

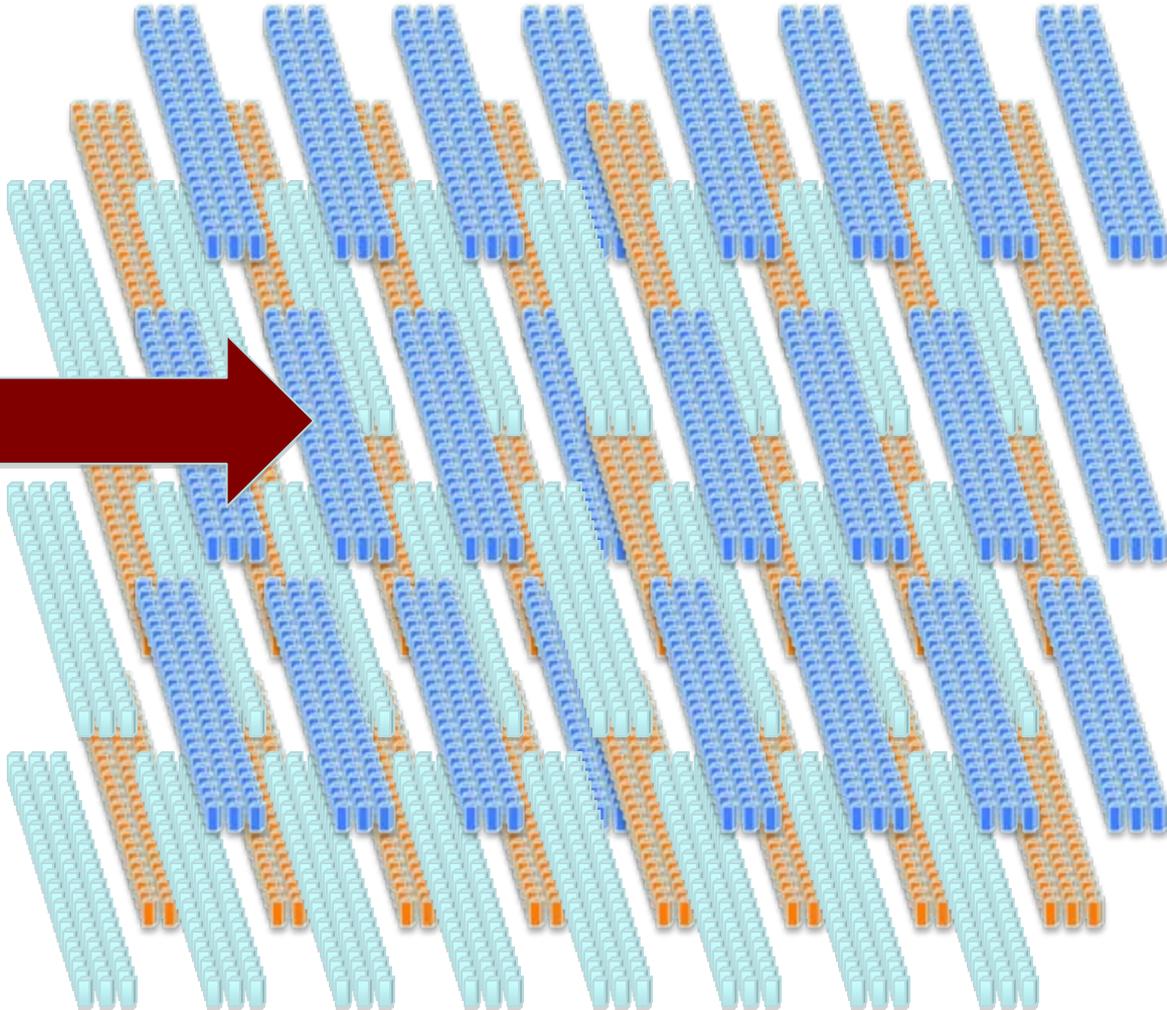
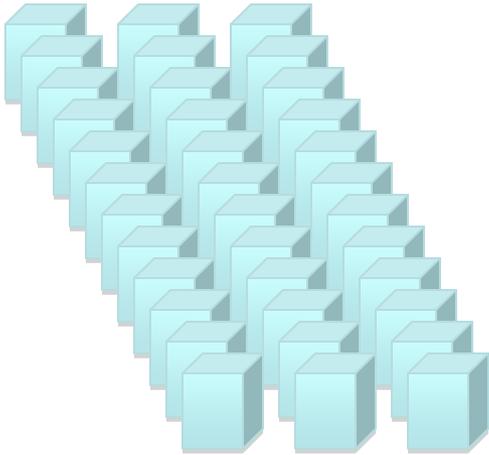
HEC FSIO Session 6: Management and RAS Talks & Roadmap

Evan Felix, Pacific Northwest National Lab

August 2009

Management and RAS

Today



- Manage
- Monitor
- Automatically healing
- Report failures
- Manage power

Current R&D Gaps 2008

- Automated problem analysis and modeling
- Formal failure analysis and tools for storage systems
- Improved scalability
- Power consumption and efficiency
- Reliability, and degraded performance in HEC systems

2009 R&D Gaps from Breakout

- Manual and Automated problem analysis and action
 - Lack of standard API for health reporting from all parts of stack including storage devices of all flavors
- **Understanding failures we face** (Should this be moved to Measurement & Understanding)
 - New storage technologies
 - Subtle problems are hard to find
 - Silent data corruption
- Formal failure analysis and tools for storage systems
- **Proactive health methods**
 - Compensating for changing reliability over time
- Improved scalability (management of all parts of stack, including data collection and fusion)
- Power consumption and efficiency
 - Entire storage path power (should this be in Next Gen I/O or M&U)
- Reliability, and degraded performance in HEC systems
 - Carry on in the face of failure even if degraded

Notable Comments

- Failure of Processors and memory was deferred until the 2009 National HPC Workshop on Resilience
- Large Test Beds are needed for M&RAS also.
- “We were really, really, really, really, really far away before; now we can remove a single really.”
- “The amount of power taken by storage is laughable compared to compute.”
- We need an place for researching the multiple ‘Trade-Off’s we face in building large systems
- “Still want human interactive tools, not a completely automated system “

Management and RAS Results

