

	2024-3				2024-4				2024-5				2024-6				2024-7				2024-8														
	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35									
Calamari	Calamari																																		
UI: Admin - Password Retrieval																																			
Remove plugin system from the cthulhu																																			
Fix SELinux for Calamari on RHEL7																																			
Calamari should use actual OSD devices on (...)																																			
remove workaround for grains in relax_salt_perms.sls																																			
Use FQDN for ceph_deploy AND ServerName in (...)																																			
asok handling code in minion module should (...)																																			
mon counters incorrectly calculated when (...)																																			
Gather and store historical cluster state																																			
Expose historical cluster state in REST API																																			
Backend support for account recovery by email																																			
Security for graphite connections																																			
Efficient latest-value querying for graphite (...)																																			
Diamond collector should use librados directly																																			
Per-pool PG details and state counters																																			
Manage pool snapshots																																			
Respect GPG auth when installing salt+diamond																																			
On-demand load of full PG detail																																			
Internationalizable events and request descriptions																																			
Proxy support in calamari bootstrap																																			
Add some examples for bulk DELETE and PATCH (...)																																			
403 Errors on API return HTML																																			
Document expected response codes in REST (...)																																			
/api/v2/key entries attaching extra metadata (...)																																			
Collect ceph version per service. expose (...)																																			
Show a friendly message when apache is running (...)																																			
Metadata tags on UserRequests																																			

<p>Metadata tags on events</p> <p>Bulk requests (many requests in one)</p> <p>Give events machine-readable type identifier</p> <p>Make UserRequests persistent</p> <p>Consume salt presence events for server upness</p> <p>Use public salt interfaces only</p> <p>Avoid generating "OSD went down" events during (...)</p> <p>Generate postgres account credentials at (...)</p> <p>Remove workaround for saltstack #11919</p> <p>Remove workaround for saltstack #11928</p> <p>Communication failure between salt-master (...)</p> <p>Annotations API</p> <p>Epic: Cache Tier Pool CRUD</p> <p>Epic: Erasure Code CRUD</p> <p>Epic: RBD Management</p> <p>Log spam: "Re-opening connection to salt-master" (...)</p> <p>Cleaner handling of CRUSH maps with non-hostname-named (...)</p> <p>Add cache tier object to cthulhu</p> <p>Expose cache tier object CRUD in rest-api</p> <p>Handle clusters where FSID isn't set in ceph.conf</p> <p>CRUSH: Add CRUD to calamari api for crush_rules</p> <p>Update django to fix known security flaws</p> <p>RFE for calamari-ctl initialize: show Calamari (...)</p> <p>Remove log handling from cthulhu</p> <p>sense SELinux and warn about it</p> <p>Reserve blue text for clickable items</p> <p>Link to product doc during setup</p> <p>Terminology Cleanup</p> <p>Buttons and tooltips</p> <p>Drop down arrows for drops</p> <p>Clean up Delete Pool warning dialog</p> <p>CRUSH: Filter retrieval of crush nodes for (...)</p> <p>LOCALE: Allow the user to select the locale</p> <p>LOCALE: Support localization for 'Manage'</p>	
--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--

CRUSH: Allow weight of bucket nodes be set (...)	
CRUSH: Report Not modified on the crush_node (...)	
LOCALE: add placeholders to displayed text (...)	
LOCALE: add placeholders to displayed text (...)	
LOCALE: create a non-sense language locale (...)	
LOCALE: create howto documentation for new (...)	
LOCALE: Document that locale is chosen based (...)	
LOCALE: check that features in changelog (...)	
LOCALE: Upgrade Angular JS from 1.2 to 1.3 (...)	
BE: resolve a host to crush node, decorate (...)	
remove workaround in rest-api Validating-serializer (...)	
rename "calamari-clients" packages to "romana"	
calamari-server postinstall should use supervisor (...)	
Host info pop-up window is empty in Calamari (...)	
CentOS: Calamari UI uses lsb_distrib_description (...)	
calamari_common salt_wrapper sucks for diagnosability	
The cthulhu should have the thread clear the (...)	
Add Time Filtering to the events target.	
RGW type in /server endpoint shows ambiguous (...)	
Include more health info in /mon targets.	
Include all Addition and Removal of Cluster (...)	
calamari stats broken on infernalis	
Ceph support Plan tasks or Timing task Function	
rest-api: api/v1/space returned a history (...)	
Diamond: One of my OSD has no performance (...)	
Diamond: when pull out the disk first, then (...)	
cthulhu: When i change a host's hostname, (...)	
Add rbd and radosgw-admin support to /cli (...)	
update pg_num of pool will return 500error,	
Error 500 :: Internal server error :: ubuntu (...)	
Calamari dashboard :: can't connect to the (...)	New 0%
Calamari Rest API: "api/v2/cluster/<fsid>/cli (...)	Calamari -1.3-rc1 80%
Calamari - 1.3-rc1	New 0%
installing Salt 2014.1.10 via calamari-bootstrap (...)	

