



An Introduction to the Open Grid Services Architecture

Robert Piotter,
piotter@hlrs.de
Distributed System Group
High Performance Computing Centre
Stuttgart (HLRS)
<http://www.hlrs.de/organization/ds/>

Presentation Outline



- Open
 - Organisations and Standards
- Grid
 - Definitions and Virtual Organisations
- Services
 - Services and Resources
- Architecture
 - Bringing it all together

In Short



- The “Open” in OGSA
 - The process by which the architecture is defined is open to all.
 - Define and build on open standards.
- The “Grid” in OGSA
 - From requirements that are derived from use cases to a proposed standard for Grids.
- The “Services” in OGSA
 - Resource virtualization through Grid Services.
- The “Architecture” of OGSA
 - Service Oriented Architecture (SOA) built on Web services

Open...



- Open
 - Organisations and Standards
- Grid
 - Definitions and Virtual Organisations
- Services
 - Services and Resources
- Architecture
 - Bringing it all together

What and Who is OGSA?



- The OGSA is an effort to document and define standards, best practices and implementation guidelines for Grid technologies
- OGSA is the GGF's flagship architecture and the blueprint for industry standard grid computing (Dave Berry, Research Manager, NeSC, European Grid Technology Days 2005)
- OGSA Contributors:
 - Industry: Fujitsu, IBM, HP, NEC, Hitachi, Platform, Univa, ...
 - Academia: UK e-Science, CERN, Argonne, ISI, ...

OGSA History



- Announced at GGF4 (2002/2)
- OGSA Working Group created (2002/9)
- OGSA Usecase document released at GGF10 (2004/3)
- OGSA was declared as GGF's flagship architecture at GGF10 (2004/3)
- OGSA Version 1.0 released at GGF11 (2004/6, update 2005/1)
- Current working version is 1.5

GGF – The Global Grid Forum



- “The Global Grid Forum (GGF) is the community of users, developers, and vendors leading the global standardization effort for grid computing.”
- Founded in 1998 at the Supercomputing conference
- Merged with the European Grid Forum and the Asia-Pacific Grid Forum in 2000
- Among the members are:
 - U.S. Department of Energy, NASA, UK e-Science Programme/DTI
 - National Institute of Advanced Industrial Science and Technology (AIST), Japan
 - IBM, Intel, Microsoft, Sun, hp, ...

The Globus Alliance



- “The Globus Alliance is an international collaboration that conducts research and development to create fundamental Grid technologies.”
- Members include:
 - Argonne National Laboratory (University of Chicago)
 - EPCC (University of Edinburgh)
 - The National Center of Supercomputing Applications NCSA, ...
- Develops the open source software Globus Toolkit
- Now available in Version 4.0

OASIS



- "OASIS is a non-for-profit, global consortium that drives the development, convergence and adoption of e-business standards."
- Members of OASIS are
 - Vendors, users, academics and governments
 - Organisations, individuals and industry groups
- Best known for e-business standards that address real world business requirements, including
 - UDDI, SAML, ebXML, WS-Security, WSRP, WS-Reliability, ...
- Host of key Grid standards including:
 - WSDM, WSRF, WS-N

OGSA - Competitors?



