Cheatography

Ceph Cheat Sheet

by Eagle1992 (Eagle1992) via cheatography.com/195419/cs/40944/

ceph config	
ceph config dump	Dumps entire monitor configuration database
ceph config get {who} {option}	Dumps the configuration option stored in monitor configuration database for specific daemon or client(for example osd.0). Options können angegeben werden.
ceph config set {who} {option} {value}	specifies a configuration option in the monitor configuration database (for example ceph config set osd.0 debug_ms 20)
ceph config show {who}	shows runtime settings for a running daemon (to see all settings use: ceph config show-with-defaults)
ceph config assimilate-conf -i {input_file} -o {output_file}	ingests a configuration file from inpute file and moves any valid options into the monitor configuration database
ceph config help {option} (-f json-p-retty)	to get help for particular option (not optional).
ceph tell {who} config set {option} {value}	set temporarily other settings (for example ceph tell osd.123 config set debug_osd 20) You can also specify wildcards: osd.* (to change settings for all OSDs.)

ceph is part of Ceph, a massively scalable, open-source, distributed storage system. Please refer to the Ceph documentation at https://docs.ce-ph.com for more information.

Options	
-i infile	will specify an input file to be passed along as a payload with the command to the monitor cluster. This is only used for specific monitor commands.
-o outfile	will write any payload returned by the monitor cluster with its reply to outfile. Only specific monitor commands (e.g. osd getmap) return a payload.
--setuser user	will apply the appropriate user ownership to the file specified by the option '-o'.
--setgroup group	will apply the appropriate group ownership to the file specified by the option '-o'.

Options (cont)

Options (cont)

-c ceph.conf, --conf=ceph.conf	Use ceph.conf configuration file instead of the default /etc/ceph/ceph.conf to determine monitor addresses during startup.
--id CLIENT_ID, --user CLIENT_ID	Client id for authentication.
--name CLIENT_NAME, -n CLIENT_NAME	Client name for authentication.
--cluster CLUSTER	Name of the Ceph cluster.
--admin-daemon ADMIN_SOCKET, daemon DAEMON_NAME	Submit admin-socket commands via admin sockets in /var/run/ceph.
-s, --status	Show cluster status.
-w, --watch	Watch live cluster changes on the default 'cluster' channel
-W, --watch-channel	Watch live cluster changes on any channel (cluster, audit, cephadm, or * for all)
--watch-debug	Watch debug events.
--watch-info	Watch info events.
--watch-sec	Watch security events.
--watch-warn	Watch warning events.
--watch-error	Watch error events.
--version, -v	Display version.
--verbose	Make verbose.
--concise	Make less verbose.
-f {json,json-pretty,xml,xml-pretty,plain,yaml}, --format	Format of output. Note: yaml is only valid for orch commands.
--connect-timeout CLUSTER_TIMEOUT	Set a timeout for connecting to the cluster.
--no-increasing	--no-increasing is off by default. So increasing the osd weight is allowed using the reweight-by-utilization or test-reweight-by-utilization commands. If this option is used with these commands, it will help not to increase osd weight even the osd is under utilized.
--block	block until completion (scrub and deep-scrub only)

Ceph Cheat Sheet

by Eagle1992 (Eagle1992) via cheatography.com/195419/cs/40944/

ceph mon

ceph mon dump {int0-}	dumps formatted monmap - if integer is given, then you'll get mon-map from epoch {integer]
ceph mon add {name} {IPaddr[:-port]}	adds new monitor named {name} at {addr}
ceph mon getmap {int0-}	gets monmap (from specified epoch)
ceph mon remove {name}	removes monitor named {name}
ceph mon stat	summarizes monitor status

ceph mgr

ceph mgr dump	dumps latest MgrMap, which describes the active & standby manager daemons
ceph mgr fail {name}	will mark a manager daemon as failed, removing it from managerr map
ceph mgr module ls	will list currently enableld manager modules (plugins)
ceph mgr module enable {module}	will enable a manager modules. available modules are included in MgrMap and visible via mgr dump
ceph mgr module disable {module}	will disable an active manager module

ceph mgr (cont)

ceph mgr metadata {name}	will report metadata about all manager daemons or if the name is specified a single manager daemon
ceph mgr versions	will report a count of running daemon versions
ceph mgr count--metadata {field}	will report a count of any daemon metadata field