ice-tools needs to install ice_setup (...)	New 0%
Calamari - 1.4 backlog	Calamari -1.4 backlog 9%
EPIC: as an administration I want to (...)	New 0%
refine SMART error reporting to have (...)	New 0%
SMART, support for other RAID controllers	New 0%
expose events in the API when smart (...)	New 0%
decorate OSD API endpoint with smart (...)	New 0%
ceph.py assumed no '-' in ceph cluster (...)	New 0%
admin in romana is unused. remove it.	New 0%
calamari crush location needs to deal (...)	New 0%
calamari virtualenv applies SWIG flag (...)	New 0%
Add RGW data to endpoints	New 0%
Calamari - 1.3.1	
UI reports "Cluster Updates Are Stale. (...)	New 0%
dashboard widgets IOPs and Usage are (...)	
Ceph	Ceph
set up a no-atomic-ops gitbuilder	
mon: check for slow performing osds	
track object states to inform error injection/testing	
msgr workloads	
Save kernel core dumps on all of our test (...)	
install + configure package everywhere	
The ceph command should time out	
osd: add opaque 'class <name> <foo>' cap (...)	
rados ager	
osd: close read hole	
mon: simple dm-crypt key management	
log: stop logging in statfs reports usage (...)	
Add perf tools to the kernel debian package (...)	
osd: localized reads (from replicas) reordered (...)	
auth startup reports "ObjectNotFound" when (...)	
Show usable space via CLI/API	
formatter: send output to ostream	
EC: add mechanism for mon to detect and whitelist (...)	

<p>deprecate rados_pool_create</p> <p>What's the policy on URL stability for public (...)</p> <p>erasure-code: last stripe is not truncated</p> <p>Epic: Kerberos/LDAP Support</p> <p>related to integrity of objects</p> <p>erasure-code: use LRU to cache decoding matrix (...)</p> <p>osd: ec plugin leak</p> <p>msgr: encode osd epoch in nonce to avoid (...)</p> <p>Add support for collecting usage information (...)</p> <p>erasure-code: better LRC strategy</p> <p>ceph CLI: Help is missing -k keyring option</p> <p>Cache warmup and ejection</p> <p>annotate config options</p> <p>ceph should fail when an argument is of the (...)</p> <p>erasure-code: ErasureCode::encode overhead (...)</p> <p>osd: hard cap on PGs per OSD</p> <p>libgoogle-perftools4: tcmalloc performance (...)</p> <p>osd: store scrub error state in kv store; (...)</p> <p>osd: add scrub result query interface</p> <p>librados: add ability to read a specific (...)</p> <p>rados: add ability to read a specific replica/shard (...)</p> <p>librados, osd: per-object scrub operation</p> <p>librados: improve get_objects and get_position (...)</p> <p>rados: new import from pipe/file</p> <p>librados: set user_version operation</p> <p>option to disable erasure code plugin version (...)</p> <p>ceph: 'daemon list' command to enumerate (...)</p> <p>add ceph-report-parsing</p> <p>osd: an all/mostly delete pg log stalls during (...)</p> <p>mon: finer-grain Paxos::is_readable() / PaxosService::is_readable()</p> <p>mon: refactor OSDMonitor map methods to use (...)</p> <p>mon: move OSDMonitor osd_epoch value into (...)</p> <p>msgr: optionally encrypt messenger session</p> <p>add stop_scrub command for ceph</p>	
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--

