Ceph Management and Monitoring with Dashboard V2

Lenz Grimmer
Engineering Team Lead
SUSE Enterprise Storage

openATTIC History

2011: openATTIC founded by it-novum

–"Traditional" storage (NAS/SAN): SMB, NFS, iSCSI

2014: Added initial Ceph support

Feb. 2016: Collaboration with SUSE started

Nov. 2016: openATTIC project and team joins SUSE

2017: openATTIC 3.x focuses on Ceph Management/Monitoring exclusively

–Salt/DeepSea

–Prometheus

–Grafana

Dashboard v1 Overviev

- Added in Ceph Luminous
- Ceph health status, logs, performance metrics
- List of nodes, OSDs
- RBD images, mirroring status, iSCSI daemon status
- Python Backend (CherryPy)
- Javascript UI (Rivets.JS)
- New information added after Luminous
  - RGW details
  - MON list
  - Perf counters
  - Config settings browser

Dashboard v1 Limitations

- "read-only" - no management functionality
- No built-in authentication system
- Limited functionality of Rivets.JS to create a "real app"

- Intentions to evolve it into a full-blown management and monitoring web UI
  - Wishlist: http://pad.ceph.com/p/mimic-dashboard
  - See "dashboard in mimic" on ceph-devel - https://marc.info/?l=ceph-devel&m=151376737304374
  - Sage during the Dec. 2017 CDM call: https://youtu.be/YNfp_4S7mYE?t=28m37s

Dashboard v2 History

Jan. 2018: Initial discussions with Sage and John

POC of a Ceph Mgr Dashboard converted to Angular

- **http://pad.ceph.com/p/ceph-dashboard-angular-prototype**
- **https://github.com/tspmelo/ceph/tree/ceph-dashboard-angular/**

Feb. 22nd 2018: Dashboard v2 development branch created

- **https://github.com/openattic/ceph/tree/wip-mgr-dashboard_v2**

Milestone 1 (Dashboard v1 feature parity) merged on 2018-03-06

- **https://github.com/ceph/ceph/pull/20103**

Dashboard v2 Overview

- Modular Python backend (CherryPy), RESTful API
- WebUI (Angular 5), inspired by / derived from openATTIC UI
- Basic username/password authentication
- All features of Dashboard v1 from current master branch

