

Grid Computing With FreeBSD

USENIX ATC '04: UseBSD SIG
Boston, MA, June 29th 2004

Brooks Davis, Craig Lee
The Aerospace Corporation
El Segundo, CA
{brooks,lee}@aero.org

<http://people.freebsd.org/~brooks/papers/usebsd2004/>

Outline

- What is Grid Computing?
- Examples of Grid Computing
- The Globus Toolkit
- Porting the Globus Toolkit

Grids are **HOT**

- Oracle 10^g
- Sun Grid Engine
- Apple's Xgrid
- Other Grid supporters:
 - HP
 - IBM
 - Microsoft
 - Platform Computing

What are Grids?

- Many different concept
 - Globus Toolkit based computation or data grids
 - Sun Grid Engine grids
 - P2P networks
- This talk focuses on HPC Grids

What is an HPC Grid?

- An analogy with the power grid
 - Transparent access to resources
- Allows users to form virtual organizations (VOs) among people from different organizations
- VOs share resources to solve problems
 - Clusters, Supercomputers, Sensors, Instruments, Data Collections, etc

Grid Computing Software

- The Globus Toolkit
- Condor-G
- BOINC – Berkeley Open Infrastructure for Network Computing
- Sun Grid Engine

Global Grid Forum (GGF)

- Standards body modeled after the IETF.
- Meets three times a year around the world.
- Standards
 - Open Grid Services Infrastructure (OGSI)
 - Web Services Resource Framework (WSRF) in conjunction with OASIS
 - Family of composable specifications

