



RED HAT STORAGE ON CISCO UCS-S SERIES

Taco Scargo
Red Hat
Senior Specialist Solution Architect
February 7th, 2017

An aerial photograph of a coastline. The water is a deep orange color, and the land is a lighter orange. A breakwater made of large, rectangular concrete blocks extends from the land into the water. The sky is a pale, hazy blue. The overall image has a warm, orange-toned aesthetic.

STORAGE IS CHANGING

SOFTWARE-DEFINED STORAGE

THE DATA EXPLOSION



WEB, MOBILE, SOCIAL MEDIA, CLOUD
Our digital assets have grown exponentially due to web scale services like Facebook, Flickr, Snapchat, YouTube, and Netflix.



VIDEO ON-DEMAND SERVICES
Rapid growth of video on-demand has culminated in 50% of households using this service.

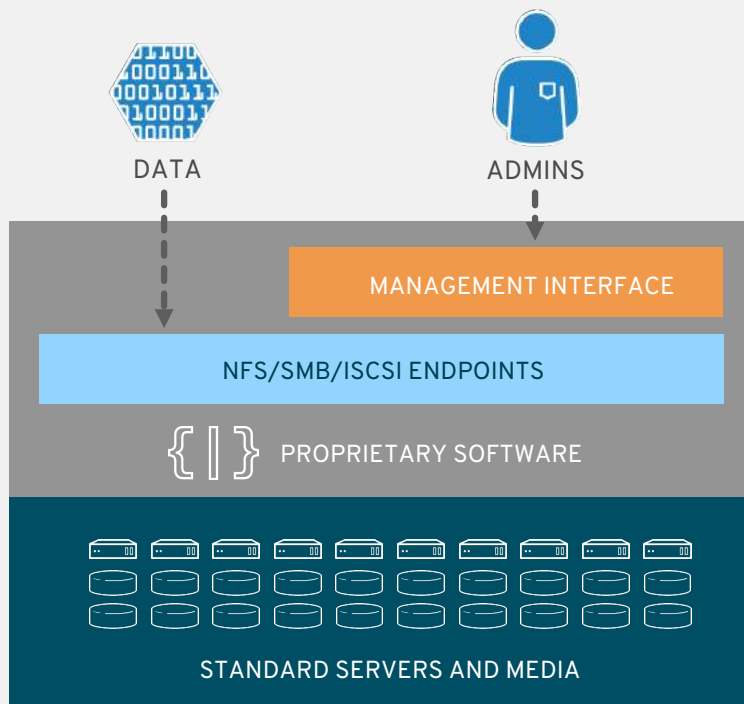


MEDIA AND ENTERTAINMENT INDUSTRIES
A staggering amount of content is created during today's optimized production processes.



MEDICAL INDUSTRY
Medical imaging needs are vast, and regulatory requirements can be demanding.

APPLIANCES AREN'T ENOUGH



THE TRADITIONAL APPROACH TO STORAGE



Complexity hidden from end users, along with flexibility




Vendor lock-in leads to pricing premium



Price premium over constituent components is difficult to sustain

PUBLIC CLOUD STORAGE ISN'T ENOUGH



CONVENIENT BUT LIMITED



Complexity still hidden from end users, pay-as-you-go pricing

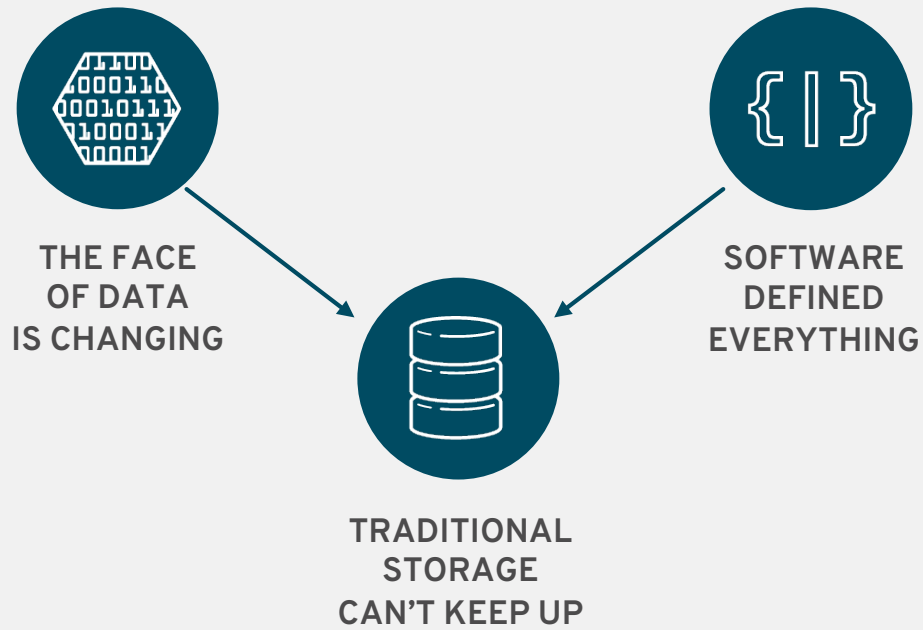


Fastest-growing segment of IT storage budgets



Mostly built with proprietary software

MARKET DYNAMICS



FLEXIBILITY IS CRUCIAL



THE INDUSTRY IS RETHINKING STORAGE



38% of IT decision makers report inadequate storage capabilities as one of their top three weekly pain points

