WOW – That's Freaky Quick!

Early Adopter Observations of Flash in the Datacenter

Scott Read

Montana Dept. of Transportation

About the Montana Dept. of Transportation

- Helena-Primary DC
- Miles City- 2nd DC
- 2x10Gb MPLS fiber
- Latency 8ms 342 miles
- 250TB Capacity per DC
- 40TB of flash capacity
- 3 Storage admins
- 18 Hypervisors

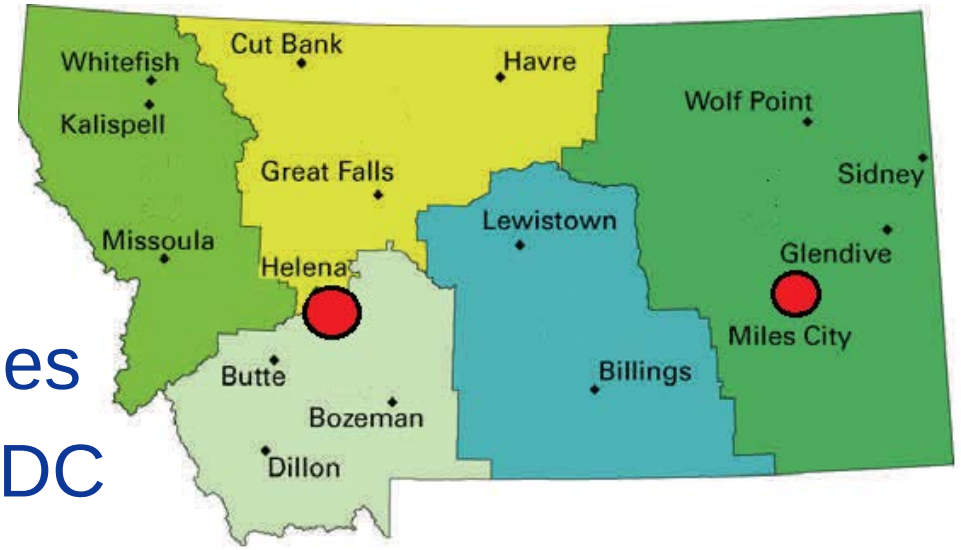

Adoption of Flash at MDT

- Adopting flash allows:
  - Replacement of heritage large spinning-disks
  - Future-proofing storage environment
  - Increased DB and Hypervisor performance
  - Resolves storage as a bottleneck for file services
  - Ties in our endpoint flash initiatives

Spring 2013 - Onboarding SLC

- SLC tier of storage activated
  - Added SLC flash to an existing array
  - Synthetic testing showed excellent performance potential.
  - Initial test DB workloads perform very well
  - Bandwidth capacity exceeds expectations
  - Limited amount of SLC prevents use with workloads outside of databases