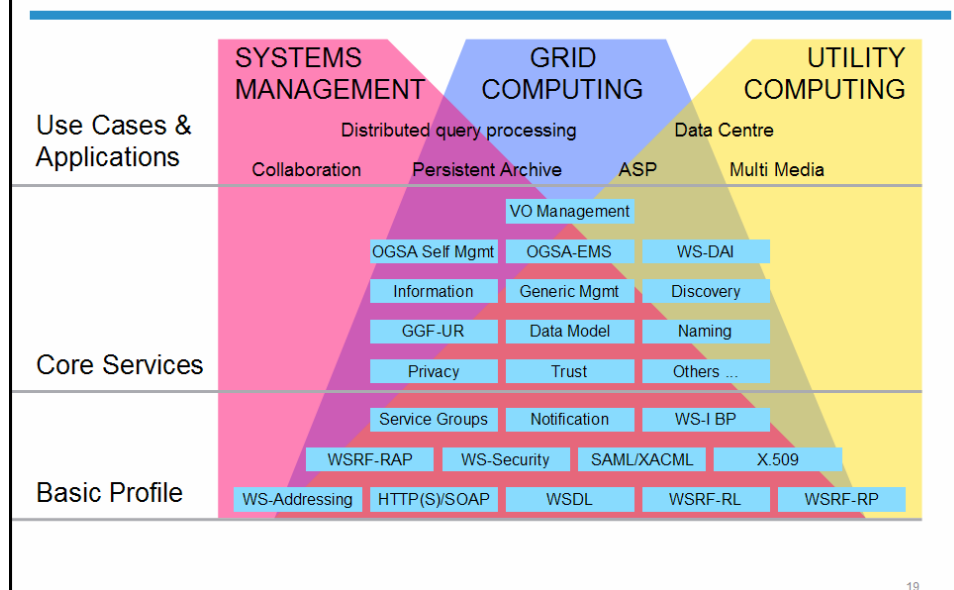
- The Enterprise Grid Alliance
 - "The Enterprise Grid Alliance is a consortium formed to develop enterprise grid solutions and accelerate the deployment of grid computing in enterprises."
 - "EGA has a pragmatic approach. Our goal is interoperable solutions for grid computing that meet enterprise needs."
 - www.gridalliance.org
- Members include:
 - Oracle, Hewlett-Packard, NEC, Sun, Fujitsu-Siemens, ...

OGSA and Standards



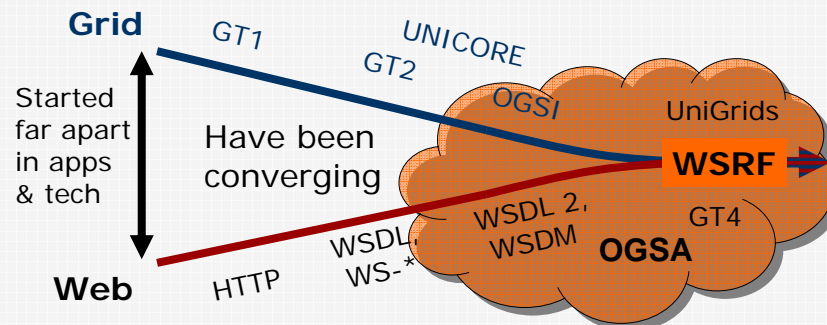
- Open standards:
 - Encourage interoperability
- What Standards?
 - Web Services provide a common infrastructure with industry support in a distributed and heterogeneous environment. (Dave Berry, European Grid Technology Days 2005)
- Normative reference specifications:
 - WS-I Basic profile 1.1 & Basic security profile 1.0
 - WS-Addressing
 - WS Resource Framework & WS Notification
 - WS-Security

OGSA Specifications Landscape



19

Convergence of Grid and Web Services



- The definition of WSRF means that Grid and Web communities can move forward on a common base
- Support from major WS vendors

Grid...



- Open
 - Organisations and Standards
- Grid
 - Definitions and Virtual Organisations
- Services
 - Services and Resources
- Architecture
 - Bringing it all together

Grid – Where are we today?



- Grid computing is an evolving discipline
- Grids exist mainly in academia
- Grids have been used inside enterprises (so-called “intra-Grids”)
 - This simplifies operation and management
- Standards are needed to bring the parts together
- Standards are still evolving

What is the Grid?



