

RED HAT
SUMMIT

Red Hat Storage Console 3

Unified Management of Ceph and Gluster

Jeff Applewhite
Principal Product Manager
May 3, 2017

Agenda

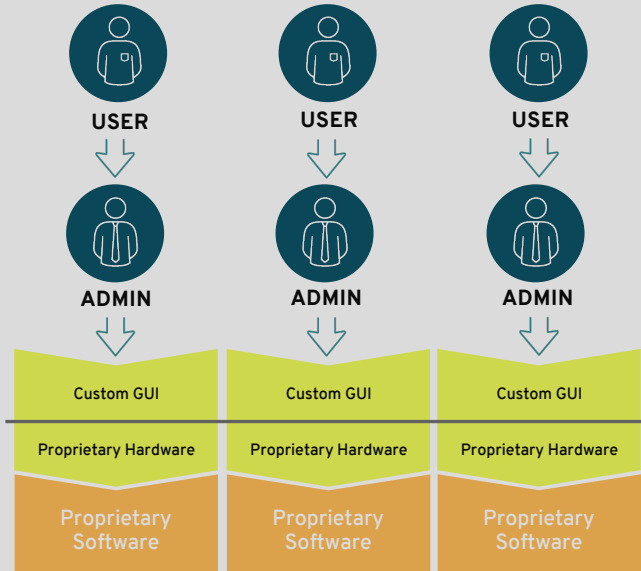
May 3, 2017
Red Hat Summit

- Introduction
- Why Red Hat Storage?
- What's new in the upcoming release of Storage Console?
- Architectural Overview
- Business Value
- Features
- Console UI
- Roadmap
- QA

Introduction: Why Red Hat Storage?

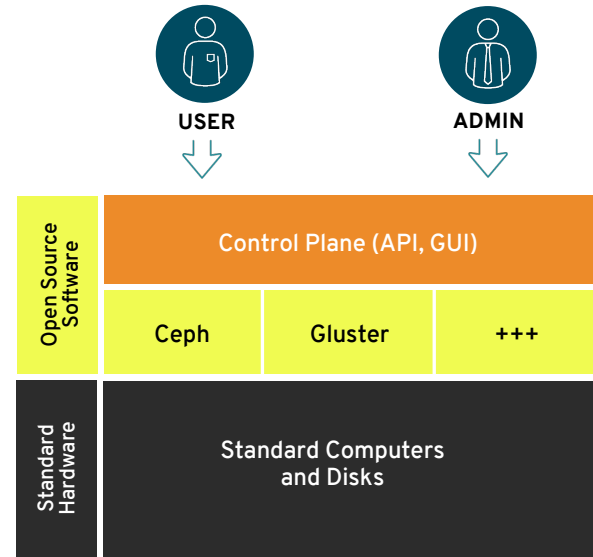
STORAGE IS EVOLVING

TRADITIONAL STORAGE Complex proprietary silos



RED HAT IS LEADING

OPEN, SOFTWARE-DEFINED STORAGE Standardized, unified, open platforms



A RISING TIDE

Software-Defined Storage is leading a shift in the global storage industry, with far-reaching effects.


“By 2020, between 70%-80% of unstructured data will be held on lower-cost storage managed by SDS.”

Innovation Insight: Separating Hype From Hope for Software-Defined Storage

“By 2019, 70% of existing storage array products will also be available as software-only versions.”

Innovation Insight: Separating Hype From Hope for Software-Defined Storage

SDS-P MARKET SIZE BY SEGMENT



Source: IDC

GARTNER MAGIC QUADRANT

Figure 1. Magic Quadrant for Distributed File Systems and Object Storage



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.

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>

Storage Console 3

Red Hat Storage Console: What is it?

- Unified graphical manager for Red Hat Ceph and Gluster storage.
- Based on the PatternFly UI framework for enterprise web applications.
- Focuses on:

Monitoring Dashboard	Fault Alerts
Task management and reporting	Host details, graphs, and network
Storage Provisioning	Cluster Creation / Import

What's New In RH Storage Console 3?

Fall 2017

Gluster 3.2
Support!