Miscellaneous

ceph tell mon.<id> quorum enter|exit	Cause a specific MON to enter or exit quorum.
ceph quorum-_status	Reports status of monitor quorum.
ceph report {<tags> [<tags>...]}	Reports full status of cluster, optional title tag strings.
ceph status	Shows cluster status.
ceph tell <name (type.id)> <comma-nd> [options...]	Sends a command to a specific daemon.
ceph tell <name (type.id)> help	List all available commands.
ceph version	Show mon daemon version
ceph fs dump	get MDS Map
ceph balancer status	get status of ceph balancer

Miscellaneous (cont)

ceph balancer off / on	enable / disable ceph balancer
ceph balancer mode crush--compat or upmap	set balancer mode to crush-compat or upmap (default)
ceph -s / --status	see actual ceph status
ceph df detail	shows data usage in raw storage and pools
ceph-volume lvm list or ceph device ls	shows all disks
ceph crash ls / ls-new	shows all mgr module crash dumps (or only list new crash dumps with ls-new)
ceph crash info {crashid}	shows exact information for crashdump with specific crashid
ceph crash archive-all	archive all crash dumps
ceph crash rm {crashid}	removes crash dump with specific id
ceph crash stat	List the timestamp/uuid crashids for all newcrash info.
ceph crash prune {keep}	Remove saved crashes older than ‘keep’ days. {keep} must be an integer.

By **Eagle1992** (Eagle1992)
cheatography.com/eagle1992/

Published 23rd October, 2023.
Last updated 23rd October, 2023.

Miscellaneous (cont)

ceph crash archive {crashid}	Archive a crash report so that it is no longer considered for the RECENT_CRASH health check and does not appear in the crash ls-new output (it will still appear in the crash ls output).
crushtool -d {compiled-crusmap-file} -o {output_decomp-crushmap-file}	decompile crushmap (ceph osd getcrushmap -o {file}) to readable format. Now you can open it with every common texteditor (vim, nano, vi) or read with cat / less
crushtool -c {modified-crush-map-filename} -o {modified-compiled-crushmap-file}	recompile crushmap after modifying to output file (-o)
ceph osd setcrushmap -i {modified-compiled-crushmap-file}	set new crushmap from file

ceph pg

ceph pg debug unfound_objects_exist|degraded_pgs_exist	shows debug info about pgs.
ceph pg deep-scrub <pgid>	starts deep-scrub on <pgid>.

ceph pg (cont)

ceph pg dump {all|summary|sum|delta|pools|osds|pgs|pgs_brief} [{all|summary|sum|delta|pools|osds|pgs|pgs_brief...]}	shows human-readable versions of pg map (only 'all' valid with plain).
ceph pg dump_json {all|summary|sum|delta|pools|osds|pgs|pgs_brief} [{all|summary|sum|delta|pools|osds|pgs|pgs_brief...]}	shows human-readable version of pg map in json only.
ceph pg dump_pools_json	shows pg pools info in json only.
ceph pg dump_stuck {inactive|unclean|stale|undersized|degraded [inactive|unclean|stale|undersized|degraded...]} {<int>}	shows information about stuck pgs.
ceph pg getmap	gets binary pg map to -o/stdout.
ceph pg ls {<int>} {<pg-state> [<pg-state>...]}	lists pg with specific pool, osd, state
ceph pg ls-by-osd <osdname (id|osd.id)> {<int>} {<pg-state> [<pg-state>...]}	lists pg on osd [osd]
ceph pg ls-by-pool <poolstr> {<int>} {<pg-state> [<pg-state>...]}	lists pg with pool = [poolname]
ceph pg ls-by-primary <osdname (id|osd.id)> {<int>} {<pg-state> [<pg-state>...]}	lists pg with primary = [osd]
ceph pg map <pgid>	shows mapping of pg to osds.

ceph pg (cont)

ceph pg repair <pgid>	starts repair on <pgid>.
ceph pg scrub <pgid>	starts scrub on <pgid>.
ceph pg stat	shows placement group status.

https://docs.ceph.com/en/quincy/rados/operations/pg-states/
https://docs.ceph.com/en/latest/rados/troubleshooting/troubleshooting-pg/

ceph osd

ceph osd blocklist add {EntityAddr} {<float[0.0-]>}	add {addr] to blocklist
ceph osd blocklist ls	show blocklisted clients
ceph osd blocklist rm {EntityAddr}	remove {addr] from blocklist
ceph osd blocked-by	prints a histogram of which OSDs are blocking their peers
ceph osd new {<uuid>} {<id>} -i {<params.json>}	To create a new OSD or recreate a previously destroyed OSD with specific id. Please look up Documentation if you're planning to use this command.
ceph osd crush add <osdname (id|osd.id)> <float[0.0-]> <args> [<args>...]	adds or updates crushmap position and weight for <name> with <weight> and location <args>.
ceph osd crush add-bucket <name> <type>	dds no-parent (probably root) crush bucket <name> of type <type>.

