

Challenges in Data Intensive Computing

Rob Ross

Mathematics and Computer Science Division
Argonne National Laboratory

ross@mcs.anl.gov



Abstract (in short)

- Some HEC sites have been pushing for **globally accessible, secure, high-performance parallel file system** (GASPFs) that is directly accessible by
 - parallel applications on HEC systems
 - other applications on other on-site systems (e.g. visualization, analytics)
 - wide-area users (e.g. scientists working remotely)
- Nobody has made this happen
- Meanwhile, three **distinct interaction models have emerged**
 - POSIX (or POSIX-like) and structured I/O reign supreme on HEC systems
 - Search capabilities lessening role of directory hierarchy on the desktop
 - Storage management systems abstracting away hierarchy in the wide-area

Questions

- **How would you characterize the models of interaction with HEC storage systems?**
 - On the HEC machine?
 - At a desktop or other system on-site?
 - Remotely?
- **Given these models of interaction, how compelling is the vision of a globally accessible file system?**
 - What parts of the vision match with your models of interaction?
 - What is missing from the vision?
 - What is unnecessary?
- **What are the biggest technical challenges facing the implementation of your vision?**
 - Are there technical challenges that seem insurmountable?
 - What technologies can we adopt from other areas?
 - What should we be focusing on now?

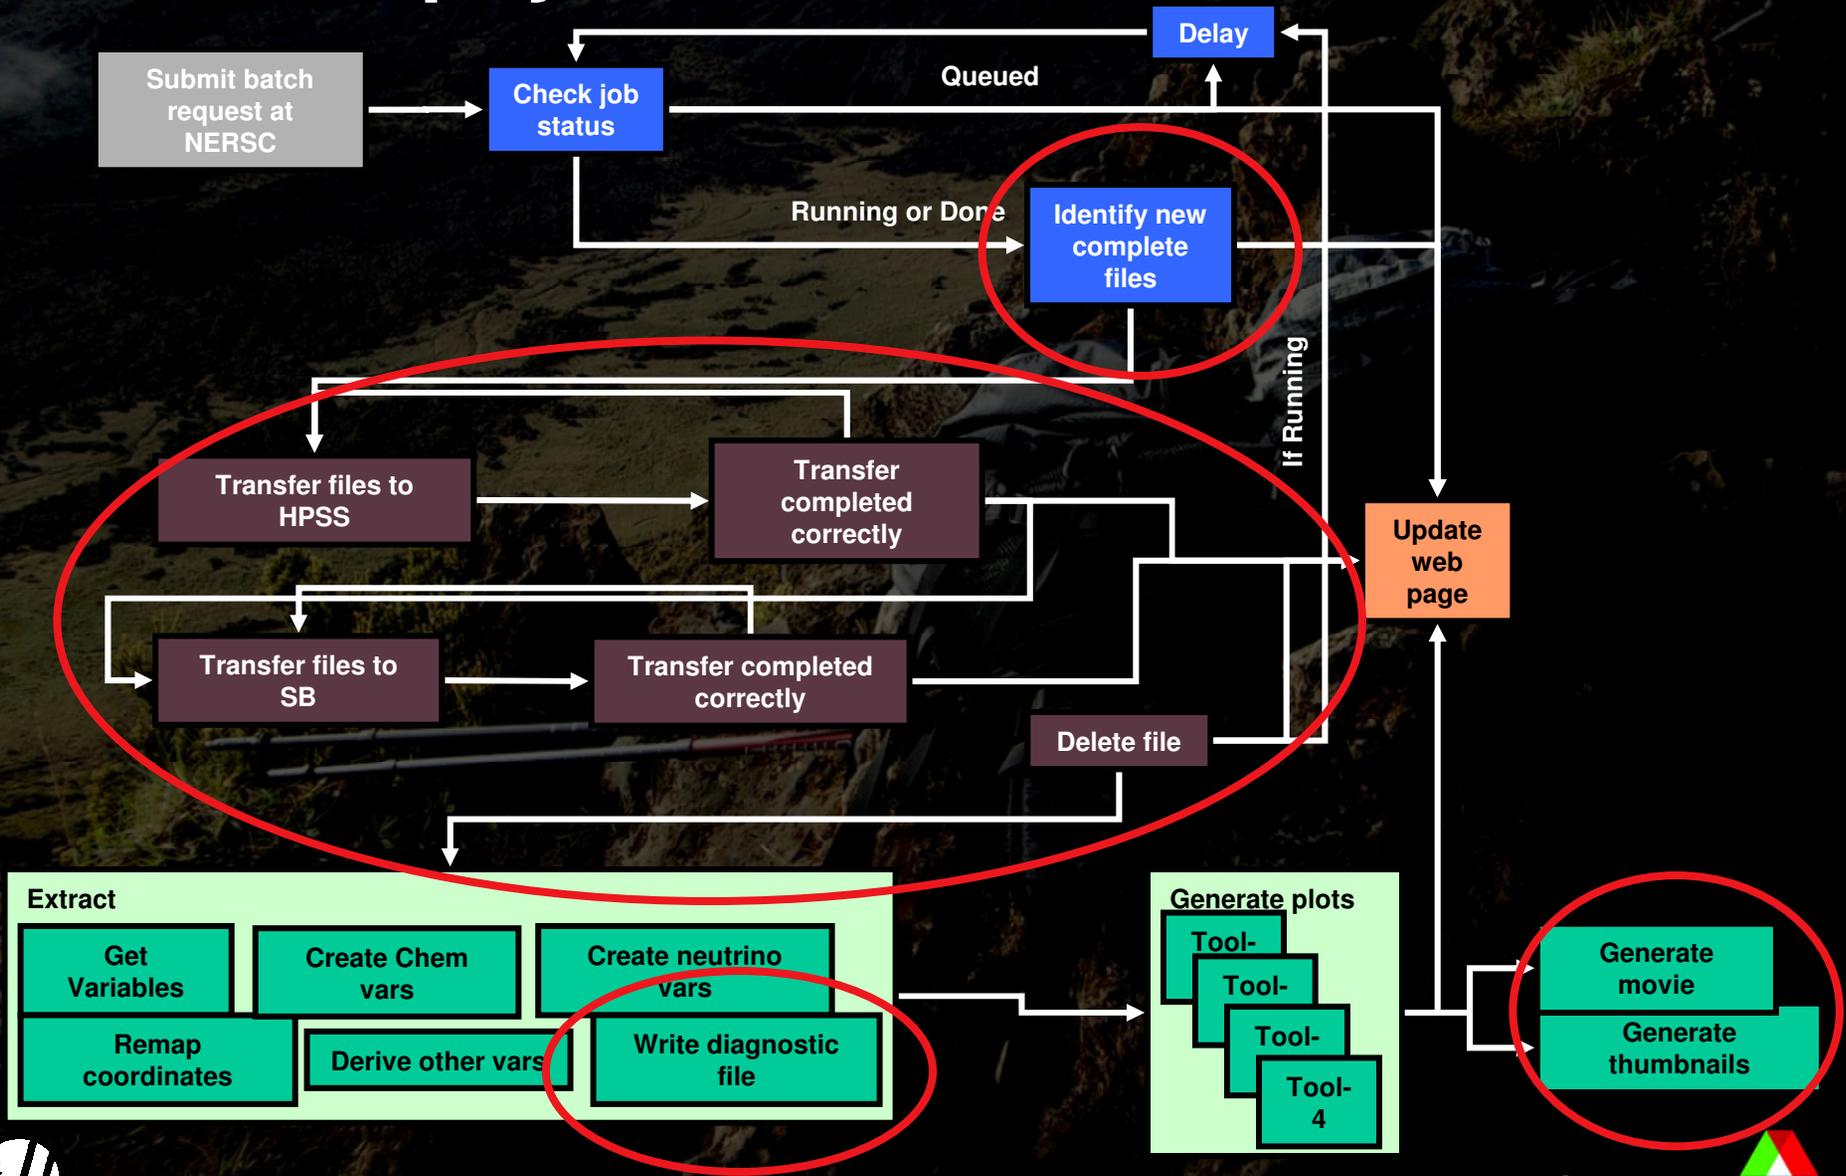
Panelists

- Bill Blake, Netezza
- Matt O’Keefe, Cray
- Jim Tomkins, Sandia
- David White, Sandia
- Pete Wyckoff, Ohio Supercomputer Center
- Lee Ward, Sandia

My model of the scientific process

- Scientists are increasingly mobile
 - Collaborations
 - Working off-site