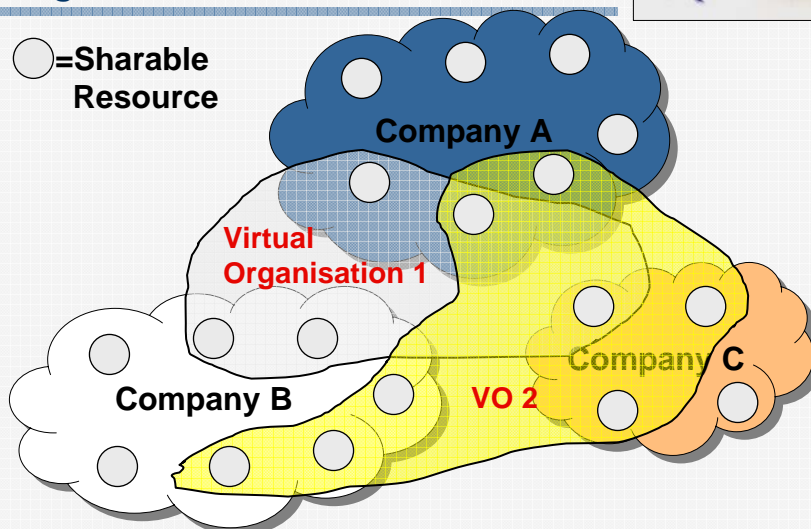
- The **metacomputer** is a network of heterogeneous, computational resources linked by software in such a way that they can be used as easily as a personal computer. (L. Smarr and C.E. Catlett, “Metacomputing”, Communications of the ACM 35/6, 45-52, 1992.)
- A **computational grid** is a hardware and software infrastructure that provides dependable, consistent, pervasive, and inexpensive access to high-end computational capabilities. (Ian Foster & Carl Kesselman, “The Grid”, MK, 1998.)
- “**Grid Problem**”, which we define as flexible, secure, coordinated resource sharing among dynamic collections of individuals, institutions, and resources. (Ian Foster, Carl Kesselman and Steven Tuecke, “The Anatomy of the Grid”, ..., 2001.)
- “**A Grid provides an abstraction for resource sharing and collaboration across multiple administrative domains...**” (Source: NGG Expert Group, 16 June 2003 “European Grid Research 2005-2010)

Working Definition



- Grid: System for coordinated resource sharing to achieve a non trivial task.
- All the definitions have more or less been functional definitions.
- But besides making it real, making money is important to some people.
- Money means contracts, needs security.
- Business means planning, management, coordination.
- How can Grids support all this?

Virtual Organisation (VO) - High Level view



What is a Virtual Organisation?

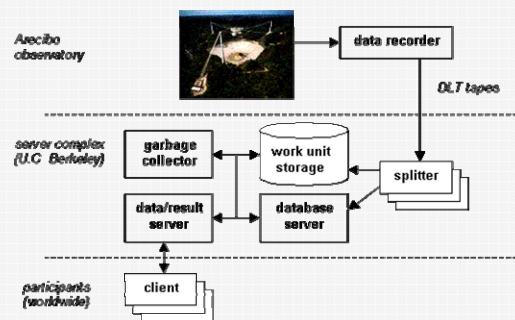


- A temporary or permanent coalition of geographically dispersed individuals, groups, organisational units or entire organisations that pool resources, capabilities and information to achieve common objectives.
 - Can provide services and thus participate as a single entity in the formation of further Virtual Organisations.
 - Enables the creation of recursive structures with multiple layers of "virtual" value-added service providers.
- The parties that form a virtual organization are typically part of a larger enterprise network of which a selection of partners is made.
- Participation in the enterprise network indicates disposition to work together in a future market opportunity.

SETI@home – A Special Purpose Computational Grid Application



- SETI@home: An Experiment in Public-Resource Computing (2002)
- By David P. Anderson, Jeff Cobb, Eric Korpela, Matt Lebofsky, Dan Werthimer



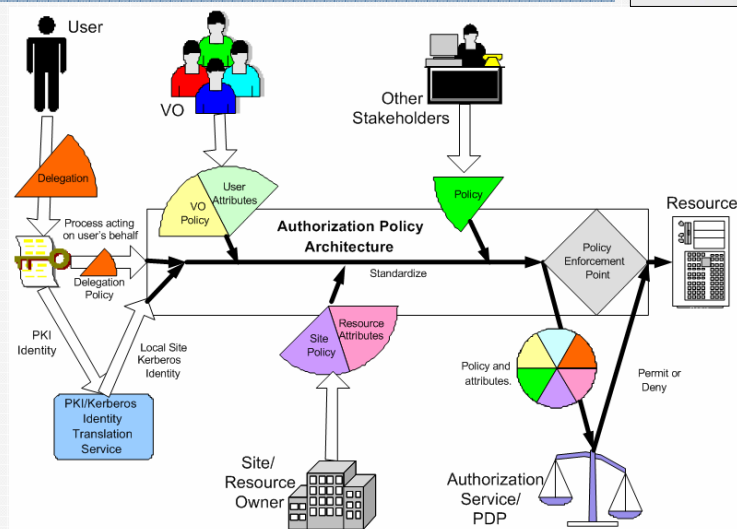
- <http://seticlassic.ssl.berkeley.edu/cacm/cacm.html>
- Computing Power: 27,000 GFLOPS

SETI@home – A Virtual Organisation?



