

SNIA VDBENCH Rules of Thumb

Steven.A.Johnson@oracle.com June 24, 2013 Colorado Springs







Phases test

- Pre-fill
- Warm up
- Hot band test
- 4 corners test
 - > Small block random read and write
 - > Large block sequential read and write



- Large block prefill and Sequential work loads are significantly different from the small block workloads
- The optimal number of streams needs to be thought through
- Generally speaking the number of streams and threads will be a function of the number of connections to the controller
- Goal is to achieve the maximum bandwidth (GBs/second)



SAS 4 lane x 600 MB/Sec

15K rpm drive, ~200 MB/sec sustained Calculate the number of active drives necessary to saturate Controller

Controller bw / Drive bw = 4,000/200 = -20 drives Set streams = 20, number of threads to 40 or 60 to keep drives busy

Example 2- Medium system Green Storage Initiative **JBODs** Host HBA Controller Four 4x600 MB SAS lanes SAS 4 lane x 600 MB/Sec ~ 2,000 MB/sec / port SAS 4 lane x 600 MB/Sec Controller total bandwidth 8,000 MBs/SEC

7.2K rpm drive, ~150 MB/sec sustained

Calculate the number of active drives necessary to saturate Controller

Controller bw / Drive bw = 8,000/150 = -54 drives Set streams = 50, number of threads to 100 or 150 to keep drives busy



- Small block Hot Band and random work loads are driven by the number of drives in the system and response time
- These workloads have a cap of 20 ms response times
- Goal is to maximize IOPs with out exceeding 20 ms response time
- Need to take into account the drive service times
 - 7.2K RPM drives = \sim 13 ms. 20ms / 13 = \sim 2 threads per drive
 - 10K RPM drives = ~ 6.8 ms. 20ms / 6.8 = ~ 3 threads per drive
 - 15K RPM drives = \sim 5.5 ms. 20ms / 5.5 = \sim 4 threads per drive



144 15K rpm drives, \sim 4 threads/drive = 4*144 = 576 threads

Example 2- Medium system Green Storage Initiative **JBODs** Host HBA Controller Four 4x600 MB SAS lanes SAS 4 lane x 600 MB/Sec ~ 2,000 MB/sec / port Controller total bandwidth 8,000 MBs/SEC

288 7.2K rpm drive, \sim 2 threads/drive = 2*288 = 576 threads



- Each one of these examples is just a starting point
- New technology and specific device properties will may require you to move up or down the number of starting threads (YMMV)
- SSDs are a real challenge. Very fast devices may require a significant number of active threads
- Adjust based on response over 30 ms (reduce the number of threads)
- Adjust based on controller not hitting it's limits (increase threads for more concurrency)

Sample Scatter plot with 4 hot bands, random access and sequential streams





Common problems – uneven phy load







Questions

Steven.A.Johnson@oracle.com

Draft version 6/13/2013