



GRID ENGINE WORKSHOP

Introduction

Fritz Ferstl

Director Sun Grid Engineering & HPC on Solaris



Grid Engine Workshop Agenda

- Community status, new features in 6.2, TACC, Grid Engine XML, Grid Engine future plans
 - > #1+2: Tue, May 13 – 10:30 AM - 3:00 PM
- Service Domain Management, Multi-Clustering, Accounting & Reporting, HA & Visualization with SGE, PluS Advance Reservation
 - > #3+4: Wed, May 14 – 8:30 AM – noon
- Voice of Community BoF
 - > #5: Wed, May 14 – 3:15 – 4 PM PM



Grid Engine Tutorials & Trainings

- Using new features in SGE 6.2
 - > Tue, May 13 – 4:30 - 6:00 PM
- Service Domain Manager (Hedeby) Installation
 - > Wed, May 14 – 1:30 - 3:00 PM
- Multi-Cluster Accounting
 - > Thu, May 15 – 8:30 – 10:00 AM
- Managing Grid Engine Clusters
 - > Fri, May 16 – All Day

Other Grid Engine Related Talks

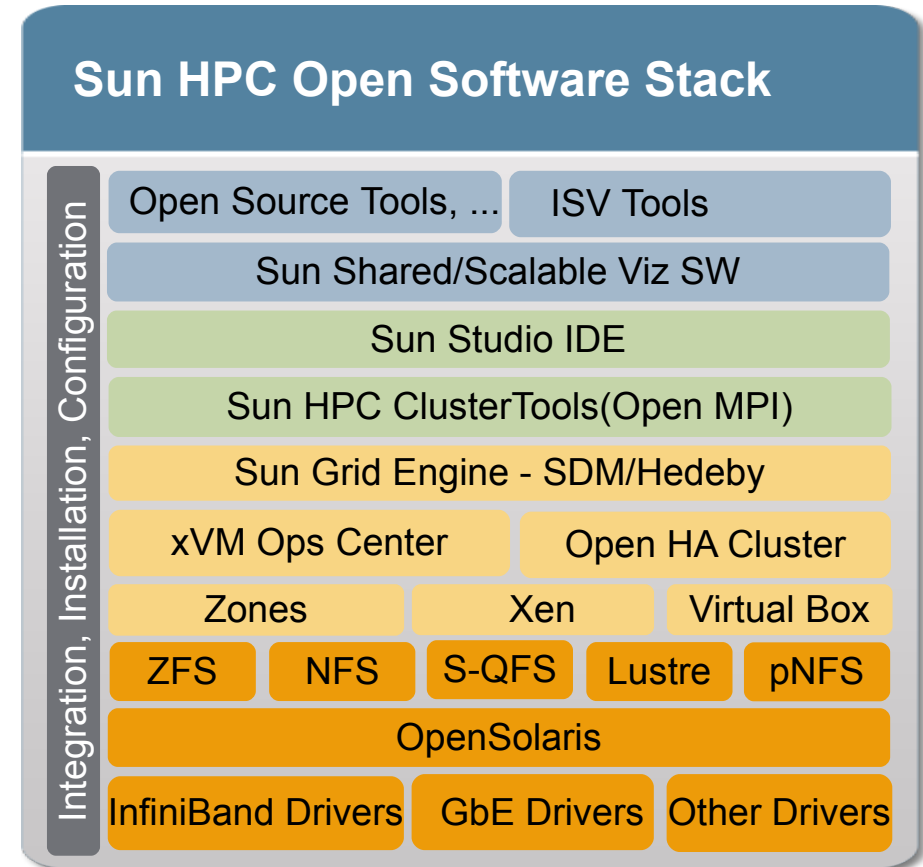
- Sun HPC Cluster Stack, Dan Templeton
 - > Tue, May 13 – 11:30 AM - noon
- Grid Engine in EDA, Grid Heating, LSF vs. Grid Engine
 - > Tue, May 13 – 4:30 - 6:00 PM
- Virtualization: Grid Engine & EC2, OpenNEbula
 - > Wed, May 14 – 1:30 – 3:00 PM
- Lustre and OpenSolaris Storage
 - > Thu, May 15 – 9:00 – 9:30 AM