ceph osd (cont)

ceph osd crush create-or-move <osdname (id|osd.id)> <float[0.0-]> <args> [<args>...]	creates entry or moves existing entry for <name> <weight> at/to location <args>.
ceph osd crush dump	dumps crush map.
ceph osd crush link <name> <args> [<args>...]	links existing entry for <name> under location <args>.
ceph osd crush move <name> <args> [<args>...]	moves existing entry for <name> to location <args>.
ceph osd crush remove <name> {<ancestor>}	removes <name> from crush map (everywhere, or just at <ancestor>).
ceph osd crush rename-bucket <srcname> <dstname>	renames bucket <srcname> to <dstname>
ceph osd crush reweight <name> <float[0.0-]>	change <name>'s weight to <weight> in crush map.
ceph osd crush reweight-all	recalculate the weights for the tree to ensure they sum correctly
eph osd crush reweight-subtree <name> <weight>	changes all leaf items beneath <name> to <weight> in crush map
ceph osd crush rm <name> {<ancestor>}	removes <name> from crush map (everywhere, or just at <ancestor>).
ceph osd crush rule create-erasure <name> {<profile>}	creates crush rule <name> for erasure coded pool created with <profile> (default default).

ceph osd (cont)

ceph osd crush rule create-simple <name> <root> <type> {firstn|indep}	creates crush rule <name> to start from <root>, replicate across buckets of type <type>, using a choose mode of <firstn|indep> (default firstn; indep best for erasure pools).
ceph osd crush rule dump {<name>}	dumps crush rule <name> (default all).
ceph osd crush rule ls	lists crush rules.
ceph osd crush rule rm <name>	removes crush rule <name>
ceph osd crush set <osdname (id|osd.id)> <float[0.0-]> <args> [<args>...]	set with osdname/osd.id update crushmap position and weight for <name> to <weight> with location <args>.
ceph osd crush show-tunables	shows current crush tunables.
ceph osd crush tree	ows the crush buckets and items in a tree view.
ceph osd crush unlink <name> {<ancestor>}	unlinks <name> from crush map (everywhere, or just at <ancestor>).
ceph osd df {plain|tree}	shows OSD utilization
ceph osd deep-scrub <who>	initiates deep scrub on specified osd.

ceph osd (cont)

ceph osd down <ids> [<ids>...]	sets osd(s) <id> [<id>…] down.
ceph osd dump	prints summary of OSD map.
ceph osd find <int[0-]>	find osd <id> in the CRUSH map and shows its location.
ceph osd getcrushmap	gets CRUSH map.
ceph osd getmap	gets OSD map.
ceph osd getmaxosd	shows largest OSD id
ceph osd in <ids> [<ids>...]	sets osd(s) <id> [<id>…] in.
ceph osd lost <int[0-]> {--yes-i-really-mean-it}	marks osd as permanently lost. THIS DESTROYS DATA IF NO MORE REPLICAS EXIST, BE CAREFUL.
ceph osd ls	shows all OSD ids.
ceph osd lspools	lists pools
ceph osd map <poolname> <objectname>	finds pg for <object> in <pool>.
ceph osd metadata {int[0-]} (default all)	fetches metadata for osd <id>.
ceph osd out <ids> [<ids>...]	sets osd(s) <id> [<id>…] out.

ceph osd (cont)

ceph osd ok-to-stop <id> [<ids>...] [--max <num>]	checks whether the list of OSD(s) can be stopped without immediately making data unavailable. That is, all data should remain readable and writeable, although data redundancy may be reduced as some PGs may end up in a degraded (but active) state. It will return a success code if it is okay to stop the OSD(s), or an error code and informative message if it is not or if no conclusion can be drawn at the current time.
ceph osd pause	pauses osd.
ceph osd perf	rints dump of OSD perf summary stats.
ceph osd force-create-pg <pgid>	forces creation of pg <pgid>.
ceph osd pool create <poolname> {<int[0-]>} {<int[0-]>} {replicated|erasure} {<erasure_code_profile>} {<rule>} {<int>} {--autoscale-mode=-<on,off,warn>}	creates pool.