HA, Stateless,
Distributed
Architecture

New upstream
Tendrl project
and code base!

Architectural Overview

Architecture

RHS-C 3 is completely redesigned and based on the upstream Tendrl project

<http://www.tendrl.org>

<http://github.com/Tendrl>






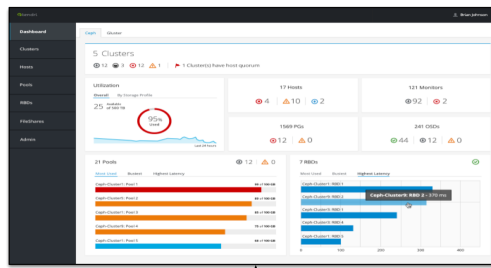


Architecture

The design consists of a core which is a trio of

- Node agent (generic service, runs on all nodes)
- Central Store (etcd, HA, connected to all nodes)
- And a stateless, HA, documented API service





Architecture

The design also consists of the Performance Monitoring module:

- Graphite: Store and graph metrics
- Carbon cache: Makes receiving & storing data efficient
- CollectD to feed generic host data to Graphite
- Custom Tendrl modules to feed Ceph and Gluster specific data to Graphite



Business Value

Makes distributed SDS storage easier!

Operational consistency and efficiency

(streamlined provisioning, enhanced discovery, and management via the integrated service dashboard).

Easily install or import Ceph 2.x or Gluster 3.2 storage “on demand”

Get important information about storage utilization to help troubleshoot and diagnose issues (noisy neighbor, flaky disks, network bottlenecks)

Business Value

Achieve comprehensive visibility with a unified view of Ceph and Gluster storage infrastructure in a single console.

Proactively monitor and manage health, performance, and capacity utilization and gain operational intelligence at scale.

Receive alerts for operational issues requiring intervention.

Red Hat Storage Console 3: Features

Feature Overview

Install, manage and monitor any SDS Storage (Future proofed by design)

Graph real-time data from the Ceph Calamari API and Gluster-D services to manage & monitor health, performance and capacity

Provisioning: (RADOS pools, RBDs, Gluster Bricks and Volumes, etc.)

Alert on operational issues: (OSD state, Cluster state, Failed Drive..)

Cluster and per node utilization, IOPs, performance, statistics

Features: **RED HAT**[®] CEPH STORAGE

Integrated Ceph Dashboard (utilization, IOPs, trends, issues)

Graphical installation of Ceph (via Ceph-Ansible)

Import existing Ceph 2.x clusters

Discover and add hosts, create cluster

Expand cluster (add mon's/osd's)

OSD add/remove/set state

Create / resize storage pools and RADOS Block Devices

Features: **RED HAT®** GLUSTER STORAGE

Integrated Gluster Dashboard (utilization, IOPs, trends, issues)

Graphical installation of Gluster (via gDeploy / Ansible)

Import existing Gluster 3.2+ clusters

Discover and add hosts, create trusted pool

Expand cluster

Add/Remove bricks

Supports Distributed replicated and EC volumes

Console UI

Dashboards..

- Dashboard
- Clusters
- Hosts
- Pools
- RBDs
- FileShares
- Admin

Ceph Gluster

5 Clusters

📉 12 🟡 3 🛑 12 ⚠️ 1 | 🚩 1 Cluster(s) have lost quorum

Raw Storage Utilization

25 Available of 500 TB

95% Used

Last 24 hours

51 Hosts

🛑 4 ⚠️ 10 🟢 2

15 Monitors

📉 1 🛑 2

1569 PGs

🛑 12 ⚠️ 0

241 OSDs

🟢 44 📉 12 ⚠️ 0

21 Pools

📉 12 ⚠️ 0

Most Used

Pool Name	Usage
Ceph-Cluster1: Pool 1	99 of 100 GB
Ceph-Cluster5: Pool 2	89 of 100 GB
Ceph-Cluster1: Pool 3	85 of 100 GB
Ceph-Cluster9: Pool 4	75 of 100 GB
Ceph-Cluster1: Pool 5	68 of 100 GB