aio gets EPERM when update-grub runs
 generic plugin loader framework
 ceph-disk: suppress non fatal partx errors
 Standardize Ceph removal commands
 should not "rm -rf /etc/ceph" when uninstalling (...)
 erasure-code: measure bandwidth usage
 erasure-code: add cross datacenter example (...)
 erasure-code: add chance_test_backfill_full (...)
 erasure-code: lrc k/m/l sanity check must (...)
 replace snprintf(s,...); x = strlen(s); with (...)
 document include/xlist.h
 replica read: do proper locking on replica (...)
 "ceph pool set-quota max_bytes" fails to (...)
 libmailrados: Mailbox storage on RADOS
 cmake
 add an option for ceph_test_rados to have (...)
 mon: cull / scrub unused crush rules when (...)
 ceph-post-file key should only be readable (...)
 scrub/repair: persist scrub results.
 should make librocksdb static link default (...)
 Ceph CLI + python bindings restrict string (...)
 Write Developer Guide
 src/test/cephtool-test-mon.sh Hooked allocator (...)
 OSD`s doesnt start after reboot
 SubsystemMap.h: 62: FAILED assert(sub < m_subsys.size())
 ceph-osd systemd unit files incomplete
 crush: deprecate tree buckets
 RFE: expose OP_OMAP_RMKEYRANGE in librados (...)
 ceph-osd: some flags are not documented in (...)
 ceph-disk deactivate fails with --verbose, (...)
 catch std::bad_alloc and die with helpful (...)
 qa/benchmark: Explore using Filebench for (...)
 add an option to force a down osd to be marked (...)
 Document old design

implement new design
rados: support omap (leveldb) and locator (...)
pkgconfig files for libraries
crush: create a visualizer for crush maps
rados: support asynchronous replication of (...)
Efficient access in distributed data centers
make Ceph build faster, take less RAM to (...)
Minimize recovery time of a degraded cluster
figure out how to best set the heartbeat (...)
log: detect dup messages
Automatically determine routing data based (...)
add ability to make ceph ignore io errors (...)
qa: add test to compile a simple program (...)
recovery from partially partitioned networks
Decode exception descriptions don't show (...)
Journal entries should record last known (...)
OpTracker: support hierarchies of request (...)
create and automate scenario based testing
Deep scrub CPU limit behavior
Backport option to disable deep scrub to (...)
document src/objclass/objclass.h
Some clone errors aren't repaired
objclass: maintain global namespaces
repair should handle snapset/clone discrepancies
Objecter: support multi-read-from-replica
pipeline large object recovery
cephtool: sanitize extra args before configuring (...)
expand upgrade tests to be able to test downgrade
packaging, admin_socket: create ceph group, (...)
Objecter,OSD: make it easy to determine when (...)
make osd crush placement on startup handle (...)
backfill_toofull => try backfilling another (...)
ceph-rest-api authentication
Ability to set up/down/in/out based on CRUSH (...)

Allow partial dump of pg statistics
keyring leading spaces
Add a new Info state to ceph health
rest-api: support commands requiring 'w' (...)
build: error out on inappropriate use of (...)
Hinted recovery
qa: test encoding semantics, not just being (...)
Make units of measurement constant and definable (...)
remove dead code from OSDMonitor::should_propose
Per CRUSH type mon_osd_down_out_interval
Illogical defaults in Python bindings: conffile=None (...)
better wording when authentication fails
Create a health severity between OK and WARN
OpTracker output doesn't include op size (...)
Cancel RADOS Bench writes and still do reads
please clarify copyright and the license
investigate having the messenger (or dispatch (...)
ceph: display tier relationships visually
shorten window of highest risk during recovery
include "ChangeLog" and/or "NEWS" files to (...)
add nonce/sequence number associating the (...)
better error reporting from ceph-disk
leveldb: initial inquiries on how we can (...)
Improve threading for ObjectCacher
Uniqueness between two or more CRUSH ruleset (...)
request for tools/commands to see hits/misses (...)
erasure-code: cache minimum-to-decode results
erasure-code: optimize IsReadablePredicate (...)
admin socket: allow multiple pending requests
Crush optimization for unbalanced data/pg (...)
mon: audit read-only commands on a lower (...)
OSD: block until we have the same map as (...)
osd: prioritize recovery of OSDs with most (...)
stop backfill when osd becomes too full