- Do participants of SETI@home have a contract?
- How are the participating entities managed?
- Comparing a special purpose application with an architecture would not be fair.
- So: no comparison, just a starting point.
- A key Grid concept is only vaguely visible: The Virtual Organisation which provides the grounds for mutual agreements among the partners of that virtual organisation.

Policy Decision and Enforcement



Source: OGSA Version 1.0

The Business Need for Virtual Organisations – Engineering Process




- A large number of people are involved in the lifecycle of a complex engineering process.
- They work in:
 - A large number organisations from single-person specialist consultancies to major corporations
 - Agencies, such as regulatory authorities and government bodies.
- Design, production, marketing and support of a product are no longer the responsibility of only one company.
- Information is distributed and disparate across diverse applications in the same or different organisations
- Collaborators are at the same time competitors
- Key Performance Indicators:
 - Improve information access
 - Reduce the time between design sourcing & engineering lifecycles
 - Reduce time for application integration

Structural View of VO



■ VO-topologies



- Self-management
 - Manage operation and evolution with limited or no intervention from VO partners' administrators or parties outside the VO
- Scalability
 - Ability to be realised in different scales depending on objectives of and the kind of partners involved in a VO
 - Ability to achieve objectives that VO partners could not achieve individually – added value via collaboration
- Security
 - Ability to satisfy agreed confidentiality, integrity, availability and accountability requirements

The Lifecycle View of a VO



- **Identification**
 - Identify capabilities of the needed services
 - Use them for service discovery
 - Select services and also organisations
- **Formation**
 - Set up a relationship enabling the interoperation of the services
 - Generate workflow descriptions
 - Establish the orchestration/aggregation service
- **Operation**
 - Services within the VO can communicate
 - The generated workflows get executed
 - Control mechanisms supervise the operation
- **Dissolution**
 - Clean up everything from the Formation phase

What have VOs to do with Grids?



- Grid systems and applications aim to integrate, virtualize, and manage resources and services within distributed, heterogeneous, dynamic "virtual organizations"
 - Foster, I., Kesselman, C. and Tuecke, S. The Anatomy of the Grid: Enabling Scalable Virtual Organizations. International Journal of Supercomputer Applications, 15 (3). 200-222. 2001
- Doing business needs the VO concept.
- What is the difference to SETI@home?
- No business, no VO?

SETI@home VO



- VO goal
 - Defined by the University of California
- Identification
 - You somehow hear about SETI
- Formation
 - You download the code and register
- Operation
 - Your client downloads data and analyzes it
 - The analyzed data is sent to the server
- Dissolution
 - You may uninstall SETI@home at any time
 - You do not have any obligation to run the program.

Services...



- Open
 - Organisations and Standards
- Grid
 - Definitions and Virtual Organisations
- Services
 - Services and Resources
- Architecture
 - Bringing it all together

Virtualizing Resources



- Virtualized resources are the building blocks of the Virtual Organisation
- Services that virtualize resources are called Grid Services
- Grid services are stateful Web services
- Resources can be:
 - CPU cycles, hard disks, information, licenses, OS processes , ...

The "Grid Service"

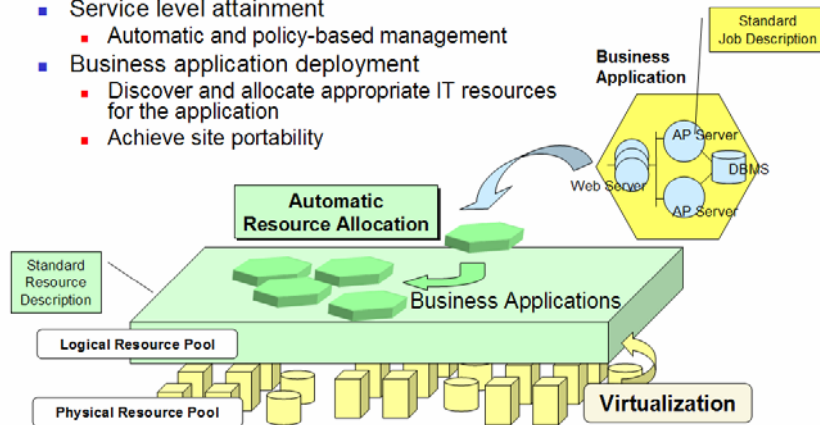


- Adding state to a Web service is not enough, to support VOs we need more:
- The combination of a Web service and one or more resources with specified interfaces & behaviors, including
 - Registration & Discovery
 - Authentication/Authorization
 - Notification
 - Concurrency
 - Manageability (including support for accounting)
 - Explicit State

