Building Blocks for the Grid

Chris Smith, VP of Standards, OGF

NCOIC
Dec 2007
St. Petersberg, FL

Overview

- A bit about OGF
- Grid Building Blocks
- Standardization
- A Future for Grid Standards

History

http://www.ggf.org

- Birthed in high performance computing community in 1998
- Merged with European & Asian grid efforts in 2001
- 52 groups and 81 documents

http://www.gridalliance.org

- Birthed in enterprise data center community in 2004
- Raised awareness of grids in “end-user” organizations
- 5 groups and several important documents

Merger completed, June ’06; OGF Launched September of ‘06

Open Grid Forum

Events & Forums	Community Practice	Industry Standards
Bring communities together to share, innovate, workshop and outreach	*Leverage expertise & experience of the community to enable successful building and operating of grids*	*Align with/influence other SDOs and/or develop specifications that lead to interoperable software standards*

Technical Strategy & Roadmap

- Defines the overall OGF technical strategy for the development of standards
- Three-year timeframe from 2006 to 2010
- Describes the output of the OGF standards working groups as well as the requirements that drive them.
- Technical Strategy Document
  - Focus Areas
  - Goals
  - Alignment Process
  - Use Cases
  - Roadmap

http://www.ogf.org/documents/GFD.113.pdf

Broad and Narrow Grids

- The concept allows for clearer communication encompassing all aspects and points of view on Grids.

- A **Broad Grid** is any collection of services
- A **Narrow Grid** is defined by both the technologies used, as well as the application focus. Some examples:
  - OGSA compliant Enterprise compute cluster
  - Web 2.0 Sensor Net linked to Google maps

Narrow Grids of Focus

- Collaboration Grids
  - Multiple institutions, secure, widely distributed, VOs
  - Collaborative agreements & commercial partnerships
  - Financial Model: Increase overall revenue
- Data Centre Grids
  - Centralized management of multiple platforms
  - Aggregation of enterprise resources and applications
  - Financial Model: Reduce Total Cost of Ownership (TCO)
- Cluster Grids
  - Networks of Workstations, Blades, etc.
  - Cycle scavenging, Homogeneous workload
  - Financial Model: Lower marginal costs

Grid Building Blocks

- Grids are made up of many services, aggregated into narrow grids of application focus
  - applications are “composite”; made up of many of these grid services
- Allows developers to collect the “best of breed” of technologies for addressing their application requirements
- Blocks fit much better together when the interfaces are well defined

One Size Does Not Fit All

- Depending on your role in an organization, you might have different requirements for interacting with the Grid
  - resource consumers are interested in how to access the Grid
  - resource providers are interested in managing the Grid infrastructure and meeting SLAs
- In terms of broad and narrow Grids, a particular physical infrastructure could be represented as different narrow Grids, depending on your role!

Grid Resource Consumers

- Focused on how to use the Grid
  - Represent the demand side of the supply and demand equation
  - APIs and Protocols for accessing services
  - Describing the application requirements
  - Interested in establishing SLAs
  - Concerned with value for money
- Interested in
  - Managing their workload
  - Application performance
  - Data access via caching, transfer, databases
  - Application portability (run my work in a local grid, or in a computing utility)
- Aligned with work going on in OGF

Grid Resource Providers

- Focused on how to provide Grid services
  - Represent the supply side of the supply and demand equation
  - Monitoring the resources in the Grid
  - Interested in managing and fulfilling SLAs
  - Want to effectively utilize the resources in the Grid
  - Isolating the users of the Grid from each other
- Interested in
  - Managing the infrastructure
  - Overall system performance
  - “Real Time Infrastructure”
- Aligned with work going on in DMTF

Platform Solution Stack: Consumers vs Providers

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Image © 2007 Platform Computing