Demo / Screencast / Screen Shots

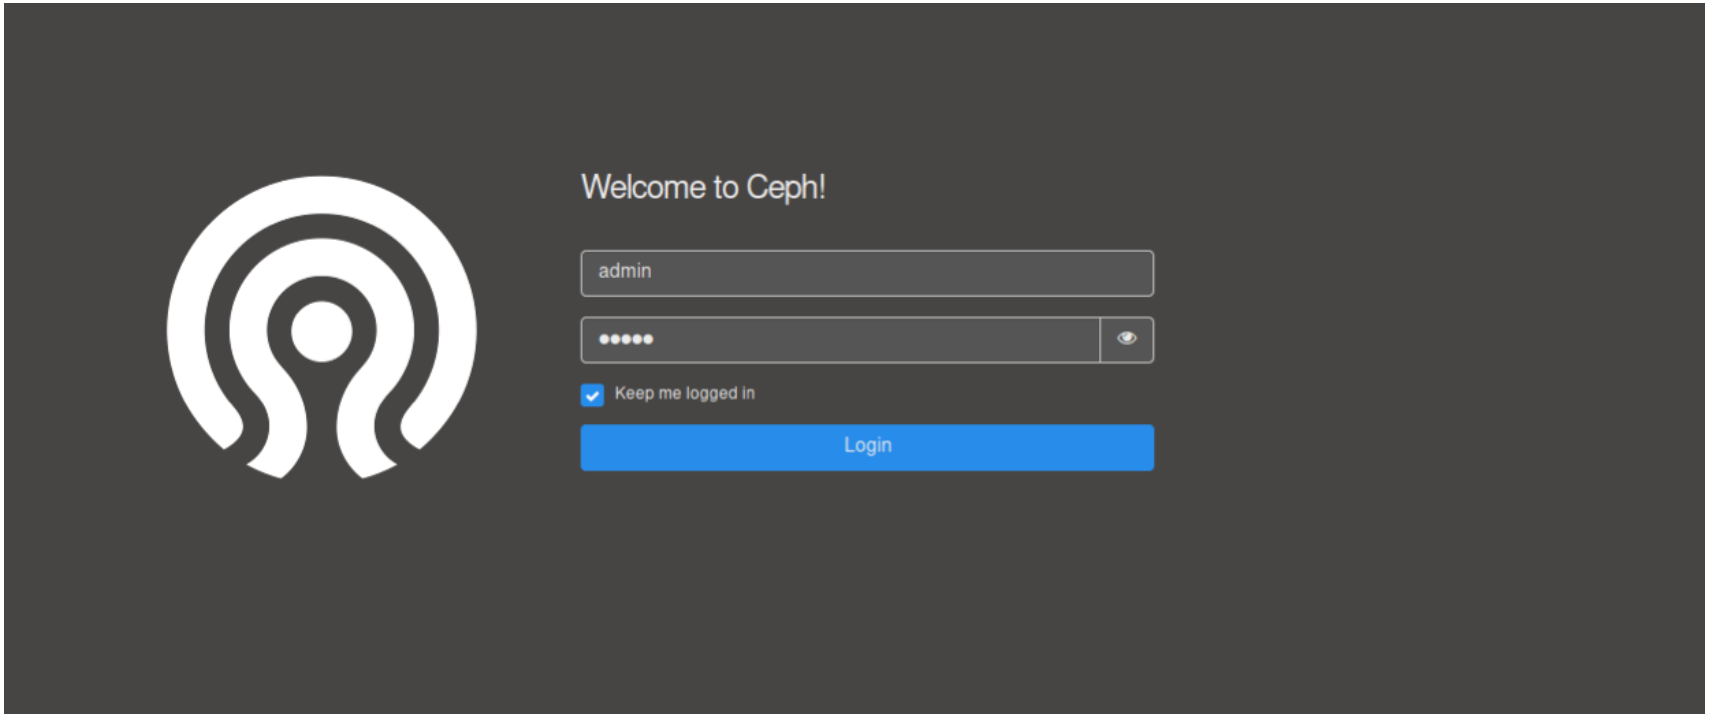
Welcome to Ceph!
admin
Keep me logged in
Login

ceph

Dashboard | Cluster | Block | Filesystems | Object Gateway

Health

Overall status: HEALTH_OK

MONITORS
3 (quorum 0, 1, 2)

OSDS
3 (3 up, 3 in)

METADATA SERVERS
1 active, 2 standby

MANAGER DAEMONS
active: x

Usage

259

Objects

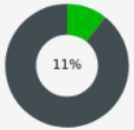

Raw capacity
(3.20GiB used)

Usage by pool

Pools

Name	PG status	Usage	Read	Write
cephfs_data_a	8 active+clean	0B / 8.89GiB	0 0 ops	0 0 ops
cephfs_metadata_a	8 active+clean	2.19KiB / 8.89GiB	0 0 ops	0 0 ops
.rgw.root	8 active+clean	1.09KiB / 8.89GiB	0 0 ops	0 0 ops
default.rgw.control	8 active+clean	0B / 8.89GiB	0 0 ops	0 0 ops
default.rgw.meta	8 active+clean	2.53KiB / 8.89GiB	0 0 ops	0 0 ops
default.rgw.log	8 active+clean	0B / 8.89GiB	0 0 ops	0 0 ops
rbd	8 active+clean	133B / 8.89GiB	0 0 ops	0 0 ops

Logs

Cluster log | Audit log

```
2018-02-23 12:00:00.000204 [INF]  overall HEALTH_OK
2018-02-23 11:58:49.835010 [INF]  daemon mds.a is now active in filesystem cephfs_a as rank 0
2018-02-23 11:58:49.746042 [INF]  daemon mds.a assigned to filesystem cephfs_a as rank 0 (now has 1 ranks)
```

