



Ceph Performance: Recent Testing and Findings LIGHTNING VERSION

Mark Nelson

Ceph Community Performance Lead

mnelson@redhat.com

11/16/2015


Ceph Memory Allocator Testing

We sat down at the 2015 Ceph Hackathon and tested a CBT configuration to replicate memory allocator results on SSD based clusters pioneered by Sandisk and Intel.

Memory Allocator	Version	Notes
TCMalloc	2.1 (default)	Thread Cache can not be changed due to bug.
TCMalloc	2.4	Default 32MB Thread Cache
TCMalloc	2.4	64MB Thread Cache
TCMalloc	2.4	128MB Thread cache
Jemalloc	3.6.0	Default jemalloc configuration

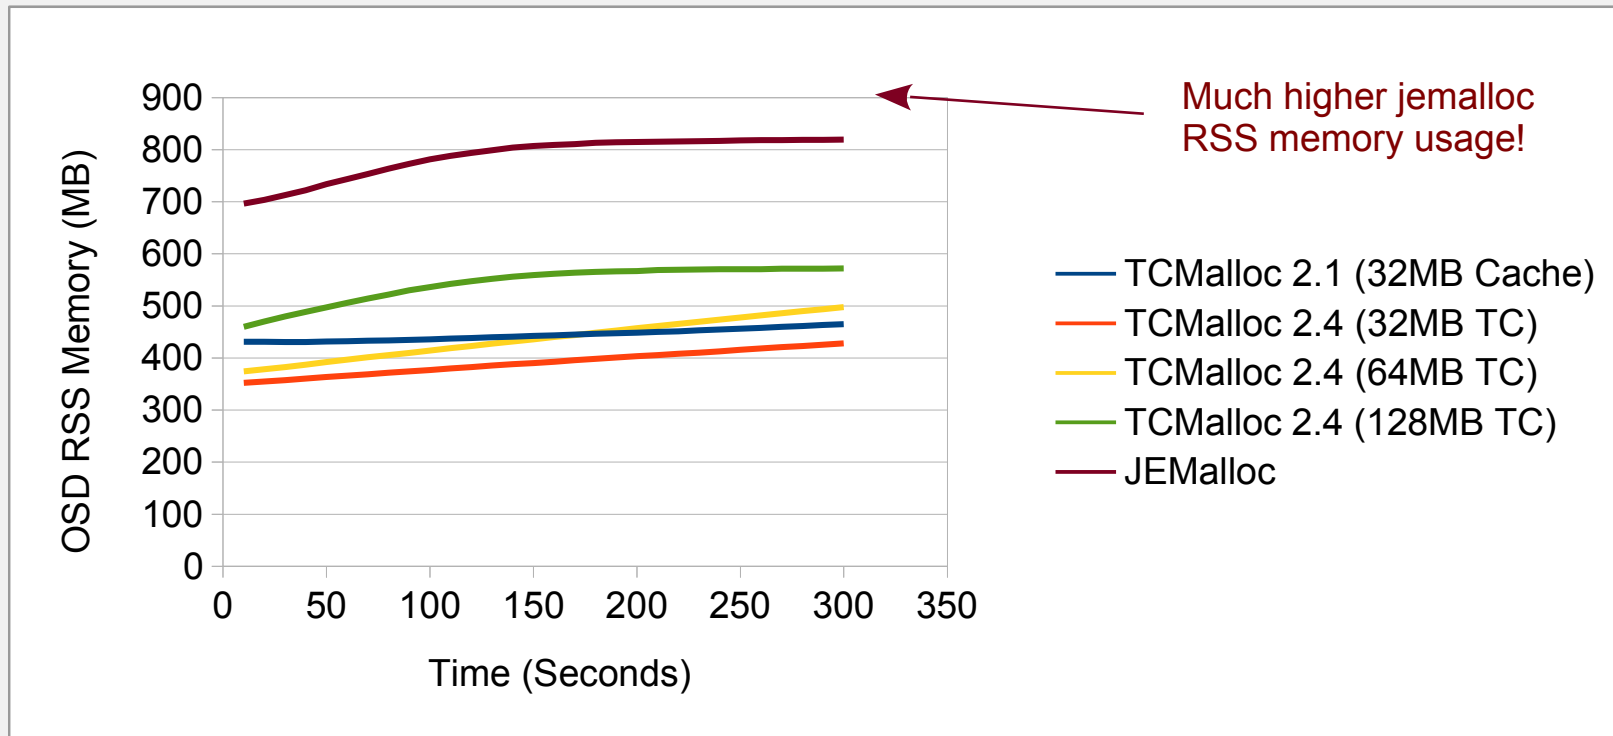
4KB Random Write IOPS

Ceph Memory Allocator Tests Using CBT



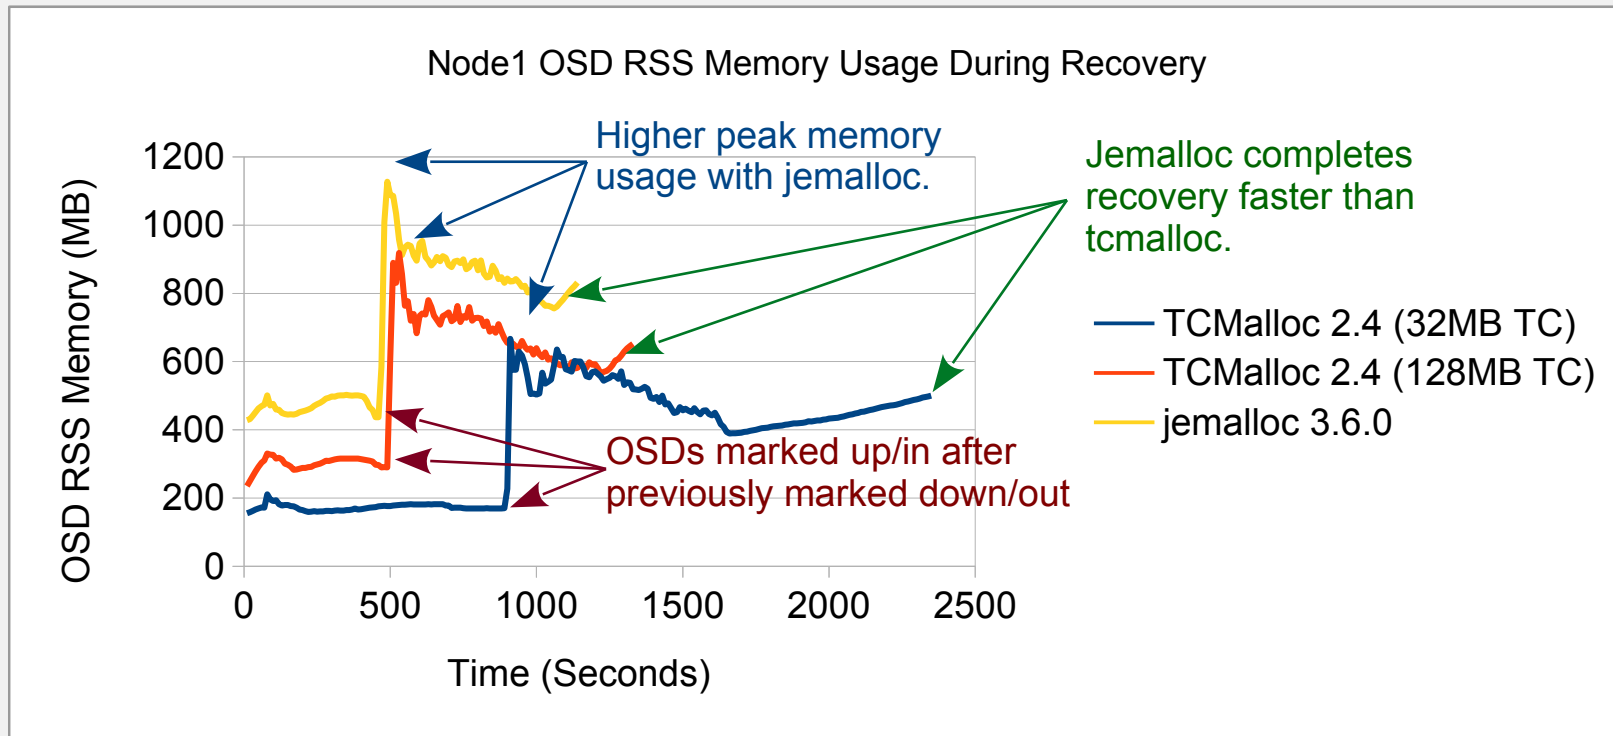
4KB Random Write OSD RSS Memory