ceph osd (cont)

ceph osd pool delete <poolname> {<poolname>} {--yes-i-really-really-mean-it}	deletes pool. (DATA LOSS BE CAREFUL!)
ceph osd pool get <poolname> size|min_size|pg_num|pgp_num|crush_rule|write_fadvise_dontneed	gets pool parameter <var>
eph osd pool get <poolname> all	to get all pool parameters that apply to the pool's type:
ceph osd pool get-quota <poolname>	obtains object or byte limits for pool.
ceph osd pool ls {detail}	list pools
ceph osd pool mksnap <poolname> <snap>	makes snapshot <snap> in <pool>.
ceph osd pool rename <poolname> <poolname>	renames <srcpool> to <destpool>.
ceph osd pool rmsnap <poolname> <snap>	removes snapshot <snap> from <pool>.

ceph osd (cont)

ceph osd pool set <poolname> size|min_size|pg_num|pgp_num|crush_rule|hashpspool|nodelete|nopgchange|nosizechange| hit_set_type|hit_set_period|hit_set_count|hit_set_fpp|debug_fake_ec_pool|target_max_bytes|target_max_objects|cache_target_dirty_ratio| cache_target_dirty_high_ratio| cache_target_full_ratio|cache_min_flush_age|cache_min_evict_age| min_read_recency_for_promote|write_fadvise_dontneed|hit_set_grade_decay_rate| hit_set_search_last_n <val> {--yes-i-really-mean-it}	sets pool parameter <var> to <val>.
ceph osd pool set-quota <poolname> max_objects|max_bytes <val>	sets object or byte limit on pool.
ceph osd pool stats {<name>}	obtain stats from all pools, or from specified pool.
ceph osd repair <who>	initiates repair on a specified osd.

ceph osd (cont)	
ceph osd reweight-by-pg {<int[100-]>} {<pool-name> [<pool-name...]} {--no-increasing}	reweight OSDs by PG distribution [overload-percentage-for-consideration, default 120].
ceph osd reweight-by--utilization {<int[-100-]> {<float[0.0-]> {<int[0-]>}}} {--no-increasing}	reweights OSDs by utilization. It only reweights outlier OSDs whose utilization exceeds the average, eg. the default 120% limits reweight to those OSDs that are more than 20% over the average. [overload-threshold, default 120 [max_weight_change, default 0.05 [max_osds_to_adjust, default 4]]]
ceph osd rm <ids> [<ids>...]	removes osd(s) <id> [<id>…] from the OSD map.

ceph osd (cont)	
ceph osd destroy <id> {--yes-i--really-mean-it}	marks OSD id as destroyed, removing its cephx entity's keys and all of its dm-crypt and daemon-private config key entries. This command will not remove the OSD from crush, nor will it remove the OSD from the OSD map. Instead, once the command successfully completes, the OSD will show marked as destroyed. In order to mark an OSD as destroyed, the OSD must first be marked as lost.
ceph osd purge <id> {--yes-i--really-mean-it}	performs a combination of osd destroy, osd rm and osd crush remove.

ceph osd (cont)	
ceph osd safe-to-destroy <id> [<ids>...]	checks whether it is safe to remove or destroy an OSD without reducing overall data redundancy or durability. It will return a success code if it is definitely safe, or an error code and informative message if it is not or if no conclusion can be drawn at the current time.
ceph osd scrub <who>	initiates scrub on specified osd.
ceph osd set pause|-noup|nodown|noout|-noin|nobackfill| norebalance|norecover|noscrub|nodeep-scrub|notieragent	sets cluster-wide <flag> by updating OSD map. The full flag is not honored anymore since the Mimic release, and ceph osd set full is not supported in the Octopus release.
ceph osd setcrushmap	sets crush map from input file.
ceph osd setmaxosd <int[0-]>	sets new maximum osd value.

ceph osd (cont)	
ceph osd set-require--min-compat-client <version>	enforces the cluster to be backward compatible with the specified client version. This subcommand prevents you from making any changes (e.g., crush tunables, or using new features) that would violate the current setting. Please note, This subcommand will fail if any connected daemon or client is not compatible with the features offered by the given <version>.
ceph osd stat	prints summary of OSD map.
ceph osd tree {<int[0-]->}	prints OSD tree.
ceph osd unpause	unpauses osd.
ceph osd unset pause|-noup|-nodown|-noout|-noin|-nobackfill|norebalance|-norecover|noscrub|-nodeep-scrub|notieragent	unsets cluster-wide <flag> by updating OSD map.