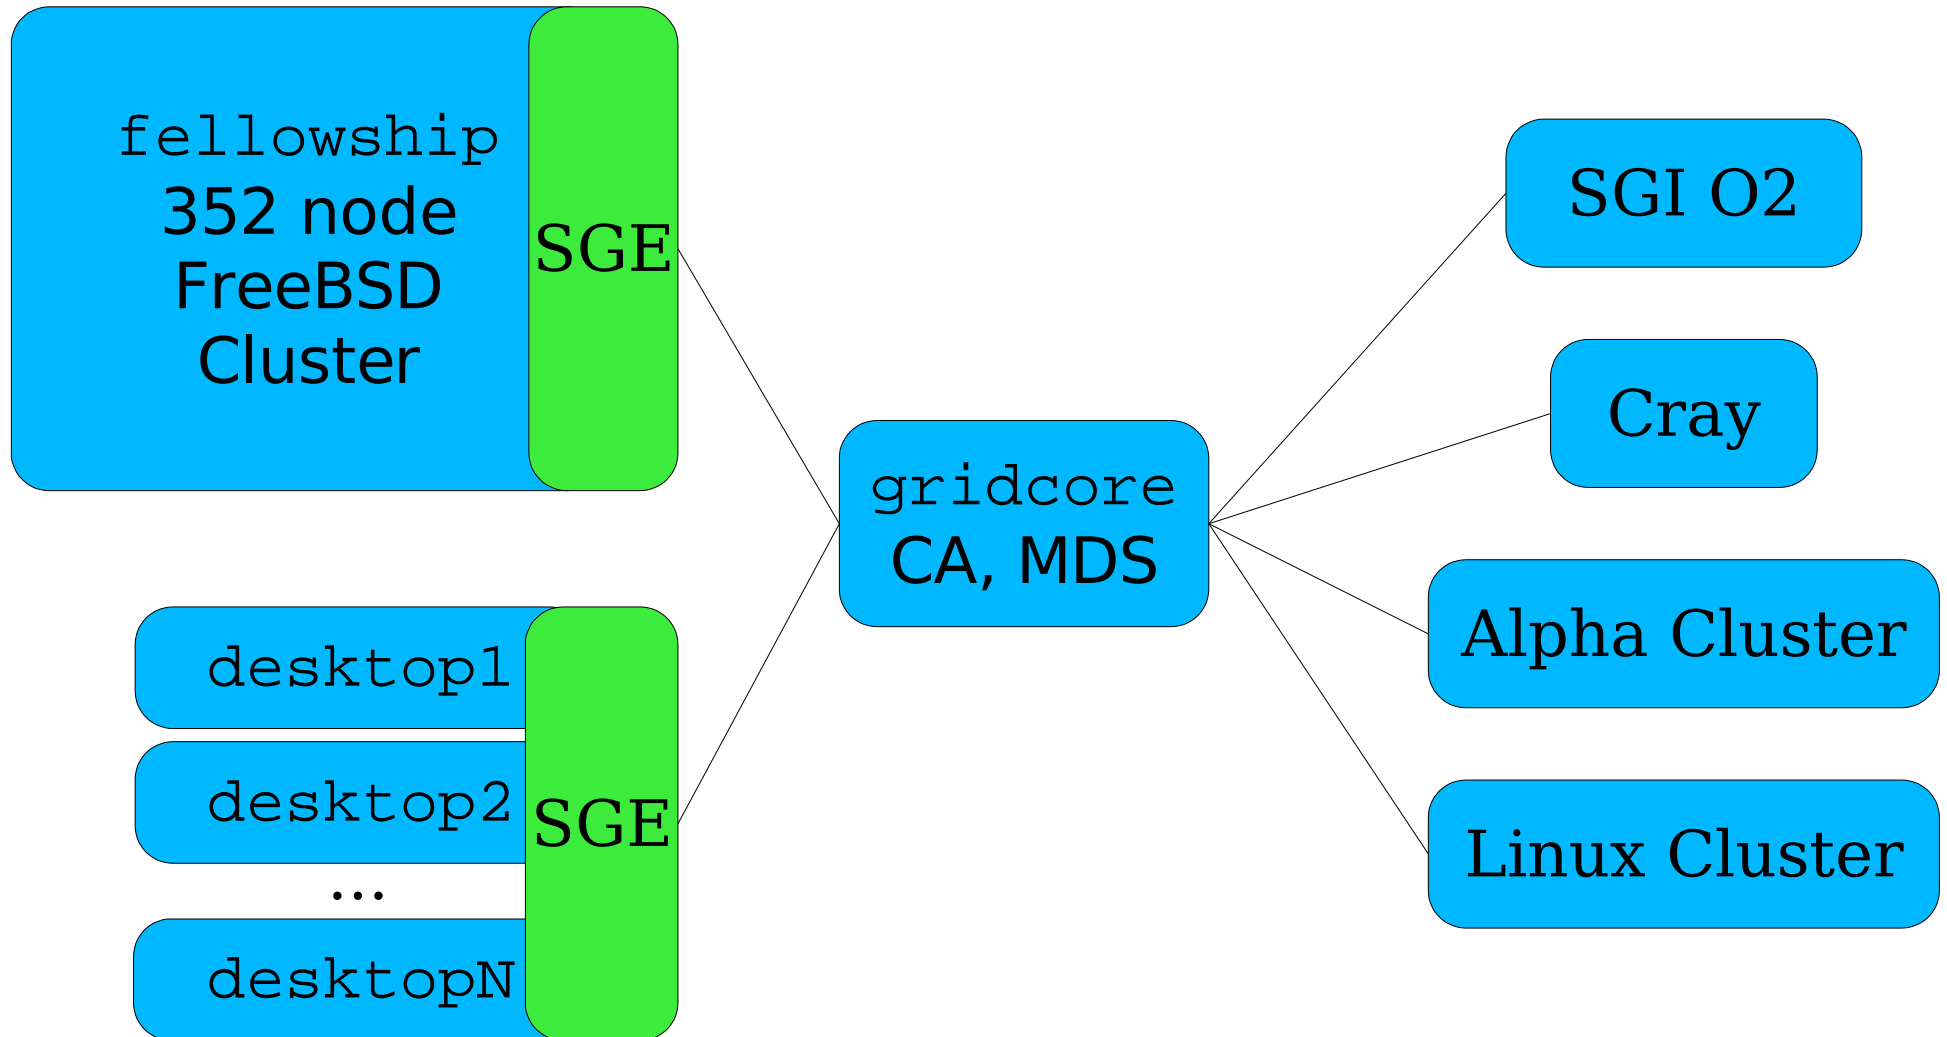
Examples of Grids

- TeraGrid – US based research grid
 - >20TFLOPS and >1petabyte
- PPDG – Particle Physics Data Grid
- EGEE – Enabling Grids for E-science in Europe
- FreeBSD Package Grid
- AeroGrid

