

Benchmarking the PetaCache

Benchmarking the PetaCache

Methodology

SCCS has set up a farm in IR2 for us to use while benchmarking the Petacache. As this farm is shared with BaBar simulation generation, we need to schedule the usage for serious tests. For tool development, there should be no conflict.

The farm consists of 40 nodes of dual processor, dual core AMD Opteron(tm) 275 nodes running RHEL 5.5. (AMD nominal speed = 4x2200MHz). (Note: These systems seem to be using dynamic power management to reduce their current cpu speed, so /proc/cpuinfo may show values as low as 1000 MHz.) The node names are `ir2farm{01,02...40}`. Home directories are AFS, and SLAC credentials transfer.

Benchmarking Tools

- Open Source/source available
 - [IOzone](#)
 - [FileBench](#)

FileSystem Benchmarking Projects

- Project FSBench (SUNY-SB)
 - [Main Page](#)
 - [Portal](#)
 - [A Nine Year Study of File System and Storage Benchmarking](#)
 - [Notes on a Nine Year Study of File System and Storage Benchmarking](#)
- Parallel I/O Benchmarking Consortium (ANL)
 - [Main Page](#)
 - [List of Benchmarks](#)

Discussions of Flash

- Sun: Flash/ZFS
 - [Sun News](#)
 - [ZDNet](#)
 - [ACM:Longer article](#)
- Usenix FAST-10 Proceedings
 - [DFS: A FileSystem for Virtualized Storage](#) (Princeton/Fusion-io)
Discussion of filesystem design using commercial flash drives (PCI-EX based). Design is an actual minimal filesystem, managed by the kernel. Loadable kernel module required – thus, superuser privs and kernel access needed. These filesystems might possibly be able to be exported via NFS.
Benchmarks discussed include:
 - Micro benchmarks (write/read rate using Iozone (q.v.))
Problems with their method: only two threads pounding on it.
 - Application benchmarks: Most not useful for us. 3/5 are either mem-mapped IO or database like.
 - [Extending SSD Lifetimes with Disk-Based Write Caches](#) (Slides) (Toronto/Microsoft)
 - [Write Endurance in Flash Drives: Measurements and Analysis](#) (Slides) (Northeastern)

Standalone Papers

- [Benchmarking Parallel I/O Performance For a Large Scale Scientific Application on the Teragrid](#)
- [File System Workload Analysis For Large Scale Scientific Computing Applications](#)
- [PostMark: A New File System Benchmark](#)

Other

- [SLAC mstore](#)
- [HP SFS](#)
- [Parallel IO Examples & Benchmark Codes](#) (Dartmouth)
- [Open MPI](#)
- [Flash Core Set \(FCS\)](#) (SLAC)