New: OpenSolaris/Solaris HPC Distro

<http://opensolaris.org/os/community/hpcdev/>

<http://opensolaris.org/os/project/hpc-stack/>

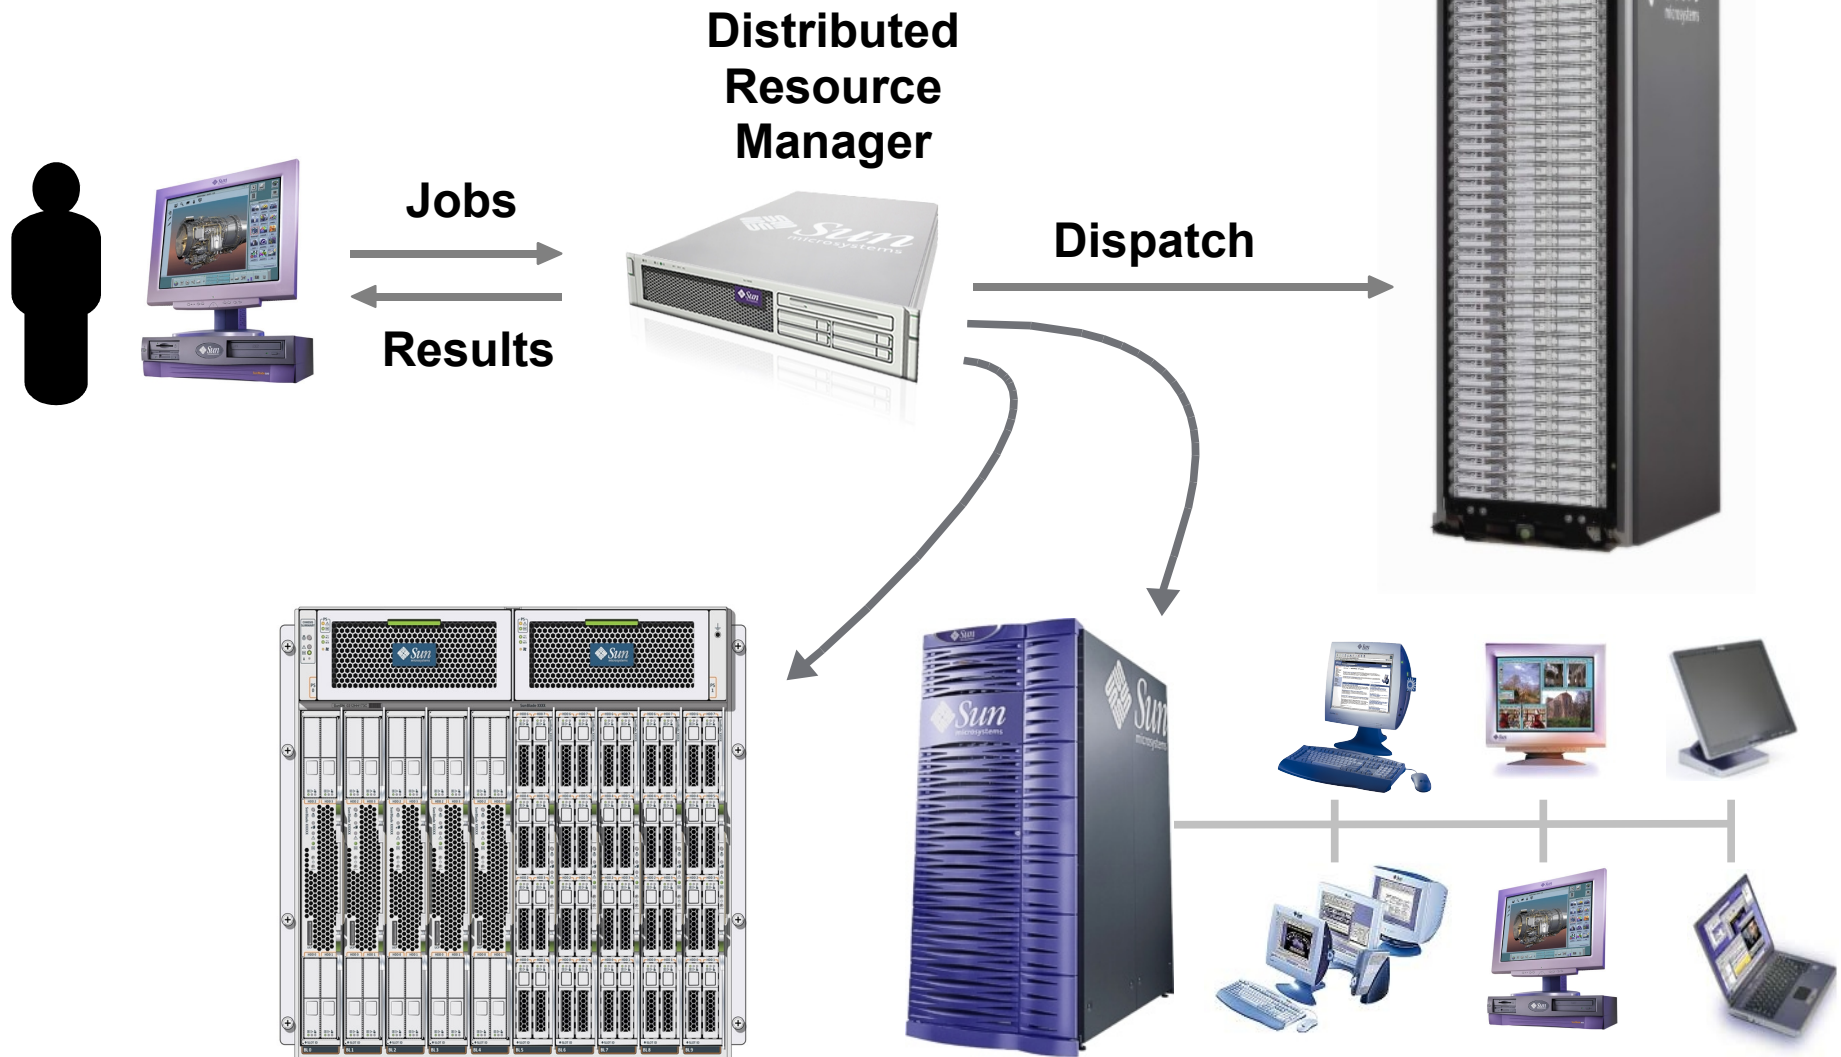
- HPC Developer Community & HPC Stack Project
- Objective:
 - > Build a tightly integrated, open & low-barrier-to-entry stack
- Two life lines – two platforms Solaris 10 & OpenSolaris*
 - > HPC Developer Preview
 - Centered around IDE
 - Seamless developer experience
 - > HPC vertical distros
 - Deployment stack targeted at specific vertical market segment needs
 - Integrated with OpenSolaris and Solaris 10
 - Also available unbundled