<p>rework bufferlist::*aligned* functions</p> <p>ceph df reports % of global size used instead (...)</p> <p>Issues found by Clang</p> <p>do not leak ceph-disk activate lock to the (...)</p> <p>try to make address sanitizer work</p> <p>forbid erasure code profile modifications (...)</p> <p>make it harder to remove cluster data pools</p> <p>When replaced a disk, the CRUSH weight of (...)</p> <p>gtest: update, move to submodule</p> <p>create selinux policies for ceph-mon, ceph-osd, (...)</p> <p>osd: mark pg and use replica on EIO from (...)</p> <p>Cluster unreachable after restart</p> <p>ECUtil::HashInfoRef should have a NONE value</p> <p>OSD level metadata mismatch handling</p> <p>PG metadata corruption detection and handling</p> <p>Object metadata mismatch detection and handling</p> <p>Create new git repo for json_spirit</p> <p>librados: mark old objects_begin interface (...)</p> <p>double resource for setting up ceph-deploy</p> <p>mon,osd: long-term non-clean PGs prevent (...)</p> <p>Per-pool PG recovery prioritisation</p> <p>Permanently mark daemons and clusters which (...)</p> <p>OSD Keeps Going Down</p> <p>Initial pg_num choice is static, different (...)</p> <p>rados command line gives no access to striper (...)</p> <p>OSDs get marked as down in docker</p> <p>cache pool: try not to keep around very old (...)</p> <p>Schedule scrubbing by considering PG's last_scrub_timestamp (...)</p> <p>Handle delete log entries in merge_log out (...)</p> <p>replicas need to track unstable objects to (...)</p> <p>CLI error numbers are not described anywhere</p> <p>mon add check to verify compatible version (...)</p> <p>Generic support for plugins installation (...)</p> <p>invalid syntax in Apache config file examples</p>	
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--

docs: federated rgw setup needs to be trimmed (...)

docs: federated rgw setup keyring section (...)

docs: federated rgw setup create-pools section (...)

docs: federated rgw setup create pools section (...)

docs: federated rgw setup create pools section (...)

docs: federated rgw setup uses Debian-based (...)

docs: federated rgw setup assumes that the (...)

docs: federated rgw setup doesn't warn about (...)

docs: federated rgw setup needs to normalize (...)

docs: federated rgw setup assumes user is (...)

docs: federated rgw setup is not explicit (...)

docs: federated rgw setup should post example (...)

Improve scrubbing throughput

throttle objects and bytes / second for scrub/deepscrub

ceph health warn if pgs are not being scrubbed (...)

LMDB key/value backend for Ceph

New feature : Ceph should warn if the kernel (...)

move ceph-deploy(8) man page into ceph-deploy (...)

ceph-disk breaks when using non-existing (...)

mons: add switch to never delete pools, only (...)

make sure notcmalloc builds continue working

erasure-code: hitchhiker plugin

test: verify that expected feature bits are (...)

add a command to the ceph tool which returns (...)

Elector: throttle election attempts from (...)

Certain configuration options not safe for (...)

improve OPTION macro to add a check function (...)

add HELP command to ceph interactive shell

pg a copy_from->pg b peered (blocked on backfill (...)

OSD: ensure that filestore queue_transaction (...)

msgr: throttle max open sockets

mon: automatically blacklist clients after (...)

change ceph-mon TCP port to an IANA-registered (...)

Show osd as NONE in "ceph osd map <pool> (...)