7 RBDs

🟢

Most Used

RBD Name	Usage
Ceph-Cluster1: RBD 1	88 of 100 GB
Ceph-Cluster5: RBD 2	86 of 100 GB
Ceph-Cluster1: RBD 3	78 of 100 GB
Ceph-Cluster9: RBD 4	75 of 100 GB
Ceph-Cluster1: RBD 5	59 of 100 GB

Ceph Clusters (Multi-Cluster View)

Dashboard

Ceph **Gluster**

2 Clusters

👇 1 🌐 1 🛑 4 ⚠️ 1

Utilization

25 Available of 500 TB



15 Hosts

🛑 2 | ⚠️ 6 | 👆 3

5M Files



27 Clients



Services

- glusterd
- smbd
- NFS
- NFS - ganesha

12 File Shares

👇 3 🗑️ 2 🗑️ 8 ⚠️ 2

Most Used

Gluster-Cluster1: File Shares 1	99 of 100 GB
Gluster-Cluster2: File Shares 2	89 of 100 GB
Gluster-Cluster2: File Shares 3	85 of 100 GB
Gluster-Cluster9: File Shares 4	75 of 100 GB
Gluster-Cluster2: File Shares 5	68 of 100 GB

36 Bricks

🛑 2 | 👇 12 | ⚠️ 0

Most Used

Gluster-Cluster1: Brick 1	99 of 100 GB
Gluster-Cluster2: Brick 2	89 of 100 GB
Gluster-Cluster2: Brick 3	80 of 100 GB
Gluster-Cluster9: Brick 4	75 of 100 GB
Gluster-Cluster2: Brick 5	68 of 100 GB

Gluster Trusted Pools (Multi-Cluster View)



Clusters

Hosts

File Shares

Pools

RBDs

Admin

Tasks


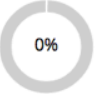




Clusters

Name

Name

Import Create

2 Clusters










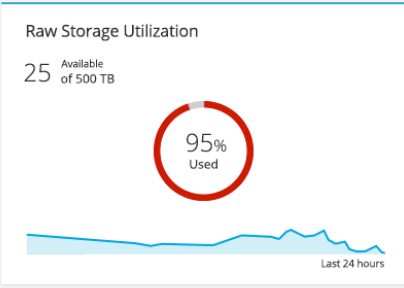
				Hosts	Pools	Alerts	
	MyCeph2	 0% 86.5 MB of 40.0 GB used		5	7	NA	
	MyGluster	 No Data Available		3	NA	NA	

Click on <cluster name> to see the cluster object details

- Dashboard
- Clusters**
- Hosts
- Pools
- RBDs
- FileShares
- Admin

Clusters > MyCeph2





Overview Hosts Pools RBDs OSDs Configuration Events



Ceph Dashboard






Clusters > MyGluster

Overview Hosts File Shares Bricks Configuration Events



Services

glusterd	✓
smbd	✓
NFS	✓
NFS - ganesha	✓



Gluster Dashboard

- Clusters
- Hosts
- File Shares
- Pools
- RBDs
- Admin
- Tasks

Tasks > CreateCluster: ecf9a9a1-bf1e-49d7-b777-32ea2c11a696

Task history

Time Submitted: 19 Apr 2017

Status: ✔ Finished

Events

info	Processing Job ecf9a9a1-bf1e-49d7-b777-32ea2c11a696
info	Running Flow tendrl.flows.CreateCluster
info	Created SSH setup job 882101ba-fa00-434b-b8d4-15a0fb200d73 for node cb819b40-ec42-40a8-bef5-797e4a304dd2
info	Created SSH setup job a5e71a91-0121-4018-80f4-e2c1dc959c28 for node 20047220-9edf-4916-9cc7-c4d86257f8d4
info	Processing Job a5e71a91-0121-4018-80f4-e2c1dc959c28
info	Running Flow tendrl.flows.SetupSsh
info	JOB[a5e71a91-0121-4018-80f4-e2c1dc959c28]: Finished Flow tendrl.flows.SetupSsh
info	Created SSH setup job 941671e2-a06f-4b2e-be04-2e25b34390d1 for node 563498f2-29a7-4684-b27f-17ebfafa0d32
info	Processing Job 882101ba-fa00-434b-b8d4-15a0fb200d73
info	Running Flow tendrl.flows.SetupSsh
info	Created SSH setup job e5fb2a47-0bc5-4dac-8572-0c24a24a1390 for node d173b0a3-7c54-4314-830f-2f4263875aab
info	JOB[882101ba-fa00-434b-b8d4-15a0fb200d73]: Finished Flow tendrl.flows.SetupSsh
info	Created SSH setup job 4381246d-719b-4272-b5db-dbf597b3072 for node 96764bef-e956-458b-bdd1-a3ca88620597
info	SSH setup completed for all nodes in cluster bfcea228-e01c-473a-9493-045cb4862c16