AeroGrid

- Internal grid project at Aerospace Corp.
- Testbed moving towards production
- Up and running
 - Certificate authority
 - Monitoring and Discovery Service
 - A few small servers
 - Primitive access to Fellowship cluster

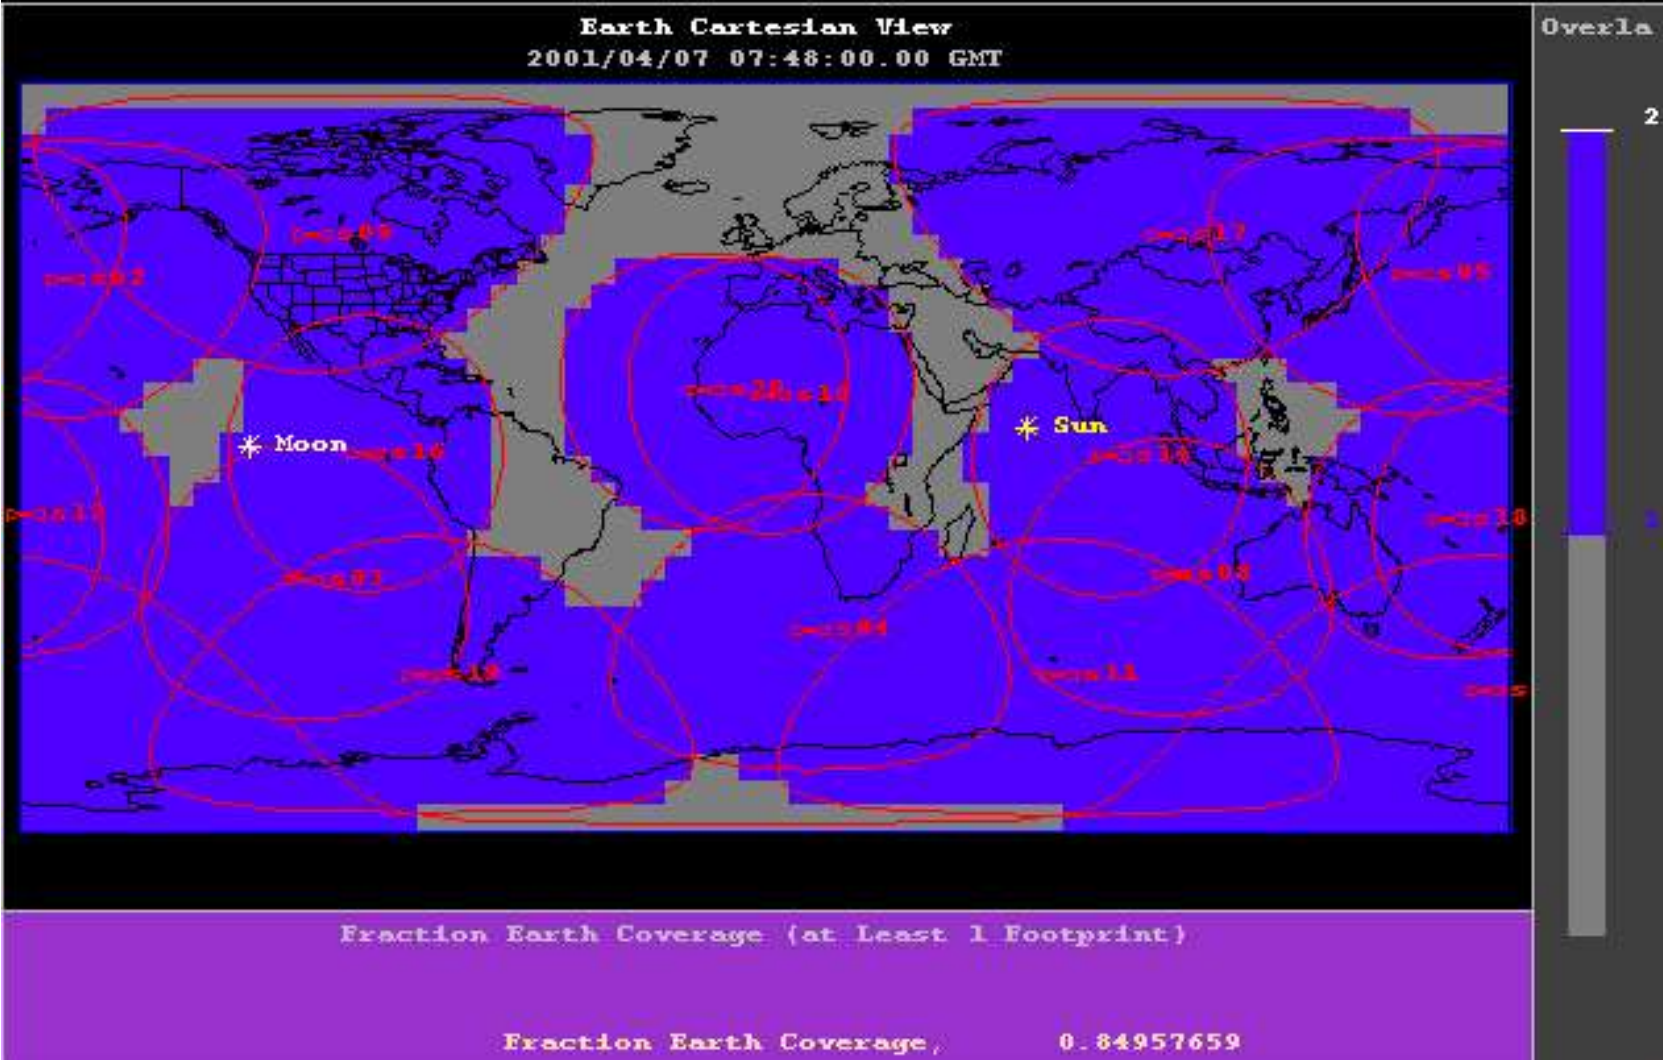
AeroGrid in the Future



AeroGrid in the Future

- Multiple clusters and other computing resources such as SGI and Cray systems
- Portals for easy job submission
- Grid backends to desktop apps
 - SOAP – Satellite Orbit Analysis Program

SOAP



The Globus Toolkit

- The de facto standard for grid computing infrastructure
- Set of orthogonal services
 - Single sign-on, encryption, file transfer, job submission, resource location
- Developed by the Globus Alliance (formerly the Globus Project)

Globus Toolkit Versions

- GT1 – All C, obsolete
- GT2 – Mostly C (some Java interfaces), no longer actively developed, widely deployed
- GT3 – Combined C and Java based services (Implements OGSI and classic (GT 2.4) services), widely deployed
- GT4 – Web Services (WSRF) and C, under development (3.9.1 available)

Globus Platform Support

- GT2
 - Supported on AIX, HP-UX, Irix, Linux, and Solaris
 - Expected to run on most POSIX-like OSes
 - Over 60 packages
- GT3
 - All GT2 components
 - Java based Web Services

Globus Services: GSI

- Grid Security Infrastructure
- Single sign-on access to resources across multiple independent administrative domains
- Per-host (and per-service) mapping of X.509 distinguished names and user accounts
- Credential delegation via proxy certificates

Globus Services: GRAM

- Globus Resource Allocation Manager
- Interfaces between users or meta-schedulers and local resources managers
- Resource managers include
 - Sun Grid Engine (SGE)
 - The Portable Batch System (PBS)
 - Condor
 - `fork()`

Globus Services: GridFTP

- FTP protocol extensions
 - GSI security on control and data channels
 - Parallel transfers
 - Partial transfers
 - Third-party transfers