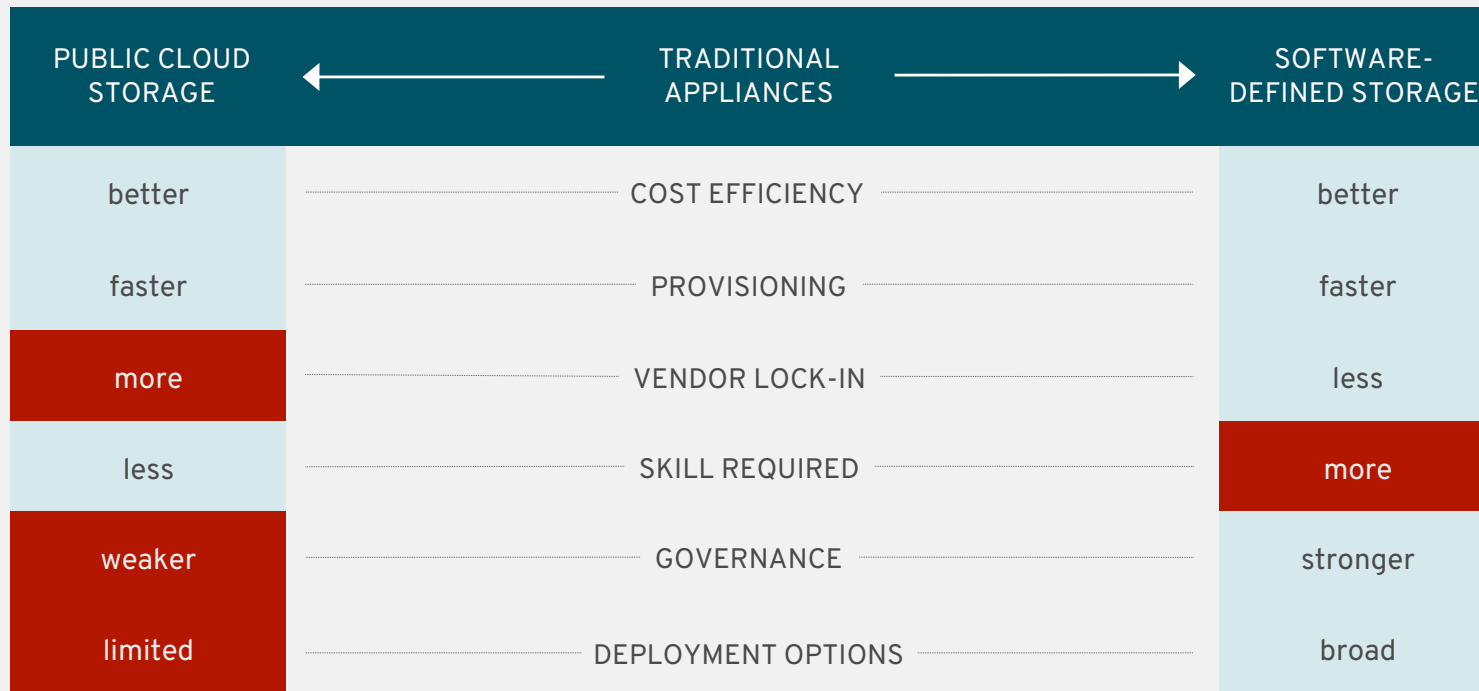


70% of IT decision makers admit that their organization's current storage can't cope with emerging workloads



98% of IT decision makers believe a more agile storage solution could benefit their organization

DISRUPTION IN THE STORAGE INDUSTRY




An aerial photograph of a coastline. The water is a deep orange color, and the land is a lighter orange. A breakwater made of large, rectangular concrete blocks extends from the land into the water. The text "WHAT IS SOFTWARE-DEFINED STORAGE?" is overlaid in white, bold, sans-serif font in the center of the image.


WHAT IS SOFTWARE-DEFINED STORAGE?

SERVER-BASED STORAGE

The use of software and standard hardware to provide services traditionally provided by single-purpose storage appliances (similar to server virtualization, which uses software to emulate servers), providing increased agility and efficiency



VIRTUALIZED STORAGE SCALES BETTER



STORAGE ORCHESTRATION

The ability to provision, grow, shrink, and decommission storage resources on-demand and programmatically, providing increased control and integration of storage into a software-defined data center

Web Console

A browser interface designed for managing distributed storage

API

A full API for automation and integration with outside systems

Command Line

A robust, scriptable command-line interface for expert operators

PROVISION

INSTALL

CONFIGURE

TUNE

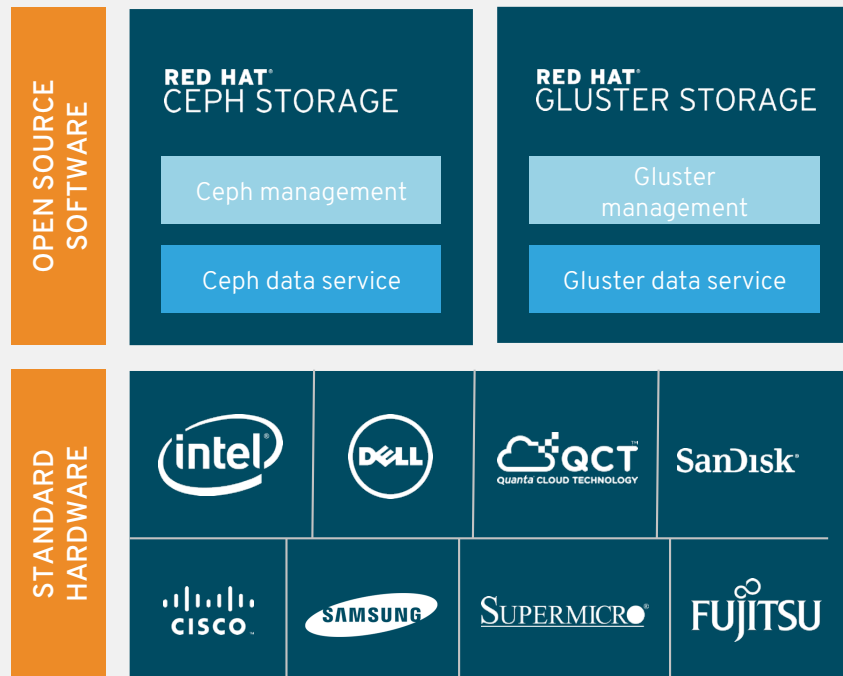
MONITOR

Full life cycle management for distributed, software-defined data services



RED HAT STORAGE

THE RED HAT STORAGE PORTFOLIO



Share-nothing, scale-out architecture provides durability and adapts to changing demands