Platform EGO Building Blocks

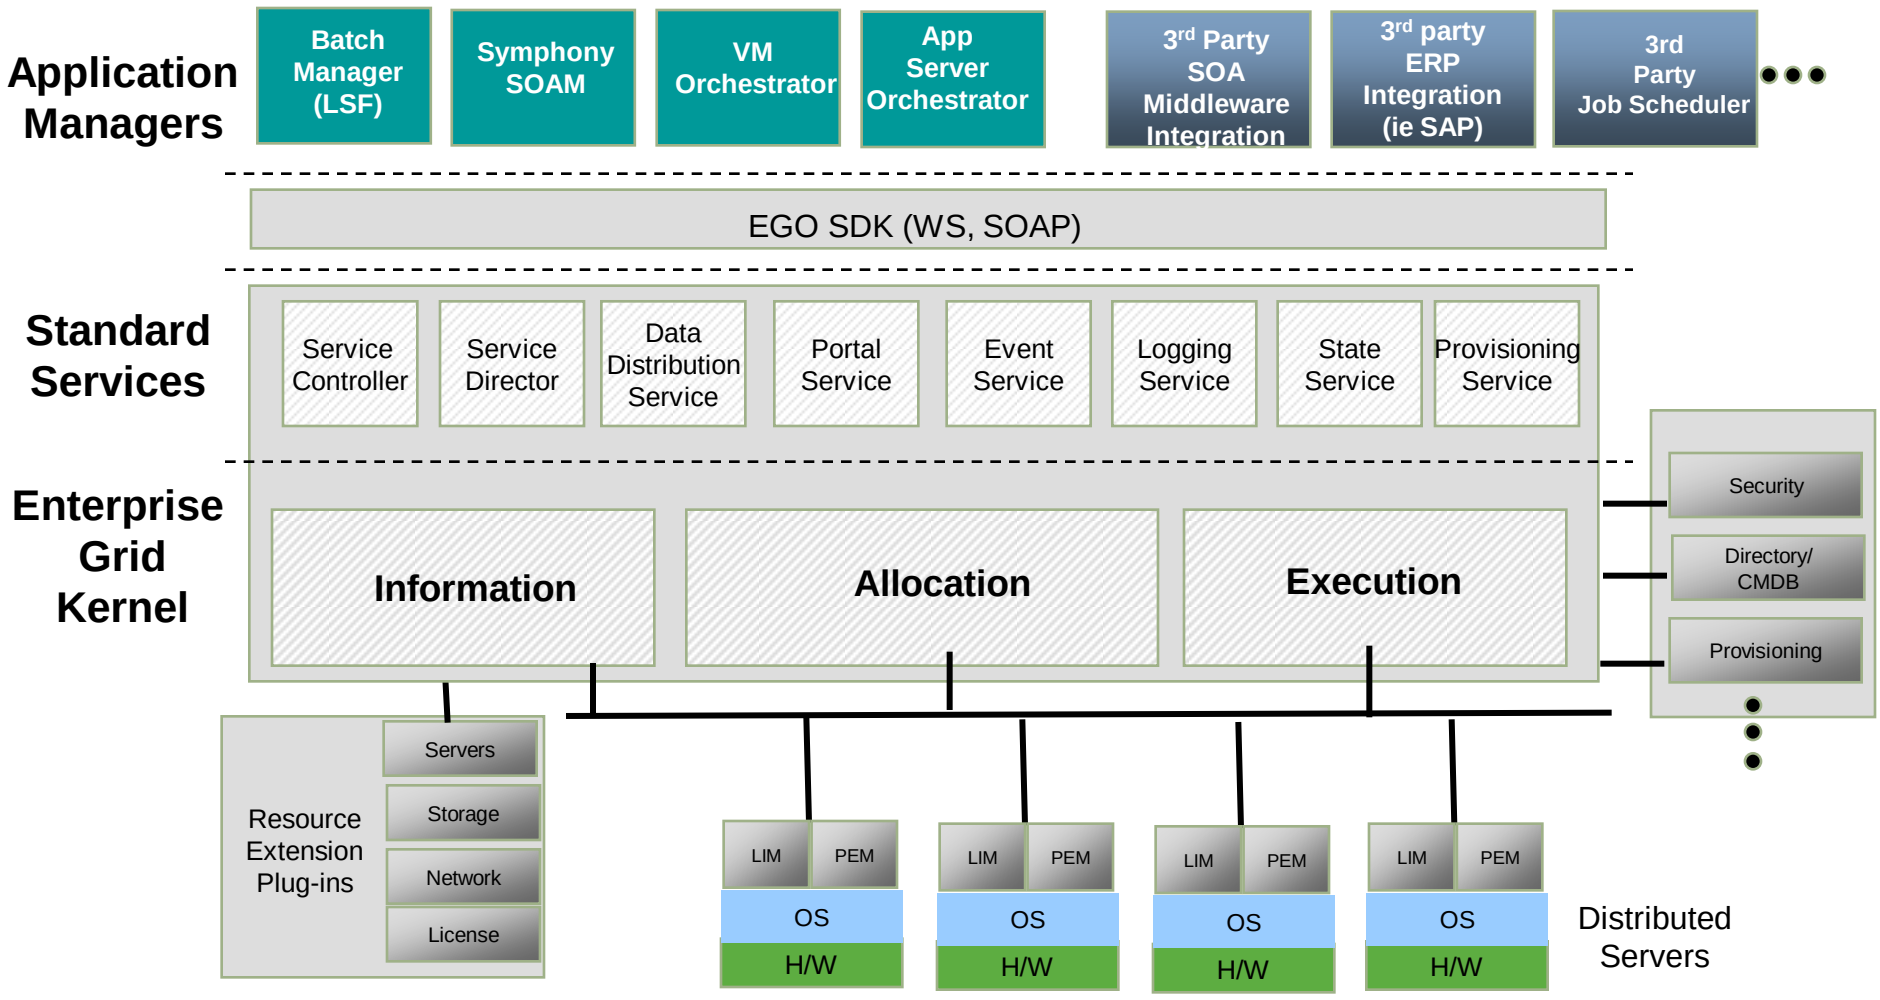

Image © 2007 Platform Computing

Why Standardize?

- What is the value of implementing standards?
- For vendors
  - meet customer demand for interoperability
- For developers
  - leverage the expertise of other developers
  - offer a choice of tools and platforms in order to speed implementations
  - only need to support one integration interface
- For end-users
  - reduce the costs and risks of adopting grid technology
  - get insight into the best practices of the industry at large

Why Not Standardize?

- Standards are not always appropriate
- A technology might be “too new”
  - you stifle innovation with standardization, which focuses on commonality
- A technology might be very niched
  - defacto standards will emerge in this case