Clusters

Hosts

File Shares

Pools

RBDs

Admin

Tasks

Hosts

Name

Name ↓↑

9 Hosts

		Storage	CPU	Memory	Cluster	Role	Alerts	
	dhcp43-44.lab...	 4% 1.6 GB of 35.5 GB used	 39% 39 % used	 24% 217.9 MB of 990.7 MB used	...	OSD Host	0	⋮
	dhcp43-45.lab...	 8% 1.6 GB of 17.2 GB used	 21% 21 % used	 13% 251.1 MB of 1.8 GB used	...	Peer	0	⋮
	dhcp43-46.lab...	 4% 1.6 GB of 15.5 GB used	 39% 39 % used	 13% 218.2 MB of 990.7 MB used	...	Peer	0	⋮
	dhcp43-47.lab...	 8% 1.5 GB of 17.2 GB used	 20% 20 % used	 13% 252.2 MB of 1.8 GB used	...	Peer	0	⋮
	dhcp43-54.lab...	 11% 1.9 GB of 15.5 GB used	 45% 45 % used	 24% 238.1 MB of 990.7 MB used	...	Monitor	0	⋮
	dhcp43-55.lab...	 8% 1.6 GB of 17.2 GB used	 32% 32 % used	 22% 218.2 MB of 990.7 MB used	...	OSD Host	0	⋮

Clusters

Hosts

File Shares

Pools

RBDs

Admin

File Shares

Name

Name



No File Shares Detected

If no File Shares are detected after creating or importing a Gluster cluster, reconfigure the File Shares correctly.

- Clusters
- Hosts
- File Shares
- Pools**
- RBDs
- Admin

Pools

Name

Name ↕

Create

7 Pools

Name	Storage	Replicas	OSD	PGs	Quotas	Alerts	
Jeffallinone In Ceph Replicated	0% 0 % used	3	NA	128	Disabled	NA	⋮
JeffBPool In Ceph Replicated	0% 0 % used	3	NA	128	Disabled	NA	⋮
Jpool In Ceph Replicated	0% 0 % used	3	NA	192	30%	NA	⋮
Ju-Ec-Pool In Ceph Erasure Coded	0% 0 % used	EC Profile 2+1	NA	32	Disabled	NA	⋮
Jupool In Ceph Replicated	0% 0 % used	2	NA	64	Disabled	NA	⋮
MyPool In Ceph Replicated	0% 0 % used	3	NA	129	Disabled	NA	⋮
Rbd In Ceph Replicated	0% 0 % used	3	NA	128	Disabled	NA	⋮

- Edit
- Grow PGs
- Remove

Now Let's Create a Ceph Storage Pool...

- Clusters
- Hosts
- File Shares
- Pools**
- RBDs
- Admin
- Tasks

Pool > Create Pool

Create Pool

Basic Settings >

Advance Configuration

Quotas

Name

Pools to Create demo-pool

Type

Replicas

Cluster

PG Count Optimize for production developments
 Use the [PG Calculator](#) to estimate PGs required for production development. Although PG's can be increased at any time later, you cannot decrease the PG's. Using an optimal PG value is always encouraged as increasing PG's after the fact can cause performance degradation during redistribution of data in the cluster.

