

Ceph

	2024-1					2024-2					2024-3					2024-4					2024-5					2024-6				
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26				
Ceph	<div><div></div><div>Ceph</div></div>																													
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 perftools 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 (...)																														
deprecate rados_pool_create																														
What's the policy on URL stability for public (...)																														
erasure-code: last stripe is not truncated																														
Epic: Kerberos/LDAP Support																														
related to integrity of objects																														
erasure-code: use LRU to cache decoding matrix (...)																														
osd: ec plugin leak																														
msgr: encode osd epoch in nonce to avoid (...)																														
Add support for collecting usage information (...)																														

<p> erasure-code: better LRC strategy  ceph CLI: Help is missing -k keyring option  Cache warmup and ejection  annotate config options  ceph should fail when an argument is of the (...)  erasure-code: ErasureCode::encode overhead (...)  osd: hard cap on PGs per OSD  libgoogle-perftools4: tcmalloc performance (...)  osd: store scrub error state in kv store; (...)  osd: add scrub result query interface  librados: add ability to read a specific (...)  rados: add ability to read a specific replica/shard (...)  librados, osd: per-object scrub operation  librados: improve get_objects and get_position (...)  rados: new import from pipe/file  librados: set user_version operation  option to disable erasure code plugin version (...)  ceph: 'daemon list' command to enumerate (...)  add ceph-report-parsing  osd: an all/mostly delete pg log stalls during (...)  mon: finer-grain Paxos::is_readable() / PaxosService::is_readable()  mon: refactor OSDMonitor map methods to use (...)  mon: move OSDMonitor osd_epoch value into (...)  msgmgr: optionally encrypt messenger session  add stop_scrub command for ceph  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 (...) </p>	
--	--

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 doesn't 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, (...)  
 Please support windows Server 2016  
 ceph\_test\_rados dumped core, Error: finished (...)  
 strongly warn against custom cluster names  
 install-deps.sh fails in Debian/stretch  
 Teuthology tests using the thrasher are hitting (...)  
 udev is not declared as a dependency, fails (...)  
 RDMA has an error in handle\_connection func  
 Doc: There are no recommendations about OS (...)  
 backfill targets' stats.sum.num\_bytes get (...)  
 Debian Package Repositories - Missing packages (...)  
 Ubuntu Bionic package index missing most (...)  
 disable "pgs not deep-scrubbed in time" message  
 ceph-mon Illegal instruction (core dumped)

<p><b>Allow User Contributed Notes in the documentation</b></p> <p><b>ceph_argparse: binary flags cannot have detached (...)</b></p> <p><b>SELinux denials in ceph-daemon test</b></p> <p><b>Ceph code doesn't have namespace</b></p> <p><b>[clslua] Why clslua_map_get_keys returns (...)</b></p> <p><b>Monitor lookup in DNS needs to be case-insensitive</b></p> <p><b>unittest_confutil crashes on NON-ASCII chars</b></p> <p><b>[ceph-mon]ceph-mon crash when modify pool's (...)</b></p> <p><b>catch std::bad_alloc and die with helpful (...)</b></p> <p><b>qa/benchmark: Explore using Filebench for (...)</b></p> <p><b>add an option to force a down osd to be marked (...)</b></p> <p><b>Document old design</b></p> <p><b>implement new design</b></p> <p><b>rados: support omap (leveldb) and locator (...)</b></p> <p><b>pkgconfig files for libraries</b></p> <p><b>crush: create a visualizer for crush maps</b></p> <p><b>rados: support asynchronous replication of (...)</b></p> <p><b>Efficient access in distributed data centers</b></p> <p><b>make Ceph build faster, take less RAM to (...)</b></p> <p><b>Minimize recovery time of a degraded cluster</b></p> <p><b>figure out how to best set the heartbeat (...)</b></p> <p><b>log: detect dup messages</b></p> <p><b>Automatically determine routing data based (...)</b></p> <p><b>add ability to make ceph ignore io errors (...)</b></p> <p><b>qa: add test to compile a simple program (...)</b></p> <p><b>recovery from partially partitioned networks</b></p> <p><b>Decode exception descriptions don't show (...)</b></p> <p><b>Journal entries should record last known (...)</b></p> <p><b>OpTracker: support hierarchies of request (...)</b></p> <p><b>create and automate scenario based testing</b></p> <p><b>Deep scrub CPU limit behavior</b></p> <p><b>Backport option to disable deep scrub to (...)</b></p> <p><b>document src/objclass/objclass.h</b></p> <p><b>Some clone errors aren't repaired</b></p>	
--	--

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  
rework bufferlist::\*aligned\* functions  
ceph df reports % of global size used instead (...)  
Issues found by Clang  
do not leak ceph-disk activate lock to the (...)  
try to make address sanitizer work  
forbid erasure code profile modifications (...)  
make it harder to remove cluster data pools  
When replaced a disk, the CRUSH weight of (...)  
gtest: update, move to submodule  
create selinux policies for ceph-mon, ceph-osd, (...)  
osd: mark pg and use replica on EIO from (...)  
Cluster unreachable after restart  
ECUtil::HashInfoRef should have a NONE value  
OSD level metadata mismatch handling  
PG metadata corruption detection and handling  
Object metadata mismatch detection and handling  
Create new git repo for json\_spirit  
librados: mark old objects\_begin interface (...)  
double resource for setting up ceph-deploy  
mon,osd: long-term non-clean PGs prevent (...)  
Per-pool PG recovery prioritisation  
Permanently mark daemons and clusters which (...)