Globus Services: MDS

- Monitoring and Discovery Service
- Based on LDAP
- Two parts
 - Grid Resource Information Service (GRIS): provides information about a resource
 - Grid Index Information Service (GIIS): aggregates data from GRISs for search
- Separate GRISs are removed in GT3 as all OGSI services act as their own GRIS

Globus Services

- Replica Location Service
- OGSi-compliant database service
- GSI-OpenSSH

Grid Packaging Tools (GPT)

- Supports patching and building from source like ports/pkgsrc plus binary package creation and management
- Flavors indicate compiler, 32 vs. 64-bits, debugging, and threading
- Multiple flavors of a package may be installed in a single installation
- No automatic dependency installation
 - Sets of packages shipped as bundles

The FreeBSD Ports Collection

- Collection of Makefiles and patches to build third-party software
- Easy mechanism for installing software
 - `cd <portdir>; make install clean`
- Ports are used to build packages
 - Assuming the license allows this
- Over 11,000 ports producing over 10,000 packages

Globus Ports

- `misc/gpt` – The Grid Packaging Tools
- `misc/gpt31` – Needed for GT2
- `misc/globus2` – Globus Toolkit 2.4.3
 - GRAM, GridFTP, SimpleCA (no port) working
 - MDS requires additional patches to scripts

Porting GPT

- GPT installs standard Perl modules in non-standard locations
 - Port depends on ports of these modules
- GPT installs a modified `Archive::Tar` module
 - Modifications do not appear necessary and the port version is used instead
- Build and install are a single step

Porting The Globus Toolkit 2.4

- Globus distributed as package bundles
- Package bundles conflict
 - These conflicts also represent distfile bloat
- Current port installs all bundles (except replica manager)
 - Can't patch this way
- Investigating building from individual packages

Issues With Separate Ports

- Multiple flavors of each package need to be installed
 - `globus_core` needs to be thread and non-threaded
- Some files conflict between flavors
 - Primary version of executables from only one flavor
 - Different bundles will choose different flavors so the choice can not be hard coded!

Porting The Globus Toolkit 3/4

- Not done yet
- Builds on GT2 port
- Adds new dependencies
 - Java
 - Apache Ant
 - Junit
 - Tomcat
- Even more bundled

Improving Globus on FreeBSD

- Improve GT2 port(s)
- Add GT3 and GT4 ports
- FreeBSD specific enhancements
 - `sendfile()` in GridFTP?

Enhancing FreeBSD for Globus

- NSS support for grid map files
 - Grid map file maintenance is a pain and the NSS infrastructure is powerful and extendable
- Continued research in high-performance I/O (network and disk)

Conclusions

- Grids are an important part of the future computing landscape
- FreeBSD is a viable platform for Grid Computing
 - The Globus Toolkit 2 runs today
 - The Globus Toolkit 3 and 4 should be portable, but are not yet ported
- Help wanted with ports

High Performance BSD BOF at SC'04

- What: HPC BSD BOF
- When: SC2004, November 6-12
- Where: Pittsburgh, PA
- Why: Support HPC BSD!

Grid Computing With FreeBSD

USENIX ATC '04: UseBSD SIG
Boston, MA, June 29th 2004

Brooks Davis, Craig Lee
The Aerospace Corporation
El Segundo, CA
{brooks,lee}@aero.org

<http://people.freebsd.org/~brooks/papers/usebsd2004/>

Background Slides

What is Grid?

- *A computational grid is a hardware and software infrastructure that provides dependable, consistent, pervasive, and inexpensive access to high-end computational capabilities. —Ian Foster & Carl Kesselman 1999*

What is a Grid?

- *A Grid is a hardware and software infrastructure that provides dependable, consistent, and pervasive access to resources to enable sharing of computational resources, utility computing, autonomic computing, collaboration among virtual organizations, and distributed data processing, among others. —Wolfgang Gentzsch 2003*

Grid Computing With FreeBSD

USENIX ATC '04: UseBSD SIG
Boston, MA, June 29th 2004

Brooks Davis, Craig Lee
The Aerospace Corporation
El Segundo, CA
{brooks,lee}@aero.org