Ceph Memory Allocator Tests Using CBT



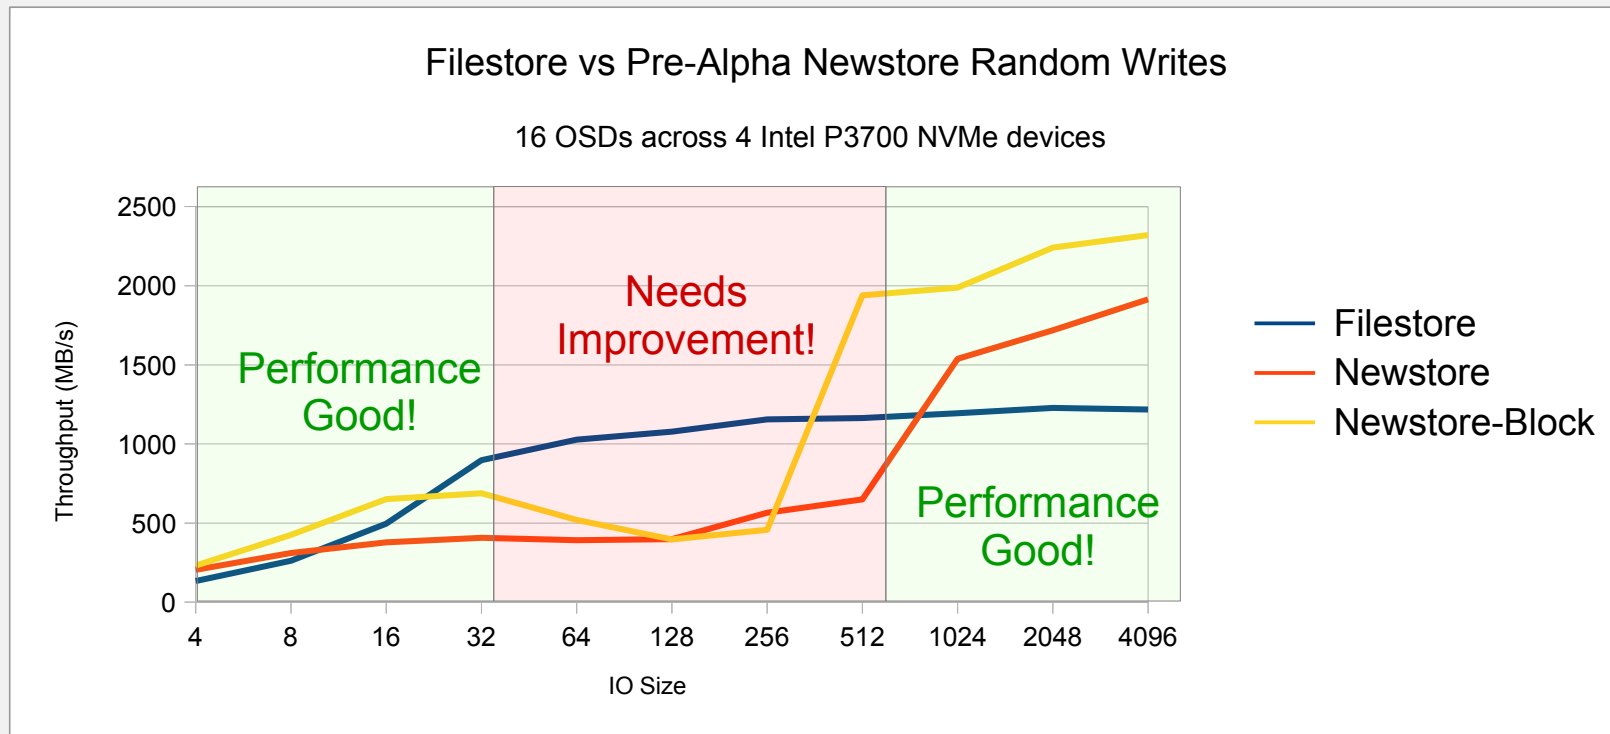
4KB Random Write OSD RSS Memory

Ceph Memory Allocator Recovery Tests Using CBT



Pre-Alpha Newstore Teaser:

Let's make Ceph even faster!





redhat.

THANK YOU



plus.google.com/+RedHat



facebook.com/redhatinc



linkedin.com/company/red-hat



twitter.com/RedHatNews



youtube.com/user/RedHatVideos