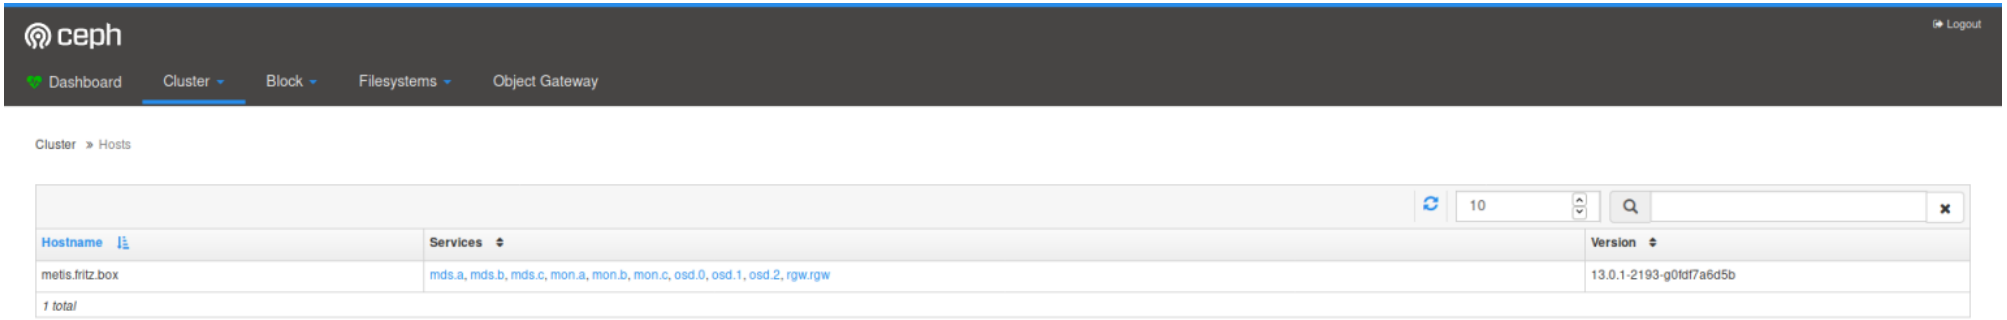
ceph
Logout
Dashboard
Cluster
Block
Filesystems
Object Gateway
Cluster » Hosts
10
Hostname
Services
Version
metis.fritz.box
mds.a, mds.b, mds.c, mon.a, mon.b, mon.c, osd.0, osd.1, osd.2, rgw.rgw
13.0.1-2193-g0fdf7a6d5b
1 total

