HEC-IWG
File Systems and I/O R&D Workshop

Scott A. Brandt, Carlos Maltzahn
Darrell D. E. Long, Ethan L. Miller
Storage Systems Research Center
University of California, Santa Cruz

UCSC Storage Systems Research Center

Systems-oriented storage research center focusing on *storage algorithms, architectures, and systems*

Research Challenges:
- Huge capacity and scalability
- Performance
- Security
- Portability
- New storage technologies

Research Thrusts:
1. Peta-scale object-based storage
2. New storage technologies
3. Archival storage
4. Predictive/adaptive techniques
5. Secure storage

SSRC Features:
- High degree of collaboration among faculty, students, visitors, sponsors
- Significant educational component
- Diversified support
- Close cooperation with sponsors□□□

SSRC Sponsors:
- National Labs: LLNL, LANL, SNL
- National Science Foundation
- HP, IBM, Microsoft, Veritas, Intel, Network Appliance, Overland Storage, Hitachi, Engenio, …□□

SSRC File Systems and I/O Research (1)

- High-Performance Object-based Storage (DOE)
  - Goals: 20+ Petabytes, 1 TB/s, billions of files, 10,000+ clients, **and** good general-purpose performance
  - Research: Metadata mgmt., OSD FSes, reliability, data distribution, security, interconnects, QoS, location-aware processing
  - Functional prototype (Ceph)
  - Open source (very soon)
- New Storage Technologies (NSF × 2)
  - Architectures and algorithms for new non-volatile memory/storage devices
  - MEMS: Model, power mgmt., sched., layout, arch., reliability
  - MRAM: Metadata architectures, LiFS, on-line compression
  - FLASH: Performance/power mgmt.

SSRC File Systems and I/O Research (2)

- Archival Storage: Deep Store (NSF)
  - Goals: Efficient scalable on-line write-mostly data storage
  - Research: Differential compression, similarity detection, representation, secure deletion
- Predictive/Adaptive Techniques (NSF)
  - Adaptive: Cache mgmt., disk spindown, file lifetime prediction
  - Predictive: Prefetching, refetching
- Secure Storage (DOE)
  - Secure object-based storage
  - Formal models for storage system security
- New UCSC/LANL Institute for Scientific Data Management
  - Researching infrastructure for scientific data management
    - Real-time data collection → Storage → Information management
  - Research, graduate education, focused MS degree

Areas that need to have more focus

- Enhanced metadata
  - Interested: LANL/LLNL/SNL, Industry, NSF
- File system scalability/evolvability
  - Interested: LANL/LLNL/SNL, Industry
- Digital preservation/archiving
  - Interested: Industry, NSF
- Relaxing FS semantics to be more HPC friendly
  - Interested: LANL/LLNL/SNL
- Predictive/adaptive methods
  - Interested: UCSC

Recommendations

- Greater focus on file system metadata
  - Search, context, relationships, usage patterns
- Relax/change FS semantics
  - Greater parallelism
  - Location-aware filtering/processing
- More flexible storage models/systems
  - File system scalability and evolvability
  - Prediction/adaptation
  - Virtualization
- Long-term storage issues
  - Archiving/preservation, reliability, security