

NNSA Tools Users View “From the Trenches”

ASRC Exascale Tools Workshop

October 13, 2011

Dave Montoya

Reference to ASC Workshop on R&D Challenges for
HPC Simulation Environments, Working Group on
Exascale Tools, March 23-24, 2011

Who are Users?

- Application Developers/ Support Analysts
- System Integrators / Developers
- System Administrators / Support

What they have seen coming

- Increased scale
- Multi-node parallelism
- Memory Hierarchy
- Heterogeneity
- Fault Resiliency

Developers that both add the science to the application and need to make it work in new HPC environments. Don't like change and would rather focus on the science. Want to focus on the application and would like the environment to just work.

Issues:

- Feel that we are still working on effective Petascale Tools
- Understanding Application Interaction with Environment; runtime, architecture, infrastructure. Need tools to analyze for further optimization.
- Exascale environment features new level of complexity; cores, threads, heterogeneity, deeply hierarchical
- Where to go? Balance new science vs. new programming models
- Large integrated codes, very complex, hard to change

The development process toward an exascale ecosystem. This is the ability for expertise at HPC sites to have the ability to assess/develop requirements and work with vendors that integrate with the Open Source communities to develop capability. - Tools

Issues:

- Need ability to assess balance of Architecture/Infrastructure/Application(s)
 - memory, bandwidth, message rates, runtime, scheduling, network, data infrastructure, etc.
- HPC Environment development (how to structure a tools friendly environment): Vendor interaction, coordination, collaboration
 - Open Source Community – develop standards and coordination
 - Vendors – platform, infrastructure, tools. Hesitant to cooperate.

This is the group that needs to make the environment work regardless of the design and development effort. Less time to focus (things have to be fixed NOW!), have limited tools and are facing scale transition with limited tools to support.

Issues:

- Day to Day care the environment and feeding of jobs to be run
 - Need throughput and successful results
- Troubleshooting problems with users and environment
 - Environment needs tools that are integrated to provide a picture that can be assessed
 - Tools need same scale and data reduction/filtering capabilities that application tools need
- Ensure health and optimum performance
 - A monitoring and probing infrastructure that supports sustainability

- Tools for application development (debugging, correctness, performance)
 - Wide spectrum: memory, power, locality, resilience, ...
 - Environment information and feedback
 - Static analysis tools for code evaluation
- Tools for System Software (SSW) to evaluate/integrate the exascale stack itself
 - SSW, I/O, Network, File systems, Scheduler, ...
 - Integration of stove piped systems into an integrated whole
 - Need to get away from ad-hoc tools, need whole system solution
- Tools for System Environment Troubleshooting (debugging, health, performance)
 - Wide spectrum: computing cores, network, data sub-systems, run-time ...
 - Monitoring tools for health, optimization

- Shared infrastructure for measurement, data gathering and presentation
 - Online analysis, data aggregation, shared across the system stack
- HW and SW APIs / information exchange/collaboration with vendors
 - **** **This is a Big One** *****
 - APIs that we want to wrap and monitor
 - Introspection APIs (HW and SW)
 - Guidance for other system components (targeted, information isolation)
 - APIs exposing information from the users
 - Need the ability for tools to be tightly integrated with the overall environment
- Post-mortem, online, in situ and batch tools
 - Resources for testing/validation of the system (incl. tools)

- Scalability of measurement, analysis, and presentation
 - Incl. new metrics: memory, power, ...
- Turning information into insight
 - Despite flood and complexity of data from billions of threads
- Quick design of prototype tools for new scenarios
 - Agile development to keep up with PMs
 - Need them early, enable specialized tools in this and other areas
- Getting right interfaces with the right abstractions
 - To SSW, HWA, Apps, Libraries, Runtimes, Compilers, ...
- Dealing with new programming methodologies
 - Heterogeneous systems/architectures (HW and SW)
 - Coupled systems and applications
- Resiliency for tools and tool infrastructures