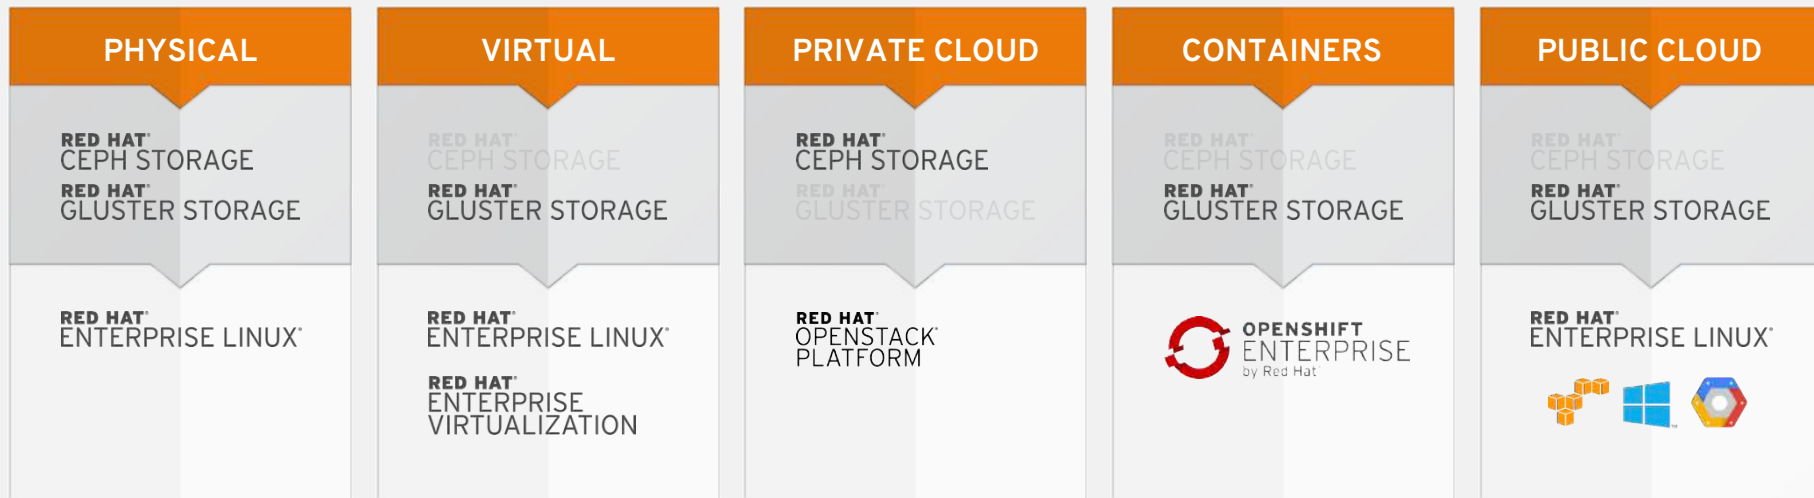
Self-managing and self-healing features reduce operational overhead

Standards-based interfaces and full APIs ease integration with applications and systems


Supported by the experts at Red Hat

HOW STORAGE FITS

RED HAT® STORAGE



GARTNER MAGIC QUADRANT



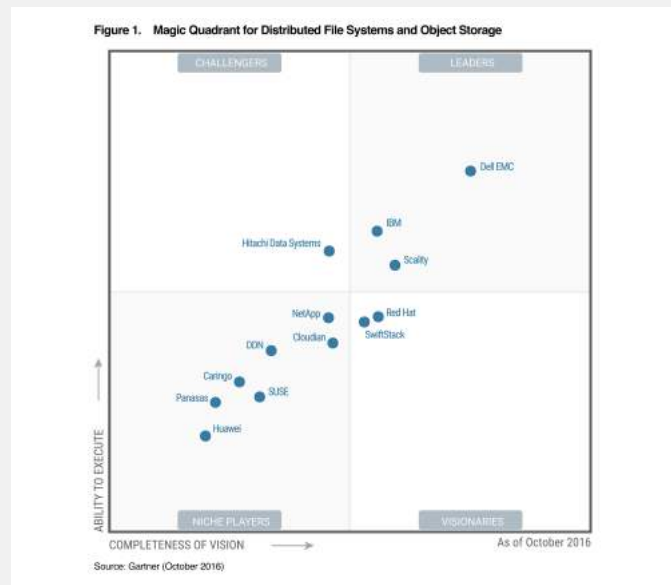
This graphic was published by Gartner, Inc. as part of a larger research document and should be evaluated in the context of the entire document. The Gartner document is available upon request at <https://engage.redhat.com/gartnermagic-quadrant-storage-s-201610121525>

Red Hat Storage recognized as a **Visionary** by Gartner in their first Magic Quadrant for Distributed File Systems and Object Storage.

Red Hat Storage positioned furthest and highest in both Completeness of Vision and Ability to Execute in the Visionaries quadrant.

Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

GARTNER MAGIC QUADRANT



This graphic was published by Gartner, Inc. as part of a larger research document and should be evaluated in the context of the entire document. The Gartner document is available upon request at <https://engage.redhat.com/gartnermagic-quadrant-storage-s-201610121525>

Gartner LICENSED FOR DISTRIBUTION

Magic Quadrant for Distributed File Systems and Object Storage

Published: 20 October 2016 ID: G00307708
 Analyst(s): Julia Palmer, Arun Chandrasekaran, Raj Bala

Summary
 Bi-modal approaches that demand infrastructure agility and scalability are stimulating market adoption of scale-out file and object storage products. This research helps I&O leaders assess the key attributes, vision and execution prowess of distributed file systems and object storage market vendors.