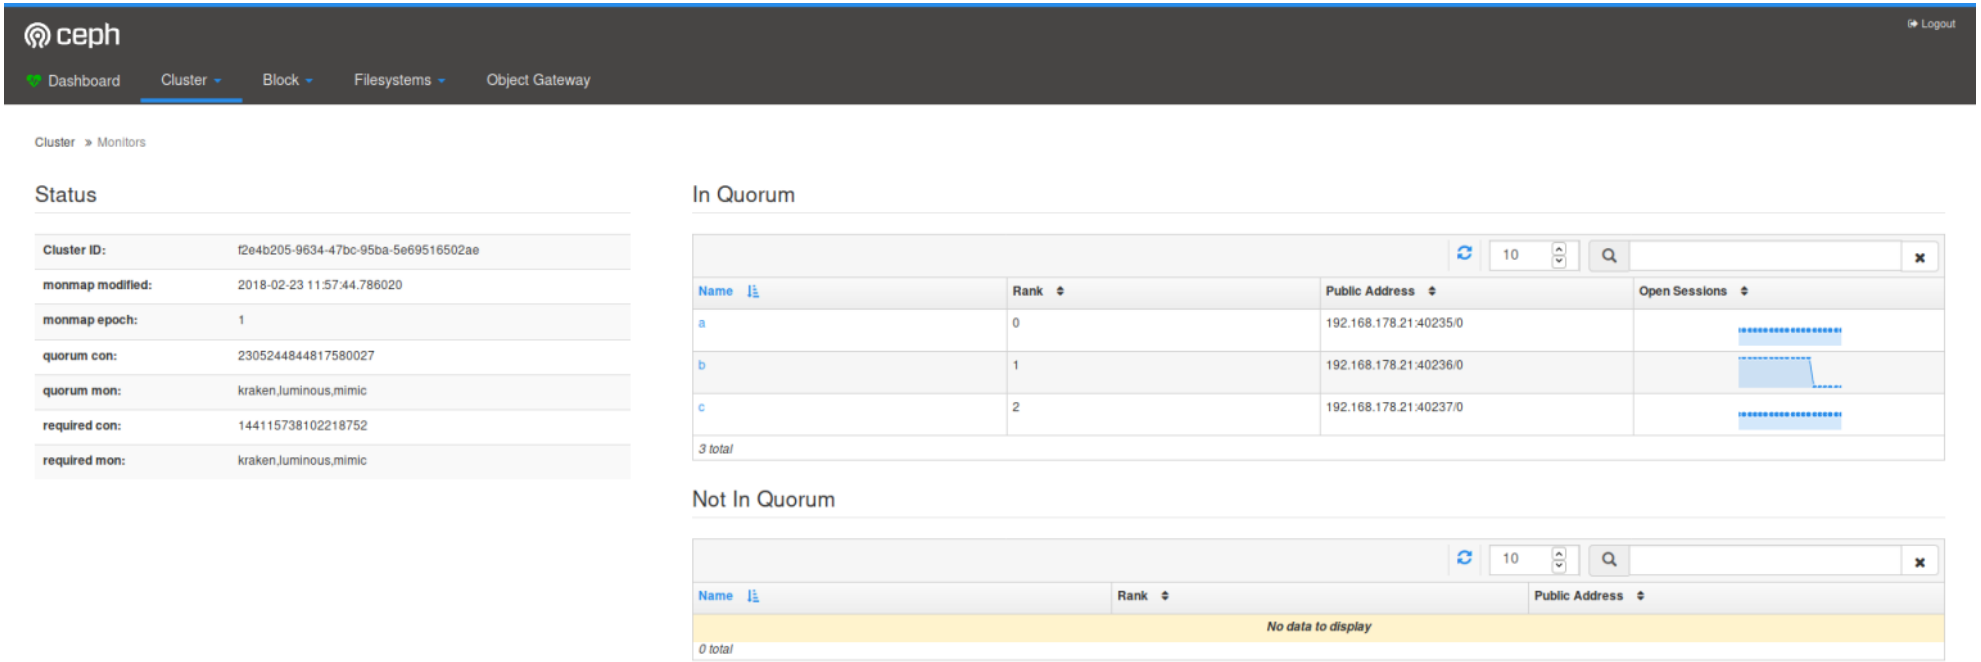
ceph
Logout
Dashboard
Cluster
Block
Filesystems
Object Gateway
Cluster » Monitors
Status
Cluster ID: f2e4b205-9634-47bc-95ba-5e69516502ae
monmap modified: 2018-02-23 11:57:44.786020
monmap epoch: 1
quorum con: 2305244844817580027
quorum mon: kraken,luminous,mimic
required con: 144115738102218752
required mon: kraken,luminous,mimic
In Quorum
10
Name
Rank
Public Address
Open Sessions
a 0 192.168.178.21:40235/0
b 1 192.168.178.21:40236/0
c 2 192.168.178.21:40237/0
3 total
Not In Quorum
10
Name
Rank
Public Address
No data to display
0 total

