



Performance Evaluation of RAID6

Yan Li, Roland Ibbett, Nigel Topham and Tim Courtney[‡]

School of Informatics
University of Edinburgh
[‡]Xyratex, UK

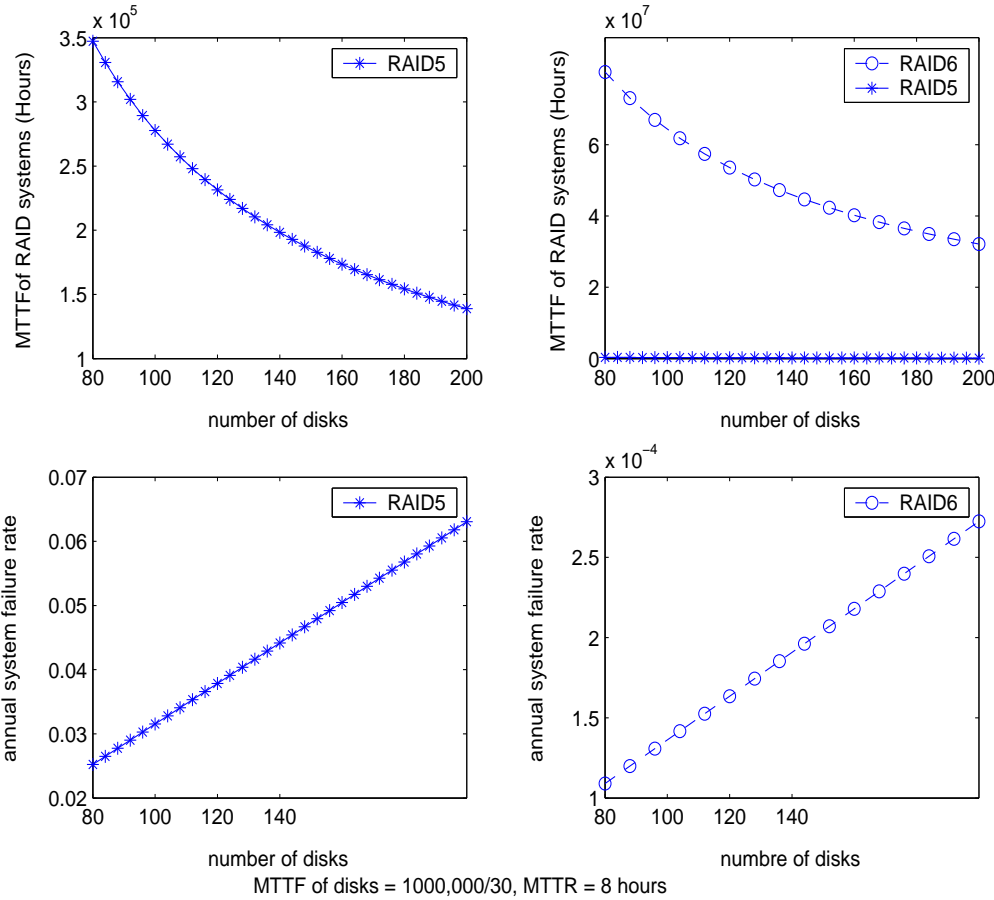


Project Motivation

- Increasing the number of disks results in higher disk failure possibility.
- The larger disk capacity leads to a longer rebuilding time, resulting higher bit error possibility during the reconstruction.
- The end-of-life issue.
- Double disk failure protection scheme RAID6 is becoming more and more necessary.



Reliability of systems over 5 years old





Project Goal

- Evaluate the performance of RAID6 under the SPC Benchmark-1 (SPC-1) workload through simulation and explore the design space.
 - The performance of RAID6 under fault-free operation.
 - The performance of degraded mode.
 - Find out how to allocate resources between rebuilding and serving of incoming requests during the rebuilding stage so that it needs the shortest time for the system to return to normal operation mode




Performance of RAID6

- We focus on the RAID6 algorithms that store the redundant information separately from the data in each stripe, like Read-Solomon, EVENODD and RDP. $8+P+Q$ data layout is used.
- Effect of RAID controller processing time.
- Optimum size of the stripe unit under SPC-1 workload
- The effect of increasing cache size and number of disks and their interactive effects.