Strategic Planning Assumption
 By 2021, more than 80% of enterprise data will be stored in scale-out storage systems in enterprise and cloud data centers, up from 30% today.

Market Definition/Description
 Scale-out storage systems are growing fast and becoming a popular platform to tackle the unabated growth of unstructured data. With data growth exceeding 40% year over year in many enterprises, infrastructure and operations (I&O) leaders are looking for extensible storage products that can address an increasing number of use cases with lower acquisition and operational costs. Enterprises are demanding features and capabilities prevalent in big data cloud infrastructures, such as self-healing and ease of management. Software defined storage (SDS), deployed on commodity hardware, is emerging as a threat to extensible controller-based (ECB) storage arrays in environments with a steep growth of unstructured data. New and established storage vendors are continuing to develop scalable storage clustered file systems and object storage products to address cost and scalability limitations in traditional, scale-up storage environments.

Gartner defines distributed file systems and object storage as software and hardware solutions that offer object and/or scale-out file technology to address requirements for unstructured data growth and based on "shared nothing architecture." A shared nothing architecture is a distributed computing architecture in which each node is independent and self-sufficient, and there is no single point of contention across the system (see the Acronym Key and Glossary terms section).

Distributed file system storage uses a single parallel file system to cluster multiple storage nodes together, presenting a single namespace and storage pool to provide high bandwidth for multiple hosts in parallel. Data is distributed over multiple nodes in the cluster to deliver data availability and resilience in a self-healing manner, and to provide high throughput and capacity linearly.

Object storage refers to devices and software that house data in structures called "objects," and serve clients via RESTful HTTP APIs, such as Amazon Simple Storage Service (S3) and OpenStack Swift.

Access the report at <https://engage.redhat.com/gartnermagic-quadrant-storage-s-201610121525>


An aerial photograph of a coastal breakwater structure. The breakwater is composed of numerous large, rectangular concrete blocks arranged in a staggered pattern. The water is a deep blue, and the sky is a pale, hazy blue. The entire image has a semi-transparent orange overlay. In the bottom right corner, there is a small red hat logo with the text 'redhat' next to it.

RED HAT CEPH STORAGE

RED HAT CEPH STORAGE

Distributed enterprise storage platform

- Proven for large-scale, modern workloads
- Open, **massively-scalable**, software-defined
- Flexible, scale-out architecture on clustered commodity x86 hardware
- Efficient, unified storage platform (object, block, file) in one **self-managing, self-healing** platform with **no single point of failure**
- Integrated, easy-to-use management console



