fibre Channel, FCoE
- ✓ Hyperconverged (HCI)
- ✓ Storage protocols (block, file, object)
- ✓ Virtualized storage
- ✓ Software-defined storage

- ◆ The material contained in this presentation is copyrighted by the SNIA unless otherwise noted.
- ◆ Member companies and individual members may use this material in presentations and literature under the following conditions:
 - ◆ Any slide or slides used must be reproduced in their entirety without modification
 - ◆ The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- ◆ This presentation is a project of the SNIA.
- ◆ Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- ◆ The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Agenda

- History of Object Storage
- What is Object Storage?
- Use cases
- Market needs
- Standards

➤ What is Object Storage?

- ◆ Object storage provided three key characteristics from the start:
 - **Shared access to data**
 - **Heterogenous computing**
 - **Dynamic scaling without interruption**
- ◆ Object storage, as a definition, can be: A storage system that manages and manipulates data storage as distinct units, called objects *

The History of Object Storage

➤ Object storage originated in the late 1990s:

- ◆ Seagate specifications from 1999
- ◆ Introduced some of the first commands and how operating system effectively removed from consumption of the storage

› <https://www.t10.org/ftp/t10/document.99/99-341r0.pdf>

› https://pdfs.semanticscholar.org/bcd1/97cb0f8544b651289dfdb95efd0b1fd70753.pdf?_ga=2.24472250.1666047995.1580239914-1595228147.1580239914

➤ One common name carries through most of it

- ◆ Dr. Garth Gibson, formerly of CMU
- ◆ Dr. Gibson participated in one of the first research collaborations <http://www.cs.cmu.edu/~garth/>




Object Storage

➤ How can we visualize object storage?

Image Source Caringo

<https://www.caringo.com/blog/back-basics-object-storage>

Block




Client Via OS

Fixed Sys Attributes

Transactional Data

Performance

File




Client Via OS

Fixed Sys Attributes

Shared Changing File

Access, Single Site

Object



Client is App

Custom Metadata

Shared Semi-Static File

Scalable, Multi-Site

➤ How does Object Storage work?

- ◆ Application Connection via REST APIs
- ◆ The networking is different, made for the Internet!
- ◆ It is *effectively* unlimited when viewed by Internet standards
- ◆ Introduces metadata for additional object information
- ◆ A new model for permissions and I/O operations are introduced
 - › **PRO TIP:** Security for public cloud object storage resources is a significant obligation; the hyperscale public cloud is a constant security threat target.

➤ Storage is still storage: Many flavors and variations

- ◆ Tiers, Consumption and retrieval-based capabilities via APIs
- ◆ Special Capabilities:
 - › Governance & Compliance object lock (WORM)
 - › Soft Delete
 - › Immutable
- ◆ Cold/Archive
 - › Expedited (1-5 minutes)
 - › Standard (3-5 hours)
 - › Bulk (5-12 hours)
- › **PRO TIP:** Consume carefully. Different products or tiers of object storage have different pricing and impact usability.

- What is a good use case for Object Storage?
 - ◆ Content distribution is a common example (Think Netflix)
- What is **NOT** a good use case for Object Storage?
 - ◆ Running operating systems or highly transactional, low-latency, immediately consistent data
- Further information in a SNIA webcast: **Transactional Models & Their Storage Requirements**
 - ◆ <https://www.snia.org/educational-library/webcast-transactional-models-their-storage-requirements-2017>

➤ Object storage is not just for the Internet!



- ◆ Microsoft Azure, Amazon Web Services and other clouds lead the offerings for hyperscale public clouds.
- ◆ There are also scores of on-premises and hybrid solutions that provide object storage within ones own data centers.
 - › Caringo, Cloudian, Ceph, Dell EMC ECS, Hitachi HCP, Huawei FusionStorage, Minio, NetApp StorageGrid, OpenStack Swift, Scality RING, Wasabi and more!
- ◆ Object storage is **durable**: Erasure coding may be used as data durability technique for massive scale storage systems.