Workqueue timeouts and the filestore sync_entry (...)
 mon: add ability to gradually remove an osd
 OSD crashes under cached cluster benchmark
 ceph-cli: let "ceph report" support compressed (...)
 erasure code plugin for tornado code
 data is no well-distributed among osds of (...)
 Does the ceph debian package really need (...)
 removing cached objects doesn't quickly remove (...)
 init scripts: try to detect when a daemon (...)
 Messenger: use setsockopt to set sk priority (...)
 Implement a minimal access keyring for read-only (...)
 Support for building on Alpine Linux
 Please publish rados-java artifacts to central (...)
 Why copy a 3T image without any data takes (...)
 PGMonitor::get_health get ops' blocked the (...)
 FileSystem become read-only after over 6 (...)
 create the upgrade test suite for gmt and (...)
 Make "rados bench" support json output
 Translate fadvise flag NOCACHE into DONTNEED (...)
 OSDMonitor refactor
 compile warning on gcc-4.8.3
 automatically set journal file name as osd.{osd-num} (...)
 ceph-disk: will not work with an software (...)
 osd: prioritize more degraded PGs for recovery (...)
 Add flag to delete a pool and its cache tier.
 ceph-disk: use systemd-escape instead of (...)
 ceph-authtool(8) doesn't mention namespace
 ceph-objectstore-tool: add feature to adjust (...)
 OSD assert with 'Current_Pending_Sector'
 On osd startup we should log all configuration (...)
 librados: c + python api for list_snaps()
 scrub/repair: add librados APIs
 add checksum for the decode/encode
 AsyncMessenger: Support memory pool per worker

Async support for Ceph-fio-objectstore
 objectstore_bench did not support newstore
 osd assert is too crudely, maybe notify mon (...)
 Create can_discard_op_initial_decode and (...)
 ceph-disk list not support ID_PART_ENTRY_* (...)
 EC overwrites
 Activation support for EC Overwrites
 EC overwrites -- cleanup of ReplicatedPG::do_op (...)
 EC overwrites: PGBackend needs to be (...)
 PGBackend: add a way for the Backend (...)
 ECBackend: delay writes which cause (...)
 ECBackend cache extents with unapplied (...)
 EC Overwrites: create temporary testing (...)
 EC Overwrites: ECBackend TPC support
 update ec overwrites implementation (...)
 update ec overwrites implementation (...)
 EC Overwrites: Phase 2: Peering support (...)
 EC Overwrites: Phase 2: PGLog support (...)
 performance counters for ECbackend cache
 EC Overwrites: work out implications (...)
 updated ec overwrites to support fadvise
 ec overwrites write zeroes efficiently
 EC Overwrites: Missing Test Case: aligned (...)
 EC Overwrites: Missing Test Case: multiple (...)
 EC Overwrites: Missing Test Case: zero
 degrade a log message to a default level
 scripts/ceph-release-notes does not handle (...)
 add backfill targets to osd stats and a way (...)
 Show the total pg on each osd.
 the method 'Message::decode_message' is too (...)
 Show the pool quota info on ceph df detail (...)
 Crush when get immediate parent quickly reply.
 Command to override Incomplete PGs
 scrub/repair: repair corrupted/missing objects

<p>unsafe handle_conf_change() methods</p> <p>Write software version to data store after (...)</p> <p>ceph-crush-location manpage missing</p> <p>Installation (Quick) mount disk question</p> <p>mon: report pool deletion in "ceph -w"</p> <p>Can't add new osd an existing cluster.</p> <p>move rados-classes from /usr/lib to /usr/lib/ceph</p> <p>Make SOCKET_BACKLOG configurable</p> <p>Improve ceph-disk error message when there (...)</p> <p>osd: try to flush/evict object from cache (...)</p> <p>NEON error when building ceph on ARM v7 targets</p> <p>uniquely identify /tmp build directories</p> <p>Enhance KeyValueDB to include atomic operator (...)</p> <p>CAS Pools</p> <p>Tiering: Object Redirects</p> <p>Faster Peering -- primary remember infos/missings (...)</p> <p>Add --wait-for-complete flag to ceph pg scrub (...)</p> <p>general pg inconsistency report</p> <p>choose randomly among 3 different filestore_merge_threshold/filestore_split_multiple (...)</p> <p>modify ceph-qa-suite to use tc to inject (...)</p> <p>unit files missing after upgrade 10.2.1</p> <p>Use power crc32-vpmsum to accelerate crc32c (...)</p> <p>OS Recommendations do not handle distribution-included (...)</p> <p>messenger only uses IPv4 address of hostname</p> <p>print out error on failure to load OSDMap (...)</p> <p>SMR Benchmarking Scripts and libzbc code (...)</p> <p>AIO version of selfmanaged_snap_create/selfmanaged_snap_remove</p> <p>ceph-objectstore-tool shouldn't fully lock (...)</p> <p>Lost xattr from all osd's</p> <p>Take pipe::connect() returned errno to rados_connect() (...)</p> <p>usability: allow decreasing pg_num when pg_num (...)</p> <p>Include more detail in perf dump</p> <p>rados put: use the FULL_TRY flag to report (...)</p> <p>rados CLI's "rmomapkey" and "setomapval" (...)</p>	
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--