BUSINESS BENEFITS

OPEN SOURCE

No proprietary lock-in, with a large commercial ecosystem and broad community

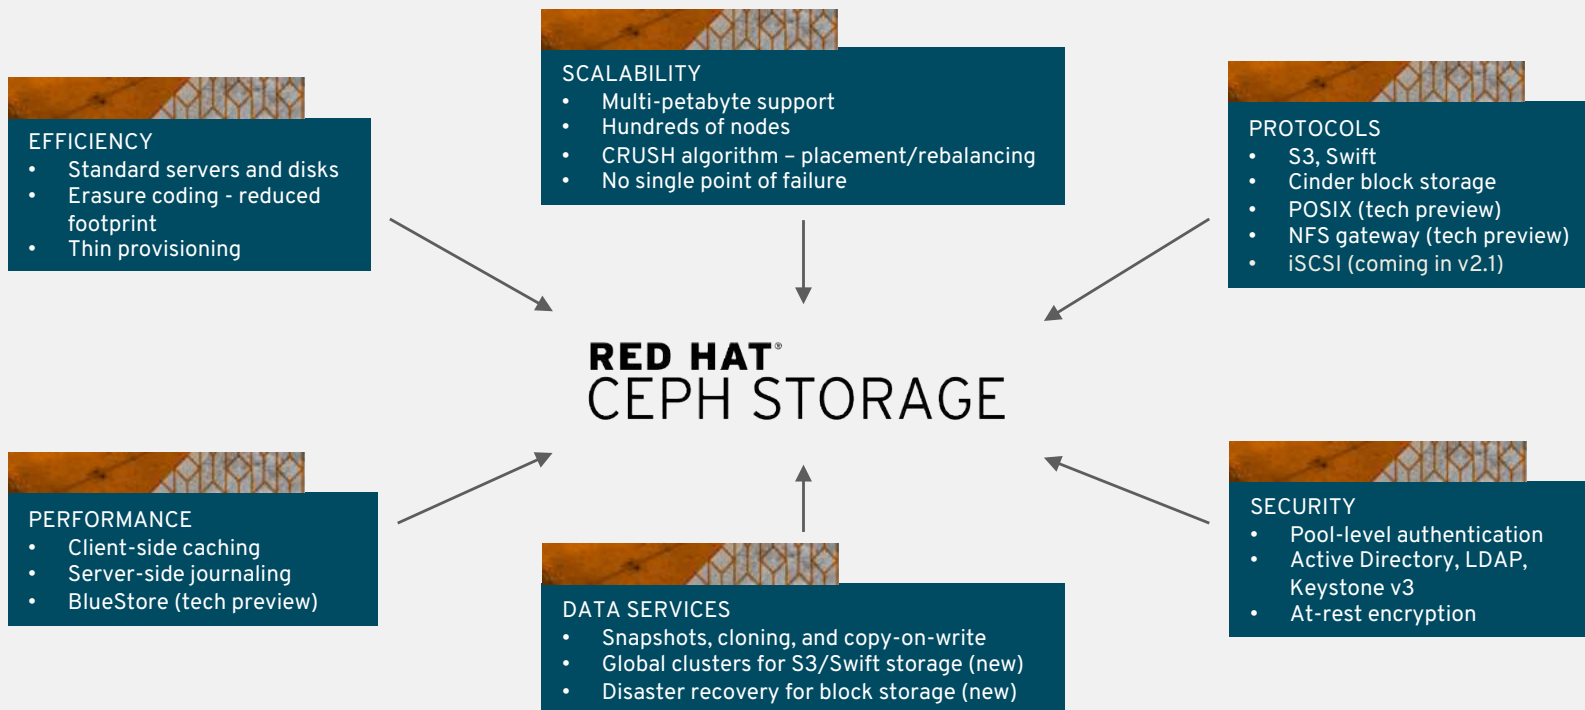
PEACE OF MIND

Over a decade of active development, proven in production and backed by Red Hat

LOWER COST

More economical than traditional NAS/SAN, particularly at petabyte scale

CORE PRODUCT FEATURES



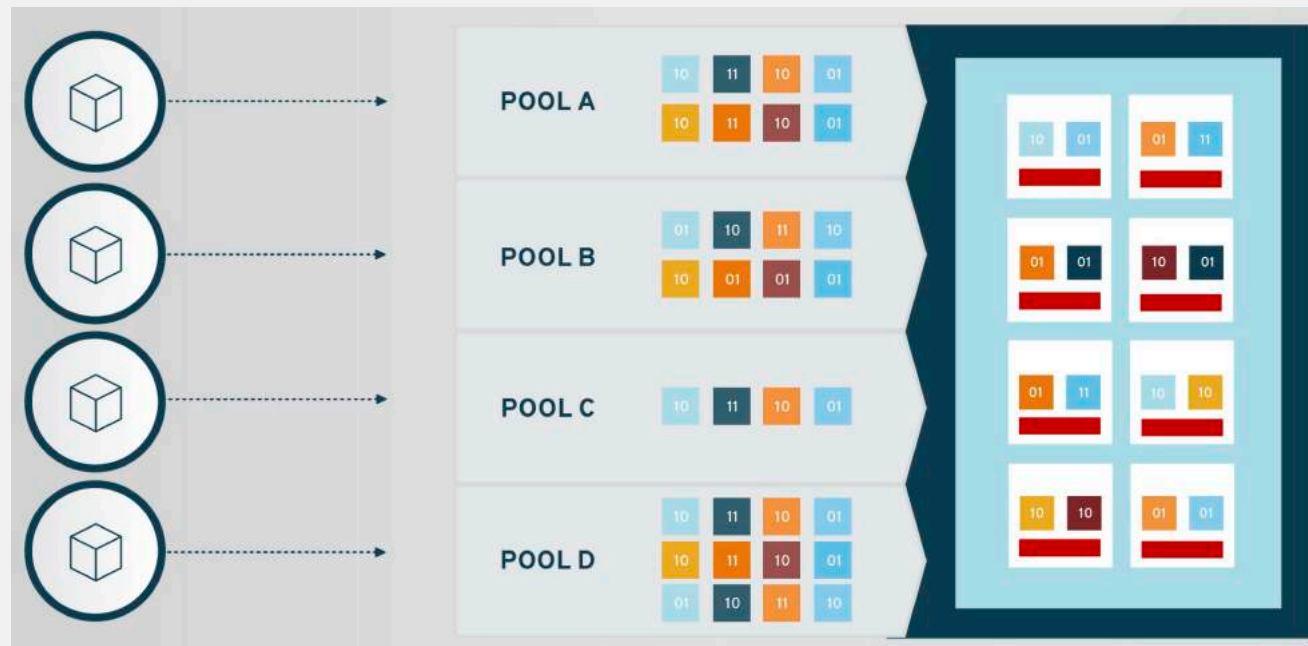
“Run storage the way Google does”



An aerial photograph of a large, complex concrete breakwater structure in the ocean. The structure consists of numerous interlocking concrete blocks arranged in a grid-like pattern, extending from the shore into the water. The water is a deep blue, and the sky is a pale, hazy blue. The overall scene is captured from a high angle, looking down at the structure. A semi-transparent red hat logo watermark is visible in the bottom right corner of the image.

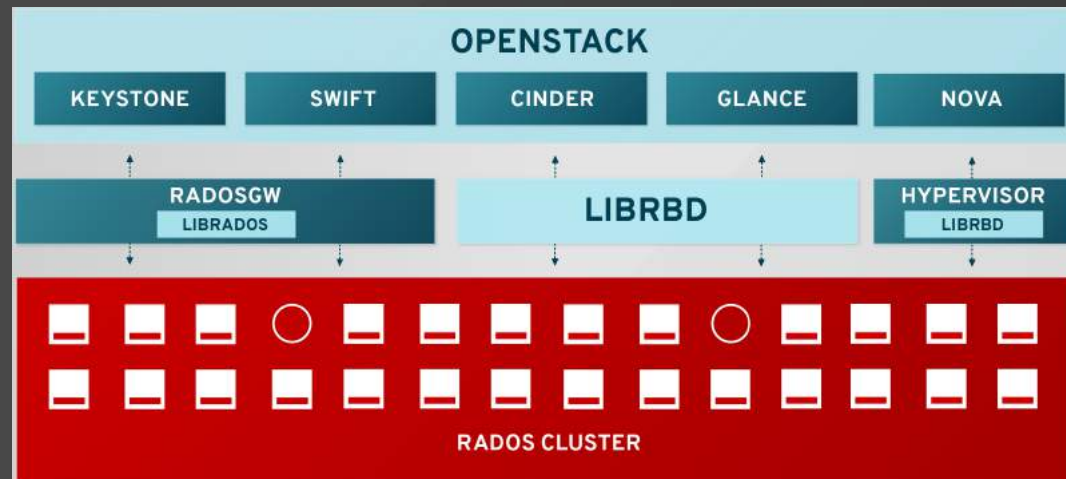
RED HAT STORAGE TARGET WORKLOADS

STRENGTH: MASSIVELY DISTRIBUTED DATA




USE CASE: OPENSTACK

- Allows for instantaneous parallel creation of VMs at massive scale
- Integrates easily and tightly with OpenStack Cinder, Glance, Nova, Keystone, and Manila
- Offers instant backup capabilities
- Provides persistent object, file, and database storage for applications