Performance Counters

mon.a

Name	Description	Value
mon.election_call	Elections started	0 /s
mon.election_lose	Elections lost	0 /s
mon.election_win	Elections won	0 /s
mon.num_elections	Elections participated in	0 /s
mon.num_sessions	Open sessions	5
mon.session_add	Created sessions	0 /s
mon.session_rm	Removed sessions	0 /s
mon.session_trim	Trimmed sessions	0 /s
paxos.accept_timeout	Accept timeouts	0 /s
paxos.begin	Started and handled begins	0.6 /s
paxos.begin_bytes	Data in transaction on begin	4.12k /s
paxos.begin_keys	Keys in transaction on begin	1.8 /s
paxos.begin_latency	Latency of begin operation	1.23G
paxos.collect	Peon collects	0 /s
paxos.collect_bytes	Data in transaction on peon collect	0 /s
paxos.collect_keys	Keys in transaction on peon collect	0 /s
paxos.collect_latency	Peon collect latency	0
paxos.collect_timeout	Collect timeouts	0 /s
paxos.collect_uncommitted	Uncommitted values in started and handled collects	0 /s
paxos.commit	Commits	0.6 /s
paxos.commit_bytes	Data in transaction on commit	4.05k /s
paxos.commit_keys	Keys in transaction on commit	2.4 /s
paxos.commit_latency	Commit latency	1.06G
paxos.lease_ack_timeout	Lease acknowledgement timeouts	0 /s
paxos.lease_timeout	Lease timeouts	0 /s
paxos.new_pn	New proposal number queries	0 /s

Ceph中国社区

IT大咖说 知识共享平台

Cluster » OSDs

10

Host	ID	Status	PGs	Usage	Read bytes	Writes bytes	Read ops	Write ops
metis	1	exists,up	56	1.14G / 10.8G			12.2 /s	0 /s
metis	0	exists,up	56	1.14G / 10.8G			8.4 /s	0 /s
metis	2	exists,up	56	1.14G / 10.8G			0.2 /s	0 /s

1 selected / 3 total

osd.1

Attributes (OSD map) | Metadata | Performance counter | Histogram

cluster_addr	192.168.178.21:6806/1557
down_at	0
heartbeat_back_addr	192.168.178.21:6807/1557
heartbeat_front_addr	192.168.178.21:6808/1557
id	1
in	1
last_clean_begin	0
last_clean_end	0
lost_at	0
osd	1
primary_affinity	1
public_addr	192.168.178.21:6805/1557
state	exists, up
up	1
up_from	9
up_thru	22
uuid	10e75fb1-1ef4-44b8-8136-b01e29d08333
weight	1

Service: any Level: basic

Name	Description	Type	Level	Default	Tags	Services	See_also	Max	Min
cluster_addr	cluster-facing address to bind to	entity_addr_t	basic	-	network	osd			
err_to_graylog	send critical error log lines to remote graylog server	bool	basic	false			log_to_graylog log_graylog_host log_graylog_port		
err_to_stderr	send critical error log lines to stderr	bool	basic	false true					
err_to_syslog	send critical error log lines to syslog facility	bool	basic	false					
fsid	cluster fsid (uuid)	uuid_d	basic	00000000-0000-0000-0000-000000000000	service	common			
host	local hostname If blank, ceph assumes the short hostname (hostname -s)	std::string	basic		network	common			
log_file	path to log file	std::string	basic	/var/log/ceph/$cluster-$name.log			log_to_stderr err_to_stderr log_to_syslog err_to_syslog		
log_graylog_host	address or hostname of graylog server to log to	std::string	basic	127.0.0.1			log_to_graylog err_to_graylog log_graylog_port		
log_graylog_port	port number for the remote graylog server	int64_t	basic	12201			log_graylog_host		
log_stop_at_utilization	stop writing to the log file when device utilization reaches this ratio	double	basic	0.97			log_file	1	0
log_to_graylog	send log lines to remote graylog server	bool	basic	false			err_to_graylog log_graylog_host log_graylog_port		
log_to_stderr	send log lines to stderr	bool	basic	true false					
log_to_syslog	send log lines to syslog facility	bool	basic	false					
mds_cache_memory_limit	target maximum memory usage of MDS cache This sets a target maximum memory usage of the MDS cache and is the primary tunable to limit the MDS memory usage. The MDS will try to stay under a reservation of this limit (by default 95%; 1 - mds_cache_reservation) by trimming unused metadata in its cache and recalling cached items in the client caches. It is possible for the MDS to exceed this limit due to slow recall from clients. The mds_health_cache_threshold (150%) sets a cache full threshold for when the MDS signals a cluster health warning.	uint64_t	basic	1073741824		mds			