 - Visiting exotic locales
- Wherever they are, want to manage jobs
 - Check status of running jobs, start new ones
 - Move data from one site to another
 - Perform post-processing
 - View results of post-processing
- **This is mostly a workflow problem...**
(When did I drink that Kool-Aid?)

An astrophysics workflow



Very little data makes it to the scientist

- Simulations can create hundreds of TB of data
 - Checkpoints
 - Data needed for visualization
 - Often write-once
- Only a subset is ever read again
 - Post-processing on certain variables and time-steps
 - Only certain files are ever touched
- Post-processing output is viewed by scientist
 - Scientist doesn't directly access anything else

Why don't we need a global file system?

- Workflow systems automate movement of data where it is needed
 - Dedicated tools for wide-area transfer exist and should be more efficient than a FS for moving data
- No need for all data to be visible all places
 - Workflow system makes final results available conveniently
 - Scientist never looks at raw files in most cases
- No I/O semantics challenges at this scale
 - Only coarse grain movement between systems

Challenges in data intensive computing

- **Data flow management in workflow tools**
 - Tools like Kepler still maturing
- **File systems and wide-area data movement**
 - Teach wide-area tools about parallel I/O and PFSeS
 - Sub-setting in wide-area transfer tools
- **Metadata management**
 - Integrating search into all file systems (?)
 - Automatically storing provenance
(metadata describing how file was created)
 - Relationship between PFSeS, archives, and other online data resources



On to the panelists!

Please save questions until after all the panelists have had a chance to speak.