Breakout 2 Gaps and Needed Breakthroughs

Session Chair: Lee Ward (SNL)

Translator: Rob Ross (ANL)

Breakout Room 2: Sessions

- Metadata
- Quality of Service
- Measurement and Understanding
- Next Generation I/O Stacks

- Format
  - Slide of topics from breakout session
  - Slide of top ranked issues (score in parenthesis)

- On the score
  - 1st = 5 points, 2nd = 4, 3rd = 3, ...

Metadata: Brainstorming

- Commodity devices, reliability, and metadata
  - Silent data corruption on commodity storage an increasing risk
- Name space operations at massive scale (was ls-lR)
  - Layout improvements, recursive tree walk, indexing of name space
  - Consistency of name space
- Structured data formats
  - Lack of knowledge of data formats at the file system layer
  - Structured and semi-structured data
- Interoperable data formats for use between HEC sites
  - Possibly covered by Grid projects (SRB, SRM, etc.)
- Content-related metadata
  - Automatic generation, search and reference based on content-related metadata
  - Well-defined semantics/usage models for this metadata for sharing purposes
- Creates and deletes per second
- Correctness in face of asynchrony

Metadata: Ranked Gaps and Needed Breakthroughs

Top 3 of 7:

- Correctness in face of asynchrony (109)
- Content-related metadata (100)
  - Automatic generation, search and reference based on content-related metadata
  - Well-defined semantics/usage models for this metadata for sharing purposes
- Creates and deletes per second (98)

#2 for Gov't, #1 for Industry:

- Commodity devices, reliability, and metadata (90)
  - Silent data corruption on commodity storage an increasing risk

Honorable mention:

- Name space operations at massive scale (was ls-lR) (91)
  - Layout improvements, recursive tree walk, indexing of name space
  - Consistency of name space

Quality of Service: Brainstorming

- Defining QoS as it relates to HEC
  - Different than QoS in real-time viz, data capture, multimedia, etc.
  - Evaluating/Benchmarking QoS
- Robust Adaptability
  - Still not done, robustness is a new emphasis
  - Feedback mechanisms (likely to be seen in protocols as well)
- Hard, real-time, end-to-end, QoS
  - Implication of errors on this capability
  - In enterprise environment
- Security and availability as relate to QoS (other dimensions than perf.)
- Policy management for QoS
  - Authorization for granting QoS
  - Policy control over robust adaptability
  - Computing bounds (e.g. utility curves from networking)
- Next-generation networks as building blocks for storage QoS
- QoS on commodity hardware
- Software component interfaces for composition of QoS applications

QoS: Ranked Gaps and Needed Breakthroughs

Top 3 of 8:

- Robust Adaptability (105)
  - Still not done, robustness is a new emphasis
  - Feedback mechanisms (likely to be seen in protocols as well)
- Defining QoS as it relates to HEC (103)
  - Different than QoS in real-time viz, data capture, multimedia, etc.
  - Evaluating/Benchmarking QoS
- Policy management for QoS (90)
  - Authorization for granting QoS
  - Policy control over robust adaptability
  - Computing bounds (e.g. utility curves from networking)

A distant 4th:

- Hard, real-time, end-to-end, QoS (57!, but #1 for gov't and industry!)
  - Implication of errors on this capability
  - In enterprise environment

Measurement and Understanding: Brainstorming

- Understanding system workload in enterprise environment
  - Multi-application, aging, reconfiguration/redistribution
  - Models of these systems
- Applying cutting edge visualization/analytics to large-scale I/O traces
  - Identifying causal dependencies in traces
  - Scalable analysis tools
- Implication of order of magnitude increase in disk heads (10^4-10^5 range)
  - Simulation of systems of this scale
- Testbeds for I/O Research
  - Scalability, security, etc.
- Standards for HEC I/O benchmarks
  - Relevance and coverage of HEC I/O benchmarks
  - Accounting for aging and other known issues in repeatability
  - Interchangeability of trace data
- Measurement techniques
  - White box, black box, ...

M&U: Ranked Gaps and Needed Breakthroughs

Top 3 of 6:

- Understanding system workload in enterprise environment (110)
  - Multi-application, aging, reconfiguration/redistribution
  - Models of these systems
- Standards for HEC I/O benchmarks (101)
  - Relevance and coverage of HEC I/O benchmarks
  - Accounting for aging and other known issues in repeatability
  - Interchangeability of trace data
- Testbeds for I/O Research (90, #1 for gov't)
  - Scalability, security, etc.

Honorable mention:

- Applying cutting edge viz/analytics to large-scale I/O traces (88)
  - Identifying causal dependencies in traces
  - Scalable analysis tools

Next Generation I/O Stacks: Brainstorming

- Architectures using 10^4-10^5 storage components
- Hybrid architectures leveraging emerging storage tech. (e.g. MRAM, NAND)
- Virtualized storage devices
- System-wide power management
- Introspective devices (high-level queries)
- Self-assembling/reconfiguring/healing storage components
  - Autonomics (mechanisms and policies)
  - Interfaces for guiding
- Underlying FS abstractions
  - Next-gen VFS interface
  - Alternative naming and organization schemes
  - Possible convergence of DB and FS technologies, transactional FSes
- Application I/O interfaces
  - HEC application interfaces for data access
  - I/O Middleware
- Extensible file systems and interfaces for extensibility
- Quality assurance for FS implementations
  - Tools, languages, etc. for enabling

NGIO Stacks: Ranked Gaps and Needed Breakthroughs

Top 4 of 10:

- Underlying FS abstractions (111)
  - Next-gen VFS interface
  - Alternative naming and organization schemes
  - Possible convergence of DB and FS technologies, transactional FSes
- Self-assembling/reconfiguring/healing storage components (81)
  - Autonomics (mechanisms and policies)
  - Interfaces for guiding
- Architectures using 10^4-10^5 storage components (78)
- Hybrid architectures leveraging emerging storage tech. (73)
  - e.g. MRAM, NAND

Interesting anomaly:

- Application I/O interfaces (62, 2nd by gov't ranking)
  - HEC application interfaces for data access
  - I/O Middleware