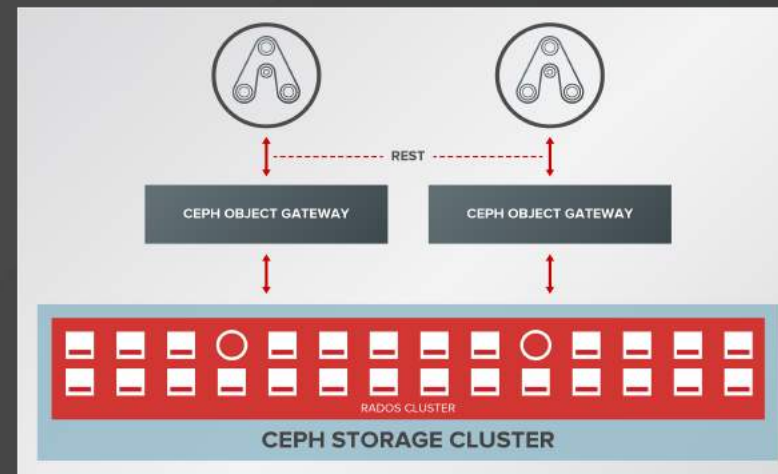


STRENGTH: PETABYTE SCALE



USE CASE: OBJECT STORAGE

- Stores unstructured data at web scale, using standard hardware
- Works with standard object APIs for a wide range of compatibility
- Spans multiple regions with no single point of failure
- Ideal for active archives, big data archives, and content libraries



STRENGTH: DEPLOYMENT FLEXIBILITY



HOSTS



VIRTUAL
MACHINES




PRIVATE
CLOUD



PUBLIC
CLOUD



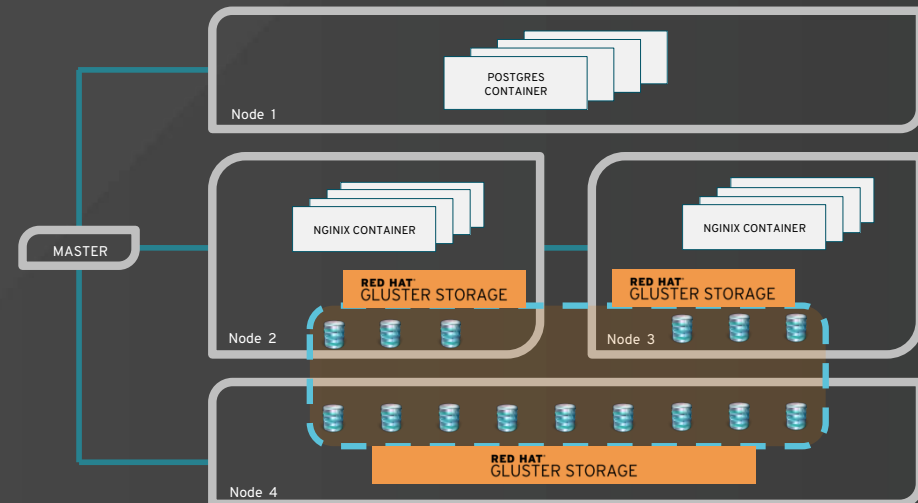
HYBRID
CLOUD



CONTAINERS

USE CASE: CONTAINER-NATIVE STORAGE

- Deploys storage alongside applications in containers
- Lowers TCO by increasing utilization of resources
- Unifies container and storage orchestration
- Allows for rapid adjustments to compute/storage ratio



OUR TARGET WORKLOADS

CONTAINERS

Container-ready storage
Container-converged storage

CLOUD

OpenStack VM storage
OpenStack database storage

OBJECT STORAGE

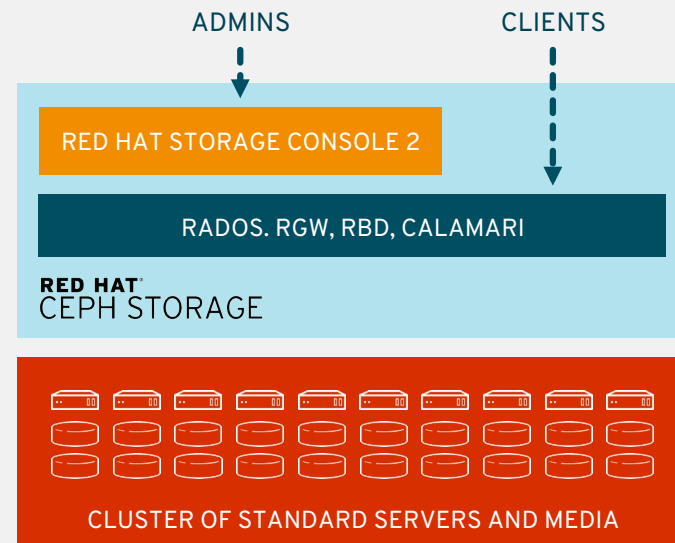
Rich media/active archives
Storage as a Service



RED HAT STORAGE CONSOLE

RED HAT STORAGE CONSOLE

- Easy to use graphical interface to manage storage cluster life cycle
- Ansible-based deployment tools for installation, importation, and granular configuration from CLI or GUI
- Monitoring and troubleshooting with statistical and graphical information about cluster components