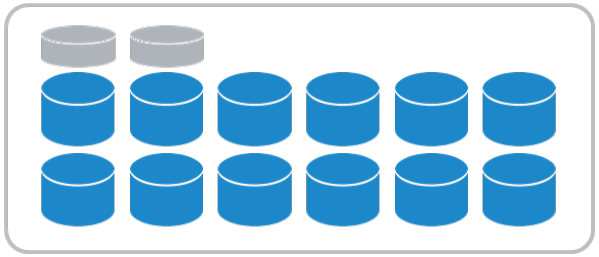

SLC Performance

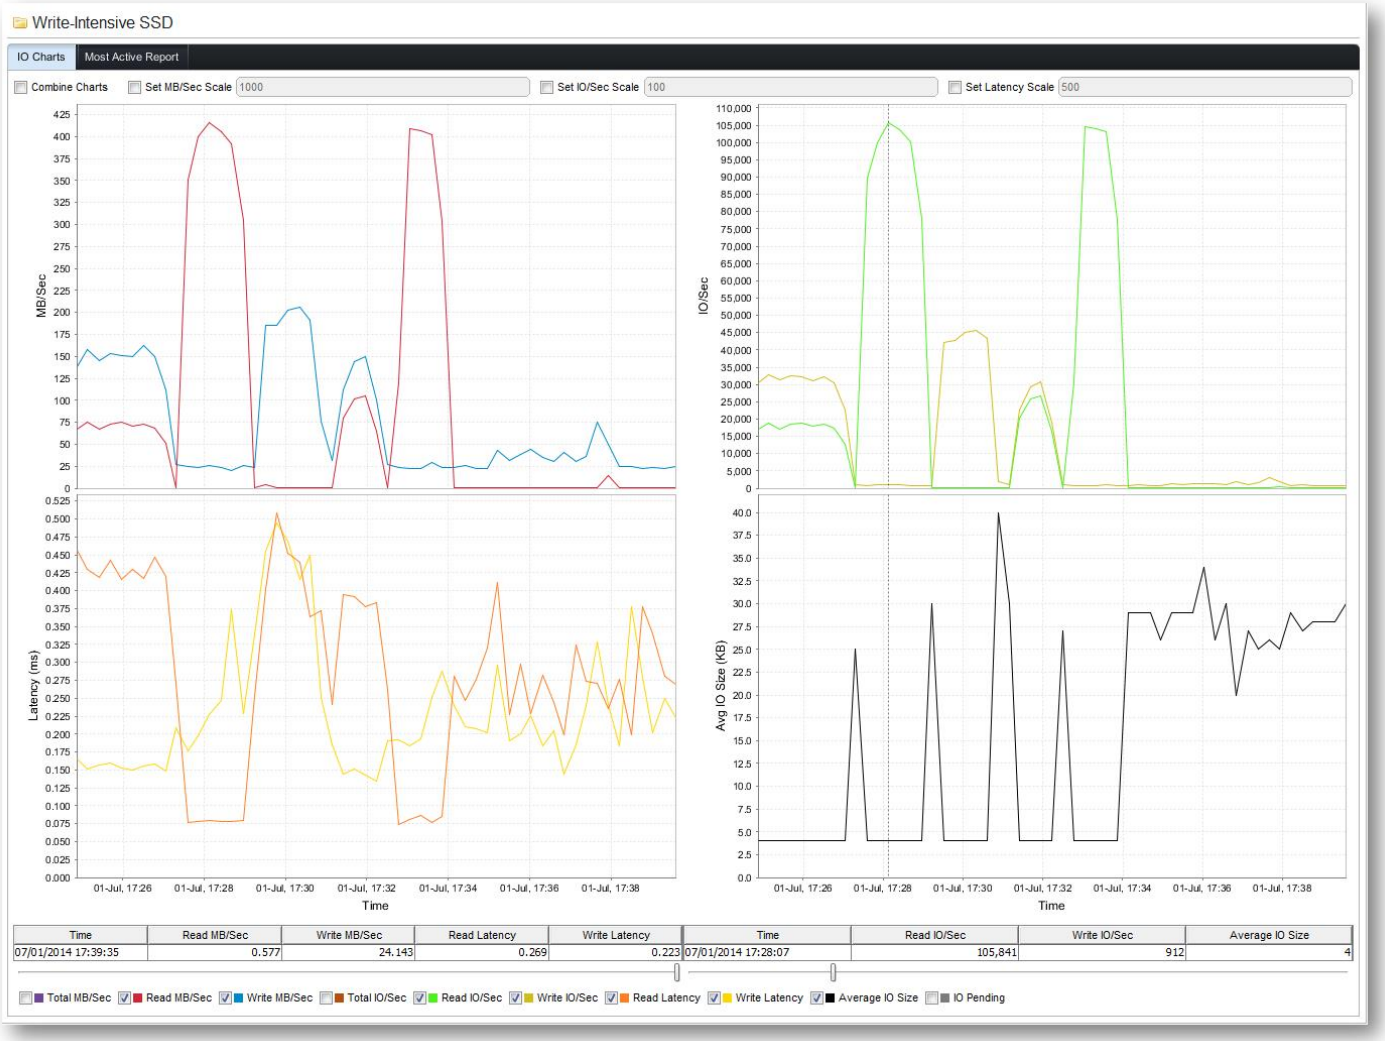

Spring 2014 - Onboarding MLC

- Conversion of primary array to flash optimized
  - Added MLC drives to existing array
  - Commissioned new page pool to support the SLC and new MLC storage
  - Added HDDs to hold the cold data
  - Process was completely non-service disrupting

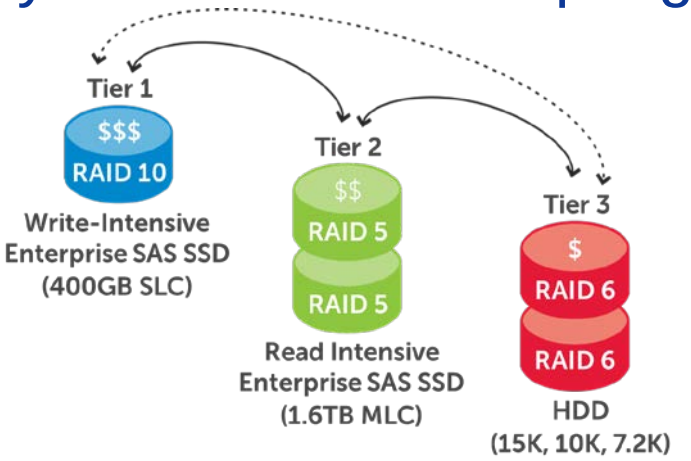

Observations

Loading data from 15K to Flash optimized

- Volumes transparently migrated
- SAS loops were able to sustain 22,000 Gb/sec (2.75GB/sec)
- 45,000 sustained write IOPs with room to spare
- Observed backend latency average was 2ms.
- Re-located .5TB in ~15 minutes
- Typical result when migrating our 15k volumes

Migration to Flash from 15k disks

Migration to Flash from 7k disks

Observations with VMWare

- VMWare activities
  - In testing, removing a large snapshot from a VM is no longer a time consuming event.
  - Creating volumes inside the workloads are quick (seconds instead of minutes)
  - VMWare Data Recovery performance is very stable
  - Storage vMotion is painless.
    - Movement of a 200GB VMDK between DataStores in less than 4 minutes