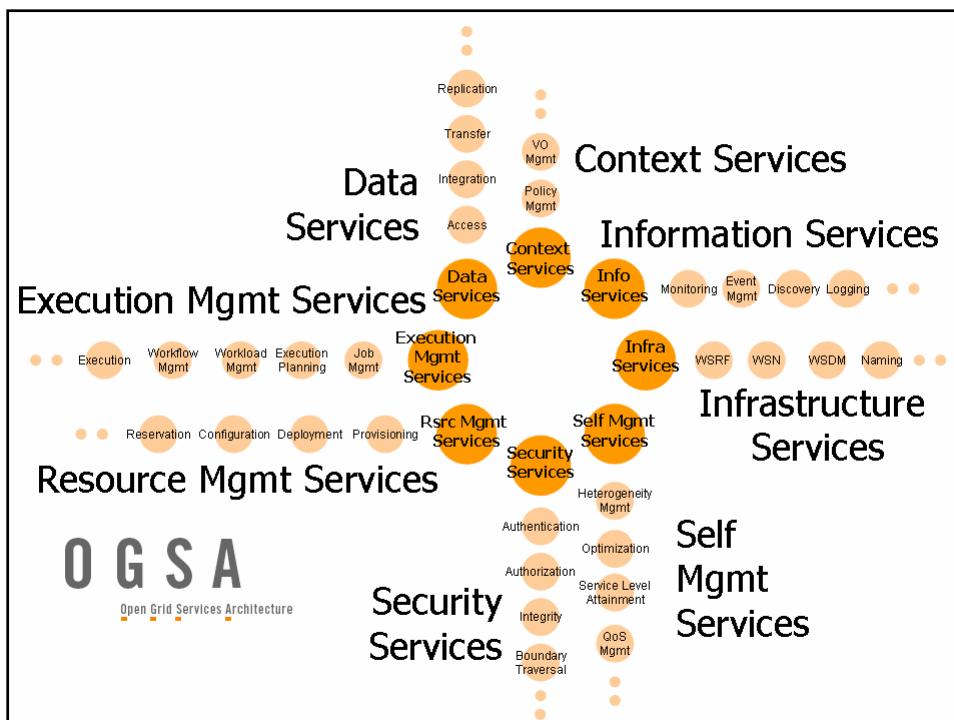
The Business Grid Project



- Resource virtualization
 - To handle distribution and heterogeneity
- Service level attainment
 - Automatic and policy-based management
- Business application deployment
 - Discover and allocate appropriate IT resources for the application
 - Achieve site portability



Source: OGSA Rising, by Hiro Kishimoto and Ian Foster.



OGSA Services



- **Execution Management Services:** Capability to instantiate and manage, to execute units of work. Units of work can be OGSA applications or legacy (non-OGSA) applications.
- **Data Services:** Capabilities for movement, access and update of data resources. A data resource is any entity that can act as a source or sink of data, such as files, streams, databases, and catalogues.
- **Resource Management:** Capabilities for management of the resources themselves, management of the resources on Grid, and management of the OGSA infrastructure.
- **Security:** Capabilities for authentication, identity mapping, authorization, credential conversion, audit and secure logging, and privacy.
- **Self-Management:** Self-configuring, self-healing and self-optimizing mechanisms for resources.
- **Information Services:** Naming, discovery, message delivery, logging, and monitoring capabilities.