<p><b>OSD Keeps Going Down</b></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> <p>docs: federated rgw setup needs to be trimmed (...)</p> <p>docs: federated rgw setup keyring section (...)</p> <p>docs: federated rgw setup create-pools section (...)</p> <p>docs: federated rgw setup create pools section (...)</p> <p>docs: federated rgw setup create pools section (...)</p> <p>docs: federated rgw setup uses Debian-based (...)</p> <p>docs: federated rgw setup assumes that the (...)</p> <p>docs: federated rgw setup doesn't warn about (...)</p> <p>docs: federated rgw setup needs to normalize (...)</p> <p>docs: federated rgw setup assumes user is (...)</p> <p>docs: federated rgw setup is not explicit (...)</p> <p>docs: federated rgw setup should post example (...)</p> <p>Improve scrubbing throughput</p> <p>throttle objects and bytes / second for scrub/deepscrub</p> <p>ceph health warn if pgs are not being scrubbed (...)</p> <p>LMDB key/value backend for Ceph</p> <p>New feature : Ceph should warn if the kernel (...)</p> <p>move ceph-deploy(8) man page into ceph-deploy (...)</p> <p>ceph-disk breaks when using non-existing (...)</p> <p>mons: add switch to never delete pools, only (...)</p> <p>make sure notcmalloc builds continue working</p> <p>erasure-code: hitchhiker plugin</p>	
---	--

<p>test: verify that expected feature bits are (...)</p> <p>add a command to the ceph tool which returns (...)</p> <p>Elector: throttle election attempts from (...)</p> <p>Certain configuration options not safe for (...)</p> <p>improve OPTION macro to add a check function (...)</p> <p>add HELP command to ceph interactive shell</p> <p>pg a copy_from-&gt;pg b peered (blocked on backfill (...)</p> <p>OSD: ensure that filestore queue_transaction (...)</p> <p>msgmgr: throttle max open sockets</p> <p>mon: automatically blacklist clients after (...)</p> <p>change ceph-mon TCP port to an IANA-registered (...)</p> <p>Show osd as NONE in "ceph osd map &lt;pool&gt; (...)</p> <p>Workqueue timeouts and the filestore sync_entry (...)</p> <p>mon: add ability to gradually remove an osd</p> <p>OSD crashes under cached cluster benchmark</p> <p>ceph-cli: let "ceph report" support compressed (...)</p> <p>erasure code plugin for tornado code</p> <p>data is no well-distributed among osds of (...)</p> <p>Does the ceph debian package really need (...)</p> <p>removing cached objects doesn't quickly remove (...)</p> <p>init scripts: try to detect when a daemon (...)</p> <p>Messenger: use setsockopt to set sk priority (...)</p> <p>Implement a minimal access keyring for read-only (...)</p> <p>Support for building on Alpine Linux</p> <p>Please publish rados-java artifacts to central (...)</p> <p>Why copy a 3T image without any data takes (...)</p> <p>PGMonitor::get_health get ops' blocked the (...)</p> <p>FileSystem become read-only after over 6 (...)</p> <p>create the upgrade test suite for gmt and (...)</p> <p>Make "rados bench" support json output</p> <p>Translate fadvise flag NOCACHE into DONTNEED (...)</p> <p>OSDMonitor refactor</p> <p>compile warning on gcc-4.8.3</p> <p>automatically set journal file name as osd.{osd-num} (...)</p>	
--	--



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

<p>EC Overwrites: Missing Test Case: aligned (...)</p> <p>EC Overwrites: Missing Test Case: multiple (...)</p> <p>EC Overwrites: Missing Test Case: zero</p> <p>degrade a log message to a default level</p> <p>scripts/ceph-release-notes does not handle (...)</p> <p>add backfill targets to osd stats and a way (...)</p> <p>Show the total pg on each osd.</p> <p>the method 'Message::decode_message' is too (...)</p> <p>Show the pool quota info on ceph df detail (...)</p> <p>Crush when get immediate parent quickly reply.</p> <p>Command to override Incomplete PGs</p> <p>scrub/repair: repair corrupted/missing objects</p> <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>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>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><a href="http://www.irchelp.org/irchelp/new2irc.html">http://www.irchelp.org/irchelp/new2irc.html</a> (...)</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>	
--	--