Standard APIs vs Protocols

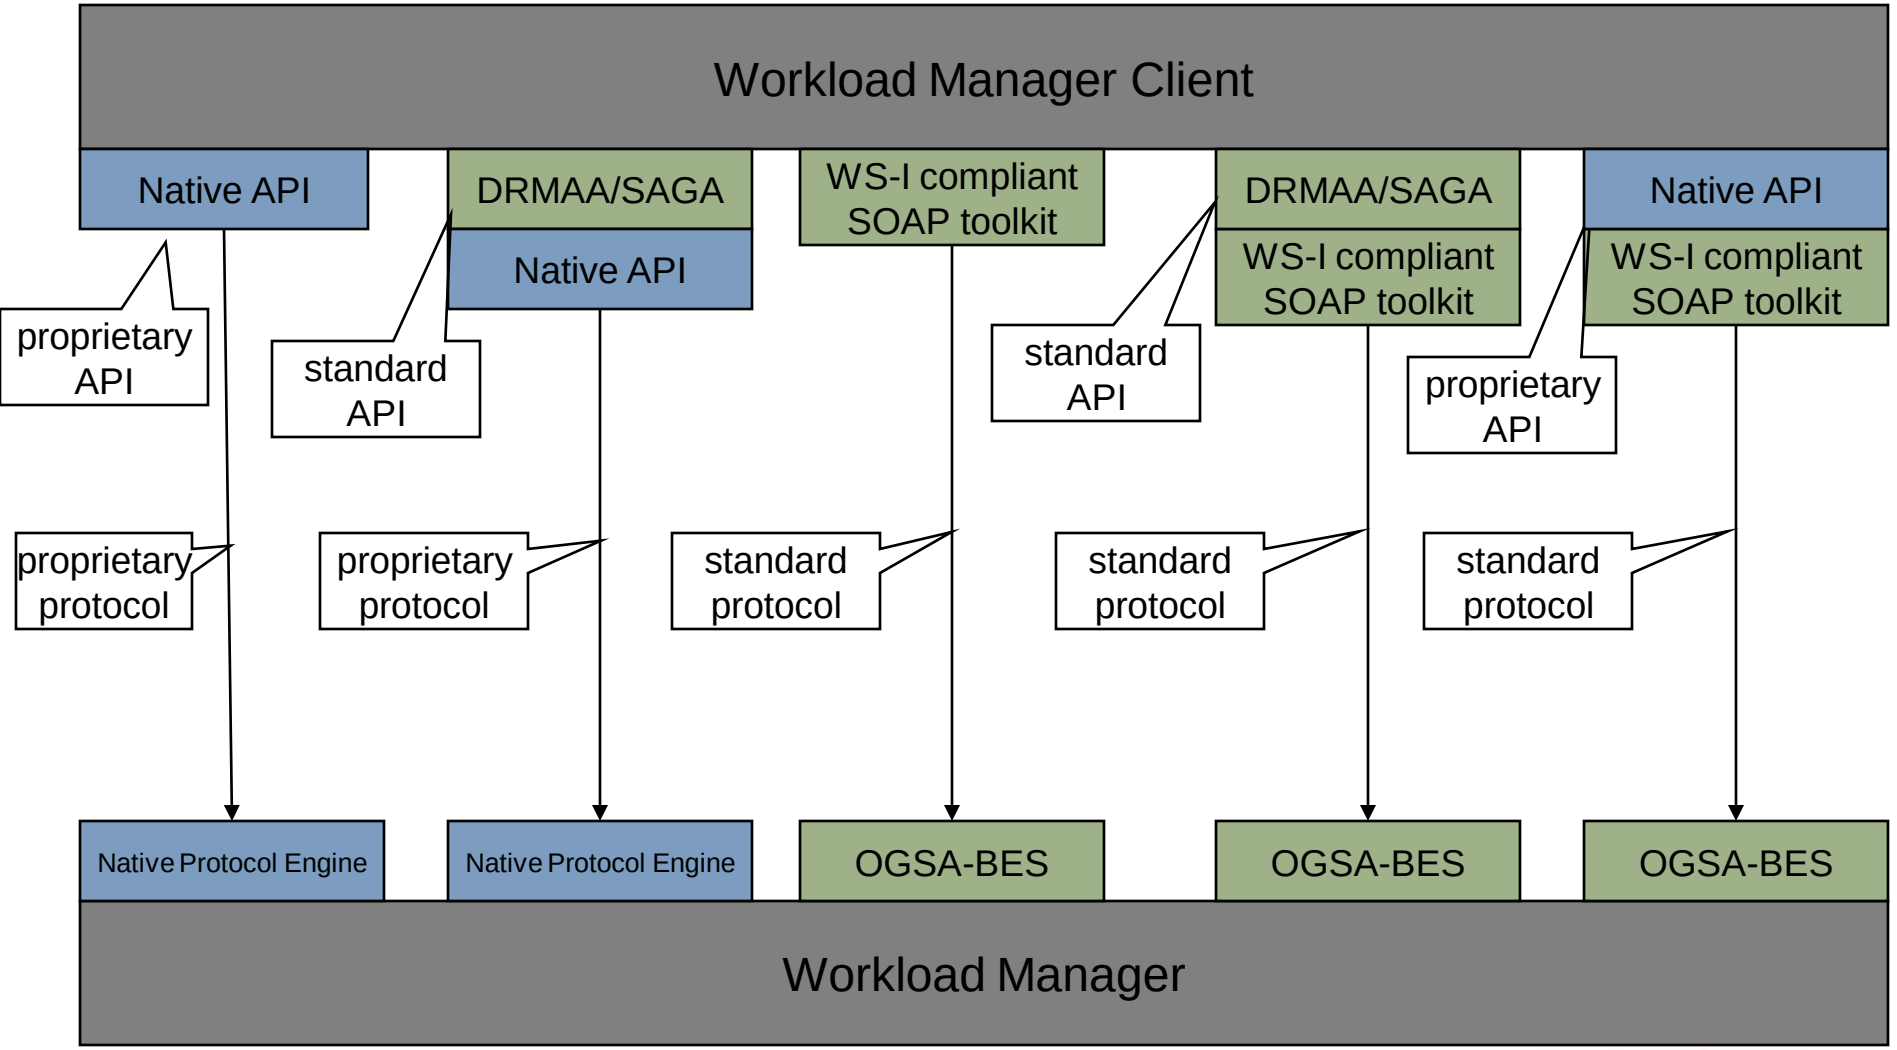

Enabling Technology for Consumers (1)

Workload Management

- Schedulers/brokers are typically used to mediate consumer access to resources
- Consumers submit descriptions of work using interfaces exposed by the resource provider
- API specs include DRMAA and SAGA
- Protocol specs include OGSA-BES and the HPC Basic Profile
- Work is described using JSDL
- Application environments described using CDL
- Resources are exposed using GLUE

Enabling Technology for Consumers (2)

Data Access

- Need to allow applications to access data sets for processing
- These might be local to the computing resources, or moved in from another location
  - transfer protocols, distributed caches and replica catalogs are all available technologies
- API specs include SAGA
- Protocols include Byte-IO, DAIS and GridFTP
- Resources exposed using GLUE

Enabling Technology for Providers (1)

- Virtualization
  - Provides a means for isolating user applications
  - Used to manage application requirements (for specific OSes, machine types, etc, etc)
  - OVF is a key new specification for describing virtual environments
    - composed of VM images and meta-data about the virtual appliance, including meta-data for checking integrity, etc
    - can describe collections of VMs
    - any virtualization middleware supporting the standard should be able to run an image described by OVF
    - http://www.dmtf.org/events/idf/OVF_Press_Release.pdf

Enabling Technology for Providers (2)

- Systems Management
  - Use of the CIM and various profiles for modeling the Grid
  - Use protocols such as WBEM or WS-Management to manage Grid components
  - The new CIM System Virtualization Model from DMTF instruments virtual environments for management
    - composed of a number of new profiles that describe how pools of resources are aggregated into virtual machines
    - http://www.dmtf.org/standards/published_documents/DSP2013_1.0.0.pdf

Information Models

- Consumers and providers needs to be able to communicate requirements and describe infrastructure to each other
- Standardized information models are the key to enabling this communication
- Allows consumers to describe their requirements
- Allows providers to describe the infrastructure available to a consumer

OGF & DMTF Alliance

- Intention of the alliance is to cross-leverage each organization's domain strength
  - DMTF focus on distributed management infrastructure and information/data modeling of resources
  - OGF focus on the deployment, operation and access to resources in a Grid
- OGF has been contributing CIM schema extensions based on Grid requirements
  - CIM Job extensions
  - GLUE information model to describing Grid environments to consumers