- Dashboard
- Clusters
- Hosts
- Storage >
- Admin >

1 Clusters

✓

Utilization

84.1% Used

138.7 GB of 164.8 GB used

Last 24 hours

Most Used Storage Profiles

default	84% Used
---------	----------

8 Hosts

✓ + 1

3 Monitors

✓

16384 PGs

✓

Most Used Pools

rados1	40% Used
rbdPool	88% Used
rbd2	0% Used
ecPool	0% Used

4 Pools

✓

15 OSDs

✓

11759 Objects

✓

System Performance

Last 1 hour

CPU Memory 3.1 k IOPS

Dashboard

Clusters

Hosts

Storage

Admin

138.7 GB of 164.8 GB used

Last 24 hours

8 Hosts

1

3 Monitors


16384 PGs

4 Pools

15 OSDs

11759 Objects

Most Used Pools



System Performance

Last 1 hour

CPU



Memory



3.1 k IOPS




53663.3 KB/s Throughput



0.3 ms Latency



WHY RED HAT STORAGE CONSOLE 2?



Easily install Red Hat Ceph Storage 2 in less than an hour and ensure best practices



Provision storage in seconds and auto-expand when disks are added



Proactively monitor and manage to gain operational intelligence



Receive alerts for operational issues requiring intervention



**UCS S3260 + RED HAT CEPH
SOLUTION**

CEPH IS NOT JUST SCALE-OUT CAPACITY

IOPS OPTIMIZED

NVMe SSD in SLED chassis

High IOPS / GB
Smaller, random IO
Read / write mix

Use Case: MySQL



THROUGHPUT OPTIMIZED

SSD, HDD in standard / dense chassis

High MB/s throughput
Large, sequential IO
Read / write mix

Use Case: Rich Media



COST/ CAPACITY OPTIMIZED

HDD in dense / ultra-dense chassis

Low cost / GB
Sequential IO
Write mostly

Use Case: Active Archives



RULES FOR DESIGNING CEPH CLUSTERS SIZE

RED HAT®
CEPH STORAGE

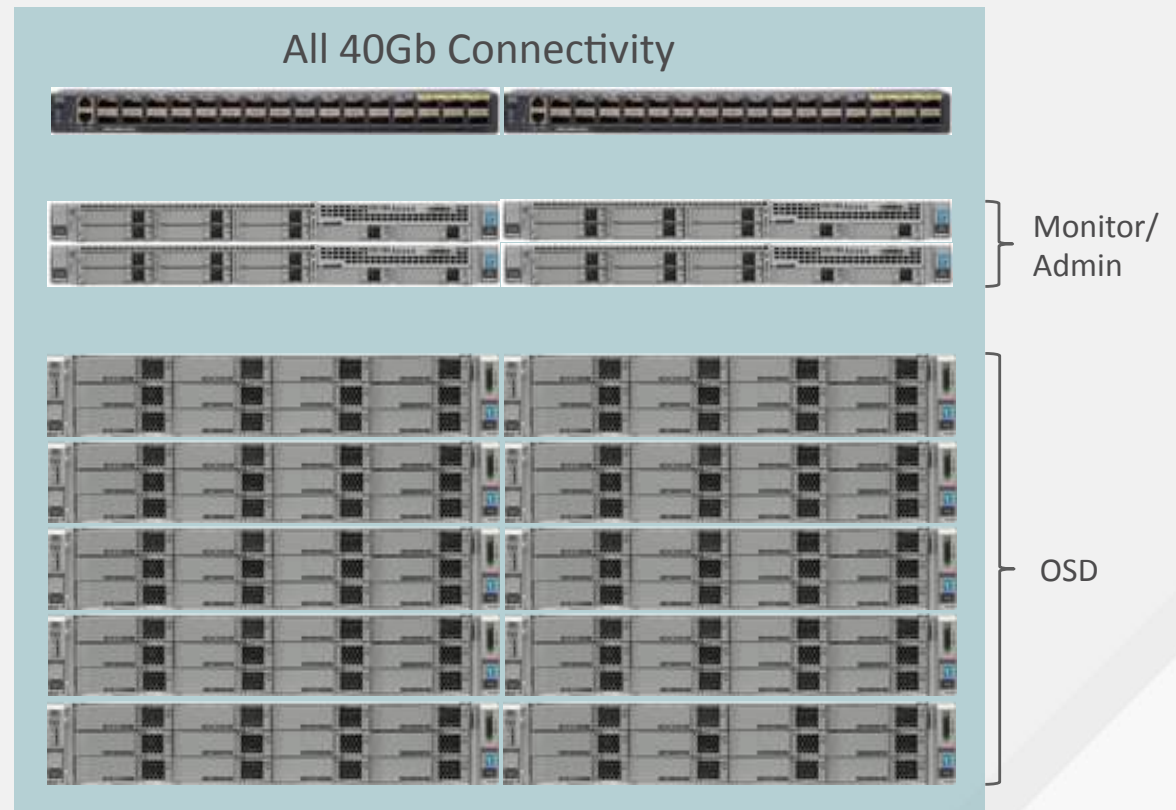


- Supported Ceph clusters can start very small with just 3 Monitor Nodes and 3 OSD Nodes but for
 - IOPS- and Throughput-intensive workload
 - Minimum recommended 10 OSD nodes
 - Capacity-intensive workloads
 - Minimum recommended 7 OSD nodes
 - Mixed workloads
 - Minimum recommended 10 OSD nodes

IOPS-INTENSIVE WORKLOAD

UCS DESIGN EXAMPLE

- **Parts Overview**
 - 2 x FI 6332
 - 4 x C220 M4S
 - 3 x Monitor + 1 x Admin
 - 10 x C240 M4L
 - Each OSD Node
 - 4-6 x SSD
 - 6-8 x HDD
- **Data Protection → 3 x Replication**
 - Usable Capacity up to 260 TB



