

Next Generation I/O and Visualization Requirements for In Situ Processing

BoF at SC
November 15, 2011

Jay Lofstead



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.





Limiting Factor

- **Science workflows involve a LOT of IO to store intermediate data**
- **Exascale makes petascale systems worse**
 - 1000x computation
 - 100x memory
 - 10x IO bandwidth
- **How do we do 10x the IO 10x as often without enough raw bandwidth?**



Proposed Solution

- **Eliminate disk-based storage as much as possible**
 - Flash is not a solution (cost/speed for bulk ops)
 - Many parallel file systems inadvertently limit saturating file system (e.g., Lustre and 160 OSTs)
- **Move to all online, interactive simulations**
 - Set parameters, watch and explore in real time
 - Store parameters and provenance for reproducibility



What Needs To Change?

- **Simulations must become slaves to the viz/analysis**
 - eSiMon a good first step at what it might look like
- **Resilience must be rethought**
 - No longer overload analysis output with resilience technique
- **Assembly of online workflows must be independent of science source code**



Current Progress

- **File systems community exploring alternatives to POSIX for metadata and data storage**
- **Data staging efforts figuring out data storage/processing requirements**
- **IO APIs (particularly ADIOS) abstracting away all but the start/write/stop processing affording redirection at runtime**



Abstract

- Next generation exascale systems will require HPC to revolutionize the concept of I/O, data analysis, and visualization. HPC is ready for this revolution, where we can build tools, which are easy-to-use, sustainable, scalable, and portable. This BOF will focus on generating new ideas to design new algorithms and middleware to aid scientist in high performance, data intensive computing. Leadership class computing has allowed application scientists to solve problems that would have been impossible 5 years ago. Much of the focus of the community has gone into creating scalable science to work efficiently on these petascale platforms. One consequence has been that the data generated from these simulations has ballooned to unprecedented volumes, opening up new challenges and opportunities. Simulations of this scale need to incorporate new techniques to meet these challenges which requires tools that can help experts to make real-time decisions of what to do with the data.
- 5 minute presentation