Example: Building Blocks of OGSA

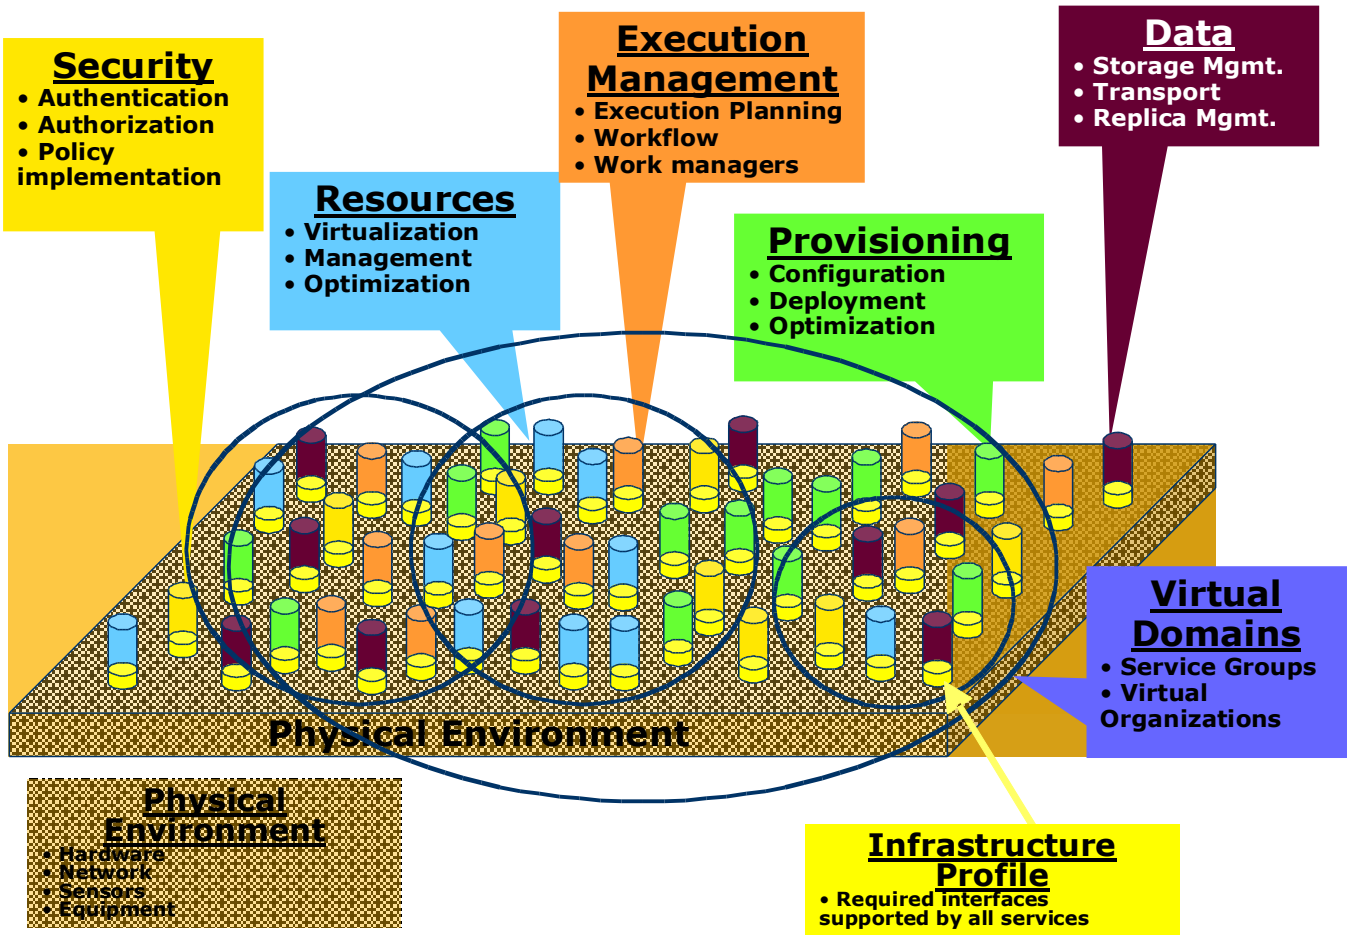

A Future for Grid Standards

- What might be the interesting areas of standardization in the next decade?
  - Programming Models
  - Application Environments
  - Affinity of Data and Computing Environments
  - Application, Data and Resource Security and Privacy

Programming Models

- Applications will be enabled to adapt themselves to the actual execution resource automatically or by modified configuration
- The cost of “porting” to new Grid APIs/models is very expensive
- Programming models will have the concepts of parallel execution and/or data built-in
  - the runtime will be abstracted, so the units of execution/data can run on single CPI, multi-core, or in a distributed environment

Application Environments

- Grids and other execution environments will be enabled to adapt themselves to application specific requirements communicated by standardized descriptions of application environments
- The execution environment is modified to suit the needs of the application
  - virtualization is one enabler of this approach
  - this will be typical of computing utilities

Data/Compute Affinity

- Standards for describing an application's Data requirements in order to express data/compute affinity, which can be communicated between Storage elements, Computing elements and Applications in order to achieve the desired balance of performance and costs

Application Security and Privacy

- Standardized description of a resource providers policies with respect to the risk associated with running a misbehaving application, and the risk associated with divulging an application's private data
  - providers want to make sure applications don't cause security breaches
  - consumers want to ensure that their data and activities stay confidential and private
  - Consumers can "pay more" for risky applications
  - Providers can "charge more" for more secure/private environments

Questions?

