

HEC FSIO Session 5: Next Generation I/O Architectures (Small and Unaligned I/O) Gaps Roadmap

**John Bent and Rob Ross
August 2008**

2007 Next Generation I/O Architectures Gap Area

Area	Researcher	FY 07	FY 08	FY 09	FY 10	FY 11	FY 12	Rankings
Understanding file system abstractions - File system architectures	Choudhary	█	█	█				 Good work, but much of research is in infancy. A small portion ready for commercialization.
	Dickens	█	█	█				
	Maccabe/Schwan	█	█	█				
	Reddy	█	█	█				
	Shen	█	█	█				
	Thain	█	█	█				
	Wyckoff	█	█	█				
	SciDAC – PDSI PNNL	█	█	█	█			
Understanding file system abstractions - naming and organization	Bender/Farach-Colton	█	█	█				 Very hard problem. More researchers need to attack this problem.
	Thain	█	█	█				
	Tosun	█	█	█				
	Zhang/ Jiang	█	█	█				
	SciDAC – SDM	█	█	█				
	SciDAC - PDSI	█	█	█				
Self-assembling, Self-reconfiguration, Self-healing storage components	Ganger	█	█	█				 Good work being done, but it's a hard problem that will take more time to solve.
	Ligon	█	█	█				
	Ma/Sivasubramaniam/ Zhou	█	█	█				
	SciDAC - PDSI	█	█	█				
	SciDAC - SDM	█	█	█				
Architectures using 10 ⁶ storage components	Ligon	█	█	█				 Very little work being done here for a very near term problem. Simulators will/must play a role here
	PNNL	█	█	█				
Hybrid architectures leveraging emerging storage technologies	Gao	█	█	█				 Big potential reward, but very little work being done in the HPC area.
	PNNL	█	█	█				
HEC systems with multi-million way parallelism doing small I/O operations	Choudhary	█	█	█				 Good initial research; needs to be moved into testing. More fundamental solutions being pondered including non-volatile solid state store.
	Dickens	█	█	█				
	Gao	█	█	█				
	FASTOS – I/O Forwarding	█	█	█				

 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

Next Generation I/O Architectures (Small and Unaligned I/O) Gaps

- Understanding file system abstractions – file system architectures
- Understanding file system abstractions – naming and organization
- Self-assembling, Self-reconfiguration, Self-healing storage components
- Architectures using 10^6 storage components
- Hybrid architectures leveraging emerging storage technologies
- HEC systems with multimillion way parallelism doing small I/O operations

Relevant Research

- Choudary - “Scalable I/O Middleware”
- Dickens - “Object Based Caching for MPI-IO”
 - Making physical layout more accurately reflect logical object view
- Gao - “High Throughput Massive I/O Storage Hierarchy for Petascale High-end Architectures”
 - I/O middleware using solid-state buffering
- IOFSL