Creating type alias for complex type definitions  
 ceph 12.2.0 health check failed osd down  
 discover ip node in ceph cluster  
 Ceph ISCSI Gateway: required rpm packages (...)  
 RDMA Direct Memory Operations  
 Ceph cluster capacity much small than it (...)  
 How to recover the EC-encoded object/file (...)  
 OSDs consumes around 50% CPU on idle cluster  
 docs: Cache Tier page has not been updated (...)  
 Key exported from auth keyring is not imported (...)  
 Infiniband send\_msg send returned error 32: (...)  
 rados object command get/put run wrong without (...)  
 make vstart fails - building ceps cluster  
 --localize-reads doc and help  
 ceph-detect-init test case failed in docker (...)  
 Add more cmake options to support building (...)  
 Inconsistent graylog option naming  
 s390x support in ceph  
 lz4 fails to compile with clang  
 fix upgrade doc section to point to release (...)  
 Address min\_size inconsistency in docs  
 --osd-objectstore is undocumented  
 Can not install LTS Ceph (Luminous) on current (...)  
 Monitor election continues for 10 minutes (...)  
 fix a race between Thread::create and Thread::set\_ioprio  
 Fix outorder between Thread::create and Thread::set\_ioprio  
 Depending on the distro a package change (...)  
 Possible rbd snapshot race with 13.2.1 causes (...)  
 debian packages seem to overdo logrotate: (...)  
 Ceph assert when enable SPDK with 64Kb kernel (...)  
 Interactive mode CLI with Python 3: Traceback (...)  
 install-deps.sh fails on centos 7.6  
 Failed to execute command: systemctl enable (...)  
 Warn the user when different version of boost (...)

<p>Missing libboost package for arm64 on ubuntu (...)</p> <p>cannot build rpm, ceph-volume fails the build</p> <p>Confusing response when changing live config (...)</p> <p>cli incorrectly reports about missing keyring (...)</p> <p>Ceph Blog - new-luminous-erasure-coding-rbd-cephfs</p> <p>Add dependencies for missing Python packages</p> <p>add mgr plugin packages to qa/packages/packages.yaml (...)</p> <p>/usr/bin/ld: cmdparse.cc.o: bad reloc symbol (...)</p> <p>vstart.sh: fail to pass params monmaptool (...)</p> <p>doc: Documentation about daemon restart with (...)</p> <p>"failed to create directory via template (...)</p> <p>Update documentation for installing monitors</p> <p>RDMA support of non-SQL qp is broken</p> <p>mimic: In function 'rte_table_hash_cuckoo_create': (...)</p> <p>Wrong order of Environment evaluation in (...)</p> <p>"MaxWhileTries: reached maximum tries (500) (...)</p> <p>Debian/Ubuntu librados2 package does not (...)</p> <p>make jemalloc/tcmalloc work.</p> <p>nautilus: Number of objects in a pool doubled (...)</p> <p>High CPU load using Ceph Nautilus in Rook</p> <p>Segmentation fault (core dumped) in upgrade:client-upgrade-kraken-luminous</p> <p>reformat release schedule page</p> <p>In the environment with EC rule and repliaction (...)</p> <p>POOL 100% USED with 17TB available raw space</p> <p>[bug] KRBD fail when ceph set wrong compression (...)</p> <p>ceph memory leak?</p> <p>Trying to upgrade from Ceph Mimic to Nautilus (...)</p> <p>ceph_test_msgr: file INSTALL cannot copy (...)</p> <p>ceph::buffer::end_of_buffer isn't caught (...)</p> <p>ceph-osd does not start correctly OSDs if (...)</p> <p>need to compile DPDK with NUMA support in (...)</p> <p>gorados - pure go client to connect RADOS (...)</p> <p>Cluster usage stats overview is out of date</p> <p>Introduction of tier local mode</p>	
--	--

<p>Need to stop ceph daemons from picking some (...) systemd restarts OSD too fast after failure</p> <p>Ceph - v0.94.11</p> <p>64pgs are in active+remapped state after (...)</p> <p>Ceph - v10.2.10</p> <p>Perfcounters: wrong latency information (...)</p> <p>Ceph - v10.2.11</p> <p>Les OSD ne redémarrent pas tous après (...)</p> <p>Ceph - v12.2.0</p> <p>failed in cmake</p> <p>Manual deployment does not work in luminous</p> <p>Negative Runway in BlueFS?</p> <p>target_max_bytes doesn't limit a tiering (...)</p> <p>Ceph - v12.2.11</p> <p>_scan_snaps no head for &lt;object&gt;</p> <p>Ceph - v12.2.12</p> <p>Rados import is potentially dangerous (...)</p> <p>Ceph - v12.2.4</p> <p>Building Ceph on armhf fails due to (...)</p> <p>Ceph - v12.2.5</p> <p>bucket policy ipdeny not in effect</p> <p>Ceph - v12.2.6</p> <p>async messenger thread cpu high, osd (...)</p> <p>Ceph - v13.2.4</p> <p>Update OS Recommendations page for Mimic</p> <p>Ceph - v14.0.0</p> <p>Command failed on smithi075 with status (...)</p> <p>nautilus: 14.0.1 build fails in fedora (...)</p> <p>upgrading ceph packages always triggers (...)</p> <p>Ceph - v14.2.2</p>	<p>New 0%</p>
---	---------------