

HEC FSIO Session 1: Measurement and Understanding Gaps Roadmap

**John Bent and Evan Felix
August 2008**

2007 Measurement and Understanding Gap Area

Area	Researcher	FY 07	FY 08	FY 09	FY 10	FY 11	FY 12	Rankings
Understanding system workload in enterprise environment	Arpaci-Dusseau							 A comprehensive tool is nowhere in sight; problem is complex.
	Reddy							
Standards for HEC I/O benchmarks	SciDAC - PDSI							 Low on agencies priorities; over simplifies problem and could drive vendors to incorrect solutions. Gap should really be replaced by release of traces, workload characterization, etc.
	SciDAC - SDM							
Testbeds for I/O Research	Ligon							 Simulators are being developed. No real testbeds being built. This problem will only get worse over time, i.e. as systems get bigger.
	Thottethodi							
Applying cutting edge visualization/analysis tools to large scale I/O traces	Reddy							 More traces are becoming available from Labs. Many opportunities to evaluate this research.
	Zadok							

- | | | |
|--|---|---|
|  Very Important |  Greatly Needs Research |  Greatly Needs Commercialization |
|  Medium Importance |  Needs Research |  Needs Commercialization |
|  Low Importance |  Does Not Need Research |  Does Not Need Commercialization |
|  Full Calendar Year Funding |  Partial Calendar Year Funding |  On-Going Work |

Measurement and Understanding Gaps

- Understanding system workload in enterprise environments
- Standards for HEC I/O benchmarks
- Test beds for I/O Research
- Applying cutting edge visualization/analysis tools to large scale I/O traces

Relevant Research

- Choudhary and Kandemir - “Scalable I/O Middleware and FS Opti...”
- **Arpaci-Dusseau - “Formal Failure Analysis for Storage Systems”**
- Chiueh - “QoS Guarantee for Scalable Parallel Storage Systems”
 - Making timing accurate parallel trace play-back tool
- **Narasimhan - “Towards Automated Problem Analysis of Large”**
- **Thottethodi - “Performance Models and Systems Optimizations...”**
 - Performance models and simulation for disk-bound applications
- **Zadok - “FS Tracing, Replaying, Profiling, and Analysis on HEC...”**
- **Mesnier - “//trace”**
- **Ligon - “Improving Scalability in Parallel File Systems for HEC”**
 - Simulation testbed
- Reddy - “Active Data Systems”
- Zhang - “Memory Caching and Prefetching to Improve I/O Perf...”
 - Distributed and hierarchical buffer cache management
- Zhang and Jiang - “Algorithms Design and Systems Implementa...”
 - Disk layout aware buffer cache management
- SciDAC PDSI and SDM