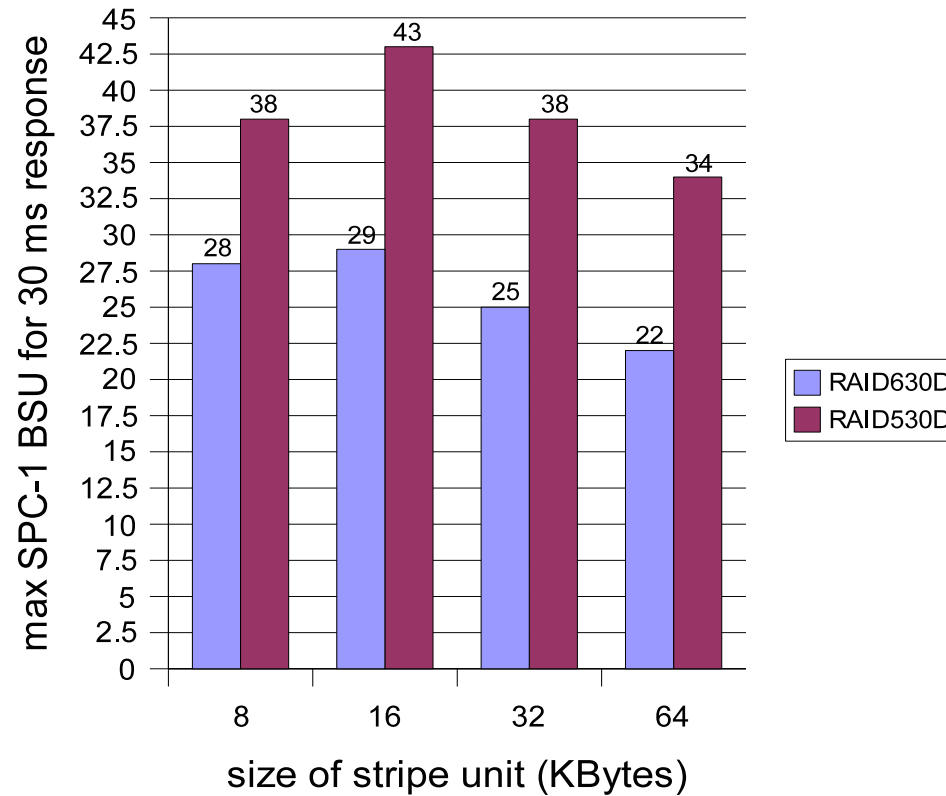


Effect of RAID controller processing time



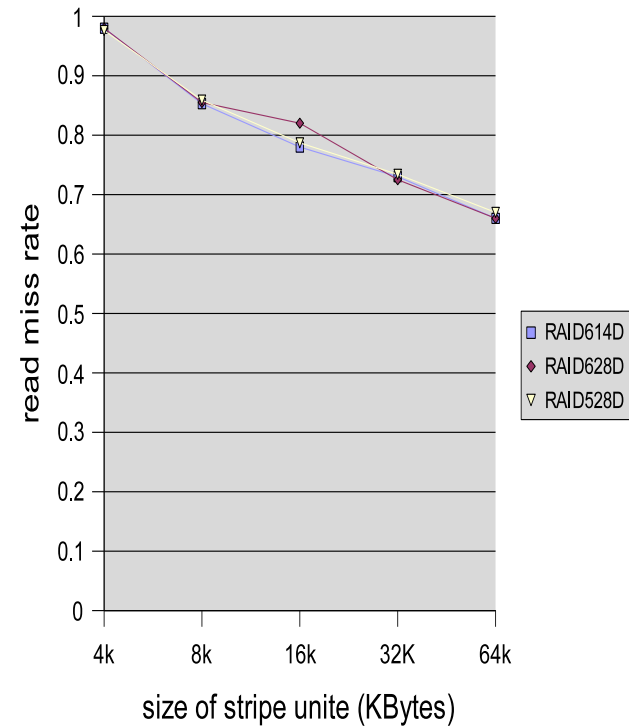



Optimum Size of Stripe Unit without Cache






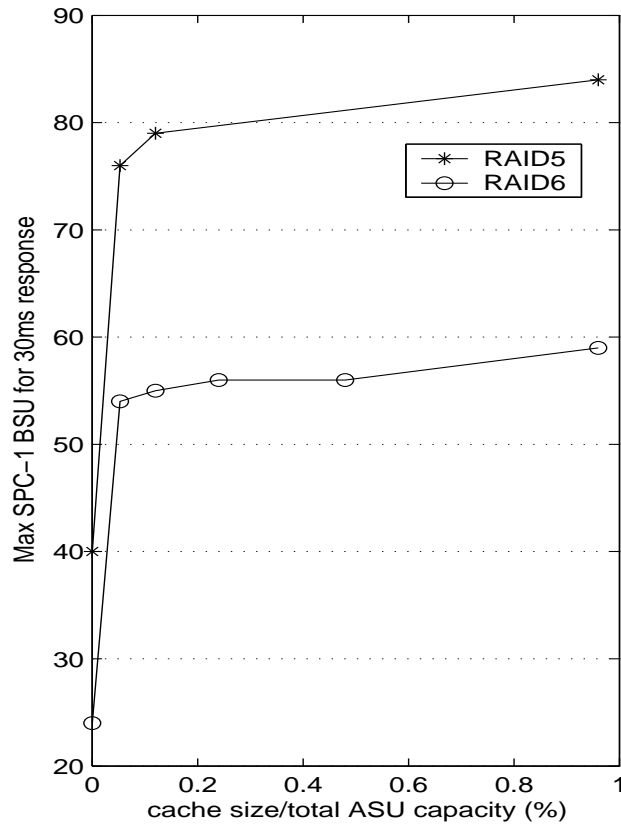
Optimum Size of Stripe Unit with Cache



Cache size = 0.12% total ASU capacity



Effect of Cache Size





Summary

- The achievable maximum SPC-1 BSU number of RAID6 is mainly decided by the disk array. The controller processing time has little effect on that.
- Among the tested values (4, 8, 16, 32, 64 kBytes), the optimum size of stripe unit for RAID6 without cache is 16 KBytes; and that of RAID6 system with cache is 32 KBytes. This is irrelevant to the number of disks in the system.
- With cache that is 0.12% of the total asu capacity, there is about 100% increase on the performance.
- With the same size of cache and asu capacity, the performance of RAID6 is about 70% of RAID5 under SPC-1 workload.



Future Work

- Test the effect of controller processing time with larger disk arrays
- Study the performance of RAID6 under degraded mode.
- The performance of RAID6 under other workload.
- Investigate the best recovering scheme.