Optimized PG Count

- Clusters
- Hosts
- File Shares
- Pools**
- RBDs
- Admin
- Tasks

Pool > Create Pool

Create Pool

- Basic Settings**
- Advance Configuration** >
- Quotas**

Data-at-Rest Encryption OFF

Owner

Minimum Replicas

< Back Next > **Create Pool** > Cancel

- Clusters
- Hosts
- File Shares
- Pools
- RBDs
- Admin
- Tasks

Pool > Create Pool

Create Pool

Basic Settings

Advance Configuration

Quotas >

Quotas

 OFF

Max Percentage used

Max number of objects

[< Back](#) [Next >](#) [Create Pool >](#) [Cancel](#)

Clusters

Hosts

File Shares

Pools

RBDs

Admin

Tasks

[Pool](#) > Create Pool

Create Pool

1 Pool to create in cluster

Name	demo-pool
Type	Standard
PG Count	128

[Continue](#)

- Clusters
- Hosts
- File Shares
- Pools
- RBDs
- Admin
- Tasks

[Pool](#) > [Create Pool](#)

Create Pool



Create Pool(s) Submitted

An task has been submitted to the background for each Pool requested. You will be notified when processing is complete and the new Pools are ready for use.

[View Task Progress](#)

Now Let's Create a Ceph RADOS Block Device...

Clusters

Hosts

File Shares

Pools

RBDs

Admin

RBDs > Create RBD

Create RBD

Basic Settings

Backing Pool

Review

1

2

3

RBD Name

demo-rbd

RBDs to Create

1

demo-rbd1

Cluster Name

ceph

39.9 GB Available

Target Size

128

GE

128 GB needed (1 RBDs at 128 GB each)

⚠ Storage needed exceeds the amount of storage available in this cluster

< Back

Next >

Cancel

Clusters

Hosts

File Shares

Pools

RBDs

Admin

RBDs > Create RBD

Create RBD

Basic Settings

Backing Pool

Review

1

2

3

1 RBD to create in cluster

Name	demo-rbd1
Target Size	128

Backing Pool

A existing pool will be created in **ceph**.

Name	Type	OSDs	Replicas	PG Count	Journal Configuration	Quotas
rbd	Replicated	NA	3	128	NA	Disabled

[< Back](#)
[Create RBDs >](#)
[Cancel](#)

Clusters

Hosts

File Shares

Pools

RBDs

Admin

[RBDs](#) > [Create RBD](#)

Create RBD



Create RBD(s) Submitted

An task has been submitted to the background for each RBD requested. You will be notified when processing is complete and the new RBDs are ready for use.

[View Task Progress](#)

- Clusters
- Hosts
- File Shares
- Pools
- RBDs**
- Admin
- Tasks









RBDs

Name

Name ↓ ↑

Create

8 RBDs

Demo-Rbd1	in	Ceph	 0% used	Backing pool rbd (Shared)	Alerts NA	⋮
JeffBsBlockDe...	in	Ceph	 0% used	Backing pool jeffBPool	Alerts NA	⋮
Jrbid1	in	Ceph	 0% used	Backing pool jpool	Alerts NA	⋮
Ju-Rbd1	in	Ceph	 0% used	Backing pool jupool (Shared)	Alerts NA	⋮
Ju-Rbd2	in	Ceph	 0% used	Backing pool jupool (Shared)	Alerts NA	⋮
MyBlockDevic...	in	Ceph	 0% used	Backing pool rbd (Shared)	Alerts NA	⋮
MyBlockDevic...	in	Ceph	 0% used	Backing pool rbd (Shared)	Alerts NA	⋮
MyBlockDevic...	in	Ceph	 0% used	Backing pool rbd (Shared)	Alerts NA	⋮











- Clusters
- Hosts
- File Shares
- Pools
- RBDs
- Admin
- Tasks

Tasks

Name
 From: 
 To: 

