CITI HEC-IWG Presentation

* Peter Honeyman
  Center for Information Technology Integration
  University of Michigan
  Ann Arbor

Current R&D at CITI

- ASC program
- GridNFS (NSF Middleware Initiative)
- Industry-sponsored projects

ASC program

- NFSv4 and parallel file systems
  - Scale and performance
    - Sessions, directory delegation, replication, pNFS, architectures for metadata service
  - Wide-area security assurance
    - Credential mappings, ACLs for multiple realms

GridNFS (NSF)

- Transparent and secure data management
  - Integrated with existing grid authentication and authorization tools
- Scalable and agile name space management
  - Establishing and controlling identity in virtual organizations
  - Specifying virtual organization data resources

Industry-sponsored

- NFSv4 for cluster coherent FS
- NFS over RDMA
- Linux NFS performance

- Sponsors include Sun, NetApp, IBM, PolyServe, EMC

Research focus

- Need more
  - Transparency in global collaboration
- Need less
  - Unsure

Near-term R&D

- Industry sponsors looking for enabling technologies
  - pNFS
  - RDMA and other offloads
  - Administrative tools
    - Migration from AFS, DFS/DCE

Longer-term R&D

- Migration and replication at global scale
  - CPU and network speedup is reshaping the memory hierarchy around fixed latencies (Amdahl's law)
- Multi-realm authorization
  - Multiple authorization domains are barely supported today

Migration and replication

- Migration and replication can hide latency
  - Memory and storage hierarchies
  - Caching
  - P2P, proxy access
  - 1st class replication
- Experience with ATLAS simulator

Research needs

- Backend protocol for migration and replication
- Protocols and tools that integrate migration and replication with other virtual organization resource management tools
- Memory hierarchy goes global
  - Automated, scheduled
- Scaling associative access to massive metadata
- Reliability vs. scale in storage over high-latency network

Multi-realm scaling

- Multiple authorization domains work today!
  - But only for long-term federation of a small number of realms.
- Breaks down for increasing number of realms
- And for the dynamic relationships of virtual organizations
- Especially when roles are taken into account

Research needs

- Tools for constructing and managing virtual organizations
- Agile resource scheduling
  - Computing plants
  - Network reservations
  - Global storage