Architecture



- Open
 - Organisations and Standards
- Grid
 - Definitions and Virtual Organisations
- Services
 - Services and Resources
- Architecture
 - Bringing it all together

OGSA Use Cases



- The Requirements for the Open Grid Services Architecture are derived from a number of use cases:
 - Commercial Data Center (CDC)
 - Severe Storm Modeling
 - Online Media and Entertainment
 - National Fusion Collaboratory (NFC)
 - Service-Based Distributed Query Processing
 - Grid Workflow
 - Grid Resource Reseller
 - Inter Grid
 - Interactive Grids
 - Grid Lite
 - Virtual Organization (VO) Grid Portal
 - Persistent Archive
 - Mutual Authorization
 - Resource Usage Service

OGSA Requirements

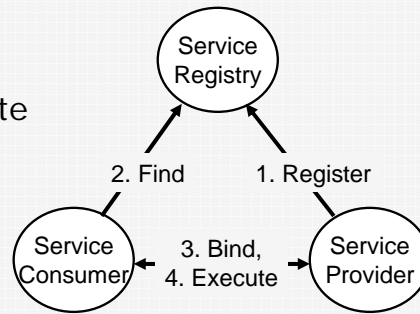


- Interoperability and Support for Dynamic and Heterogeneous Environments
- Resource Sharing Across Organizations
- Optimization
- Quality of Service (QoS) Assurance
- Job Execution
- Data Services
- Security
- Administrative Cost Reduction
- Scalability
- Availability
- Ease of Use and Extensibility

Service Oriented Architecture



- Service Providers:
 - Register their services in a Service Registry
- Service Consumers:
 - Find, bind and execute the found services
- SOAs built on Web services have a large vendor support.



Open Grid Services Architecture



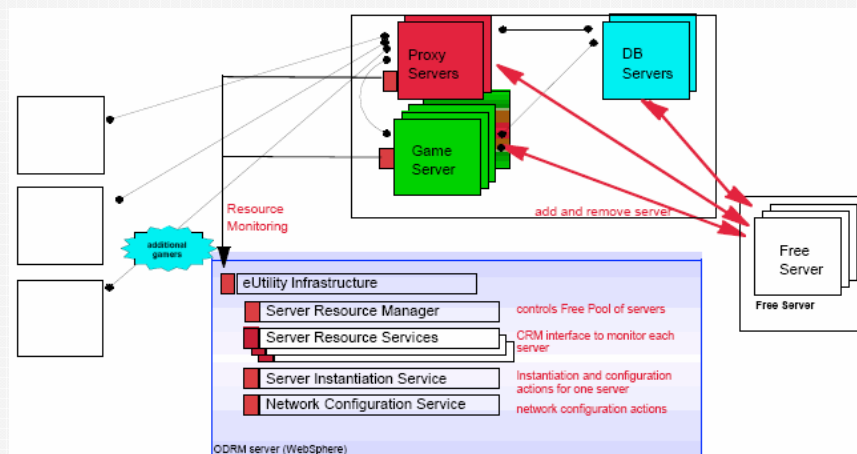
- Follow the concept of Service Oriented Architectures
- Define standard behavior and interfaces for services enabling interaction
 - Web Services offer standard interface definition mechanisms: multiple protocol bindings, multiple implementations, local/remote transparency
→ natural choice!
- System comprises (a typically few) persistent services & (potentially many) transient services
- Defined by the GGF and in discussion with the Enterprise Grid Alliance (EGA)
 - Not yet a real conclusion between these competing standardisation bodies

OGSA Use Case – Online Media and Entertainment



- Actors:
 - A consumer who consumes the entertainment content.
 - A service provider that hosts the entertainment content.
 - A publisher that offers the entertainment content.
 - A developer that consumes the entertainment content.
- Each role may consist of multiple companies
- The entertainment content consists of many different forms (e.g. movie on demand or online games).
- Different hosting capacity demands and lifecycle may be required for each form of content.