* = A Sun-supported HPC stack is also being worked on for Linux

What is Sun Grid Engine?

Dynamic Resource Management



Grid Engine & Sun Grid Engine

The ultimate Try & Buy



Grid Engine Open Source

- Full source code access; liberal license (SISSL)
- Regular courtesy binary updates
- Open development and issue tracking
- Vibrant mailing lists
- Community contributions



Sun Grid Engine

- Indemnification (also for free download from SDLC)
- Strong Support options; relief patches, on-site diagnosis, etc.
- More “muscle” to influence product direction
- Partner Solutions (UnivaUD UniCluster)
- Integrated in Sun's HPC Stack

Community Update - Overview

- Status gridengine.sunsource.net
- Projects
 - > ARCo
 - > Hedeby
 - > Testsuite
- Site Statistics
- QA Information
 - > Bug / RFE Statistics

Status gridengine.sunsource.net

- **UPDATE** ~ 700 Members
- Mailing lists
 - > <list>@gridengine.sunsource.net
 - > <list> = announce, users, dev, cvs, issues, gep, globus
- Size of content / code
 - > ~ 60MB (10 MB binary, 50MB ascii)
 - > ~ 8500 files
 - > ~ 1.5 Million lines (1 Million code + 0.5 Million documentation)
- Documentation
 - > on gridengine.sunsource.net
 - > on docs.sun.com (in transition to a Wiki)

Sub-Project ARCo

- 30 Members
- Mailing Lists
 - > announce, users, dev, issues from gridengine project
 - > cvs@arco.sunsource.net
- Size of content / code
 - > ~ 4MB (1MB binary, 3MB ascii)
 - > ~ 850 files
 - > ~ 80000 lines (code + documentation)
- Documentation
 - > on docs.sun.com (also goes to Wiki)

Sub-Project Hedeby

- 27 Members
- Mailing Lists
 - > <list>@hedeby.sunsource.net
 - > <list> = announce, users, dev, cvs, issues
- Size of content / code
 - > ~ 34MB (10MB binary, 24MB ascii)
 - > ~ 3.000 files
 - > ~ 350.000 lines (code + documentation)
- Sub(sub) Projects
 - > hedeby-config-center, hedeby-n1sm-adapter, hedeby-ge-adapter, hedeby-security, ...
- Documentation to be published on new Wiki

Sub-Project Testsuite


- 24 Members
- Mailing Lists
 - > dev@testsuite.sunsource.net
 - > cvs@testsuite.sunsource.net
 - > issues@testsuite.sunsource.net
- Size of content / code
 - ~ 13MB (0.5MB binary, 12.5MB ascii)
 - ~ 1.700 files
 - ~ 300.000 lines of code
- Documentation not yet public

