

Ceph – Slow recovery times

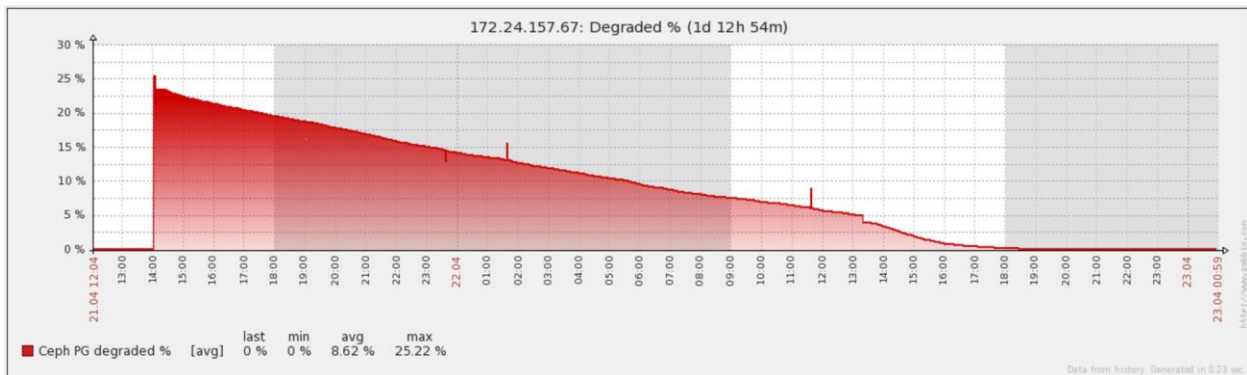
Reported by sachinkagarwal@gmail.com

Test-run

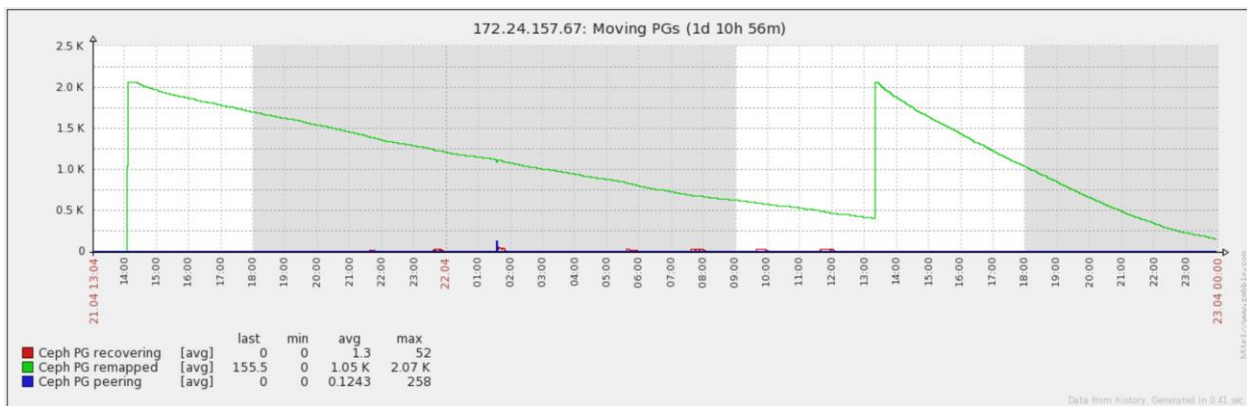
At Time = 21.04.2016, 14:00, one of 4 OSD servers was “failed” (OSD processes shutdown).

At Time = 22.04.2016, 13:20 the server was “restored”

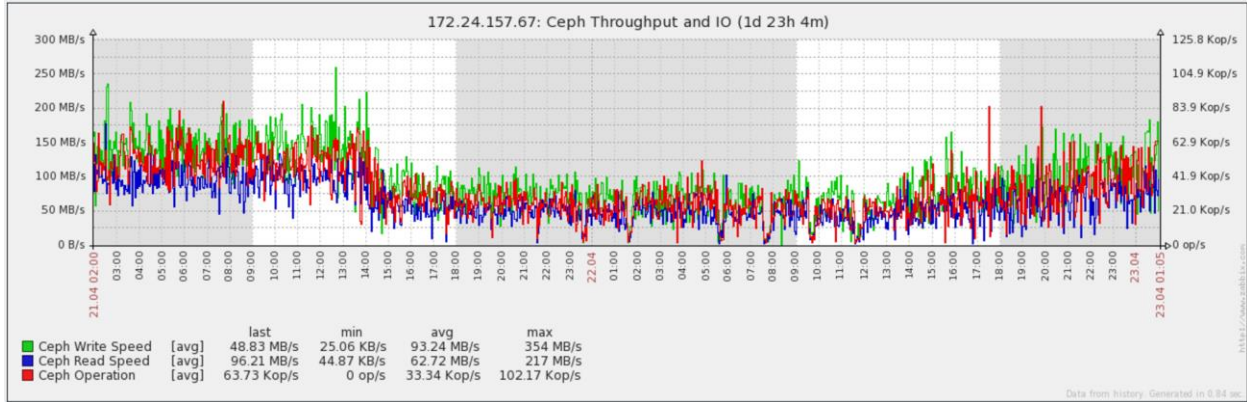
Graph 1: Number of PGs degraded (single copy)



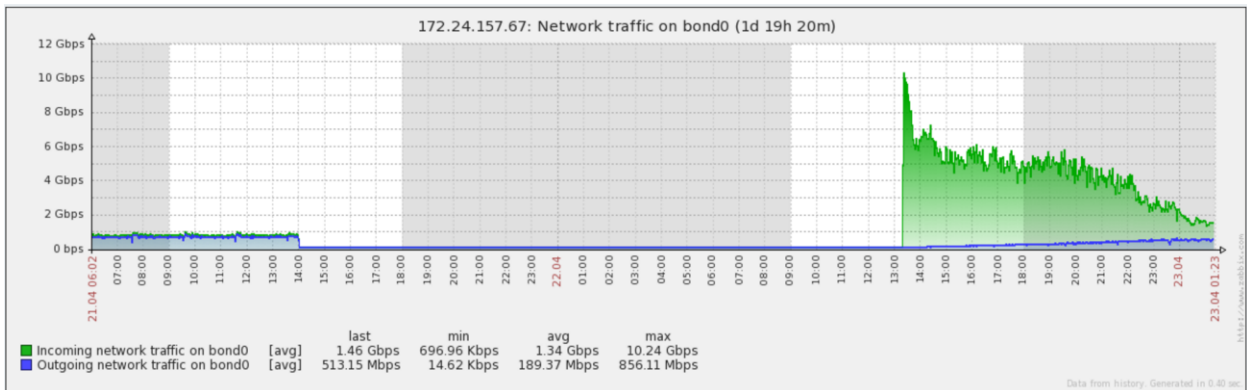
Graph 2: The corresponding “moving PGs” – note the tapering fall-offs that suggest slowdowns in the recovery as it progresses.



GRAPH 3: Cumulative Ceph client IO during the test (note the dips and recoveries)



GRAPH 4: Network traffic on the “failed” OSD server – tapers off as recovery progresses.



GRAPH 5: CPU on the “failed” OSD server. Note that the IOWait time is less than 33% at all times - so the SAS bus on the disks is not pegged.