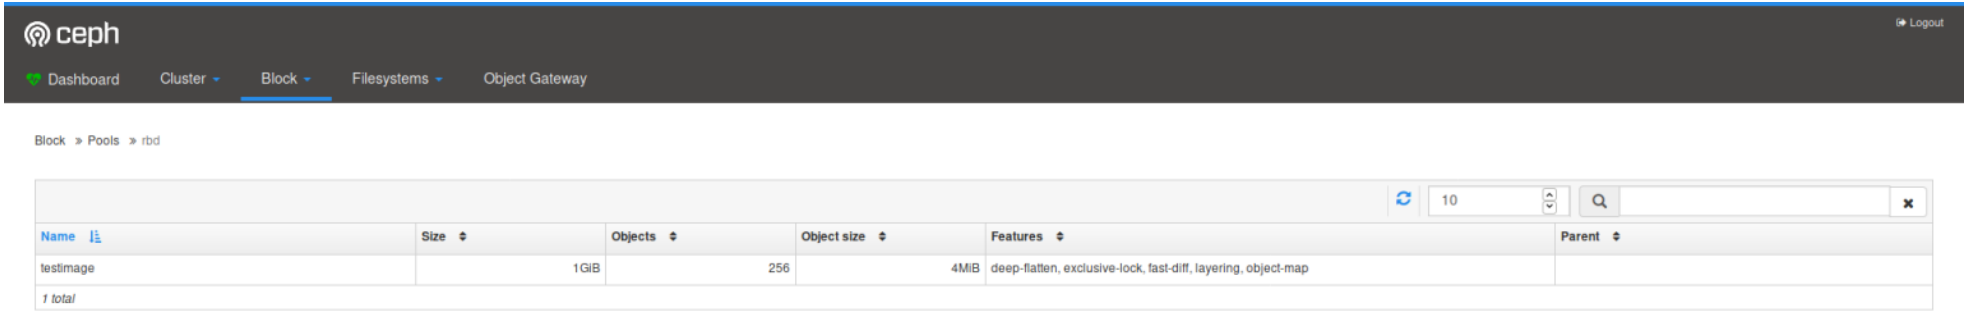
ceph
Logout
Dashboard
Cluster
Block
Filesystems
Object Gateway
Block » Pools » rbd
10
Name
Size
Objects
Object size
Features
Parent
testimage
1GiB
256
4MiB
deep-flatten, exclusive-lock, fast-diff, layering, object-map
1 total

ceph

Dashboard | Cluster | Block | Filesystems | Object Gateway

Filesystem » cephfs_a

0 Clients

Ranks

Rank	State	Daemon	Activity	Dentries	Inodes
0	active	a	Reqs: 0 /s	10	12
1 total					

Standby daemons	b, c

Pools

Pool	Type	Used	Avail	Usage
cephfs_data_a	data	0B	8.89GiB	
cephfs_metadata_a	metadata	2.19KiB	8.89GiB	
2 total				