Other Grid Engine related Projects

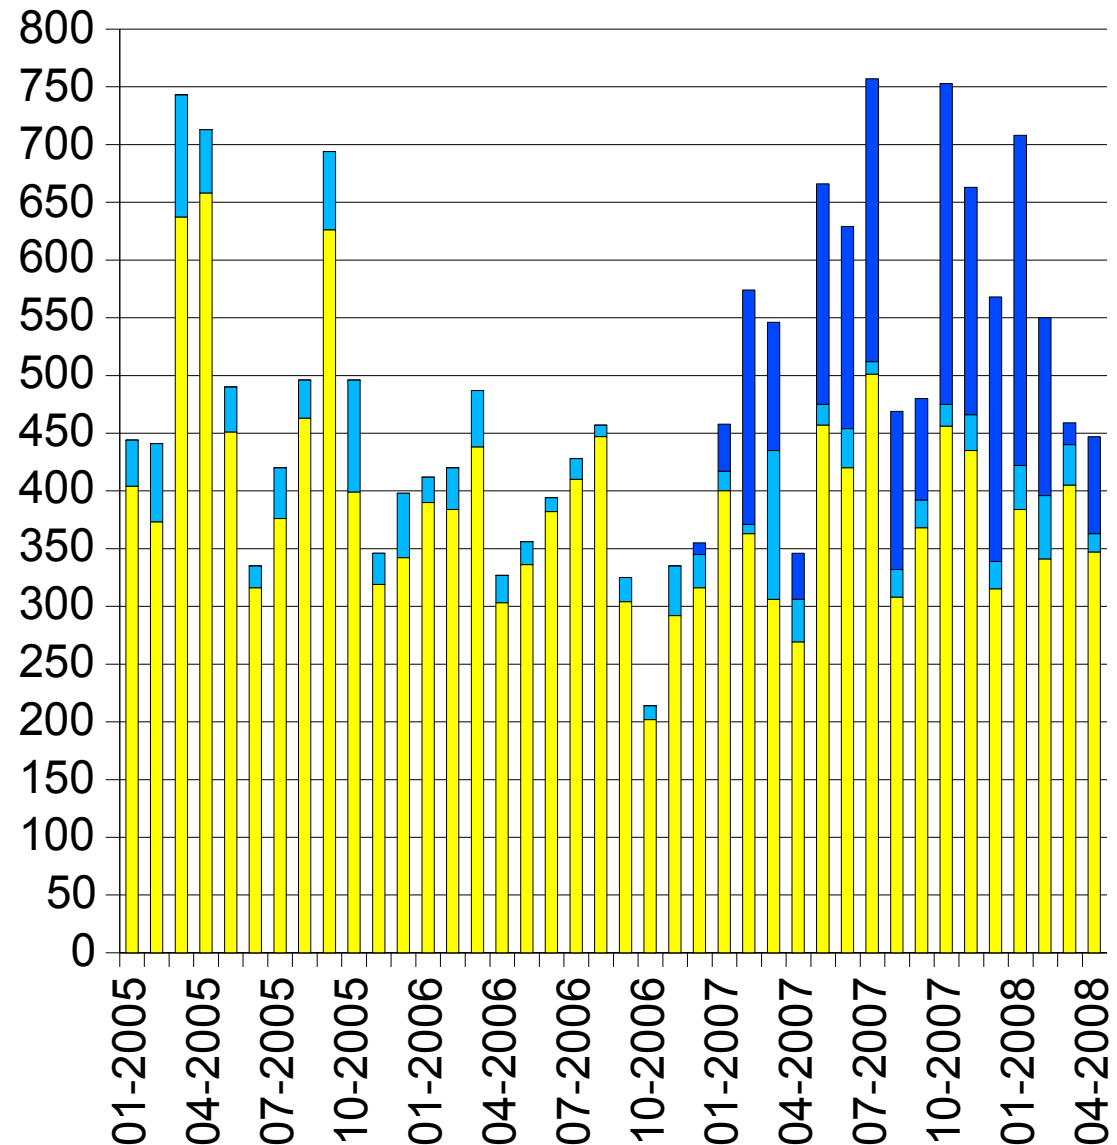
- GRID.org
 - > Open Grid stack (UniCluster Express) including Grid Engine
- drmaa4ruby.sunsource.net
 - > DRMAA binding for Ruby
- Passau Grid Engine XML Parser
 - > <http://passau.dev.java.net>
 - > Java API for accessing qstat output via JAX-B
- OpenDSP
 - > <http://sourceforge.net/projects/opendsp>
 - > Remote job management via secure DRMAA/JSDL framework
- ComputeServer
 - > <http://computeserver.dev.java.net>
 - > Framework for distributed Java applications on Grid Engine

Mailing List Traffic

posts/month




- avg pretty constant
- high traffic on users mailing list
- hedeby development discussions in opensource
- core system development discussions mostly internal



CVS Statistics

- core
 - > including Testsuite until Sep 2006
- arco (> Dec 2006)
- testsuite (> Sep 2006)
- hedeby (> Dec 2006)
- private (before OS)
 - > arco
 - > hedeby
 - > windows port

cvs checkins / month




Community Contributions

- mailing lists, lots of support
- gridengine.info (Chris Dagdigan)
- Bug reports, partly with full evaluation and patches
- development projects
 - > ssh Tight Integration (Ron Chen)
 - > FlexLM Loadsensor Integration (Mark Olesen)
 - > Array Job Dependencies (Rising Sun Pictures)
 - > others

Issues / Change Requests (May 2008)


- 2 sources of information
 - > Issuezilla (gridengine.sunsource.net)
 - > Bugster (former Bugtraq, SUN tool, more fine grained info)
- Some numbers:
 - > Total (open) 2542 (587)
 - > Bugs (open) 2197 (378)
 - > RFE's (open) 245 (209)

CR's Fixed / Introduced per Release



Open Bugs (May 2008)

Open Bugs by subcategory and priority





GRID ENGINE WORKSHOP

Introduction

Fritz Ferstl

ferstl@sun.com