OGSA - M&E Use Case



Resources and Services

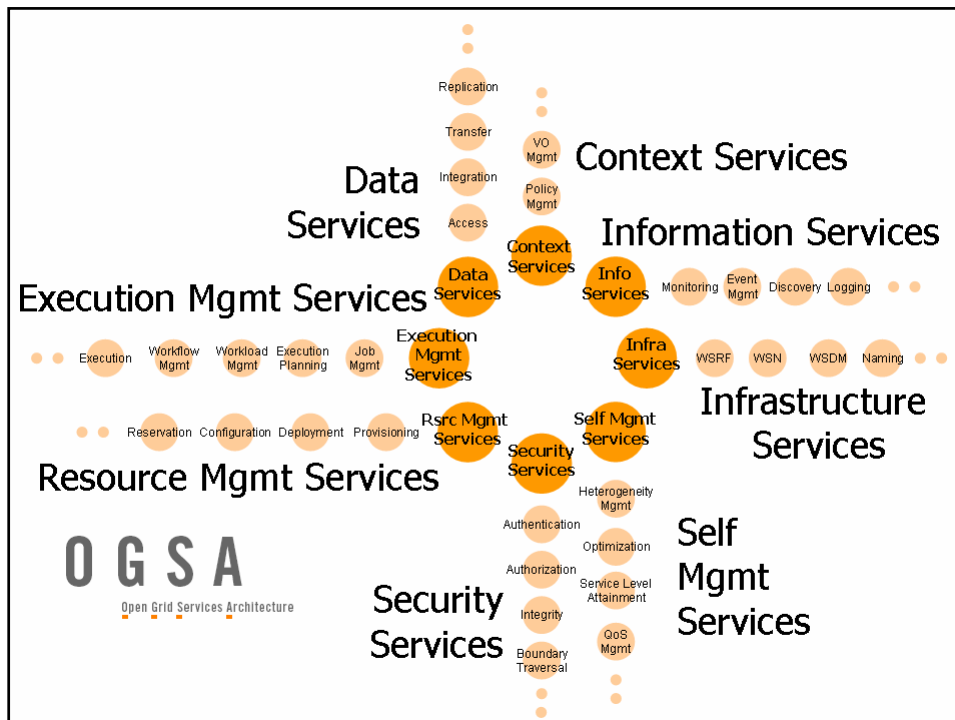


- Datacenter
 - Distributed server
 - Networked storage
 - Secure network (including multiple levels of firewalls)
 - Player consoles
- Online entertainment business
 - Security services (authentication/authorization, identity mapping, etc.)
 - Financial services (billing, rating, accounting, etc.)
 - Contracting/settlement services
 - Customer relations services (logging and data mining of user behaviors)
 - Management service (capacity management, workload management)
 - Media/Entertainment specific services (e.g. multimodal input)

The M&E Hosting Environment



- Requirements
 - Allow dynamic composition of standard pluggable components (e.g. billing service, customer relations service).
 - Be secure and trusted.
 - Have on-demand capacity (autonomic scalability according to workload), aggregation/selection of new services, and provide for integration with other companies that have needed competencies.
 - Enable new commercial business models.
 - Apply to needs of online game applications.



Execution Management Services



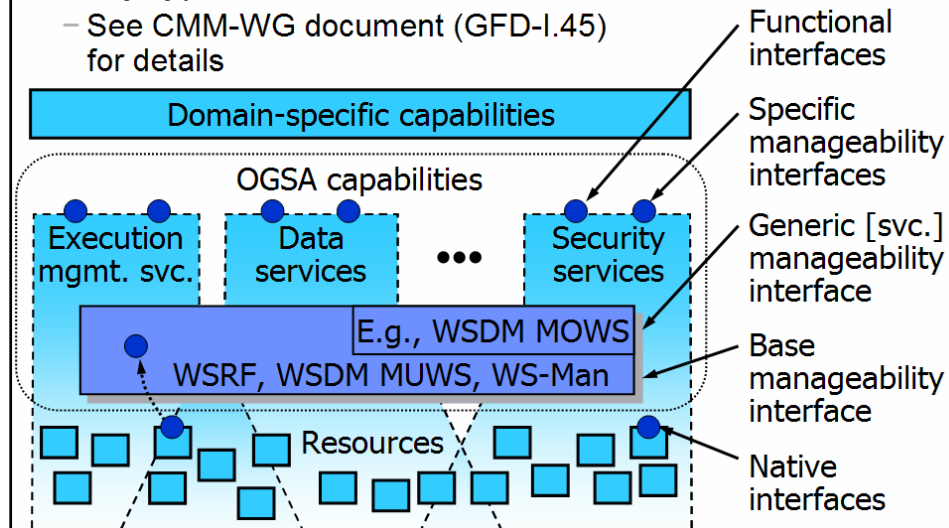
- Basic problem: provision, execute/instantiate and manage services/resources (including legacy applications) in a grid
 - Some use cases
 - start up a cache service; on-demand, utility computing ;
 - start up and manage a set of legacy applications
- Example questions to address:
 - Where can a service execute?
 - Where should the service execute?
 - Prepare the service to execute.
 - Get the service executing.
 - Manage (monitor, restart, move, etc.).

