



## MAGIC Meeting Minutes

May 1, 2013

### Attendees

|                          |                 |
|--------------------------|-----------------|
| Rachana Ananthakrishnan, | ANL             |
| Bob Bonneau              | AFOSR           |
| Scott Brim               | Internet2       |
| Rich Carlson             | DOE/SC          |
| Dan Gunter               | LBL             |
| Shantenau Jha            | LSU             |
| Kate Keahy               | ANL             |
| David Martin             | Northwestern U. |
| Barney McCabe            | ORNL            |
| Grant Miller             | NCO             |
| Von Welch                | Indiana U.      |

### Action Items

### Proceedings

This MAGIC Meeting was chaired by Rich Carlson of DOE/SC. Arthur (Barney) McCabe of Oak Ridge National laboratory (ORNL) provided a discussion of Recent Trends in Operating Systems and their Applicability to HPC. [http://www.nitrd.gov/nitrdgroups/index.php?title=Middleware\\_And\\_Grid\\_Interagency\\_Coordination\\_\(MAGIC\)#title](http://www.nitrd.gov/nitrdgroups/index.php?title=Middleware_And_Grid_Interagency_Coordination_(MAGIC)#title)

### Recent Trends in Operating Systems and their Applicability to HPC: Barney McCabe

Researchers are developing new applications and architectures to run applications such as science models faster on supercomputers. This talk addresses what is happening to enable high performance computing. A wide range of approaches has been developed including: Mach, microkernels, L4, CNK, lightweight kernels, Catamount, hypervisors, VMware, and Xen.

Microkernels provide minimal services, policy implementation, address spaces, control threads, and message passing, servers and trampoline. Hypervisors virtualize hardware to support multiple Operating Systems. Xen virtualizes the processor, MMU, and basic I/O. Lightweight operating systems include Catamount (SUNMOS, Puma/Couger and Catamount portals) and Blue Gene/L (Compute Node Kernel and I/O nodes). The Blue Gene/L I/O nodes run Linux, have storage resources, and run a separate I/O network. The compute nodes run a lightweight kernel and run a high-speed partitionable network.

Linux provides a wide range of services: libraries, development environment, and a work environment. It works on a wide range of hardware. Linux in HPC doesn't really have many devices so it, likely, provides services. Linux on Blue Gene/L already runs Linux to some extent. Catamount also uses Linux. The Quintessential Kernel (QK) is Xen and PCT is Dom 0. QK virtualization is accomplished through PCT (building address spaces, running

FOR OFFICIAL GOVERNMENT USE ONLY

c/o National Coordination Office for Networking and Information Technology Research and Development

Suite II-405 · 4201 Wilson Boulevard · Arlington, Virginia 22230

Phone: (703) 292-4873 · Fax: (703) 292-9097 · Email: [nco@nitrd.gov](mailto:nco@nitrd.gov) · Web site: [www.nitrd.gov](http://www.nitrd.gov)

contexts, and providing portals for the network). XenLinux emulates Xen hypercalls. A realistic picture is to start with XenLinux, minimizing modifications, building a wedge to provide the QK interface and the wedge can support page table construction. Using XenLinux on Catamount enables evolution by evolving the wedge, then PCT. It enables a minimal number of supported code bases.

The real goal is to run a single application per node (multiple processes, multiple threads).

### **Upcoming Meetings**

July 22-25 XSEDE Meeting, San Diego

Week of July 30-Aug 2, OGF Workshop, Miami, Florida

Week of July 30-Aug 2 Federated Cloud Workshop, Germany

November 11-15 InCommon Identity Week, Silicon Valley

OSG and XSEDE are offering a summer school to provide understanding of the principles, concepts and applications. A link to this meeting is provided on the XSEDE Web page.

### **Next MAGIC Meetings**

- June 5, 2:00-4:00, NSF, Room II-415
- July 3, 2:00-4:00, NSF, Room II-415