Storage vMotion Performance - RE-Locate 200GB in ~4 minutes

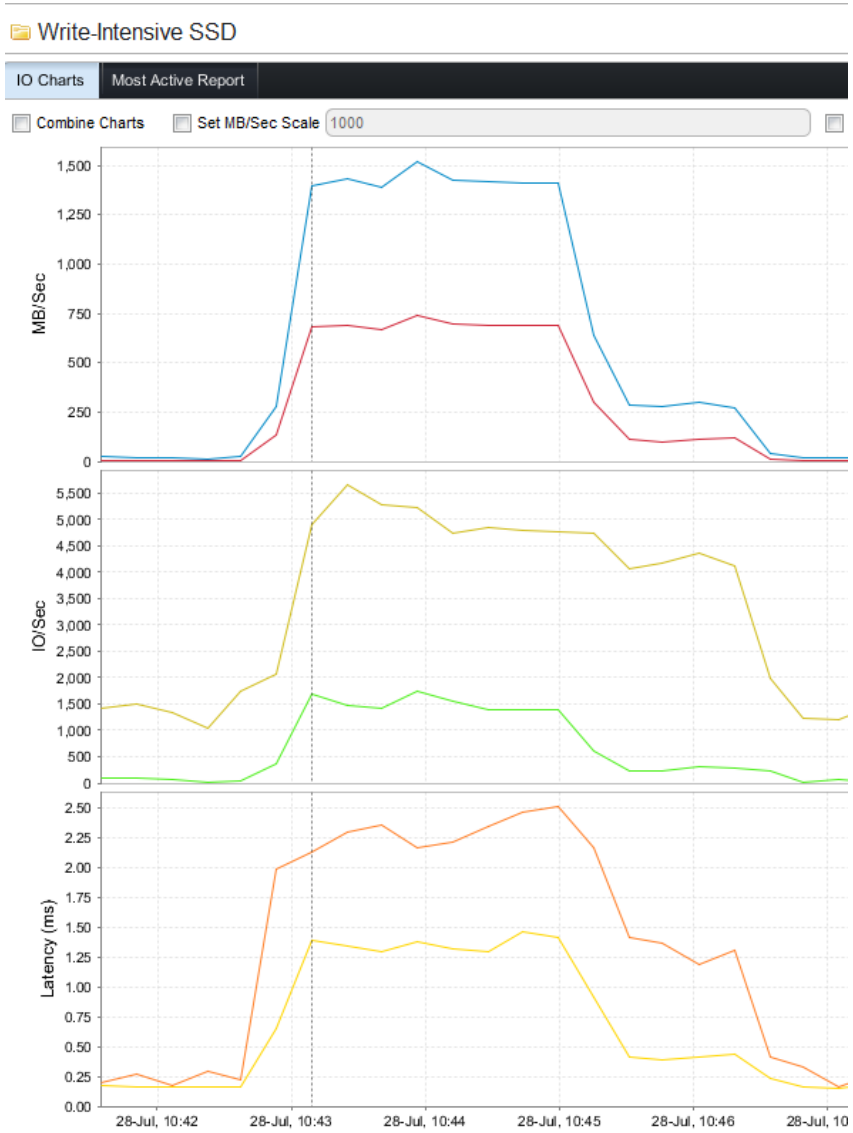

Other observations

- Charting everyday activity
  - Observing workloads is different now
    - Operations are quicker to complete
    - Real-time charting is required to observe some activities
    - Workload spikes are now more common than workload plateaus
    - Daily and weekly charts may not show performance spikes.

Other observations

- Latency
  - Now, looking at microseconds not milliseconds
- Bandwidth
  - Projects with large datasets easier to handle
    - Large aerial photos, LIDAR (Light Detection and Ranging) files, and point clouds can be manipulated to the endpoint
    - All of these files can be multiple GB each
    - Point clouds RAW size can be over 1TB

Other observations (Cont.)

- Scale Up
  - With our investment in flash, we can combine it with high-capacity storage tiers allowing the system to scale to meet the needs of our datasets.
- I/O – what do we need?
  - As we observe the real-work IOP loads, we can see array has more headroom
  - By making 60,000-100,00 IOPS available, future investments are towards capacity and bandwidth
  - 200,000 to 5,000,000 IOPS vs. large capacity, low latency storage?

Other observations (Cont.)

- Costs – Economy of flash
  - With new technologies on the way, at what point will a large system be 100% flash and be more cost effective that 7k spinning disk. Data safety and integrity are very important. Now does not seem the time to use low cost flash in the managed enterprise. When will the technology, reliability, and price combine?
- Flash is no longer a premium resource
  - Presenting flash access to all workloads should be an enterprise priority.

Flash for the future

- Flash capacity
  - How are we going to handle the new larger flash disks that are 20-40 TB in size? What kind of RAID would allow for large disks?
- Use flash to eliminate tape
  - Will the economy of flash allow for safe deep archive of large datasets.
  - Can we use this technology to safely eliminate tape?