THROUGHPUT-INTENSIVE WORKLOAD


UCS DESIGN EXAMPLE

- **Parts Overview**

- 2 x FI 6332
- 4 x C220 M4S
 - 3 x Monitor + 1 x Admin
- 6 x S3260 Dual Node
 - Each OSD Node
 - 4-6 x SSD
 - 22-24 x HDD

- **Data Protection**

- 3 x Replication
 - Usable Capacity up to 960 TB



CAPACITY-INTENSIVE WORKLOAD

UCS DESIGN EXAMPLE

- **Parts Overview**

- 2 x FI 6332
- 8 x C220 M4S
 - 3 x Monitor + 1 x Admin + 4 x RGW
- 6 x S3260 Dual Node
 - Each OSD Node
 - 28 x HDD

- **Data Protection**

- 3 x Replication
 - Usable Capacity up to 1,120 TB
- Erasure Coding
 - Usable Capacity up to 2,200 TB

All 40Gb Connectivity



The Power of UCS S3260 for Red Hat Ceph



Shared Local Resources

Designed for large object stores, high capacities and all Ceph workloads

Compute Resources

Full-Featured Modular Two-Socket Xeon Server

Stand-alone CIMC, IMC Supervisor and *UCS Managed*

Storage Resources

Massive Local Storage

Up to 600TB of dense storage in a compact 4U Form Factor that fits in a standard rack

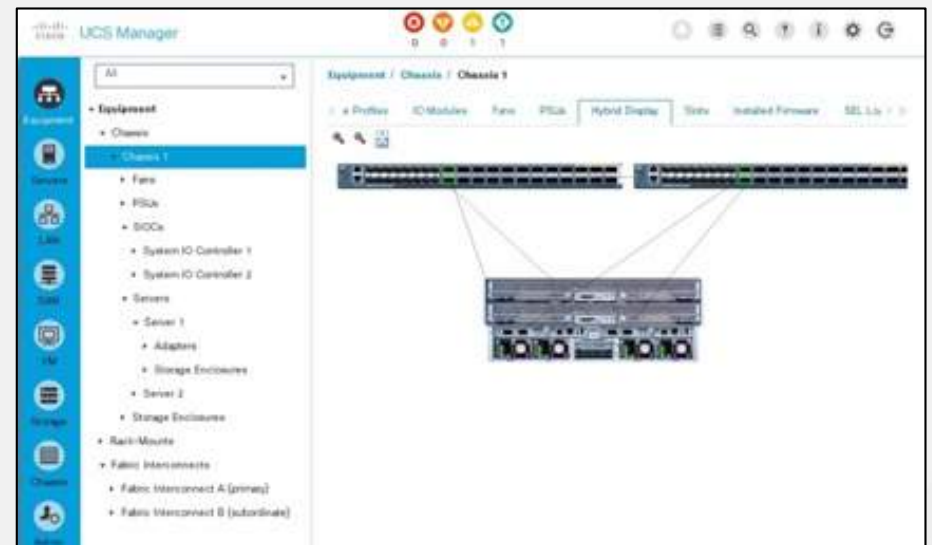
Network Resources

High I/O Bandwidth

Up to 4x 40GbE Powered by Cisco VIC


UCS S3260 – BENEFIT FOR CEPH CUSTOMERS

- Use different Ceph pools on the same hardware
 - Variety of NVMe/SSD/HDD options for S3260 available
- Grow your cluster by a few clicks
 - UCS Manager simplifies the add-on of S3260 as an OSD
- Don't get blocked by low network bandwidth
 - S3260 can easily separate Ceph network traffic with VIC technology
- Compute, Network, and Storage from Cisco – Single vendor



CEPH REFERENCE DESIGN FOR CISCO UCS S3260

- Useful for Throughput- and Capacity-intensive workloads
- UCS Managed
- 40GbE End-2-End with UCS 6332
- S3260 dual node for OSD, C220 M4 for Monitor/Admin/RGW nodes
- Easy to scale: 300TB to 15+PB per UCS Domain
- Policy Driven architecture
- Enabled for Mixed Workload



RED HAT CEPH STORAGE REFERENCE ARCHITECTURES

- [Hardware Configuration Guide](https://www.redhat.com/en/resources/red-hat-ceph-storage-hardware-configuration-guide)
<https://www.redhat.com/en/resources/red-hat-ceph-storage-hardware-configuration-guide>
- [Hardware Selection Guide for Red Hat Ceph Storage](https://www.redhat.com/en/resources/red-hat-ceph-storage-hardware-selection-guide)
<https://www.redhat.com/en/resources/red-hat-ceph-storage-hardware-selection-guide>
- [Cisco UCS C3160 High-Density Rack Server with Red Hat Ceph Storage](http://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/whitepaper-C11-735004.html)
<http://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/whitepaper-C11-735004.html>


RED HAT GLUSTER STORAGE REFERENCE ARCHITECTURES

- Large Unstructured Data Storage in a Small Datacenter Footprint: Cisco UCS C3160 and Red Hat Gluster 500-TB Solution
http://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/whitepaper_c11_734975.html
- Deploy Large Unstructured Data Storage in a Small Data Center Space: Cisco UCS C240 M4 Rack Server and Red Hat Gluster Storage 100-TB Solution
http://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/whitepaper_c11_734975.html

•

Following red hat storage

BLOG	redhatstorage.redhat.com
TWITTER	www.twitter.com/redhatstorage
FACEBOOK	www.facebook.com/RedHatStorage
YOUTUBE	www.youtube.com/user/redhatstorage
SLIDESHARE	www.slideshare.net/Red_Hat_Storage
WEB	www.redhat.com/storage





**THANK
YOU**