<p>Creating EC pool using replicated ruleset (...)</p> <p>Fix the Ceph and Ceph UT compiling error (...)</p> <p>HEALTH_ERR status is misleading</p> <p>Compiling with --without-cython seems to (...)</p> <p>ceph should warn the administrator when it (...)</p> <p>Jewel ceph-dbg packages are missing for Debian-based (...)</p> <p>Can't get pool's cache mode parameter via (...)</p> <p>Accelerate crush_do_rule with SIMD</p> <p>[systemd]: Ceph doesn't remove symlinks for (...)</p> <p>using RPC to eliminate complexity of network</p> <p>leader should recognize its own clock skew (...)</p> <p>osd: add an interface to allow an object (...)</p> <p>reweight-by-utilization needs a subtree option</p> <p>test radosgw with messenger failure injection (...)</p> <p>Support creating objects without network (...)</p> <p>osdmap calculation frequency option request</p> <p>frequent blocked request impacting Ceph client (...)</p> <p>Replace past_intervals with more compact (...)</p> <p>support shared ccache dirs across users</p> <p>ceph df should show correct pool usage</p> <p>readline isn't configured or installed in (...)</p> <p>ceph-disk may use blivet</p> <p>Add S3 lifecycle configuration validate</p> <p>Have OSD mark itself down while still running</p> <p>Configurable OSD Heartbeat packet size (MTU)</p> <p>RGW: Create users with a tenant via the admin (...)</p> <p>Where physically stored http headers for (...)</p> <p>Support lifecycle validate in s3test</p> <p>Python librados OmapIterator segfaults if (...)</p> <p>Ability to add comments in certain views (...)</p> <p>Add mon_osd_down_out_subtree_max_osd</p> <p>ceph-disk: Implement easy replace of OSD</p> <p>librados python binding reference is not (...)</p> <p>Script to find broken URLs in Ceph source (...)</p>	<p>New 0%</p> <p>New 0%</p>
--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------

<p>Output only the latest require_{release}_osds (...)</p> <p>ceph_objectstore_tool,rados tool: add commands (...)</p> <p>Merge snapshot of the cloned image back to (...)</p> <p>'journaling' feature affects write bandwidth</p> <p>If there is one FULL OSD, All IO the the (...)</p> <p>Expose snap trim speed along with client (...)</p> <p>detect broken links in documentation automatically</p> <p>Cannot remove striped object</p> <p>run check_commands.sh as part of make check</p> <p>ceph osd new -i should read from stdin</p> <p>ceph-disk should use the python random function</p> <p>mon/test: increase coverage of test-mon-msg</p> <p>give scrub a start_time property</p> <p>Change default filestore 'omap' backend to (...)</p> <p>cluster health is ok, but client IO hang (...)</p> <p>Devuan support</p> <p>Remove ceph_osd_mkfs function from src/ceph-disk/ceph_disk/main.py</p> <p>http://www.irchelp.org/irchelp/new2irc.html (...)</p> <p>Constants in code should be replaced with (...)</p> <p>Add API to cancel AIO request (for RBD and (...)</p> <p>ceph df should report %USED and MAX AVAIL (...)</p> <p>optimize some timer</p> <p>Creating type alias for complex type definitions</p> <p>ceph 12.2.0 health check failed osd down</p> <p>discover ip node in ceph cluster</p> <p>Ceph iSCSI Gateway: required rpm packages (...)</p> <p>RDMA Direct Memory Operations</p> <p>Ceph cluster capacity much small than it (...)</p> <p>Ceph - v0.94.11</p> <p>64pgs are in active+remapped state after (...)</p>	
------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--