Source: GGF

OGSA Management Framework

- Many types and levels of interfaces

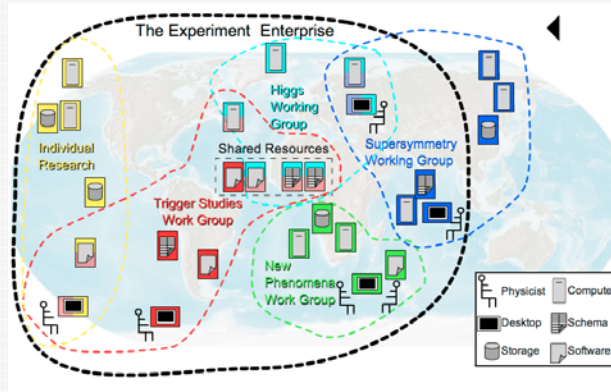
– See CMM-WG document (GFD-I.45) for details



Self Management Services

- Reduce the cost and complexity of owning and operating an IT infrastructure
- System components are
 - Self-configuring
 - Self-healing
 - Self-optimizing
- Policy- and model-based management
- Problem: Define proper means and mechanisms to implement self management.

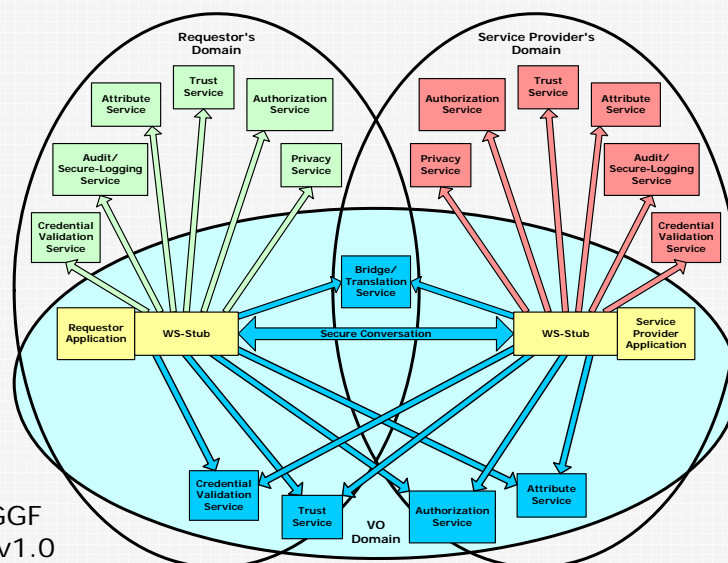
Security Services



- Authorization, Roles, and Access Privileges
 - Locally (site) managed
 - SAML and XACML Basis
 - Credential mapping provided by implementations.

Source: GGF OGSA v1.0

VO Security Services



Source: GGF
OGSA v1.0

Further information



- Standardisation activities at:
 - Global Grid Forum, OASIS Open, W3C, EGA?, IETF?
- WSRF <http://www.globus.org/wsrf> especially http://www.globus.org/wsrf/sabbah_wsrf.ppt
- WSRF Toolkits
 - Globus Toolkit 4
 - <http://www.globus.org>
 - WSRF.NET
 - <http://www.cs.virginia.edu/~gsw2c/wsrf.net.html>
- Start with OGSA e.g. at <https://forge.gridforum.org/projects/ogsa-wg>
 - OGSA (Open Grid Service Architecture) Version 1.0
 - <https://forge.gridforum.org/projects/ogsa-wg/document/draft-ggf-ogsa-spec/en>
- Look at the documents available from GRIDSTART (<http://www.gridstart.org>)
- Subscribe to a Grid newsletter such as "Primeur EnterTheGrid" or "Gridtoday" to get a feeling about the topics of the community

Fin



Gracias!