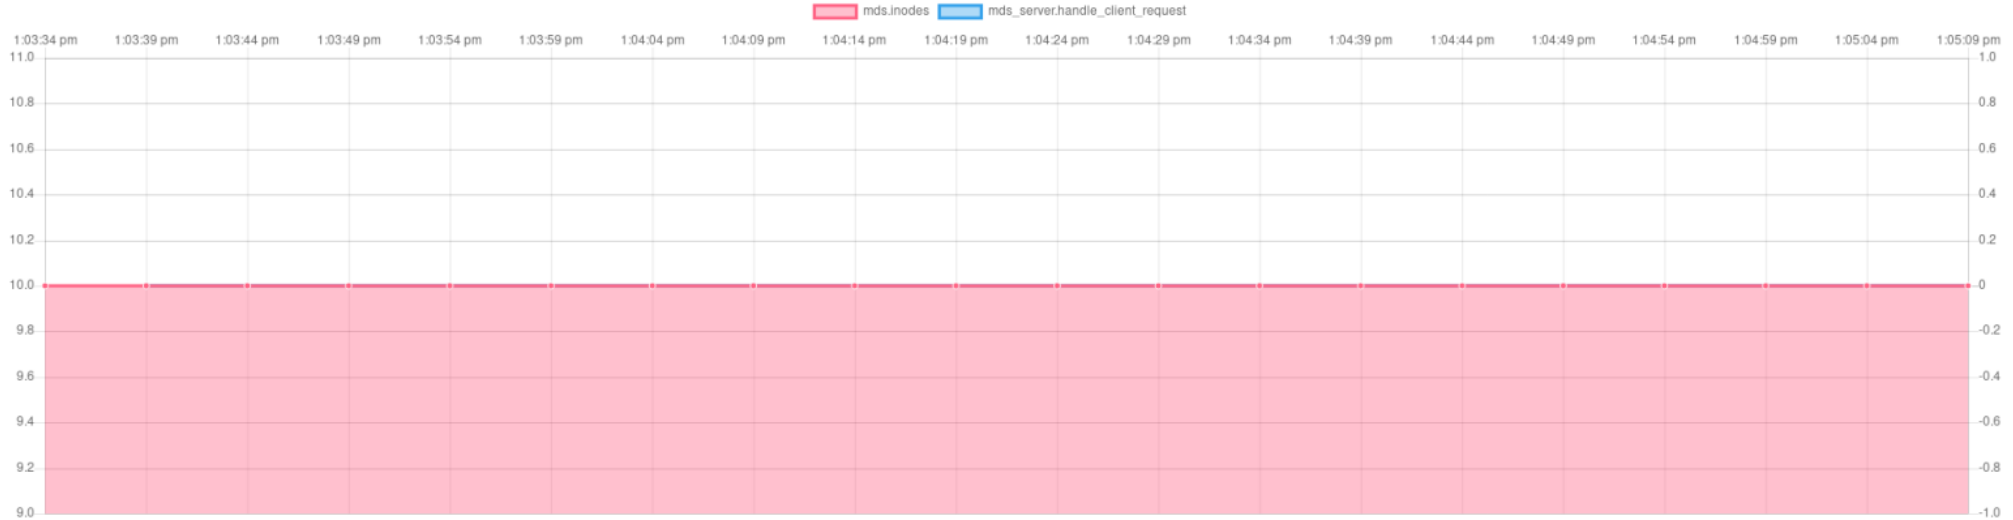

Ceph中国社区

IT大咖说 知识共享平台

Object Gateway

10

ID	Hostname	Version
rgw	metis.fritz.box	13.0.1-2193-g0fdf7a6d5b

1 selected / 1 total

Details | Performance Counters

arch	x86_64
ceph_version	ceph version 13.0.1-2193-g0fdf7a6d5b (0fdf7a6d5bba9ab1610a9983ebc49b5daaef5aac) mimic (dev)
cpu	Intel(R) Core(TM) i7-4600U CPU @ 2.10GHz
distro	opensuse
distro_description	openSUSE Tumbleweed
distro_version	20180207
frontend_config#0	civetweb port=8000
frontend_type#0	civetweb
hostname	metis.fritz.box
kernel_description	#1 SMP Tue Feb 13 17:02:01 UTC 2018
kernel_version	4.15.3-300.fc27.x86_64
mem_swap_kb	7864316
mem_total_kb	7850680
num_handles	1
os	Linux
pid	2503
zone_id	fef7b644-ca96-41ec-9739-af8118034c49
zone_name	default
zonegroup_id	891d5150-9728-466e-8a54-f01a76202d88
zonegroup_name	default

Next Steps

Reaching feature parity with openATTIC 3.x

- RBD management (create/modify/delete)
- RGW management (create/modify/delete users, keys, buckets)
- Ceph Pool management (create/modify/delete)
- SSL/TLS support
- Embedded Grafana Dashboards
- iSCSI target management (TCMU runner)
- NFS Ganesha management

Getting involved / Planning

Feedback, Patches, Bug reports welcome!

https://ceph.com/get-involved/