➤ Why does the market need Object Storage?

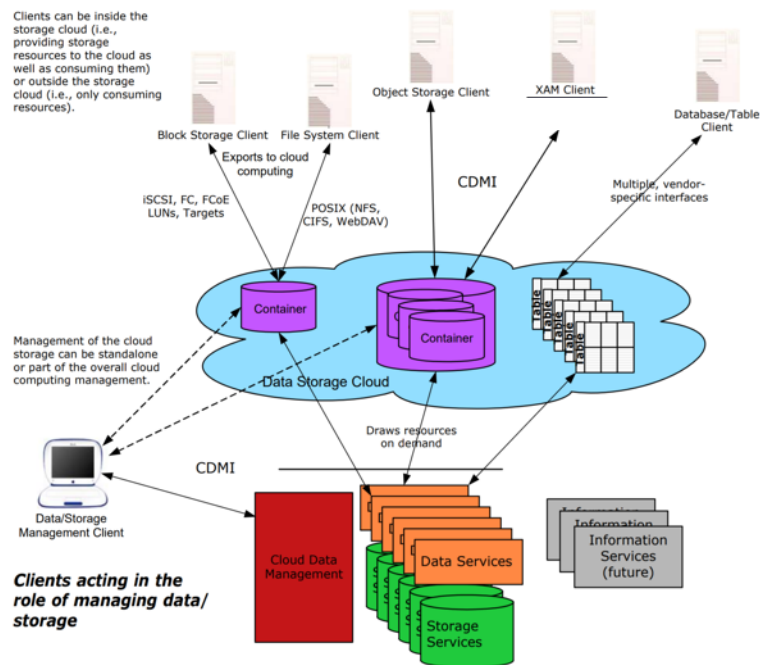
- ◆ Certain data sets were “born” for object storage: Rich Media
- ◆ Certain data sets fit well for *effectively* unlimited capacity.
- ◆ Certain application profiles fit well for Internet connectivity.
- ◆ Cloud-native applications may have “Golden Repository”.
- ◆ Object stores act as a “golden repository” for long term retention
- ◆ Certain personas like the cost model
 - › **PRO TIP:** If you don’t have a cloud economist in your organization; you may want to add one.

Object Storage and Standards

- All are RESTful API based
- S3 (Simple Storage Service), proprietary to Amazon
 - ◆ Part of a larger Amazon ecosystem
- CDMI (Cloud Data Management Interface), ISO/IEC 17826, SNIA developed
 - ◆ HTTP and JSON based
 - ◆ Provides management as well as store & retrieve
 - ◆ Provides CRUD operations
 - › create, read, update, delete
 - ◆ High levels of security through use of ACLs

CDMI Overview

- Single namespace for different datatypes
 - S3, LUNs, file systems, databases, objects, ...
- Data & management paths, capability discovery
- Use cases include data movement between different object systems
- Many commercial systems support both
- See snia.org/cloud for details on CDMI



- File vs. Block vs. Object Storage
 - ◆ SNIAVideo YouTube: <https://youtu.be/Akpu35MKzrA>
- Cloud Object Storage and the Use of Gateways
 - ◆ SNIAVideo You Tube: <https://youtu.be/EF-hsEEemJR4>
- SNIA CDMI & Cloud Technologies
 - ◆ Our website at snia.org/cloud

- Hopefully this session provided a nice overview of how object storage came to be, what it is used for and how it can be consumed. We covered
 - ◆ History of Object Storage
 - ◆ What it is
 - ◆ How it works
 - ◆ Why it matters in the market
 - ◆ Standards

Q&A

After This Webcast

- Please rate this webcast and provide us with feedback
- This webcast and a PDF of the slides will be posted to the SNIA Networking Storage Forum (NSF) website and available on-demand at www.snia.org/forums/nsf/knowledge/webcasts
- A full Q&A from this webcast, including answers to questions we couldn't get to today, will be posted to the SNIA-NSF blog: sniansfblog.org
- Follow us on Twitter @SNIANSF

Thank You