Alphabet Soup

- OGSA - Open Grid Services Architecture
- DRMAA - Distributed Resource Management Application API
- SAGA - Simple API for Grid Applications
- BES - Basic Execution Service
- CDL - Configuration Description Language
- DMTF - Distributed Management Task Force
- OVF - Open Virtual Machine Format
- CIM - Common Information Model

DRMAA

- Defines a generalized API for interacting with the job management capabilities of distributed resource management systems (DRMs) in order to facilitate integration of application programs.
  - currently there are standardized bindings for C and Java
  - non-standard bindings exist for Perl, Python and Ruby
- It's the first OGF specification to become a "full recommendation"
  - this means that it has at least 2 interoperable implementations available (it actually has more)
  - http://www.ogf.org/documents/GFD.22.pdf
- See http://drmaa.org for in depth information

JSDL

- JSDL stands for Job Submission Description Language
  - A language for describing the requirements of computational jobs for submission to Grids and other systems.
- A JSDL document describes the job requirements
  - What to do, not how to do it
- JSDL does not define a submission interface or what the results of a submission look like
  - Or how resources are selected, or …
    - The JSDL-WG is now considering its next steps.
- JSDL 1.0 is published as GFD-R-P.56
  - Includes description of JSDL elements and XML Schema
  - http://www.gridforum.org/documents/GFD.56.pdf

OGSA-BES

- The Basic Execution Service defines a Web Services interface for creating, monitoring and controlling things such as UNIX/Windows processes, service instances or parallel programs - what are called “Activities” in BES parlance
- BES defines an extensible state model for modeling the lifecycle of Activities
  - different implementations may support different sets of states and allowable transitions
- BES defines an extensible information model for the BES itself and for the Activities it manages
  - different implementations can provide extra properties on the BES container or Activities using XML extensibility
- http://www.ogf.org/documents/GFD.108.pdf

HPC Profile

- What is it?
  - A draft OGF interoperability standard for batch scheduler job submission and management
  - Defines how to combine other specifications (JSDL, OGSA-BES) to implement a particular use case scenario
- What is its value?
  - Enables the integration of HPC applications with resource managers, using Web Services
  - The HPC Profile purposefully reduces the scope of JSDL and BES in order to ensure base line interoperability by clarifying ambiguities
- Companies and organizations involved in developing the standard include
  - Platform Computing, Microsoft Corporation, University of Virginia, The Globus Alliance, Fujitsu, and many others
- http://www.ogf.org/documents/GFD.114.pdf

OGF Myths Busted (1)

Customer asks: “Do you implement the Globus standard?”

- Globus is not a standard. Globus is a package of Grid middleware. It has had fairly wide adoption in academia, and enjoys a “brand association” with Grid computing, but it is not a standard.
- Members of the Globus Alliance and some Univa UD employees participate in some of the OGF working groups where standards are defined. They contribute the same way that Platform Computing (among others) contributes, but their software interfaces are not accepted as “the standard”.

OGF Myths Busted (2)

Vendor says: “Our software is OGSA compliant”

- OGSA (Open Grid Service Architecture) is a working group at the OGF whose purpose is to define an architecture for Grid using Web services.
- They have published a number of documents, but the Open Grid Services Architecture is informational (i.e. not a specification).
- Because there is no OGSA specification to be compliant with, being OGSA compliant means nothing at this point in time.
  - although you can be compliant with some OGSA specifications, such as the OGSA Basic Security Profile
- Often people confuse OGF specifications with OGSA, due to the highlighting of OGSA activities in the past.