Show: In progress Succeeded Failed

31 Tasks

	CreateRbd Task ID: f9adca34-ad4d-4f79-b53e-1549581793c8	Submitted 28 Apr 2017 16:16:09	Completed
	CreatePool Task ID: 154e557c-bf4f-4576-9867-2227085f0c67	Submitted 28 Apr 2017 15:44:12	Completed
	CreateRbd Task ID: eae0a19b-8c7c-4262-a978-7f23f4adb9c4	Submitted 28 Apr 2017 07:45:29	Completed
	CreateRbd Task ID: 557c0af2-d2a3-4fad-b42c-8dd0ca2d35fa	Submitted 28 Apr 2017 07:45:28	Completed
	CreateRbd Task ID: 8652f92f-b73d-463b-9935-8397f089fd33	Submitted 28 Apr 2017 07:45:28	Failed
	CreatePool Task ID: 169d5fad-8400-41bc-9bda-58eb0d42d9d5	Submitted 28 Apr 2017 07:45:09	Failed
	CreatePool Task ID: 37b6a069-92db-458c-b018-afb39aecde08	Submitted 26 Apr 2017 07:02:34	Completed
	CreatePool Task ID: 7fe5a3fa-f821-4379-b824-4bca0a8a9388	Submitted 26 Apr 2017 03:52:20	Failed
	CreatePool Task ID: 8610fa19-3894-42dd-ac57-ead79ebdc739	Submitted 26 Apr 2017 03:43:06	Failed
	CreatePool	Submitted	Completed

Road Map

RHS-C ROADMAP SUMMARY

RED HAT STORAGE CONSOLE (past)

RHSC 1.0 (2013)

- Gluster Support
- Provisioning
- Monitoring, Dashboard
- EoL (support extended)

RHS-C 2.0 (Aug 2016)

- Ceph 2.0 Support
- Provisioning
- Monitoring, Dashboard
- 2 Async releases
- EoL Feb 2018

RED HAT STORAGE CONSOLE (2017 & beyond)

RHS-C 3 (early Fall 2017)

- Gluster 3.2 Support (MVP)
- Ceph 2.x Support
- RHS Console 2 Equivalency for Gluster
- Improved Dashboard

RHS-C 4 (Spring 2018)

- **Gluster** add Vol types, Tiering, Data Management, CIFS/ NFS
- Snapshots, Geo-replication
- **Ceph** RGW, CRUSH map visualization
- RH Insights & Analytics

RHS-C 5 (Fall 2018)

- Gluster: Emerging features*
- CephFS enablement
- Ceph advanced troubleshooting, RBD Top, analytics
- RH Portfolio Enablement

UPSTREAM ACTIVITY (Past & Future)

Tendr1 1.0 (Nov 14)

- Framework creation

Tendr1 2.0 (May 18)

- Ceph + Gluster MVP

Tendr1 3.0 (Aug 2017)

- Gluster EC, Gluster provisioning

Red Hat Storage Console 4

Targeted for Spring 2018

Snapshot Functionality for Gluster Volumes and Ceph RBD

Gluster: Volume types (Sharded, Striped, Tiering), Data Management, Geo-replication, CIFS/ NFS enablement

Snapshots, Ceph: Add RADOS Gate


Ceph CRUSH Map visualization

Auto re-weight? *under investigation.*

Export to EFK (Elastic Search, FluentD, Kibana)? *under investigation*

Red Hat Storage Console 4

With Red Hat Insights & Analytics



Predictive analytics

This helps you save time and money addressing problems after the fact. Red Hat Insights lets you know about problems before they affect your environment.



Reduce human error, resolve quickly: Our Storage experts prepare the solution articles.

See potential problems and issues in real-time, without the need to wait for reactive error reports and logs, then fix with ease.



Better, faster, and stronger through experience

Red Hat Storage Console 5

Targeted for Fall 2018

Gluster: Emerging features - Compression, QoS

Ceph advanced troubleshooting, “RBD Top”, analytics

Ceph CRUSH editor

CephFS Enablement

Red Hat CloudForms: Basic storage reporting, show back / chargeback, cross linking

If you have product ideas or requests

Contact me! japplewh@redhat.com

RED HAT
SUMMIT

THANK YOU



plus.google.com/+RedHat



facebook.com/redhatinc



linkedin.com/company/red-hat



twitter.com/RedHatNews



youtube.com/user/RedHatVideos

The logo consists of a red speech bubble shape pointing downwards, containing the text "RED HAT" in a smaller font above "SUMMIT" in a larger, bold font, both in white.

RED HAT
SUMMIT

LEARN. NETWORK.
EXPERIENCE
OPEN SOURCE.